Model Order Reduction with Dynamically Transformed Modes for Electrophysiological Simulations

Niklas Hornischer[ab,⋆] ⓘD

[a] Stuttgart Center for Simulation Science (SC SimTech), University of Stuttgart, Germany
[b] Institute for Applied Analysis and Numerical Simulation (IANS), University of Stuttgart, Germany

⋆ corresponding author: niklas.hornischer@simtech.uni-stuttgart.de
supervisor: Felix Huber, IANS, University of Stuttgart, Germany

Abstract: *The numerical simulation of neuromuscular processes in skeletal muscles typically imposes high computational costs, resulting in the need for techniques to reduce these costs. One submodel in these types of multiphysics multiscale simulations is the electrophysiological model, describing the propagation of an action potential (AP) in a muscle fiber. Involved in the simulation of the propagation are two traveling waves, resulting in difficulties for classical model order reduction (MOR) techniques based on linear subspaces.*

Instead, we apply the nonlinear MOR method shifted Proper Orthogonal Decomposition (sPOD) to one state variable only in order to construct dynamically transformed reduced basis functions depending on time-dependent paths spanning the adaptive reduced ansatz space.

Our numerical experiments demonstrate that the constructed reduced ansatz space can accurately capture the dynamics of two fully separated wavefronts and reduce the degrees of freedom of the whole simulation. However, it cannot represent overlapping wave parts of the hyperpolarization. The constructed reduced order model outperforms the high-dimensional full order model in terms of the computational costs while the accuracy is maintained and reaches speedup factors between 2 and 73 depending on the time discretization.

1 Introduction

A majority of tasks in human life can only be accomplished because of the complex muscular processes and interactions of and between the approximately 220 individual muscles [17]. At the core of all movements are voluntary contractions of skeletal muscles, allowing for a large variety of different movements, covering fast, persistent, finely controlled, but also exhausting actions. This diversity is based on complex interactions between the musculoskeletal and the neurological system.

Gaining insights into these systems is of immense importance, for example, in medicine and healthcare. Most incurable genetic disorders like myotonia con-

genita stem from a malfunction in the electrical processes regulating muscle contractions [23, 31]. A better understanding of the mechanisms of the neuromuscular system can improve treatments and classification, also for other muscular diseases and in post-stroke rehabilitation using patient-specific motorized prostheses [24, 33]. In the field of engineering, the development of bioinspired robots and artificial muscles, mimicking the functionality of the human skeletal muscle, requires deep knowledge about the involved neurological processes [21, 30]. Increasing the functionality and usability of such biomimetic systems can have advantages in industrial robotics and the development of exoskeletons [2, 3].

Traditionally and empirically oriented approaches in clinical studies have reached their limits regarding the level of detail and the time requirements [10]. Instead, individualized, biophysics-based simulations can be used to gain detailed insights into non-measurable processes. The numerical treatment of the underlying coupled, nonlinear partial differential equations (PDEs) necessitates large-scale computation facilities to manage the extensive computational costs required for high fidelity solutions. These computational costs are not exclusively a problem in research because of long computation times and high energy consumption. In envisioned real-life applications, in medicine for example, patient-specific simulations require low runtimes of the detailed simulations to supply fast and exact diagnostics, for example in adaptive surgery settings. To allow for execution on resource-constrained devices and the treatment of the envisioned scenarios, the computational costs and the complexity of such simulations must be reduced.

Outline and Contributions In this work, the simulation of neuromuscular processes is based on a multiphysics, multiscale model capturing the processes in the individual muscle fibers to the contraction of the whole muscle, see Section 2.1. Since the number of fibers comprising a skeletal muscle is typically high, the simulations of the individual fibers are the dominating factor in the overall computation. Simulating each fiber involves the electrophysiological model capturing the propagation of the cell membrane potential, defined as action potentials (APs). Reducing the computational complexity of this complex electrophysiological model is of utmost interest.

One promising approach to reduce the computational costs is model order reduction (MOR). Generally, in MOR, the full order model (FOM) is replaced by a suitable reduced order model (ROM) with lower computational complexity while ensuring the accuracy of the results as much as possible [4].

Recent works in the field of MOR for electrophysiological simulations involve Proper Orthogonal Decomposition (POD) and Deep Learning (DL) and show the potential of these techniques for this kind of application. In contrast to this work, these approaches were applied to cardiac electrophysiology, the simulation of the electrical activity of the heart, and not to single skeletal muscle fibers [9, 13, 14, 16].

This paper focuses on the derivation of a ROM with transformed modes, based on shifted Proper Orthogonal Decomposition (sPOD) for an electrophysiological simulation of skeletal muscle fibers. In detail, the contributions of this work are:

- The application of sPOD to construct a reduced ansatz space capturing the wave propagations.
- The use of this reduced ansatz space in a finite element approach to obtain the ROM.
- The evaluation of the reduced ansatz space and the performance of the reduced simulation.

In doing so, this work shows that the application of sPOD in the context of MOR for electrophysiological simulations is a valid approach. Even by applying the MOR approach to one state variable only, speedup factors between 2 and 73 can be reached with maintained accuracy.

The remainder of this paper is structured as follows. Section 2 deals with the necessary background of the simulation model, the construction of the FOM via the finite element method, and its simulation. Section 3 introduces the MOR approach, followed by the details of the reduced ansatz space and the resulting ROM in Section 3.4. Performance and accuracy of this MOR approach are evaluated in Section 4. Finally, Section 5 summarizes the results and gives an outlook on future work.

2 Full Order Model

Before we present the technical details of our MOR approach, we first give some background to the original full order simulation of the propagation of the APs.

The APs arrive at the neuromuscular junction located in the center of each fiber and propagate along the fiber to both ends. This propagation is triggered by biochemical processes on the subcellular scale involving ion movement through voltage-dependent ion channels. To capture these complex processes we use the overall model described in [10, 27] and extended in [26].

The model is based on two submodels, the 1D fiber activation model and the 0D subcellular model. The first one captures the propagation of the AP along a 1D fiber, whereas the second one describes the chemical reaction to this potential in the fiber cell. Hence, the 1D fiber activation model and the 0D subcellular model are strongly coupled, and we combine them in the *electro-physiological model*. The 3D mechanics of the actual muscle movement are not covered in this paper, and we refer to [26].

2.1 Electrophysiological Model

The mathematical definition of the electrophysiological model is based on [29] and described in [26]. Its derivation is mainly based on the electrical coupling between intra- and extracellular spaces of the muscle fiber and the Hodgkin-Huxley cell model, proposed in [19].

For each bounded fiber domain $\Omega = [0, \ell]$ with the muscle fiber length ℓ and simulation time $0 \leq T_0 < T$, the 1D model captures the transmembrane potential $u : \Omega \times [T_0, T] \to \mathbb{R}$ along the fiber and is a second-order PDE of the form

$$\partial_t u = \underbrace{\frac{1}{A_m C_m} \sigma_{\text{eff}}}_{=:C} \Delta u - \underbrace{\frac{1}{C_m} I_{\text{ion}}(u, y)}_{=:f(u,y)} \qquad \text{in } \Omega. \quad (1)$$

The first term in the 1D model describes the diffusion of the electrical potential along the fiber and the second one accounts for the transport of the potential, driven by the ion movements. This potential-dependent current of the ion movements I_{ion} can be described by the model of Hodgkin-Huxley [19]. The involved constants are the fiber's surface-to-volume ratio A_{m}, the capacitance of the muscle fiber membrane C_{m}, and the effective conductivity σ_{eff}.

The Hodgkin-Huxley model describes the movement of the $k > 0$ ions in terms of ion-specific conductivities $g_i(y)$ and constant Nernst potentials E_i:

$$I_{\text{ion}}(u, y) = \sum_{i=1}^{k} g_i(y)(u - E_i).$$

These conductivities depend nonlinearly on the activation probabilities of the ion channels $y : \Omega \times [T_0, T] \to \mathbb{R}^k$ and describe the pass-through of the ions through the membrane. Typically, $k \approx 50$ activation probabilities are involved in capturing the complex processes in our muscle fibers. These probabilities are defined locally on a subcellular scale, more precisely at approximately 100+ positions on the muscle fiber cell. Depending on the transmembrane voltage, the activation probabilities change over time. This process is captured by the 0D

subcellular model at every point on the fiber, formulated in the Hodgkin-Huxley model as well

$$\partial_t y = g(u, y) \qquad \text{in } \Omega. \quad (2)$$

Prescribing initial and boundary conditions as follows

$$\begin{aligned} \nabla u = 0, & \qquad \text{on } \partial\Omega, \\ u(\cdot, T_0) = u_0, \qquad y(\cdot, T_0) = y_0 & \qquad \text{in } \Omega, \end{aligned}$$

yields the final initial boundary value problem (IBVP) of the electrophysiological model. Formally, we require $u(x, \cdot), y(x, \cdot) \in C^1(]T_0, T[) \cap C^0([T_0, T])$ for all $x \in \Omega$ as well as $u(\cdot, t) \in C^2(\Omega) \cap C^0(\bar{\Omega})$ and $y(\cdot, t) \in C^0(\bar{\Omega})$ for all $t \in]T_0, T[$.

2.2 Finite Element Approach

The original full order simulation on the 1D fiber scale is based on a finite element method and is used as a reference solution of the electrophysiological model problem from Section 2.1. Constructing this finite element method results in the semi-discrete FOM, yielding a finite-dimensional approximation to the weak solution of the model problem.

For a fixed time $t \in [T_0, T]$, we seek the spatial weak solution $u(t)$ in the Sobolev space $\mathcal{Y} = H^1(\Omega)$, more precisely with $(u(t), y(t)) \in \mathcal{Y} \times \mathcal{Y}^k$ the weak forms of (1) and (2)

$$\int_\Omega \partial_t u(t) v \, \mathrm{d}x = -C \int_\Omega \left(\nabla u(t) \nabla v + \frac{1}{C} f(u(t), y(t)) v \right) \mathrm{d}x$$

$$\int_\Omega \partial_t y(t) \cdot w \, \mathrm{d}x = \int_\Omega g(y(t), u(t)) \cdot w \, \mathrm{d}x$$

must hold for all $v \in \mathcal{Y}$ and $w \in \mathcal{Y}^k$. To simplify the notation, we take $\mathcal{Y}$ as the ansatz space for $y(t)$ intentionally, even though $L^2(\Omega)$ would suffice.

We cover the spatial domain Ω by a triangulation $\mathcal{T}_h$ as a subdivision of equidistant subintervals with length h_s and use nodal basis functions $\{\psi_1, ..., \psi_n\}$ to construct the finite dimensional ansatz space V_h where we restrict ourselves to the Lagrange $\mathcal{P}_1(\mathcal{T}_h)$ space for simplicity in the following. Performing a Galerkin approximation in space for $u_h \in V_h$ with $u_h = \sum_{i=1}^{n} \bar{u}_i \psi_i$ and analogously $y_h \in V_h^k$, we arrive at the semi-discrete formulation

$$\begin{aligned} \partial_t \bar{u} &= -M_h^{-1} C A_h \bar{u} - M_h^{-1} F_h(\bar{u}, \bar{y}), \\ \partial_t \bar{y} &= (I_k \otimes M_h)^{-1} G_h(\bar{u}, \bar{y}), \end{aligned}$$

with matrices and vectors

$$\left(M_h\right)_{i,j=1}^n := \left\langle \psi_i, \psi_j \right\rangle_{L^2(\Omega)},$$

$$\left(A_h\right)_{i,j=1}^n := \left\langle \nabla\psi_i, \nabla\psi_j \right\rangle_{L^2(\Omega)},$$

$$\left(F_h(\bar{u}, \bar{y})\right)_{j=1}^n := \left\langle f(u_h(t), y_h(t)), \psi_j \right\rangle_{L^2(\Omega)},$$

$$\left(G_h(\bar{u}, \bar{y})\right)_{j=1}^n := \left\langle g(u_h(t), y_h(t)), \psi_j \right\rangle_{L^2(\Omega)}.$$

The identity matrix is denoted as $I_k \in \mathbb{R}^{k\times k}$ and $\otimes$ corresponds to the Kronecker product. The nonlinear terms from (1) and (2) are represented in the ansatz space as the maps $f : \mathscr{Y} \times \mathscr{Y}^k \to \mathscr{Y}$ and $g : \mathscr{Y} \times \mathscr{Y}^k \to \mathscr{Y}^k$. To solve for the time-dependent solution $\bar{u}$ and $\bar{y}$, we apply an explicit Euler time-stepping scheme for simplicity. For this, we partition the time interval $[T_0, T]$ into m equidistant subintervals with length h_t.

3 Model Order Reduction

The simulation of the FOM is computationally expensive due to the coupling between the two complex submodels and a large number of discretization nodes both in space and time. In this section, we present a MOR approach with the general aim to reduce the computational costs and allow for rapidly computable approximations, while still achieving a satisfactory accuracy of the FOM, building on ideas from [4].
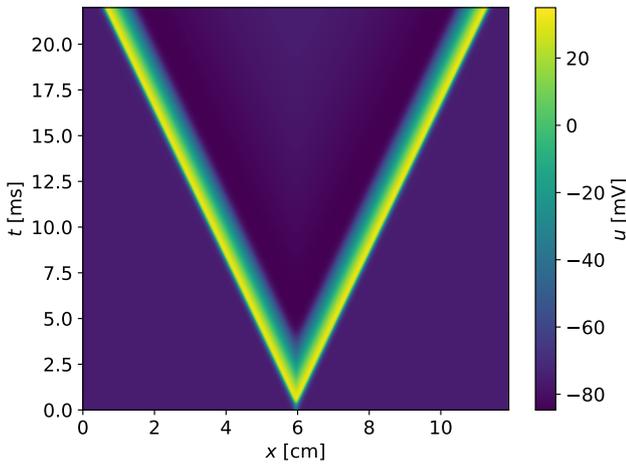
Figure 1 – Illustration of the space-time cylinder of the AP wave propagation for a muscle fiber of length 12 cm, originating in the center of the fiber. FOM simulation with the parameters from Table 1 and $h_t = 0.001$.

Our approach is based on the construction of a problem-specific, time-dependent, low-dimensional subspace of the full dimensional ansatz space and using this ansatz space to derive the ROM. More precisely the reduced ansatz space should be able to capture the dynamics of the solution, in our case the two traveling wavefronts in Figure 1. We focus on constructing and applying our MOR approach for the transmembrane voltage u for simplicity, not for the ion channel probabilities y.

3.1 Overview of MOR Approach

Since our problem is dominated by transport, we apply the sPOD introduced in [35, 36] and formalized in [5, 6, 8]. To capture the transport terms, sPOD allows the construction of multiple reduced ansatz spaces traveling in the spatial domain. But instead of moving the discretization nodes like in Moving Finite Element Methods [15, 28], the reduced basis functions $\{\varphi_i\}_{i=1}^r$ evolve over time themselves and the discretization is fixed. This can be achieved by shifting the modes over time, introducing an additional variable $p_j(t)$ for the shift amount. From the full order simulation in Figure 1 we observe that the wavefront of the transmembrane voltage travels at an approximately constant speed through the spatial domain. Following this observation, we only account for constant shift velocities of the basis functions and construct a constant-speed shifted reduced ansatz space, in contrast to [5, 8] where also non-constant shift velocities are allowed.

We use sPOD [8, 34], an adapted version of POD [18, 20, 22, 32], to obtain the reduced basis functions from the results of the full order simulation, explicitly accounting for the shift variables p_j. To obtain the ROM, we use these reduced basis functions instead while following the same Galerkin approach in comparison to the FOM.

3.2 Adaptive Reduced Ansatz Space

With the general goal to construct a ROM, we aim to approximate the solution $u(t) \in \mathscr{Y}$ by the reduced solution $\hat{u}(t) \in \mathscr{R}$ in some reduced space $\mathscr{R} \subset \mathscr{Y}$. We construct the basis functions $\{\varphi_j\}_{j=1}^r$ of the reduced space $\mathscr{R}$ as elements of the finite dimensional ansatz space V_h of the FOM to obtain continuous reduced basis functions. Since they should capture the transport quantities, we evaluate them in a shifted and time-dependent position based on unknown paths $p_j : [T_0, T] \to \mathbb{R}$ [8].

Since the path functions describe the evolution of the ansatz space and thus should capture the constant velocities we can assume that the paths $p \in H^1(T_0, T; \mathbb{R}^r)$ satisfy

$$\frac{\partial}{\partial t} p(t) = \text{const} \qquad \text{and} \qquad p(T_0) = 0,$$

i.e. $p(t) = t\bar{p}$ with $\bar{p} \in \mathbb{R}^r$.

With these preparations, we can define our modified reduced ansatz space. Let $\mathcal{T}(p_j(t))$ be a suitable transformation operator, compare [8], with linear shifts $p_j(t) = \bar{p}_j t$ with the path coefficients $\bar{p} \in \mathbb{R}^r$ for all $j = 1, \ldots, r$. We define the time-shifted reduced basis functions $\{\phi_j\}_{j=1}^r$ through

$$\phi_j(t) = \mathcal{T}(p_j(t))\varphi_j = \varphi_j(x - \bar{p}_j t)$$

and use the periodic extension proposed in [5, Ex. 4.3 and 5.2] to keep the shifted basis functions $\{\varphi_1, \ldots, \varphi_r\}$ well-defined. With these time-shifted basis functions we span the constant-speed shifted ansatz space $\hat{V}_h(t)$ with

$$\hat{V}_h(t) = \text{span}\{\phi_1(t), \ldots, \phi_r(t)\},$$

and $r = \dim \hat{V}_h(t) < \dim V_h = n$. This ansatz space travels in the spatial space Ω depending on time.

3.3 Constant-Speed Shifted Proper Orthogonal Decomposition

To use the constant-speed shifted reduced ansatz space explicitly in our ROM, we need to determine the basis functions ϕ_j. Therefore, we follow the general idea of snapshot decomposition to solve a similar minimization problem as presented in [8]. The differences are hidden in the details: Since we assume linear paths, it suffices to optimize over the constant path velocities $\bar{p}$ instead of the paths themselves. However, the constant shift velocities $\bar{p}$ must be optimized nevertheless.

The available snapshot data in our case are function samples from the FOM simulation and denoted as $u \in L^2(T_0, T; \mathcal{Y})$. In addition, let $\mathcal{L} := L^2(T_0, T; \mathbb{R}^r) \times \mathcal{Y}^r \times \mathbb{R}^r$ be the space of the optimization variables. Since the overall goal is to represent the solution u in the reduced space, the error of the representation should be as small as possible. To this end, we define the cost functional $J : \mathcal{L} \to \mathbb{R}$ as the mean squared distance to the function sample u

$$J(\hat{u}, \varphi, \bar{p}) = \frac{1}{2} \left\| u - \sum_{i=1}^r \hat{u}_i \mathcal{T}(p_i)\varphi_i \right\|^2_{L^2([T_0, T], L^2(\Omega))}. \quad (3)$$

The optimization problem then reads: Find $(\hat{u}, \varphi, \bar{p}) \in \mathcal{L}$ such that

$$(\hat{u}, \varphi, \bar{p}) \in \underset{(\hat{u}', \varphi', \bar{p}') \in \mathcal{L}}{\arg\min} J(\hat{u}', \varphi', \bar{p}') \quad (4)$$

holds. We solve this optimization problem using a quasi-Newton method for limited memory, the L-BFGS-B algorithm [12, 25]. Similar to [8], we provide the explicit discretized cost function and the corresponding gradient by using the same discretization from the FOM.

3.4 Reduced Finite Element Method

Solving the discrete optimization problem from Section 3.3, we obtain the coefficients $\bar{p}_i$ and $\bar{\varphi}$ to construct the shifted ansatz space $\hat{V}_h$. Similar to the derivation of the FOM, we insert the reduced basis representation of the Galerkin approximations $\hat{u}_h = \sum_{i=1}^r \hat{u}_i \phi_i$ and $y_h = \sum_{i=1}^n \bar{y}_i \psi_i$ into the weak formulation of the original problem. Here $V_h = \mathcal{P}_1(\mathcal{T}_h)$ is the full-dimensional discrete ansatz space of $\mathcal{Y}$ with nodal basis $\{\psi_1, \ldots, \psi_n\}$. The resulting semi-discrete ROM is thus given by

$$\partial_t \hat{u} = \hat{M}_h^{-1}(-C\hat{A}_h - \hat{N}_h)\hat{u} - \hat{M}_h^{-1}\hat{F}_h(\hat{u}, \bar{y}),$$
$$\partial_t \bar{y} = (I_k \otimes M_h)^{-1}\hat{G}_h(\hat{u}, \bar{y}),$$

with the time-dependent matrices and vectors

$$(\hat{M}_h)_{i,j=1}^r := \left\langle \phi_j, \phi_i \right\rangle_{L^2(\Omega)},$$
$$(\hat{A}_h)_{i,j=1}^r := \left\langle \nabla \phi_j, \nabla \phi_i \right\rangle_{L^2(\Omega)},$$
$$(\hat{N}_h)_{i,j=1}^r := \left\langle \partial_t \phi_j, \phi_i \right\rangle_{L^2(\Omega)},$$
$$(\hat{F}_h(\hat{u}, \bar{y}))_{j=1}^n := \left\langle f(\hat{u}_h, y_h), \phi_j \right\rangle_{L^2(\Omega)},$$
$$(\hat{G}_h(\hat{u}, \bar{y}))_{j=1}^r := \left\langle g(\hat{u}_h, y_h), \psi_j \right\rangle_{L^2(\Omega)}.$$

To assemble the reduced matrices, we need to compute the $L^2(\Omega)$-inner products with the time-shifted basis functions $\phi(t)$. The major difference in comparison to the FOM from Section 2.2 is the time dependency of the basis functions. Since the paths account for the constant velocity shift and are thus given by the reduced space, they do not occur as unknowns of the ROM thus the paths are not recalculated in the online phase.

3.5 Computation of the ROM

Solving the semi-discrete ROM from Section 3.4 involves the application of the explicit Euler scheme, similar to the FOM. This time-stepping scheme depends on the model matrices and for the FOM it is sufficient to assemble these model matrices once due to their time independence. However, the model matrices for the ROM depend on the shifted reduced basis functions and thus on the time-dependent paths. This requires their reassembly in every time step, which can be computationally expensive.

In our implementation, we evaluate the terms $\hat{F}_h$ and $\hat{G}_h$ using the nonlinear terms F_h and G_h, defined on the full-dimensional ansatz space V_h. To do so, we have to project the transmembrane voltage from the constant-speed shifted ansatz space $\hat{u}_h \in \hat{V}_h$ to the full order ansatz space $u_h \in V_h$. We call this *up-projection*

and use an L_2 projection to define it. Starting at the condition $\langle u_h - \hat{u}_h, \psi_i \rangle_{L^2} = 0$ for all $i = 1, ..., n$ we can obtain the coefficients $\bar{u}$ of the basis representation in the full order space V_h based on the coefficients $\hat{u}$ of the basis representation in the reduced space $\hat{V}_h$. This transformation is defined as

$$\bar{u} = M_h^{-1} \hat{P}_h \hat{u},$$

with the projection matrix $(\hat{P}_h)_{i,j=1}^{n,r} = \langle \phi_j, \psi_i \rangle_{L^2}$. Based on these coefficients we can easily represent a function $\hat{u}_h = \sum_{i=1}^{n} \hat{u}_i \phi_i \in \hat{V}_h$ in the space V_h through the basis representation $u_h = \sum_{i=1}^{n} \bar{u}_i \psi_i$.

Following the same procedure, we can also construct the so-called *down-projection* mapping from V_h to $\hat{V}_h$ with $\bar{P}_h = \hat{P}_h^T$ to obtain the evaluations $\hat{F}_h$ and $\hat{G}_h$. These additional steps must be taken into account in the algorithm when solving the ROM.

Typically, the evaluation of the ROM still scales with the dimension of the FOM and for that reason hyper-reduction techniques like discrete empirical interpolation method (DEIM) are commonly used to reduce the computational complexity. However, in our setting, we solely focus on the sPOD without applying hyper-reduction techniques and note that the performance of the ROM can be further improved by using these methods.

4 Numerical Experiments

In this section, we present the application of the MOR scheme defined in Section 3 to the electrophysiological model. All computations are performed on an 11th Gen Intel® Core™ i7-11700 @ 2.50GHz with 8 cores and a total DDR4 RAM capacity of 16 GB @ 3200 MHz. The implementation in *Python* is strictly sequential and the sparsity of the matrices is exploited by using CSR matrices.

The accuracy of the presented methods, especially of the constructed reduced basis in Section 4.2 and the resulting ROM in Section 4.3, is evaluated based on the relative error in the L^2-norm and the maximal error in the maximum norm. We define the global errors between a function $z \in L^2(T_0, T; L^2(\Omega))$ with $\|z\|_{L^2(T_0, T; L^2(\Omega))} \neq 0$ and its approximation z_h with $z - z_h \in L^\infty(T_0, T; L^\infty(\Omega))$

as follows

$$e_{\text{rel}}(z, z_h) := \frac{\|z - z_h\|_{L^2(T_0, T; L^2(\Omega))}}{\|z\|_{L^2(T_0, T; L^2(\Omega))}},$$

$$e_\infty(z, z_h) := \|z - z_h\|_{L^\infty(T_0, T; L^\infty(\Omega))}.$$

We refer to the local absolute error as e_{abs} and define it as

$$e_{\text{abs}}(t, x) := |z(t, x) - z_h(t, x)|,$$

for almost all $(t, x) \in [T_0, T] \times \Omega$. In addition, we denote the arithmetic mean of the absolute error over the entire spatial and time domain as $\bar{e}_{\text{abs}}$, i.e.

$$\bar{e}_{\text{abs}} := \frac{1}{T - T_0} \frac{1}{|\Omega|} \int_{T_0}^{T} \int_{\Omega} e_{\text{abs}}(t, x) \mathrm{d}x \mathrm{d}t.$$

If not stated otherwise, all the following errors refer to the transmembrane voltage state $z = u$ only, where u is the solution from Section 4.1. To compare this function with approximations in the reduced setting, we project the reduced functions to the full dimensional space

$$z_h(x, t) = \left(M_h^{-1} \hat{P}_h \hat{u}(t) \right)^T \psi(x),$$

where $\hat{u}$ are the amplitudes in the basis representation of the reduced space and ψ is the vector containing the nodal basis functions of V_h.

Furthermore, we differentiate the errors in *offline errors* and *online errors*. The *offline errors* are the errors of the reduced ansatz space based on the optimal amplitudes resulting from the optimization problem (4), which correspond to the orthogonal projection onto the (time-dependent) reduced ansatz space. These *offline errors* are used to evaluate the quality of the results from the sPOD. We denote the *online errors* as the errors based on the solution of the ROM as amplitudes. This separation allows analyzing if the errors stem from the ROM or the reduced ansatz space.

4.1 FEM Simulation

The model parameters used in all our simulations are taken from [26] and listed in Table 1. In all the simulations we use $k = 3$ as the number of voltage states and for the simulation of the FOM we use a time-step

Table 1 – Parameters of the electrophysiological simulations

ℓ [cm]	h_s [cm]	T_0 [ms]	T [ms]	C_m [μF/cm^2]	A_m [1/cm]	σ_{eff} [mS/cm]	V_{eql} [mV]	P_{Na^+}	P_{K^+}	P_L
11.9	0.01	0	22	0.58	500.0	3.828	-74.985	0.05	0.6	0.325

width of $h_t = 0.001$. To activate the propagation of the AP we initialize the voltage state by a scaled and shifted Gaussian function

$$u_0(x) = \frac{50}{2\pi} e^{-\frac{(x-\mu)^2}{2\varsigma^2}} (V_{\max} - V_{\text{eql}}) + V_{\text{eql}},$$

with $\mu = \ell/2$, $\varsigma^2 = 0.02$ cm^2 and $V_{\max} = 35$ mV for all $x \in \Omega$. Furthermore, we set the ion channel states to their equilibrium probabilities given in Table 1 with $y_0 = \left[P_{\text{Na}^+}, P_{\text{K}^+}, P_{\text{L}} \right]^T$ for all $x \in \Omega$.

The resulting voltage snapshots are visualized in Figure 1 and show the AP defined by the voltage as two wavefronts traveling through the spatial domain in both directions from its origin. The solution for all states can be seen in Figure 2, where the propagation in time is visualized by the color gradient ranging from red ($t = 0$) to green ($t = T$).
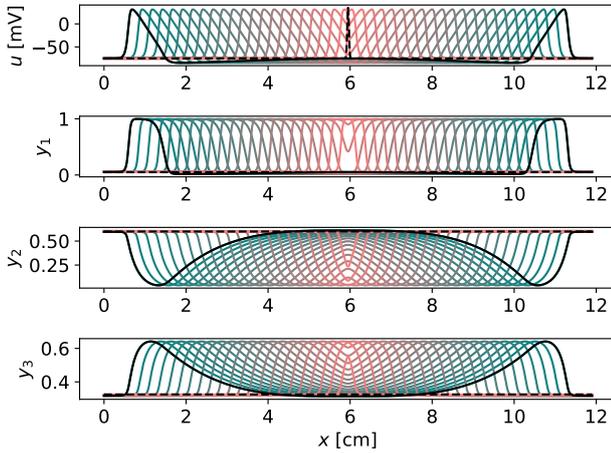
Figure 2 – States for a time range of $[0, 22]$ visualized from red to green with 1 ms between the curves. The first and last states are visualized in black.

4.2 Quality of Reduced Ansatz Space

The reduced ansatz space should capture the propagation of the voltage state in the spatial domain, and thus we only use the snapshots for the voltage to construct the reduced basis functions. To determine these reduced basis functions we numerically minimize the cost functional (3). This subsection corresponds to the offline phase and hence all errors are offline errors. The optimization problem (4) is carried out by SciPy's *minimize* function using the L-BFGS-B algorithm [37].

Based on the numerical experiments in [7], we expect it to be challenging to achieve a good approximation of the initialization of the wave, also called activation,

when using a sPOD-based ROM. For this reason, we focus on determining the reduced basis functions for the cases where the waves are separated and extract the initial values and the snapshot data from the full order simulation results for different start times T_0.

Typically, an AP consists of a narrow peak at the wavefront and a wide hyperpolarization part afterwards, characteristic for the decrease of the potential below the resting potential [1]. We furthermore differentiate between two different cases of the snapshot data. In the first one with $T_0 = 5$ ms only the wavefronts are separated but the hyperpolarization parts overlap. In the second one with $T_0 = 18$ ms the waves are almost completely separated, and the hyperpolarization parts only slightly overlap.

To reduce the offline computation time for sPOD, we do not use the whole time domain of the snapshot data for the optimization problem. Instead, we use the snapshot data until the fourth millisecond, respectively for the two cases $\left[T_0, T_0 + 4 \right]$.

Generally, we expect the cost functional of the optimization problem of sPOD to be highly non-convex. Hence, a good choice of initial values is crucial for the success of sPOD. Since we assume a constant propagation speed, we initialize the paths by linear regression over the whole time domain separately for both waves, similar to [7, Algorithm 1].

4.2.1 Completely Separated Waves

We start with the less challenging scenario based on the snapshot data for $T_0 = 18$ ms with completely separated waves and two basis functions. For the initialization, we split the snapshot data at $t = T_0$ in the center of the domain and use each half separately for different basis functions with the other half evaluated to zero. We take these values as the amplitudes in the linear combination of the basis functions $\{\psi_1, ..., \psi_n\}$ to construct the initial basis functions $\{\varphi_1, ..., \varphi_r\}$. In accordance with the theoretical construction in Section 3, these basis functions are continuous over the whole domain Ω.

We obtain a relative error of $e_{\text{rel}} = 0.405$ and a maximal absolute error of $e_\infty = 102.646$. From comparing the coefficients of the initial paths $p_i^{(0)}$ and the optimized paths p_i,

$$\bar{p}_1^{(0)} = -0.245, \qquad \bar{p}_2^{(0)} = 0.245,$$
$$\bar{p}_1 = -0.187, \qquad \bar{p}_2 = 0.187,$$

we observe that the wave is inhibited from traveling through the spatial domain to remedy the hole in the center resulting from shifting the basis functions in opposite directions to the sides.

This is a general problem because the resting potential V_{eql} is not zero, consequently, the chosen initialization is
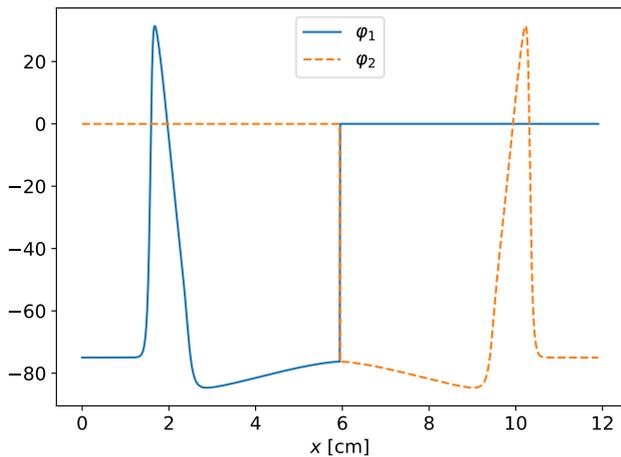
(a) $r = 2$ without transformation
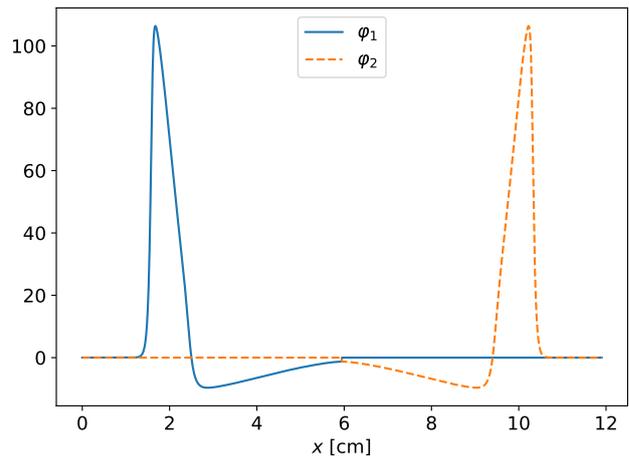
(b) $r = 2$ with transformation

Figure 3 – Spatial parts φ of reduced basis functions from sPOD for completely separated wavefronts initialized with the snapshot data.

not optimal for this case but shows an underlying problem. Shifting basis functions away from one another results in an area in the center where the basis functions evaluate to zero by construction, compare Figure 3(a).

To deal with this problem, we define the so-called equilibrium transformation that transforms the absolute voltage $u(t, x)$ to the relative voltage $u_{\mathrm{rel}}(t, x)$, such that $u_{\mathrm{rel}}(t, x)$ evaluates to zero in the equilibrium state. This transformation can easily be carried out by subtracting the resting potential from the data

$$u_{\mathrm{rel}}(t, x) = u(t, x) - V_{\mathrm{eql}}.$$

This transformation also has to be applied during the projection between the reduced and original space, and we continue under the assumption that the equilibrium transformation is applied in any necessary case.

With the transformed setting we observe a significantly smaller relative error of $e_{\mathrm{rel}} = 0.0027$ and a maximal absolute error of $e_{\infty} = 1.2607$, see additionally Table 2. A further evaluation using the absolute errors is visualized in Figure 4. It shows that the wavefronts are captured accurately; however, a small amount of the overlapping wave parts in the center cannot be represented.

We increase the number of basis functions to $r = 4$ and use the two already optimized basis functions as initialization and initialize the two additional basis functions such that they evaluate to zero with a non-zero gradient. However, this leads to no improvement due to the non-convergent behavior of the optimizer and results in unchanged basis functions. In order to increase the accuracy by adding additional basis functions, investi-

gating different approaches to the optimization problem might be necessary.

Since the cost function of the optimization problem is expected to be highly non-convex, it is likely that the optimization algorithm starts close to a local minimum by initializing the basis functions with the snapshot. To analyze this aspect, we initialize $r = 2$ basis functions such that they constantly evaluate to zero in the whole domain. These experiments result in the relative and absolute errors listed in Table 2 as the undivided case.
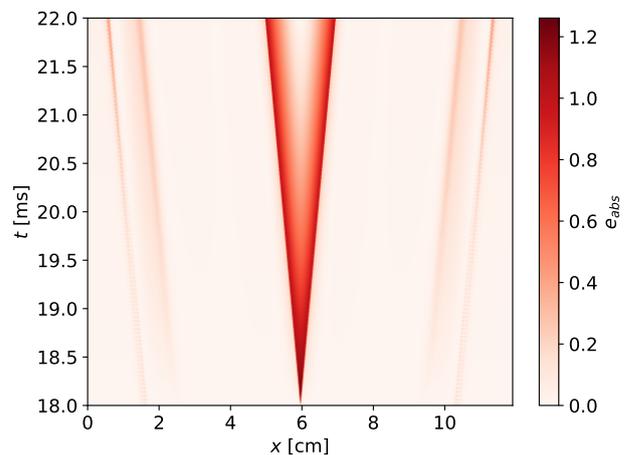
Figure 4 – Absolute error of sPOD for completely separated wavefronts with $r = 2$, initialized with snapshot data, and using the equilibrium transformation over the short time interval $[T_0, T_0 + 4]$.

From the basis functions of the reduced spaces in Figure 5(a), we conclude that the basis functions are not divided into the wavefronts. This results in wavefronts

(a) undivided domain
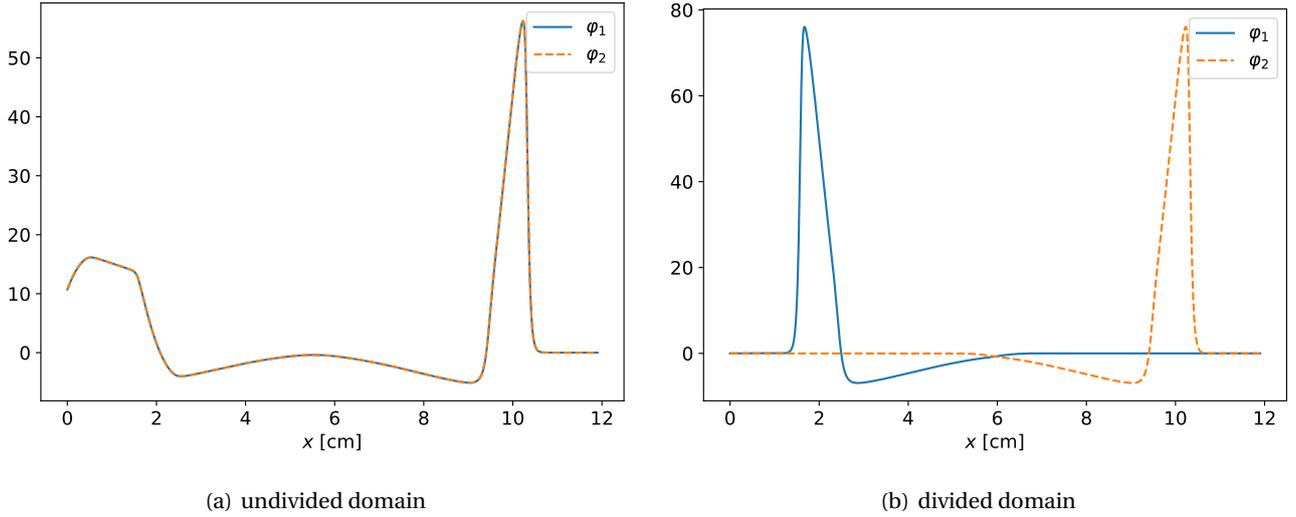
(b) divided domain

Figure 5 – Spatial parts φ of reduced basis functions from sPOD for completely separated wavefronts with $r = 2$, zero-initialized for the undivided and divided domain case.

Table 2 – Errors of sPOD for completely separated wavefronts with $r = 2$, with equilibrium shift for snapshot initialized and zero-initialized for the undivided and divided domain case

	e_{rel}	e_∞	$\bar{e}_{\text{abs}}$
Snapshot (undivided)	0.0027	1.2607	0.0762
Zero (undivided)	0.2271	106.401	7.0244
Zero (divided)	0.0031	14.2877	0.0785

that cannot be captured accurately. However, for a divided domain in Figure 5(b), the shape of the wavefront is reconstructed properly.

To analyze this problem further, we follow the approach of separately constructing basis functions for the divided domain mentioned in [7]. Therefore, we divide the snapshot data along its center axis, set the corresponding other half to the equilibrium state, construct the reduced space with one basis function $r = 1$ each for the left and right wave separately and merge the reduced spaces to obtain a merged space with $r = 2$. In that way, we obtain much smaller relative and maximal absolute errors as listed in Table 2 as well. We observe that the relative errors are similar to the results from the snapshot data initialization for $r = 2$ in Table 2. However, using the snapshot data initialization for $r = 2$ yields smaller maximal absolute errors.

Conclusions   Transforming the snapshots to the equilibrium state allows to remedy the numerical impact of shifting basis functions in opposite directions, leading to zero evaluations of the reduced ansatz space in the cen-

ter of the domain. Constructing the reduced space for both wavefronts at once with zero-initialized reduced basis functions leads to high relative and maximal absolute errors. In addition, by splitting up the wavefronts and separately constructing reduced spaces for them with zero initialization, the wavefronts can be accurately captured. We conclude that with the sPOD and the chosen optimizer, the assignment of the reduced basis functions to the two different wavefronts strongly depends on the initialization. The optimizer could converge to a local minimum due to the non-convexity of the cost functional, causing this effect.
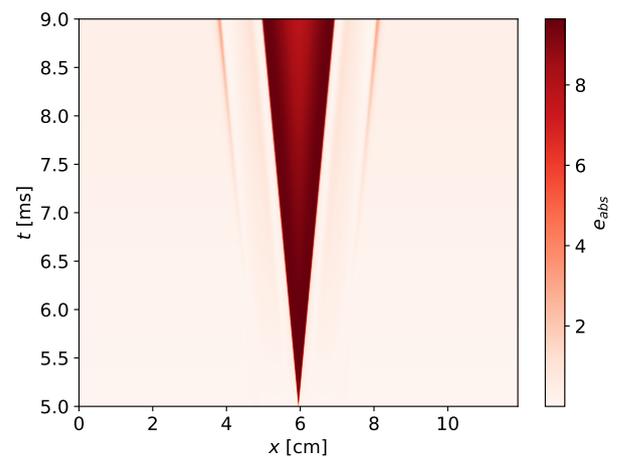
Figure 6 – Absolute error of sPOD for partly overlapping wavefronts with $r = 2$, initialized with snapshot data, over the short time interval $[T_0, T_0 + 4]$.

(a) Zoom in left peaks

(b) Zoom in center
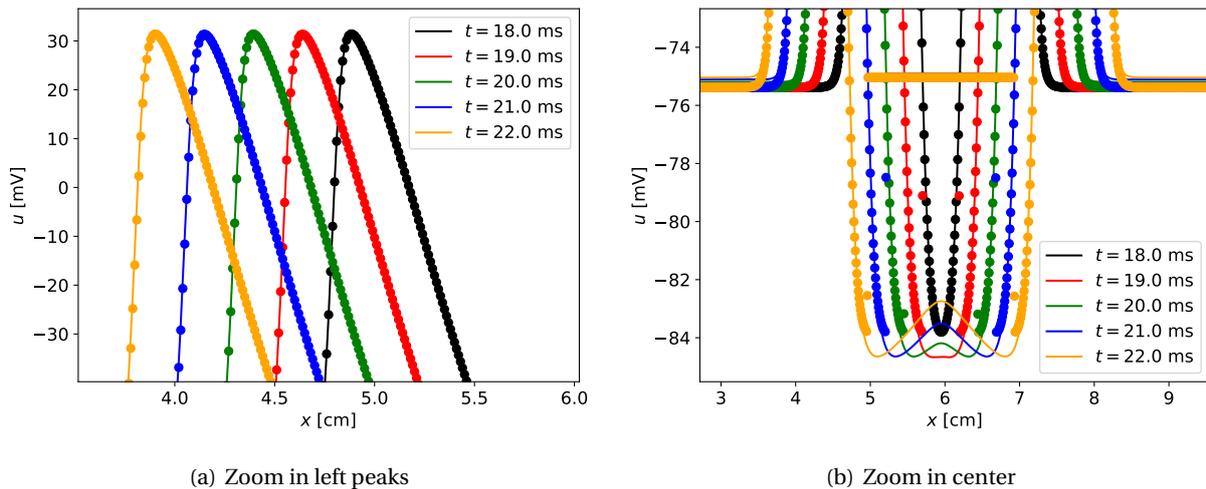
Figure 7 – Comparison of the original wavefront (solid line) and reconstructed wavefront (dots) of sPOD for partly overlapping wavefronts with $r = 2$, initialized with snapshot data for the short time interval $[T_0, T_0 + 4]$.

4.2.2 Partly Overlapping Wavefronts

We continue with the more challenging scenario of the partly overlapping wavefronts based on the snapshot data from $T_0 = 5$ ms. For an initialization with the snapshot data, we obtain larger relative and maximal absolute errors, listed in Table 3, compared to the scenario of the completely separated wavefronts.

Table 3 – Errors of sPOD for partly overlapping wavefronts with $r = 2$ for different initializations

	e_{rel}	e_∞	$\bar{e}_{\text{abs}}$
Snapshot (undivided)	0.0372	9.6404	0.9552
Zero (undivided)	0.2325	100.2865	6.8009
Zero (divided)	0.0178	14.9812	0.4595

Figure 6 shows that these larger errors mostly stem from the center part and not from the wavefront. In addition, we can see that the wavefront is captured accurately by comparing the reconstructed waves with the original waves, visualized in Figure 7. However, by zooming into the center, we observe that the hyperpolarization part of the AP is almost completely cut off and cannot be represented by the reduced space. This behavior is similar to the reduced space for the completely separated wavefronts with snapshot data initialization, where the center part is the cause of the errors but has less impact. Following that, we can deduce that with an increase in the overlap of the hyperpolarization parts of the two waves, the error in the center increases as well and information about the hyperpolarization is lost.

For the zero initialization, we obtain the relative and absolute errors listed in Table 3. Splitting the snapshot data and constructing two reduced spaces reduces both

the relative and maximal absolute error. Constructing both reduced ansatz spaces separately leads to a smaller relative error compared to the snapshot data initialization. Looking at the shape of the reduced basis functions in Figure 8 we notice that this decrease in the errors most likely stems from a better approximation of the hyperpolarization part. The initialization with snapshot data most likely got stuck in a local minimum.
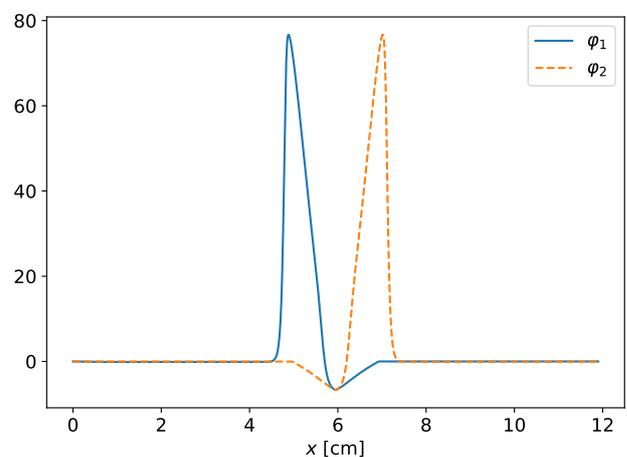
Figure 8 – Spatial parts φ of reduced basis functions from sPOD for partly overlapping wavefronts with $r = 2$ with zero initialization for the divided parts separately.

Conclusions   We conclude that in this scenario, using the split snapshot as initialization allows us to capture the non-overlapping wavefront, similar to the scenario of completely separated wavefronts. In addition, the reduced space cannot reproduce the overlapping parts of the wave, especially the hyperpolarization part can-

not be represented in the reduced space. Using a snapshot initialization is the best option if we are only interested in capturing the wavefronts. Thus, based on the presented approach, we can only approximate the full waves accurately after they are completely separated.

4.3 Simulation with the Reduced Order Model

With the constructed reduced space for the two scenarios, we can now evaluate the accuracy and efficiency of the ROM. To this end, we distinguish between the two scenarios from Section 4.2 with the different starting times $T_0 = 18$ ms and $T_0 = 5$ ms and use the reduced ansatz space with $r = 2$ reduced basis functions constructed with snapshot initialization. In the following,

Table 4 – Offline errors of the reduced ansatz space with $r = 2$ and runtime of the full order simulation

T_0 [ms]	e_{rel}	e_∞	time [s]
18	2.70×10^{-3}	1.26	39.60
5	3.72×10^{-2}	9.64	168.00

we determine the accuracy of the ROM in comparison to the FOM and compare these online errors with the offline errors in Table 4. Additionally, we use the computation time of the FOM simulations to calculate speedup factors, indicating the efficiency of the ROM. We round these speedup factors to the nearest integer and use the computation times of the FOM listed in Table 4.

4.3.1 Completely Separated Waves

We start with the first scenario ($T_0 = 18$ ms) for completely separated waves. We simulate the ROM for five time-step widths $h_t \in \{0.001, 0.005, 0.01, 0.05, 0.1\}$ and immediately observe the instability of the explicit Euler time-stepping scheme for a time-step width of $h_t = 0.1$. From the results in Table 5, we observe a slight increase in the relative errors while the computation time is drastically reduced. Even if the simulation of the ROM uses the same time-step width of the FOM, the simulation time is two times faster.

Exploiting the possibility of increasing the time-step width up to 0.05 ms, the computation time can be reduced by a factor of 73, with online errors of the same order of magnitude as the offline errors, compare Table 4. Hence, the error from the reduced ansatz space dominates the error in the simulation. We only have to accept an increase of approximately 10 % in the maximal absolute errors and an increase between 3 % and

Table 5 – Online errors and performance analysis of the ROM for completely separated wavefronts and various time-step widths

h_t	e_{rel}	e_∞	time [s]	Speedup
0.001	2.80×10^{-3}	1.39	19.41	2
0.005	2.80×10^{-3}	1.39	3.97	9
0.01	2.82×10^{-3}	1.39	2.06	19
0.05	3.32×10^{-3}	1.39	0.53	73

22 % for the relative errors depending on the chosen time-step width.
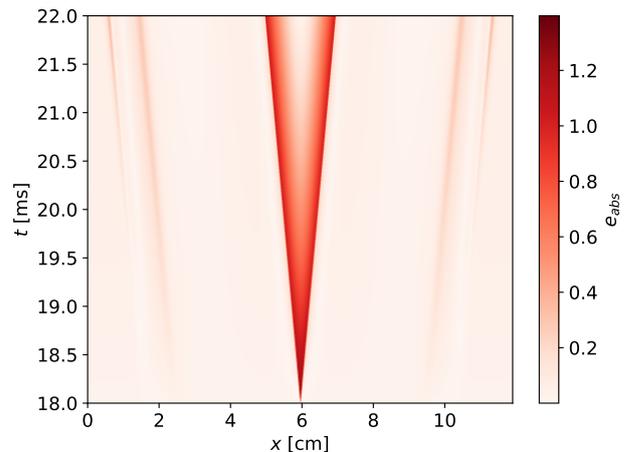
Figure 9 – Absolute error of the ROM for completely separated wavefronts simulated with $h_t = 0.001$.

To further investigate these errors, we focus on the solution of the ROM for the time-step width $h_t = 0.001$ ms. From Figure 9, we observe that the center, and thus the end of the hyperpolarization part, has a large impact on the errors even for the small time-step width and can not be represented accurately. However, since we cannot expect that the ROM can capture dynamics that the reduced space cannot represent, this effect was expected in the online phase as well since it has occurred in the offline phase. Comparing the voltage wavefronts of the ROM solution and the original wavefront in Figure 10, we note that the wavefront is captured accurately.

4.3.2 Partly Overlapping Waves

We perform the same experiments for the more challenging part with partly overlapping waves and obtain the results listed in Table 6. Comparing the online and offline errors, we observe that they have the same order of magnitude but with 50 % larger relative online error and an approximately three times larger maximal absolute online error. In addition, we can see in Figure 11

(a) Zoom in left peaks

(b) Zoom in center
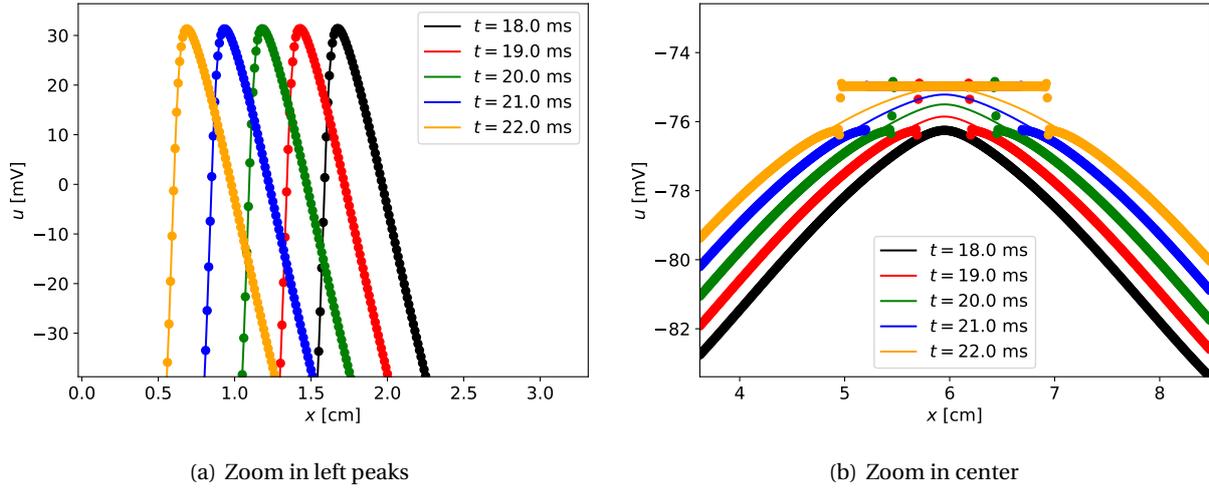
Figure 10 – Comparison of the original wavefront (solid line) and reconstructed wavefront (dots) of the ROM for completely separated wavefronts simulated with $h_t = 0.001$ over the whole time interval $[T_0, T]$.

Table 6 – Online errors and performance analysis of the ROM for partly overlapping wavefronts and various time-step widths

h_t	e_{rel}	e_∞	time [s]	Speedup
0.001	5.72×10^{-2}	23.29	78.11	2
0.005	5.72×10^{-2}	23.30	15.63	10
0.01	5.72×10^{-2}	23.32	7.88	21
0.05	5.72×10^{-2}	23.42	1.49	112

that the wavefront is not captured accurately, and the hyperpolarization part cannot be represented. Since we know that the reduced ansatz space cannot represent the hyperpolarization part, the high absolute errors in the center are no surprise. Despite the loss in accuracy,

the ROM simulation is up to 112 times faster than the FOM simulation.

Ion Channel Probability States

Since the voltage state directly affects the ion channel probability states, we analyze the impact of the errors in the voltage state on these other states as well. Therefore, we exemplarily use the ROM for completely separated wavefronts simulated with $h_t = 0.001$. From the absolute error snapshots for these states, visualized in Figure 12, we observe a small error in the center for all states y_1, y_2, y_3. In addition, a small absolute error at the wavefronts for the sodium ion channel probability y_1 is present. The error in the center is caused by the lack of

(a) Zoom in left peaks

(b) Zoom in center
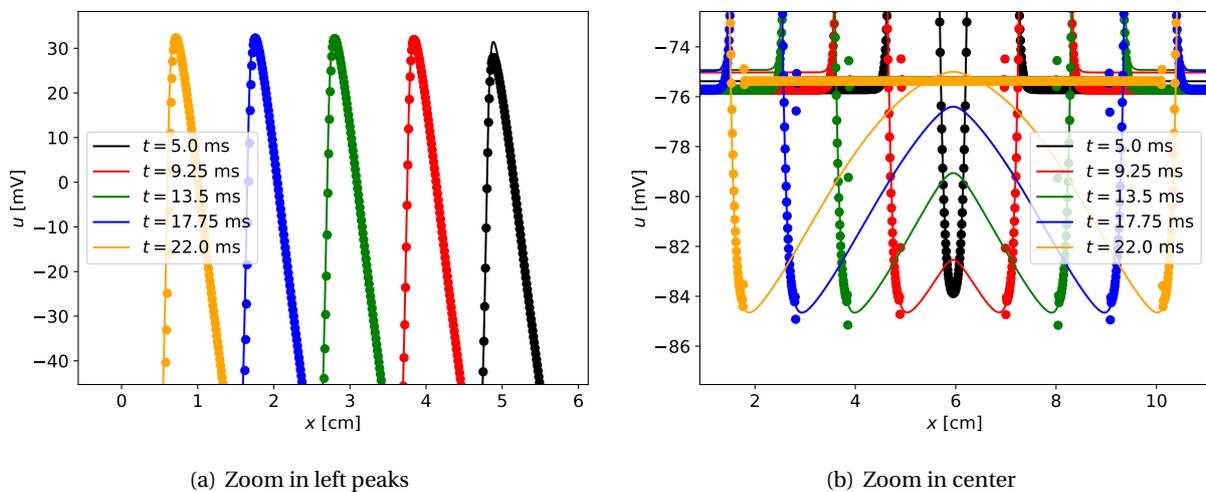
Figure 11 – Comparison of the original wavefront (solid line) and reconstructed wavefront (dots) of ROM for partly overlapping wavefronts simulated with $h_t = 0.001$ over the whole time interval $[T_0, T]$.

(a) State y_1 (b) State y_2 (c) State y_3
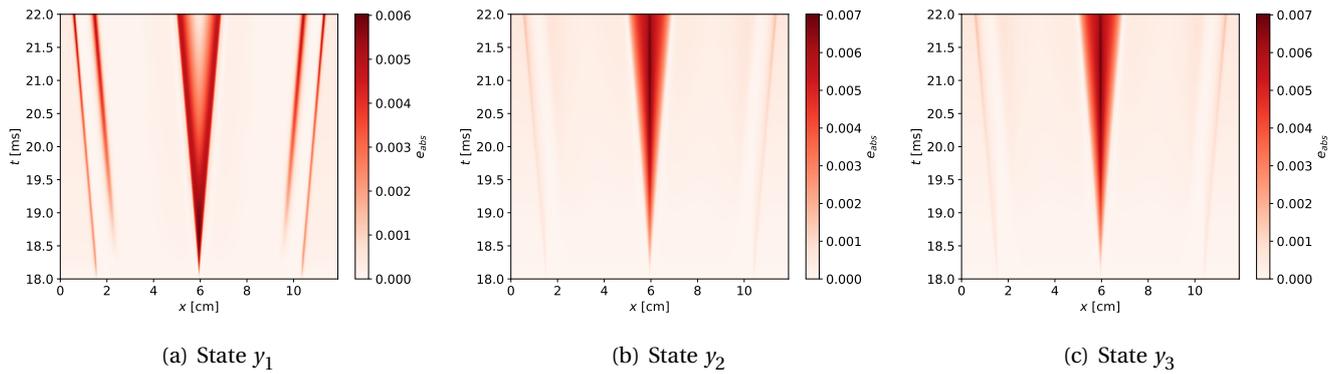
Figure 12 – Absolute error snapshots for the ion channel probability states of the ROM for completely separated wavefronts simulated with $h_t = 0.001$.

expressiveness of the reduced ansatz space for the end of the hyperpolarization in the center. Depending on the time step size these errors increase but the shape of the front profile of the traveling waves is maintained.

5 Discussion

We have applied the sPOD, which is based on dynamically transformed basis functions, to the electrophysiological simulation of the propagation of APs in muscle fibers. The reduced approximation ansatz has been based on a linear combination of transformed basis functions with constant transport velocity.

The propagation of the AP through the spatial domain originates from the center of the muscle fiber, resulting in overlapping wavefronts in the activation area. The ROM with the constructed reduced ansatz space can capture the wavefront accurately with a small relative error. However, the complete wave, especially the large hyperpolarization part, can only be captured for the completely separated waves.

We have observed a drastic reduction in the simulation time compared to the FOM computation. One important aspect is the possibility of using larger time-step widths for the time-stepping scheme resulting in fewer iterations. In that way, we have shown that applying MOR techniques in biophysical simulations for medical applications is indeed a valid approach to drastically reduce the computation time while still achieving a satisfactory accuracy of the simulation results.

Despite the achieved performance under mostly preserved accuracy, some difficulties still need to be addressed in the future. First, the activation of the AP in the center of the fiber must be integrated into the ROM. One promising approach follows [7], combining the presented ROM for the wavefront with another approach for the activation in the center. Another critical topic is the hyperpolarization part of the wave. Biologically, this part is essential for activating another AP and thus should be captured by the ROM. Since the hyperpolarization parts make up a large percentage of the complete wave and are nearly separated after 18 ms, using different reduced basis functions for the wavefronts and the hyperpolarization parts could increase the approximability. To enable predictions in the online phase using parameter settings that have not been used in the offline phase, the ROM has to be extended to paths being treated as time-dependent unknowns of the ROM, instead of being fixed. In that way, the propagation depends on the other states of the ROM, similar to the wildland fire simulation in [7]. In addition, extending the ROM to the ion channel probabilities in the electrophysiological model can also potentially increase the simulation performance further. Despite that the ion channel probabilities behave similarly to the voltage and result in a wave traveling through the spatial domain, the presented approach is not optimal due to the solely nonlinear model equations for the ion state probabilities and the additional costs of projecting between the reduced and full dimensional space. Instead, an approximation of the nonlinearity, for example, by the DEIM [11] can be adapted to the setting of the dynamically transformed basis functions, similar to [7, Chapter 3.2]. Furthermore, since the electrophysiological model is embedded in a multiscale skeletal muscle simulation, including the contraction of the muscle, the ROM can be integrated into this much more complex framework.

References

[1] AMBOSS. Resting potentials and action potentials. *Amboss Medical Encyclopedia*, 2022.

[2] Y. Bar-Cohen, editor. *Electroactive Polymer (EAP) Actuators as Artificial Muscles: Reality, Potential, and Challenges, Second Edition*. SPIE, 2004.

[3] Y. Bar-Cohen and C. Breazeal. Biologically inspired intelligent robots. In Y. Bar-Cohen, editor, *Smart Structures and Materials 2003: Electroactive Polymer Actuators and Devices (EAPAD)*, volume 5051, pages 14–20. International Society for Optics and Photonics, SPIE, 2003.

[4] P. Benner, A. Cohen, M. Ohlberger, and K. Willcox. *Model Reduction and Approximation*. Computational Science & Engineering. SIAM, 2017.

[5] F. Black, P. Schulze, and B. Unger. Projection-based model reduction with dynamically transformed modes. *ESAIM: Mathematical Modelling and Numerical Analysis*, 54(6):2011–2043, 2020.

[6] F. Black, P. Schulze, and B. Unger. Model order reduction with dynamically transformed modes for the wave equation. *PAMM*, 20(1):e202000321, 2021.

[7] F. Black, P. Schulze, and B. Unger. Efficient wildland fire simulation via nonlinear model-order reduction. *Fluids*, 6(8):280, 2021.

[8] F. Black, P. Schulze, and B. Unger. Modal decomposition of flow data via gradient-based transport optimization. In R. King and D. Peitsch, editors, *Active Flow and Combustion Control 2021*, pages 203–224, Cham, 2022. Springer International Publishing.

[9] M. Boulakia, E. Schenone, and J-F. Gerbeau. Reduced-order modeling for cardiac electrophysiology. application to parameter identification. *International Journal for Numerical Methods in Biomedical Engineering*, 28(6-7):727–744, 2012.

[10] C. P. Bradley, N. Emamy, T. Ertl, D. Göddeke, A. Hessenthaler, T. Klotz, A. Krämer, M. Krone, B. Maier, M. Mehl, T. Rau, and O. Röhrle. Enabling detailed, biophysics-based skeletal muscle models on HPC systems. *Frontiers in Physiology*, 9:816, 2018.

[11] S. Chaturantabut and D. C. Sorensen. Nonlinear model reduction via discrete empirical interpolation. *SIAM Journal on Scientific Computing*, 32(5):2737–2764, 2010.

[12] R. Fletcher. *Practical methods of optimization*. John Wiley & Sons, Chichester, England, 2 edition, 2000.

[13] S. Fresca, A. Manzoni, L. Dedè, and A. Quarteroni. Deep learning-based reduced order models in cardiac electrophysiology. *PLOS ONE*, 15(10):e0239416, 2020.

[14] S. Fresca, A. Manzoni, L. Dedè, and A. Quarteroni. Pod-enhanced deep learning-based reduced order models for the real-time simulation of cardiac electrophysiology in the left atrium. *Frontiers in Physiology*, 12, 2021.

[15] R. J. Gelinas, S. K. Doss, and K. Miller. The moving finite element method: Applications to general partial differential equations with multiple large gradients. *Journal of Computational Physics*, 40(1):202–249, 1981.

[16] J.-F. Gerbeau, D. Lombardi, and E. Schenone. Reduced order model in cardiac electrophysiology with approximated Lax pairs. *Advances in Computational Mathematics*, 41(5):1103–1130, 2014.

[17] A. M. Gilroy, B. R. MacPherson, J. Wikenheiser, M. Schuenke, E. Schulte, and U. Schumacher. *Atlas of Anatomy*. Thieme Medical, New York, NY, 4 edition, 2020.

[18] M. Gubisch and S. Volkwein. *Model Reduction and Approximation*, chapter 1: Proper Orthogonal Decomposition for Linear-Quadratic Optimal Control, pages 3–63. Computational Science & Engineering. SIAM, 2017.

[19] 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.

[20] Philip Holmes, John L. Lumley, Gahl Berkooz, and Clarence W. Rowley. *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge Monographs on Mechanics. Cambridge University Press, 2 edition, 2012.

[21] M. C. Huber, U. Jonas, and S. M. Schiller. An autonomous chemically fueled artificial protein muscle. *Advanced Intelligent Systems*, 4(4):9, 2022.

[22] M. Kahlbacher and S. Volkwein. Galerkin proper orthogonal decomposition methods for parameter dependent elliptic systems. *Discussiones Mathematicae. Differential Inclusions, Control and Optimization*, 27:95–117, 2007.

[23] A. Ketley, M. Wojciechowska, S. Ghidelli-Disse, P. Bamborough, T. K. Ghosh, M. Lopez Morato, S. Sedehizadeh, N. A. Malik, Z. Tang, P. Powalowska, M. Tanner, R. Billeter-Clark, R. C. Trueman, P. C. Geiszler, A. Agostini, O. Othman, M. Bösche, M. Bantscheff, M. Rüdiger, D. E. Mossakowska, D. H. Drewry, W. J. Zuercher, C. A. Thornton, G. Drewes, I. Uings, C. J. Hayes, and J. D. Brook. Cdk12 inhibition reduces abnormalities in cells from patients with myotonic dystrophy and in a mouse model. *Science Translational Medicine*, 12(541):eaaz2415, 2020.

[24] D. Leonardis, M. Barsotti, C. Loconsole, M. Solazzi, M. Troncossi, C. Mazzotti, V. P. Castelli, C. Procopio, G. Lamola, C. Chisari, M. Bergamasco, and A. Frisoli. An EMG-controlled robotic hand exoskeleton for bilateral rehabilitation. *IEEE Transactions on Haptics*, 8(2):140–151, 2015.

[25] D. C. Liu and J. Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1):503–528, 1989.

[26] B. Maier. *Scalable biophysical simulations of the neuromuscular system*. PhD thesis, University of Stuttgart, 2021.

[27] B. Maier, N. Emamy, A. Krämer, and M. Mehl. *Highly parallel multi-physics simulation of muscular activation and EMG*, pages 610–621. CIMNE, 2019.

[28] K. Miller and R. N. Miller. Moving finite elements. i. *SIAM Journal on Numerical Analysis*, 18(6):1019–1032, 1981.

[29] M. Mordhorst, T. Heidlauf, and O. Röhrle. Predicting electromyographic signals under realistic conditions using a multiscale chemo-electro-mechanical finite element model. *Interface Focus*, 5(2):20140076, 2015.

[30] University of Freiburg. Artificial muscles made of proteins. *ScienceDaily*, 2022.

[31] University of Nottingham. Scientists edge closer to treatment for myotonic dystrophy. *ScienceDaily*, 2020.

[32] A. Quarteroni, A. Manzoni, and F. Negri. *Reduced basis methods for partial differential equations*. UNITEXT - La Matematica per il 3+2. Springer International Publishing, Cham, 1 edition, 2015.

[33] A. Rahmatillah, O. N. Rahma, M. Amin, S. I. Wicaksana, K. Ain, and R. Rulaningtyas. Post-stroke rehabilitation exoseleton movement control using EMG signal. *International Journal on Advanced Science, Engineering and Information Technology*, 8(2):616–621, 2018.

[34] J. Reiss. Optimization-based modal decomposition for systems with multiple transports. *SIAM Journal on Scientific Computing*, 43(3):A2079–A2101, 2021.

[35] J. Reiss, P. Schulze, J. Sesterhenn, and V. Mehrmann. The shifted proper orthogonal decomposition: A mode decomposition for multiple transport phenomena. *SIAM Journal on Scientific Computing*, 40(3):A1322–A1344, 2018.

[36] P. Schulze, J. Reiss, and V. Mehrmann. Model reduction for a pulsed detonation combuster via shifted proper orthogonal decomposition. In R. King, editor, *Active Flow and Combustion Control 2018*, pages 271–286, Cham, 2018. Springer International Publishing.

[37] P. Virtanen, R. Gommers, T. E. Oliphant, M Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C J Carey, I. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.