

UCSF

UC San Francisco Previously Published Works

Title

Reconstruction of three-dimensional scroll waves in excitable media from two-dimensional observations using deep neural networks.

Permalink

<https://escholarship.org/uc/item/7vb6w2hs>

Journal

Physical Review E (statistical, nonlinear, biological, and soft matter physics), 107(1-1)

Authors

Lebert, Jan

Mittal, Meenakshi

Christoph, Jan

Publication Date

2023

DOI

10.1103/PhysRevE.107.014221

Peer reviewed



Published in final edited form as:

Phys Rev E. 2023 January ; 107(1-1): 014221. doi:10.1103/PhysRevE.107.014221.

Reconstruction of three-dimensional scroll waves in excitable media from two-dimensional observations using deep neural networks

Jan Lebert¹, Meenakshi Mittal^{1,2}, Jan Christoph¹

¹Cardiovascular Research Institute, University of California, San Francisco, USA

²University of California, Berkeley, USA

Abstract

Scroll wave dynamics are thought to underlie life-threatening ventricular fibrillation. However, direct observations of three-dimensional electrical scroll waves remain elusive, as there is no direct way to measure action potential wave patterns transmurally throughout the thick ventricular heart muscle. Here, we study whether it is possible to reconstruct simulated scroll waves and scroll wave chaos using deep learning. We trained encoding-decoding convolutional neural networks to predict three-dimensional scroll wave dynamics inside bulk-shaped excitable media from two-dimensional observations of the wave dynamics on the bulk's surface. We tested whether observations from one or two opposing surfaces would be sufficient and whether transparency or measurements of surface deformations enhances the reconstruction. Further, we evaluated the approach's robustness against noise and tested the feasibility of predicting the bulk's thickness. We distinguished isotropic and anisotropic, as well as opaque and transparent excitable media as models for cardiac tissue and the Belousov-Zhabotinsky chemical reaction, respectively. While we demonstrate that it is possible to reconstruct three-dimensional scroll wave dynamics, we also show that it is challenging to reconstruct complicated scroll wave chaos and that prediction outcomes depend on various factors such as transparency, anisotropy and ultimately the thickness of the medium compared to the size of the scroll waves. In particular, we found that anisotropy provides crucial information for neural networks to decode depth, which facilitates the reconstructions. In the future, deep neural networks could be used to visualize intramural action potential wave patterns from epi- or endocardial measurements.

Keywords

Scroll waves; excitable media; action potential waves; ventricular fibrillation; deep learning

VIII. AUTHOR CONTRIBUTIONS

JL and JC conceived the research and implemented the algorithms. JL, MM, and JC conducted the data analysis and designed the figures. JL and JC wrote the manuscript.

IX. CONFLICT OF INTEREST

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

I. INTRODUCTION

Scroll wave dynamics occur in excitable reaction-diffusion systems, termed ‘excitable media’. They are conjectured to underlie life-threatening heart rhythm disorders, such as ventricular fibrillation. In the heart, nonlinear waves of electrical excitation propagate through the cardiac muscle and initiate its contractions. The electrical waves are conjectured to degenerate into electrical scroll wave dynamics and spatio-temporal chaos via a cascade of wavebreaks during the onset of ventricular fibrillation. However, while the dynamics of scroll waves have been studied extensively in computer simulations [1–3], the direct visualization of scroll waves throughout the depths of the heart muscle remains a challenge.

Spiral wave-like action potential waves can be imaged on the heart surface during ventricular tachycardia or fibrillation using voltage-sensitive optical mapping [4–8], and the surface observations are in agreement with simulated three-dimensional scroll wave dynamics [3]. Otherwise, only few and indirect experimental evidence of scroll waves in the heart exists. Voltage-sensitive transillumination imaging was used to measure projections of scroll waves on the surface of the isolated right ventricle of porcine and sheep hearts [9–11]. The right ventricles are thinner than the left and can therefore be penetrated (~ 0.5 cm) by near-infrared light, making them semi-transparent. Consequently, it was possible to locate focal wave sources inside the volume of the right ventricle using transillumination imaging [12, 13]. More recently, it was shown that ultrasound imaging can reveal mechanical vortices in the ventricles of whole isolated porcine hearts, which co-exist with electrical vortices on the epicardial surface, suggesting that the heart’s mechanical dynamics reflect electrical scroll wave dynamics [7, 14]. However, it remains difficult to extrapolate the measured projections or surface observations of the electrical dynamics into the depths of the cardiac muscle, and eventually correlate them with mechanical measurements. Fully three-dimensional reconstructions of scroll waves were obtained in the Belousov-Zhabotinsky chemical reaction, which is a transparent excitable medium that exhibits similar but much slower wave dynamics than in the heart, using optical tomography [15–19] and magnetic resonance imaging [20]. In contrast, three-dimensional action potential waves have been directly measured in small rat and zebrafish hearts using laminar optical tomography [21] or light-sheet microscopy [22]. However, attempts to obtain three-dimensional visualizations of scroll waves inside the optically dense cardiac muscle of large mammalian hearts have not yet attained a similar quality, and better measurement and reconstruction techniques are needed.

Multiple numerical approaches for the reconstruction of scroll waves from surface observations have previously been proposed: Berg et al. [23] attempted to recover simulated scroll wave chaos from single-surface observations using a synchronization-based data-assimilation approach. However, while the approach was successful at recovering scroll wave chaos from sparse measurements within the medium, see also [24], it was not suited to extrapolate scroll wave dynamics into the three-dimensional bulk-shaped medium from surface observations. Hoffman et al. [25, 26] analyzed dual-surface observations (comparable to measuring both the epi- and endocardium) using a different data-assimilation approach, the local ensemble transform Kalman filter [27], and subsequently reconstructed simulated scroll waves successfully. The question remains if it is possible

to reconstruct truly complex three-dimensional scroll wave chaos from single- or dual-surface observations. More recently, neural networks were used to predict cardiac dynamics from sparse or partial observations with promising results [28–30]. However, the task of predicting scroll wave dynamics from surface observations using neural networks has not yet been established.

Here, we provide a numerical proof-of-principle that deep encoding-decoding convolutional neural networks, under certain conditions, can be used to reconstruct three-dimensional scroll wave dynamics from two-dimensional observations of the dynamics, see Fig. 1. We show that scroll waves can be recovered fully when the size of the waves is in the order of the thickness of the medium, or when the waves are observed in transparent anisotropic excitable media as projections. In the latter case, the waves can be much smaller and more complicated. We tested several deep convolutional neural network architectures and analyzed their reconstruction performance depending on opacity, thickness and anisotropy of simulated excitable media.

II. METHODS

We performed simulations of three-dimensional electrical and electromechanical scroll wave dynamics in bulk-shaped isotropic and anisotropic (elastic) excitable media, respectively, and used neural networks to predict the three-dimensional wave patterns from a short sequence of two-dimensional observations of the dynamics on the bulk’s surface. We distinguished ‘laminar’ scroll wave dynamics consisting of 1–3 meandering scroll waves and ‘turbulent’ scroll wave chaos.

A. Scroll Wave Dynamics in Elastic Excitable Media

We simulated electrical scroll wave dynamics in bulk-shaped excitable media of size $128 \times 128 \times d_z$ voxels with varying thicknesses or depths $d_z \in \{8, \dots, 40\}$. We used the phenomenological Aliev-Panfilov model [31] to simulate nonlinear waves of electrical excitation:

$$\frac{\partial u}{\partial t} = \nabla \cdot (\mathbf{D} \nabla u) - ku(u - a)(u - 1) - ur \quad (1)$$

$$\frac{\partial r}{\partial t} = \epsilon(u, r)(ku(a + 1 - u) - r) \quad (2)$$

The dynamic variables u and r represent the local electrical excitation (voltage) and refractory state, respectively, and are dimensionless, normalized units. Together with the term $\epsilon(u, r) = \epsilon_0 + \mu_1 r / (u + \mu_2)$, the partial differential equations describe the local excitable kinetics and diffusive dynamics. The parameters k , a , ϵ_0 , μ_1 and μ_2 , listed in Table I, influence the properties of the excitation waves. The partial differential equations were integrated

using the forward Euler method in a finite differences numerical integration scheme and used Neumann (zero-flux) boundary conditions. We simulated both isotropic

$$\mathbf{D} = D_{\text{iso}} \mathbf{I} \quad (3)$$

and anisotropic excitable media with locally varying fiber direction with diffusion coefficients for the parallel D_{\parallel} fiber direction and perpendicular $D_{\perp 1}, D_{\perp 2}$ directions as described in [1]:

$$\mathbf{D} = \begin{pmatrix} D_{11} & D_{12} & 0 \\ D_{21} & D_{22} & 0 \\ 0 & 0 & D_{33} \end{pmatrix}$$

$$\begin{aligned} D_{11} &= D_{\parallel} \cos^2(\theta(z)) + D_{\perp 1} \sin^2(\theta(z)) \\ D_{22} &= D_{\parallel} \sin^2(\theta(z)) + D_{\perp 1} \cos^2(\theta(z)) \\ D_{12} &= D_{21} = (D_{\parallel} - D_{\perp 1}) \cos(\theta(z)) \sin(\theta(z)) \\ D_{33} &= D_{\perp 2} \end{aligned} \quad (4)$$

Here, the fiber organization represents ventricular muscle tissue with muscle fibers aligned in sheets in the x - y plane and the sheet-fiber orientation rotating throughout the thickness of the bulk. $D_{\perp 1}$ is the diffusivity perpendicular to the fiber axis in the $x - y$ plane and $D_{\perp 2}$ transmurally, for simplicity we set $D_{\perp} = D_{\perp 1} = D_{\perp 2}$. We used a varying fiber angle $\theta(z)$ ranging from 0° to 90° between the top and bottom layer of the bulk for the simulation depth $d_z = 24$ voxels:

$$\theta(z) = z \cdot \Delta\theta \quad (5)$$

For the other depths, $d_z \in \{8, 12, 16, 20, 28, 32, 40\}$, we used the same $\Delta\theta$ as in the $d_z = 24$ case. The ratio between the parallel D_{\parallel} and perpendicular D_{\perp} diffusion coefficients was set to $4 : 1$. We chose $D_{\perp} = D_{\text{iso}} = 0.05$, and adapted Δt for the Euler integration such that 500 simulation time steps correspond to about $0.5 - 1.0$ scroll wave rotations. The simulation time steps required for one scroll wave rotation fluctuates and depends on the parameter values as well as an/-isotropy. To save disk space, we stored only every 80th simulation time step as one ‘snapshot’, such that 5 snapshots covered about a half to one scroll wave rotation, see also Fig. 3.

In addition to the purely electric simulations, we also simulated electromechanical scroll wave dynamics in deforming excitable media, as described in [24]. In short, we coupled a three-dimensional mass-spring damper system with hexahedral cells and tunable fiber anisotropy [32] to the electric simulation. Each cell in the mechanical part of the simulation corresponded to one cell or voxel in the electrical part of the simulation. Active tension generation in each cell was modelled using an active stress variable T_a that is directly dependent on the excitation variable u , as described in [33]:

$$\begin{aligned} \frac{\partial T_a}{\partial t} &= \epsilon(u) \cdot (k_r u - T_a) \\ \epsilon(u) &= \begin{cases} 10 & \text{if } u < 0.05 \\ 1 & \text{if } u \geq 0.05 \end{cases} \end{aligned} \quad (6)$$

The axis along which each cell exerts active contractile force could be pointed into an arbitrary direction, and, throughout the bulk, the axis alignment matched the rotating orthotropic fiber alignment already defined in the electrical part of the simulation. The mechanical parameter k_r and other parameters shown in Table I influence the magnitude of contraction and the properties of the elasticity of the mass-spring damper system. The elastic medium's boundaries were non-rigid and confined by elastic springs acting on the medium's boundary, see [24] for details. In general, the electromechanical simulation produces three-dimensional deformation patterns that are highly correlated with the electrical scroll wave chaos.

We simulated two different regimes of scroll wave dynamics: 1) a 'laminar' regime with 1–3 meandering scroll waves with wavebreaks as shown in Figs. 4 and 5 and 2) a fully 'turbulent' scroll wave chaos regime as shown in Figs. 7–9, see Table I for the respective parameter values. For both parameter regimes, we performed 125 isotropic and anisotropic simulations of electrical scroll wave chaos for different bulk depths $d_z \in \{16, 24, 32, 40\}$ voxels for the 'laminar' regime and $d_z \in \{8, 12, 16, 20, 24, 28, 32\}$ for the 'turbulent' regime. 100 simulations were used during generation of the training dataset and 25 simulations were exclusively used for evaluation. Furthermore, we performed 125 simulations of electromechanical scroll wave chaos for a thickness of $d_z = 24$, with the same split between training and evaluated dataset. The electrical and mechanical parameters were identical in each simulation. However, the initial conditions $u_{t=0}(x, y, z), r_{t=0}(x, y, z)$ were randomized and therefore different in each simulation. We used cross-field stimulation to set $u_{t=0}, r_{t=0}$ such that two scroll waves are induced at random positions $(x_1, y_1), (x_2, y_2)$. Additionally we added a small amount of Gaussian noise (standard deviation $\sigma = 0.1$) to the initial conditions $u_{t=0}, r_{t=0}$. With both parameter sets, the dynamics quickly diverged. We discarded the first 75 snapshots of each simulation (approximately 15 scroll wave rotation periods) and used the remaining 500 snapshots (approximately 100 scroll wave rotation periods) for generating the training or evaluation data, respectively. If the excitation in a simulation decayed ($u \approx 0 \quad \forall(x, y, z)$), we restarted the simulation with new initial conditions $u_{t=0}, r_{t=0}$. The training dataset consists of 20,000 randomly selected samples from the 100 training simulations, and the evaluation dataset uses 5,000 randomly chosen samples from the 25 evaluation simulations. Consequently, training and evaluation datasets were completely separate datasets. The numerical simulation was implemented in C++, the source code for the simulations is available in [24].

B. Deep Learning-based Reconstruction of Scroll Wave Dynamics

We implemented and tested deep neural networks, which each analyze a short temporal sequence of 5 subsequent two-dimensional snapshots of electrical wave patterns $\tilde{u}_i(x, y)$ to reconstruct a single fully three-dimensional snapshot $u_t(x, y, z)$ of scroll wave dynamics:

$$(\tilde{u}_1(x, y), \dots, \tilde{u}_5(x, y)) \rightarrow u_t(x, y, z). \quad (7)$$

We found empirically that 5 snapshots provide sufficient information about the dynamics, see also [28] or [35]. We also tested reconstructing the dynamics with a series of 10 snapshots and did not observe an improvement in performance compared to 5 snapshots. Consequently, we used 5 snapshots by default. Further, we found that the reconstruction accuracy does not depend on whether the network analyzes the current snapshots plus 4 snapshots sampled in the past or in the future with respect to the current snapshot, or whether 2 are sampled in the past and 2 in the future, respectively. The two-dimensional snapshots, see Figs. 2, 3 and 6, are either i) the top surface layer (single-surface mode, Fig. 2a):

$$\tilde{u}_t(x, y) = u_t(x, y, 1), \quad (8)$$

ii) both the top and bottom surface layer (dual-surface mode, Fig. 2b):

$$\tilde{u}_t(x, y) = (u_t(x, y, 1), u_t(x, y, d_z)), \quad (9)$$

iii) both the electrical wave dynamics and the mechanical displacements $\vec{d} = (dx, dy)$ which occur in corresponding electromechanical simulations from the top surface layer (single-surface mode), see Fig. 2c):

$$\tilde{u}_t(x, y) = (u_t(x, y, 1), dx(x, y, 1), dy(x, y, 1)), \quad (10)$$

or iv) a projection of all u -values along the z -direction (depth) of the bulk (Fig. 2d):

$$\tilde{u}_t(x, y) = \frac{1}{d_z} \sum_{i=1}^{d_z} u_t(x, y, i). \quad (11)$$

For the top+bottom case (iii) and the mechanical displacement case (iv), we interleaved the snapshots. The neural network architectures we chose require three-dimensional input samples ((128,128, 5) for case (i) and (ii)), whereas in cases (iii) and (iv) the sample shape is (128, 128, 5, 3) and (128, 128, 5, 2) respectively. We changed their shape to be

three-dimensional by stacking the components: e.g. in case (iv) the resulting shape is (128, 128, 10) using the top layers for the even indices and bottom frames for the odd indices ($u_1(x, y, 1), u_1(x, y, d_z), u_2(x, y, 1) \dots u_5(x, y, d_z)$).

We evaluated four different neural network architectures with basic and more intricate designs for the three-dimensional bulk prediction task (Eq. 7). While we primarily use a U-Net [36] architecture, we validate it against a generic Encoder-Decoder architecture, TransUNet [37] and MIRNet [38]. The Encoder-Decoder convolutional neural network (CNN) is similar to the architecture we previously used in [28, 35]. It consists of an encoder stage where the spatial resolution is progressively decreased, a latent space, and a decoder stage where the spatial resolution is progressively increased back to the original resolution. The encoding and decoding steps consist of three steps, in each two padded two-dimensional convolutional layers (2D-CNN) with filter size 3×3 and rectified linear unit [39] (ReLU) activation are applied, followed by batch normalization [40] and maxpooling (encoder) or upscaling (decoder), respectively. The number of filters in 2D-CNN layers in order are 128, 128, 256, 256, 512, 512, 256, 256, 128 and 128. The U-Net architecture is identical to the Encoder-Decoder CNN architecture, except that skip connections are added between the encoder and decoder stages (see [36]). The TransUNet combines the U-Net architecture with self-attention mechanisms of Transformers [41] in its the latent space. The MIRNet architecture is different from the other evaluated architectures, as it contains parallel multi-resolution branches with information exchange, as well as spatial and channel attention mechanisms [38]. It aims at maintaining spatially-precise high-resolution representations through the entire network, while simultaneously receiving strong contextual information from the low-resolution representations. The number of trainable parameters for each network architecture and training times are listed in Table II. For all neural network architectures we use the generalized Charbonnier loss function [42, 43]:

$$l(u, \hat{u}) = \sqrt{(\hat{u} - u)^2 + \epsilon^2}, \quad (12)$$

where u is the three-dimensional ground truth, \hat{u} the prediction and we choose $\epsilon = 0.001$. The Charbonnier loss function behaves like L2 loss (mean squared error) when $u \approx \hat{u}$ and like L1 loss (mean absolute error) otherwise. We evaluated the accuracy of the predictions on the evaluation datasets with the root mean squared error (RMSE) on each z -axis layer:

$$\text{RMSE}(z) = \sqrt{\frac{1}{N} \sum_{x, y, t} (\hat{u}(x, y, z, t) - u(x, y, z, t))^2}. \quad (13)$$

To validate our findings, we studied if a neural network can accomplish a simpler task than the three-dimensional prediction: estimate the depth d_z of the simulation bulk from 5 two-dimensional observations. We tested both a depth regression and a depth classification neural network, which each predict the depth d_z of the simulation bulk

$$(\tilde{u}_1(x, y), \dots, \tilde{u}_5(x, y)) \rightarrow d_z. \quad (14)$$

The depth regression network predicts the depth d_z as a continuous value, while the depth classification network predicts the depth d_z as one of $\{8, 12, 16, 20, 24, 28, 32\}$. For this task we used the encoder part from the Encoder-Decoder architecture, followed by a global average pooling layer, two dense layers with 1024 filters with batch normalization and ReLU activation, and ultimately an output dense layer with one filter (for regression) or seven filters (for classification). For the depth classification neural network we used a categorical cross-entropy loss function with a softmax activation function for the last layer, and for the depth regression network mean squared error as loss function and ReLU as activation function. The datasets for the depth estimation was generated from the bulk prediction task datasets. We used 4,000 random samples for each depth for the training dataset and 500 samples for the evaluation dataset (in total 28,000 training samples and 3,000 evaluation samples).

All networks were trained for 20 epochs using the Adam [44] optimizer with a learning rate of 10^{-3} for the bulk prediction tasks and 10^{-5} for the depth regression and classification task. We used a batch size of 32 for the Encoder-Decoder and U-Net architectures and a batch size of 4 for TransUNet and MIRNet. All neural network models were implemented in Tensorflow [45] using Keras [46]. Training and reconstructions were performed on a NVIDIA RTX A5000 graphics processing unit (GPU).

III. RESULTS

Using deep convolutional neural networks, it is possible to reconstruct three-dimensional scroll wave dynamics inside an excitable medium when the medium's thickness is not much thicker than the scroll wave, see Fig. 4 and section III A. Reconstructions become increasingly difficult in thicker excitable media or with smaller scroll waves and more complicated dynamics, see Figs. 5, 9a) and sections IIIA–IIIC. However, complicated scroll wave chaos can be reconstructed in transparent anisotropic excitable media, see Figs. 8d), 9b), 10c,f), or with dual-surface observations in thinner opaque excitable media, see Figs. 5c) and 8b) and section IIIB.

A. Medium Thickness vs. Scroll Wave Size

The size of the scroll waves with respect to the thickness of the medium determine how well scroll wave dynamics can be reconstructed in opaque media. We simulated 'laminar' scroll waves, see Figs. 4 and 5, and 'turbulent' scroll wave chaos, see Figs. 7, 8c) and 9a), in opaque bulks with varying thicknesses. While the reconstructions from single surface observations are very accurate for the 'laminar' scroll wave dynamics in the thinner bulk (with thickness $d_z = 24$) shown in Figs. 4a,b) and in Supplemental Video 1 [34], reconstructions become more and more inaccurate with increasing thickness and prediction depth. The reconstructions of the 'laminar' scroll wave dynamics in the thicker bulk (with thickness $d_z = 40$), shown in Fig. 5a,b) and in Supplemental Video 2 [34], exhibit

artifacts towards the bottom half of the bulk. Accordingly, the vortex filaments computed from the predicted scroll wave dynamics (red) deviate substantially from the ground-truth vortex filaments (gray) in the lower half of the bulk, see Fig. 5d). With increasing bulk thickness, the three-dimensional character and complexity of the wave dynamics increases, which is reflected by the dissociation of the top and bottom layers, see Figs. 6 and 7, and by the intramural alignment of the vortex filaments in Fig. 5d). The degree of dissociation and the average scroll wave size relative to the medium's thickness determine the prediction accuracy at deeper layers. In opaque excitable media, the reconstructions become increasingly difficult the deeper one aims to predict, and they do not appear to succeed deeper than one scroll wavelength, see Figs. 5b), left panel in 8b,c) and 9a).

Correspondingly, the plots in Fig. 10a,d) show how the prediction error increases with increasing depths z (with bulks depths $d_z = 8, 12, \dots, 40$) in opaque excitable media with 'laminar' and 'turbulent' scroll wave dynamics, respectively. The error increases approximately linearly and faster with thicker bulks and with 'turbulent' than with 'laminar' scroll wave dynamics. The average single-surface reconstruction error for the 'laminar' scroll wave in the thin opaque bulk ($d_z = 16$) shown in Fig. 4b) is better than 95% (note that a root mean squared error of 0.1 (RMSE) corresponds to about 0.05 mean absolute error (MAE)), whereas in thicker bulks or with 'turbulent' dynamics the prediction accuracy can become much worse. We compared the error profiles obtained with the deep learning-based reconstruction with a naive reconstruction in which the top layer is simply repeated in each following layer in Fig. S2 [34]. The naive reconstruction produces significantly steeper error curves with both 'laminar' and 'turbulent' scroll wave dynamics. Overall, the reconstruction error fluctuates moderately over time, remains small at smaller depths, and increases as the reconstruction error increases with larger depths, see Fig. 12. Supplemental Videos 1–4 [34] give an impression of the temporal stability of the reconstructions.

B. Single-Surface vs. Dual-Surface Observations

Low prediction depths in opaque excitable media can be compensated to a certain extent by analyzing both the top and bottom surface layers in dual-surface mode rather than single-surface mode, respectively. Figs. 5c) and 8b) demonstrate how the reconstruction improves in a thick bulk with the 'laminar' scroll wave and in a thin bulk with 'turbulent' scroll wave chaos, respectively. The vortex filaments (red) in the thick bulk in Fig. 5c) match the ground-truth vortex filaments (gray) much better in dual- than in single-surface mode. The plots in Fig. 10b,e) and 11 show how the profile of the reconstruction error changes in dual-surface mode. Surprisingly, the network does not appear to benefit from the additional information in dual-surface mode with scroll wave chaos: the steep linear increase in the error persists and the error at midwall is comparable to the error in single-surface mode, see Fig. 10e). Nevertheless, with larger scroll waves or thinner bulks the error slightly decreases at midwall, see Figs. 5c), 10b) and 11b).

C. Opaque vs. Transparent Excitable Media

Figs. 8 and 9 show comparisons of reconstructions obtained in opaque and transparent excitable media, respectively. The reconstructions succeed in transparent excitable media of any thickness (that we tested) under the condition that they are anisotropic, but are

difficult to obtain in thicker opaque excitable media. Fig. 8 shows cross-sections along the depth of the bulk (z -direction) whereas Fig. 9 shows layers parallel to the surface in the $x - y$ plane. Correspondingly, panels c,f) in Fig. 10 show that the reconstruction error stays small along the bulk's transmural axis in transparent anisotropic excitable media (see also additional plots in Fig. S1 with isotropy). In anisotropic transparent excitable media, the prediction error remains small (< 0.1 RMSE) with both the 'laminar' and 'turbulent' scroll wave dynamics and with various bulk thicknesses ($d_z = 8 - 32$). However, the reconstruction completely fails in isotropic transparent excitable media, see right panel in Fig. 8d) and also Fig. S1 [34].

D. Anisotropy

While ventricular muscle tissue is highly anisotropic (orthotropic muscle fiber organization), the Belousov-Zhabotinsky chemical reaction is isotropic. Both systems exhibit scroll waves, but the scroll wave morphology can be very different in anisotropic versus isotropic excitable media. In anisotropic media, scroll waves are elongated in fiber direction as they propagate faster along the fiber direction. This phenomenon can often be observed in optical mapping recordings. In the simulated anisotropic bulk, the waves are elongated differently at different depths, which is presumably why the reconstructions succeed in transparent excitable media as shown in Fig. 10c,f). By contrast, in isotropic excitable media the scroll waves are similarly shaped throughout the bulk, and therefore the network cannot distinguish scroll waves closer to the surface from scroll waves deeper in the bulk, see also discussion and Fig. S1 [34]. Anisotropy does not affect the reconstructions in opaque excitable media, as the reconstruction does not rely on depth information.

E. Analyzing Surface Deformation

If the network analyzes the mechanical deformation of the surface in addition to the excitable wave patterns visible on the same surface, the reconstruction improves slightly, see plot 'Top + Motion' in Fig. 11a). We tested this behavior with 'turbulent' scroll wave chaos, the U-Net architecture and the single-surface configuration shown in Fig. 2c), and found that the reconstruction improves slightly, but not substantially. This is a surprising finding, because deformation on the surface may also result from contractile activity further within the bulk. The reconstruction error does not rise as steeply with increasing depth as when analyzing electrics in single-surface mode alone. The reconstruction accuracy improves by roughly 25% at midwall (with a bulk thickness of $d_z = 24$ layers). We found that there was no significant difference between analyzing only two-dimensional in-plane displacements $\vec{u} = (u_x, u_y)$ with x - and y - components versus three-dimensional displacements with also a z -component, see also eq. (10) in section IIB. We cannot exclude that the latter finding is specific to our methodology and the mechanical boundary conditions that we used in the simulations.

F. Noise

Reconstructions can be performed with noise, see Fig. 13 and Supplemental Video 5 [34], if the network was previously trained with noise, similarly as described in [28] and [35]. The noise can be present in either the surface observations or projections. Fig. 13a,b) shows

reconstructions of scroll wave dynamics with noise in opaque and transparent anisotropic excitable media (thickness: $d_z = 24$ layers), respectively. In single-surface mode with the opaque media, the slope of the reconstruction error is slightly steeper with noise than without. In transparent media, the error profile remains flat and stays below 0.1 (RMSE) with noise. We tested this behavior with Gaussian noise and noise levels of up to $\sigma = 0.2$, shown in Fig. 13 (top right in each panel).

G. Network Types

We tested several deep neural network architectures (basic Encoder-Decoder, U-Net, TransUNet and MIRNet, see also section IIB) on the ‘turbulent’ scroll wave chaos prediction task and found that the prediction behavior is very similar across the different architectures, see Fig. 14. We observed that the MIRNet architecture produces the lowest reconstruction error, while Encoder-Decoder, U-Net, and TransUNet all have similar but slightly higher reconstruction errors. In opaque excitable media, the prediction error (RMSE: mean root squared error) rises linearly and steeply with increasing depth equally with all networks, see Fig. 14a) with single-surface predictions in anisotropic bulk. All networks produce the same characteristic error profile and the error saturates at depths $z > 15$ equally with all networks, while MIRNet provides a slightly lower maximal error than the other networks. In transparent excitable media with anisotropy, see Fig. 14b), all networks achieve high prediction accuracies better than 95% (note that a RMSE of 0.1 corresponds to a mean absolute error (MAE) of about 5%), while MIRNet provides the highest prediction accuracy (97% – 98%).

We primarily used U-Net in this study because TransUNet and MIRNet provided either no or incremental improvements in prediction accuracy, while requiring significantly longer to train. U-Net required about an hour to train, while being competitive with the accuracy of MIRNet, which required almost a full day to train. All results in Figs. 4–13 were obtained with the U-Net architecture, if not stated otherwise. We tested several U-Net sizes: a small model with 0.5M parameters, a medium network with 2M parameters and a large model with 8M parameters, determined that larger models perform significantly better, and subsequently used the largest model. In some circumstances U-Net and TransUNet exhibited a significantly better reconstruction performance than the Encoder-Decoder network, but we did not observe significant differences between U-Net and TransUNet.

H. Depth Estimation

It is possible to estimate the thickness or depth d_z of transparent bulks from projections of the corresponding scroll wave dynamics using either a regression or classification neural network. By contrast, it is not possible to reliably predict the thickness of opaque bulks using either approach. Fig. 15a) shows predictions obtained with a regression neural network in transparent media, which estimates the depth accurately with floating point precision (with a certain degree of uncertainty). Fig. 15c) shows a confusion matrix with depth predictions obtained with a classification neural network also in transparent media, which performs better than the regression. Out of ~ 500 predictions per thickness, only few attempts falsely classify the thickness (off-diagonal values). For both panels a) and c), predictions were made from two-dimensional observations as shown in Figs. 6c) and 15e).

In particular, Fig. 15e) shows how the contrast of the waves decreases with increasing bulk thickness as more and more waves are superimposed and the signal is averaged along the z -axis. The neural network presumably associates the bulk's thickness with the contrast.

Both the regression and classification neural networks are not able to predict the depth of opaque media correctly, see Figs. 15b,d). The classification neural network predicts random thicknesses and fails completely at correctly classifying the bulk's actual thickness. For both panels b) and d), predictions were made from two-dimensional observations as shown in Fig. 6a) or b). The data demonstrates that predicting the extend of scroll wave chaos is challenging with opacity, at least with our methodology. While scroll wave dynamics can vary qualitatively with different bulk thicknesses, in particular with thinner bulks, as shown in Figs. 6–8, there is a critical thickness beyond which the dynamics are dominated by the intrinsic excitable kinetics and are less influenced by the bulk's geometry and its boundaries, thus making depth predictions from surface observation challenging.

IV. DISCUSSION

We demonstrated that deep neural networks can be used to reconstruct three-dimensional scroll wave dynamics from two-dimensional observations of the dynamics on the surface of excitable media. Reconstructions succeed throughout opaque excitable media when the scroll wave size is not much smaller than the medium's thickness. Large scroll waves can be reconstructed even if they take on complex shapes and produce significant dissociation between the two opposing surfaces of the medium. However, multiple layers of scroll waves are challenging to reconstruct in opaque media, even when the dynamics are analyzed in dual-surface mode. Reconstructions can be performed particularly well in transparent anisotropic excitable media, in which it is possible to reconstruct also complicated scroll wave chaos far into the medium.

That encoding-decoding convolutional neural networks can reconstruct three-dimensional dynamics from two-dimensional observations is facilitated by their training on tens of thousands of similar examples. Further generalization can be achieved by diversifying the training dataset, e.g. by adding simulations with a broad range of parameters to the training data or by performing data augmentation, see also Fig. 14 in [28]. The amount of information that encoding-decoding convolutional neural networks can extract from the short sequence of snapshots to perform the 2D-to-3D prediction task is remarkable. However, it is also revealing that the dual-surface reconstruction does not provide any benefit or synergistic effects over the single-surface reconstruction. It is as if the network performs two separate reconstructions from either side. This highlights fundamental limitations of convolutional encoding-decoding neural networks in this particular application.

One interesting detail we found in transparent excitable media is that the reconstruction outcomes are very good with anisotropy, but poor with isotropy. Scroll wave chaos cannot be reconstructed at all in isotropic transparent excitable media, and reconstructions of simpler scroll wave dynamics exhibit artifacts. These findings show that anisotropy is crucial because it implicitly encodes depth. The neural network learns to associate the alignment

of the waves with the underlying fiber alignment which varies with depth. Moreover, it is able to decode this encoding even when multiple waves are superimposed in the projections, see Figs. 6c), 9b) and 15e) and Supplemental Video 3 [34]. Accordingly, the reconstructions exhibit artifacts or fail entirely in isotropic transparent excitable media as the network lacks depth information. Presumably, it would equally fail in anisotropic excitable media with uniform linearly transverse anisotropy. Unfortunately, this means that this feature cannot be exploited and is neither directly applicable to ventricular fibrillation, because the ventricular muscle is opaque, nor to the Belousov-Zhabotinsky reaction, which is transparent but isotropic.

Nevertheless, our methodology could still be used to reconstruct intramural action potential waves including scroll waves inside cardiac tissue: 1) Transillumination imaging [9–13], near-infrared optical mapping [47], or other optical techniques [21, 22], which allow imaging of action potential waves deeper inside cardiac tissue, could be used in the transilluminated right ventricle, in the atria or in small animal hearts such as mouse or zebrafish hearts. The muscle fiber architecture in the projections of the transilluminated translucent tissues could enable the depth encoding. 2) Dual-surface imaging with superficial electrode mapping or fluorescent dyes lacking the penetration depth (such as Di-4-ANEPPS) could be used to reconstruct ‘laminar’ episodes of ventricular tachycardia or fibrillation or atrial fibrillation. The wavelengths of single scroll waves or macro-reentries during ventricular arrhythmias are larger than the thickness of the right and left ventricular walls. The atria, which exhibit epi- and endocardial dissociation during atrial fibrillation [48–52], are presumably thin enough for dual-surface reconstructions to succeed. However, the training data would have to account for the complex anatomy of the atria [53] as well as the particular wave dynamics. Whether it will be possible to create ground-truth data or to train a neural network on simulated data and subsequently apply it to experimental data needs to be determined in future research. We found in previous work that the latter approach is in principle feasible, see Figs. 7 and 8 in [28]. Lastly, the deep learning-based reconstructions can be performed very efficiently within milliseconds on a graphics processing unit, and they do not require the collection of long time-series data.

We have recently used similar encoding-decoding convolutional neural networks for the prediction of electrical scroll wave chaos from three-dimensional mechanical deformation [35], as well as for the prediction of phase maps and phase singularities from two-dimensional electrical spiral wave dynamics [28]. While the networks performed very well in these applications, some of the results presented in this study, particularly the results for scroll wave chaos, are more sobering. Our study is another example of the more general notion that cardiac dynamics, and chaotic dynamics more generally, are challenging to predict [29, 54–58]. It is well known that classical deep learning approaches excel at interpolating, but do not perform well at extrapolating, which is what we aimed to do in this study. Therefore, the complete reconstruction of complicated fine-scaled scroll wave dynamics from surface observations in opaque excitable media will require more sophisticated techniques than encoding-decoding convolutional neural networks.

V. CONCLUSIONS

We demonstrated that it is possible to reconstruct three-dimensional scroll wave dynamics from two-dimensional observations using deep encoding-decoding convolutional neural networks. Reconstructions succeed under two conditions: either i) the medium is transparent and anisotropic with spatially varying anisotropy or ii) the medium is opaque and the dynamics are observed on two opposing surfaces while the scroll wavelength is not much shorter than the medium's thickness. In the future, our methodology could be used to reconstruct transmural action potential wave dynamics from epicardial or endocardial measurements.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

FUNDING

This research was funded by the University of California, San Francisco, the National Institutes of Health (DP2HL168071) and the Sandler Program for Breakthrough Biomedical Research, which is partially funded by the Sandler Foundation (to JC). The RTX A5000 GPUs used in this study were donated by the NVIDIA Corporation via the Academic Hardware Grant Program (to JL and JC).

VI. DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

- [1]. Fenton F. and Karma A, Vortex dynamics in three-dimensional continuous myocardium with fiber rotation: Filament instability and fibrillation, *Chaos: An Interdisciplinary Journal of Nonlinear Science* 8, 20 (1998).
- [2]. Clayton RH, Vortex filament dynamics in computational models of ventricular fibrillation in the heart, *Chaos: An Interdisciplinary Journal of Nonlinear Science* 18, 043127 (2008).
- [3]. Pathmanathan P. and Gray RA, Filament dynamics during simulated ventricular fibrillation in a high-resolution rabbit heart, *BioMed Research International* 2015, 10.1155/2015/720575 (2015).
- [4]. Davidenko JM, Pertsov AV, Salomonsz R, Baxter W, and Jalife J, Stationary and drifting spiral waves of excitation in isolated cardiac muscle, *Nature* 355, 349 (1992). [PubMed: 1731248]
- [5]. Pertsov AM, Davidenko R, Salomonsz JM, Baxter WT, and Jalife J, Spiral waves of excitation underlie reentrant activity in isolated cardiac muscle, *Circulation Research* 72, 631 (1993). [PubMed: 8431989]
- [6]. Winfree A, Electrical turbulence in three-dimensional heart muscle, *Science* 266, 1003 (1994). [PubMed: 7973648]
- [7]. Christoph J, Chebbok M, Richter C, Schröder-Schetelig J, Bittihn P, Stein S, Uzelac I, Fenton FH, Hasenfuss G, Gilmour RJ, and Luther S, Electromechanical vortex filaments during cardiac fibrillation, *Nature* 555, 667 (2018). [PubMed: 29466325]
- [8]. Uzelac I, Iravanian S, Bhatia NK, and Fenton FH, Spiral wave breakup: Optical mapping in an explanted human heart shows the transition from ventricular tachycardia to ventricular fibrillation and self-termination, *Heart Rhythm* (2022).
- [9]. Baxter WT, Mironov SF, Zaitsev AV, Jalife J, and Pertsov AM, Visualizing excitation waves inside cardiac muscle using transillumination, *Biophysical Journal* 80, 516 (2001). [PubMed: 11159422]

- [10]. Bernus O, Mukund KS, and Pertsov AM, Detection of intramyocardial scroll waves using absorptive transillumination imaging, *Journal of Biomedical Optics* 12, 014035 (2007).
- [11]. Mitrea BG, Wellner M, and Pertsov AM, Monitoring intramyocardial reentry using alternating transillumination, in *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (2009) pp. 4194–4197.
- [12]. Khait VD, Bernus O, Mironov SF, and Pertsov AM, Method for the three-dimensional localization of intramyocardial excitation centers using optical imaging, *Journal of Biomedical Optics* 11, 34007 (2006).
- [13]. Caldwell BJ, Trew ML, and Pertsov AM, Cardiac response to low-energy field pacing challenges the standard theory of defibrillation, *Circulation: Arrhythmia and Electrophysiology* 8, 685 (2015). [PubMed: 25772543]
- [14]. Molavi Tabrizi A, Mesgarnejad A, Bazzi M, Luther S, Christoph J, and Karma A, Spatiotemporal organization of electromechanical phase singularities during high-frequency cardiac arrhythmias, *Phys. Rev X* 12, 021052 (2022).
- [15]. Welsh BJ, Gomata J, and Burgess AE, Three-dimensional chemical waves in the Belousov–Zhabotinskii reaction, *Nature* 304, 611 (1983).
- [16]. Stock D. and Müller SC, Three-dimensional reconstruction of scroll waves in the belousov-zhabotinsky reaction using optical tomography, *Physica D: Nonlinear Phenomena* 96, 396 (1996).
- [17]. Bansagi T. and Steinbock O, Three-dimensional spiral waves in an excitable reaction system: Initiation and dynamics of scroll rings and scroll ring pairs, *Chaos: An Interdisciplinary Journal of Nonlinear Science* 18, 026102 (2008).
- [18]. Dähmlow P, Alonso S, Bär M, and Hauser MJB, Twists of opposite handedness on a scroll wave, *Phys. Rev. Lett* 110, 234102 (2013).
- [19]. Bruns C. and Hauser MJB, Dynamics of scroll waves in a cylinder jacket geometry, *Phys. Rev. E* 96, 012203 (2017).
- [20]. Cross AL, Armstrong RL, Gobrecht C, Paton M, and Ware C, Three dimensional imaging of the belousov-zhabotinsky reaction using magnetic resonance, *Magnetic Resonance Imaging* 15, 719 (1997). [PubMed: 9285812]
- [21]. Hillman EMC, Bernus O, Pease E, Bouchard MB, and Pertsov A, Depth-resolved optical imaging of transmural electrical propagation in perfused heart, *Opt. Express* 15, 17827 (2007). [PubMed: 18592044]
- [22]. Sacconi L, Silvestri L, Rodríguez EC, Armstrong GA, Pavone FS, Shrier A, and Bub G, KHz-rate volumetric voltage imaging of the whole zebrafish heart, *Biophysical Reports* 2, 100046 (2022).
- [23]. Berg S, Luther S, and Parlitz U, Synchronization based system identification of an extended excitable system, *Chaos: An Interdisciplinary Journal of Nonlinear Science* 21, 10.1063/1.3613921 (2011).
- [24]. Lebert J. and Christoph J, Synchronization-based reconstruction of electromechanical wave dynamics in elastic excitable media, *Chaos: An Interdisciplinary Journal of Nonlinear Science* 29, 10.1063/1.5101041 (2019).
- [25]. Hoffman MJ, LaVigne NS, Scorse ST, Fenton FH, and Cherry EM, Reconstructing three-dimensional reentrant cardiac electrical wave dynamics using data assimilation, *Chaos: An Interdisciplinary Journal of Nonlinear Science* 26, 013107 (2016).
- [26]. Hoffman MJ and Cherry EM, Sensitivity of a data-assimilation system for reconstructing three-dimensional cardiac electrical dynamics, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 378, 20190388 (2020).
- [27]. Hunt BR, Kostelich EJ, and Szunyogh I, Efficient data assimilation for spatiotemporal chaos: A local ensemble transform kalman filter, *Physica D: Nonlinear Phenomena* 230, 112 (2007).
- [28]. Lebert J, Ravi N, Fenton FH, and Christoph J, Rotor localization and phase mapping of cardiac excitation waves using deep neural networks, *Frontiers in Physiology* 12, 10.3389/fphys.2021.782176 (2021).
- [29]. Herzog S, Zimmermann RS, Abele J, Luther S, and Parlitz U, Reconstructing complex cardiac excitation waves from incomplete data using echo state networks and convolutional

- autoencoders, *Frontiers in Applied Mathematics and Statistics* 6, 10.3389/fams.2020.616584 (2021).
- [30]. Martin CH, Oved A, Chowdhury RA, Ullmann E, Peters NS, Bharath AA, and Varela M, EP-PINNs: Cardiac electrophysiology characterisation using physics-informed neural networks, *Frontiers in Cardiovascular Medicine* 8, 10.3389/fcvm.2021.768419 (2022).
- [31]. Aliev RR and Panfilov AV, A simple two-variable model of cardiac excitation, *Chaos, Solitons & Fractals* 7, 293 (1996).
- [32]. Bourguignon D. and Cani M, Controlling anisotropy in mass-spring systems, *Computer Animation and Simulation*, Springer, 113 (2000).
- [33]. Nash M. and Panfilov A, Electromechanical model of excitable tissue to study reentrant cardiac arrhythmias, *Progress in Biophysics and Molecular Biology* 85, 501 (2004). [PubMed: 15142759]
- [34]. See Supplemental Material at [URL will be inserted by publisher] for Supplemental Videos 1–4 and Supplemental Figures S1, S2. The Supplemental Material can also be found online at gitlab.com/janlebert/subsurface-supplementary-materials and at youtube.com/@cardiacvision under ‘Playlists’ → ‘Paper: Reconstruction of Three-Dimensional Scroll Waves in Excitable Media from Two-Dimensional Observations using Deep Neural Networks’.
- [35]. Christoph J. and Lebert J, Inverse mechano-electrical reconstruction of cardiac excitation wave patterns from mechanical deformation using deep learning, *Chaos: An Interdisciplinary Journal of Nonlinear Science* 30, 123134 (2020).
- [36]. Ronneberger O, Fischer P, and Brox T, U-net: Convolutional networks for biomedical image segmentation, in *Lecture Notes in Computer Science* (Springer International Publishing, 2015) pp. 234–241.
- [37]. Chen J, Lu Y, Yu Q, Luo X, Adeli E, Wang Y, Lu L, Yuille AL, and Zhou Y, Transunet: Transformers make strong encoders for medical image segmentation, *CoRR abs/2102.04306* (2021), 2102.04306.
- [38]. Zamir SW, Arora A, Khan S, Hayat M, Khan FS, Yang M-H, and Shao L, Learning Enriched Features for Real Image Restoration and Enhancement, in *Computer Vision – ECCV 2020*, *Lecture Notes in Computer Science*, edited by Vedaldi A, Bischof H, Brox T, and Frahm J-M (Springer International Publishing, 2020) pp. 492–511.
- [39]. Nair V. and Hinton GE, Rectified linear units improve restricted boltzmann machines (Omnipress, Madison, WI, USA, 2010) p. 807–814.
- [40]. Ioffe S. and Szegedy C, Batch normalization: Accelerating deep network training by reducing internal covariate shift, in *Proceedings of the 32nd International Conference on Machine Learning*, *Proceedings of Machine Learning Research*, Vol. 37, edited by Bach F. and Blei D. (PMLR, Lille, France, 2015) pp. 448–456.
- [41]. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser L, and Polosukhin I, Attention is all you need, in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, December 4–9, 2017, Long Beach, CA, USA, edited by Guyon I, von Luxburg U, Bengio S, Wallach HM, Fergus R, Vishwanathan SVN, and Garnett R. (2017) pp. 5998–6008.
- [42]. Bruhn A, Weickert J, and Schnörr C, Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods, *Int. J. Comput. Vis* 61, 211 (2005).
- [43]. Barron JT, A General and Adaptive Robust Loss Function, in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019) pp. 4326–4334.
- [44]. Kingma DP and Ba J, Adam: A method for stochastic optimization, in *3rd International Conference on Learning Representations, ICLR 2015*, San Diego, CA, USA, May 7–9, 2015, *Conference Track Proceedings*, edited by Bengio Y. and LeCun Y. (2015).
- [45]. Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A, Dean J, Devin M, Ghemawat S, Goodfellow I, Harp A, Irving G, Isard M, Jia Y, Jozefowicz R, Kaiser L, Kudlur M, Levenberg J, Mané D, Monga R, Moore S, Murray D, Olah C, Schuster M, Shlens J, Steiner B, Sutskever I, Talwar K, Tucker P, Vanhoucke V, Vasudevan V, Viégas F, Vinyals O, Warden P, Wattenberg M, Wicke M, Yu Y, and Zheng X, TensorFlow: Largescale machine learning on heterogeneous systems, <https://www.tensorflow.org> (2015).

- [46]. Chollet F. et al., Keras, <https://keras.io> (2015).
- [47]. Hansen BJ, Li N, Helfrich KM, Abudulwahed SH, Artiga EJ, Joseph ME, Mohler PJ, Hummel JD, and Fedorov VV, First in vivo use of high-resolution near-infrared optical mapping to assess atrial activation during sinus rhythm and atrial fibrillation in a large animal model, *Circulation: Arrhythmia and Electrophysiology* 11, e006870 (2018).
- [48]. Schuessler RB, Kawamoto T, Hand DE, Mitsuno M, Bromberg BI, Cox JL, and Boineau JP, Simultaneous epicardial and endocardial activation sequence mapping in the isolated canine right atrium., *Circulation* 88, 250 (1993). [PubMed: 8319340]
- [49]. Eckstein J, Maesen B, Linz D, Zeemering S, van Hunnik A, Verheule S, Allesie M, and Schotten U, Time course and mechanisms of endo-epicardial electrical dissociation during atrial fibrillation in the goat, *Cardiovascular Research* 89, 816 (2010). [PubMed: 20978006]
- [50]. Eckstein J, Zeemering S, Linz D, Maesen B, Verheule S, van Hunnik A, Crijns H, Allesie MA, and Schotten U, Transmural conduction is the predominant mechanism of breakthrough during atrial fibrillation, *Circulation: Arrhythmia and Electrophysiology* 6, 334 (2013). [PubMed: 23512204]
- [51]. de Groot N, van der Does L, Yaksh A, Lanter E, Teuwen C, Knops P, van de Woestijne P, Bekkers J, Kik C, Bogers A, and Allesie M, Direct proof of endo-epicardial asynchrony of the atrial wall during atrial fibrillation in humans, *Circulation: Arrhythmia and Electrophysiology* 9, e003648 (2016).
- [52]. Walters TE, Lee G, Lee A, Sievers R, Kalman JM, and Gerstenfeld EP, Site-specific epicardium-to-endocardium dissociation of electrical activation in a swine model of atrial fibrillation, *JACC: Clinical Electrophysiology* 6, 830 (2020). [PubMed: 32703566]
- [53]. Zhao J, Hansen BJ, Csepe TA, Lim P, Wang Y, Williams M, Mohler PJ, Janssen PM, Weiss R, Hummel JD, and Fedorov VV, Integration of high-resolution optical mapping and 3-dimensional microcomputed tomographic imaging to resolve the structural basis of atrial conduction in the human heart, *Circulation: Arrhythmia and Electrophysiology* 8, 1514 (2015). [PubMed: 26671938]
- [54]. Pathak J, Hunt B, Girvan M, Lu Z, and Ott E, Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach, *Phys. Rev. Lett* 120, 024102 (2018).
- [55]. Pathak J, Wikner A, Fussell R, Chandra S, Hunt BR, Girvan M, and Ott E, Hybrid forecasting of chaotic processes: Using machine learning in conjunction with a knowledge-based model, *Chaos: An Interdisciplinary Journal of Nonlinear Science* 28, 041101 (2018).
- [56]. Herzog S, Wörgötter F, and Parlitz U, Data-driven modeling and prediction of complex spatio-temporal dynamics in excitable media, *Frontiers in Applied Mathematics and Statistics* 4, 10.3389/fams.2018.00060 (2018).
- [57]. Shahi S, Marcotte CD, Herndon CJ, Fenton FH, Shiferaw Y, and Cherry EM, Long-time prediction of arrhythmic cardiac action potentials using recurrent neural networks and reservoir computing, *Frontiers in Physiology* 12, 10.3389/fphys.2021.734178 (2021).
- [58]. Shahi S, Fenton FH, and Cherry EM, A machine-learning approach for long-term prediction of experimental cardiac action potential time series using an autoencoder and echo state networks, *Chaos: An Interdisciplinary Journal of Nonlinear Science* 32, 063117 (2022).

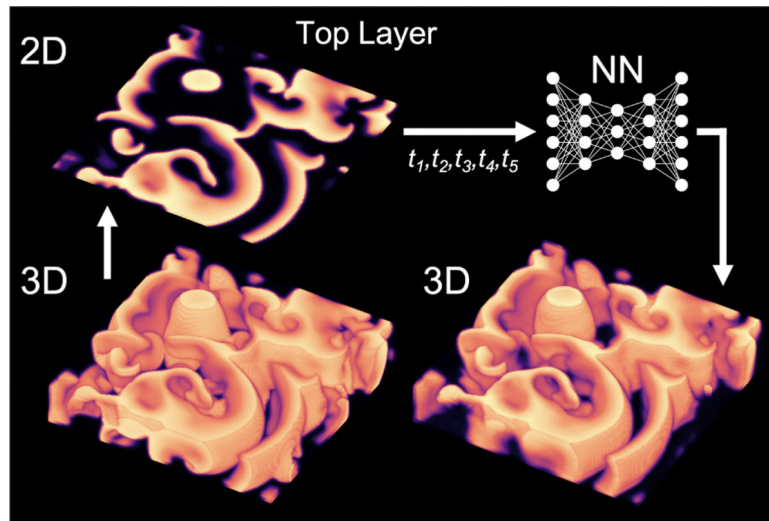
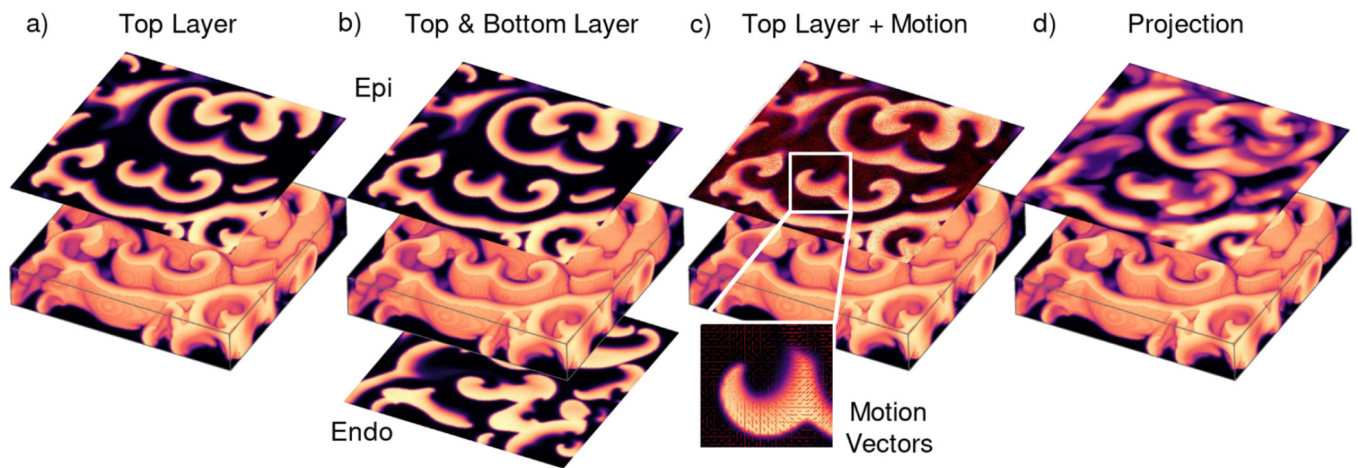


FIG. 1. Deep learning-based reconstruction of scroll wave chaos inside a three-dimensional volume from partial observations of the dynamics on its surface. Scroll wave chaos is a model for the electrophysiological dynamics underlying ventricular fibrillation. Computer simulations were performed in isotropic and anisotropic bulk-shaped excitable media. A neural network (NN) predicts scroll wave dynamics underneath the bulk's surface from a short temporal sequence (t_1, \dots, t_s) of two-dimensional observations (here shown for top layer).

**FIG. 2.**

Surface observations and projections of three-dimensional scroll wave chaos (‘turbulent’ parameter regime) in a bulk medium. a) Observation of the top surface of the bulk (layer 1) in single-surface mode. b) Observation of the top and bottom surfaces of the bulk (layers 1 and 24) in dual-surface mode. c) Observation of the electrical activity and mechanical motion on the top surface of the bulk (layer 1) in single-surface mode (red: motion vectors). In a-c), the medium is opaque and does not allow observations of the dynamics inside the medium below the top layer. d) Observation of the projection (transillumination) of the three-dimensional dynamics along the depth of the bulk. The projection is calculated by summing the values in all 24 layers along the z -axis for a specific (x, y) -coordinate and dividing the sum by the number of layers. All layers 1 – 24 are cross-sections in the $x - y$ plane.

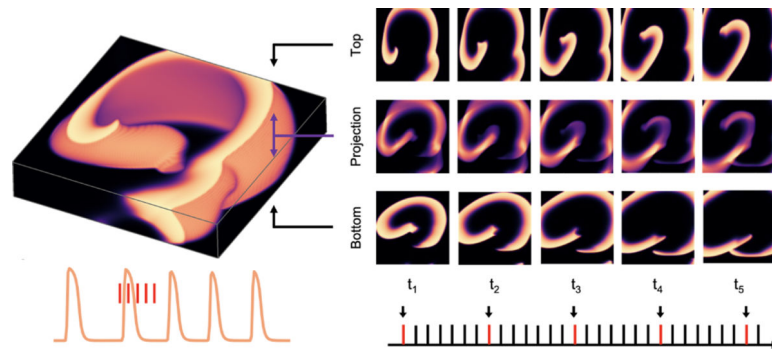


FIG. 3.

Observations of three-dimensional scroll wave ('laminar' parameter set) on the top and bottom surfaces of an opaque medium, and in the projection of the full dynamics in a transparent medium in the bulk's z -direction (depth). In each case, the neural network analyzes a short sequence of 5 snapshots, which are sampled at discrete times (red) over the period of the scroll wave from the simulation data, see section II B for details. In the simulations, one rotational period corresponds to about 500 simulation time steps.


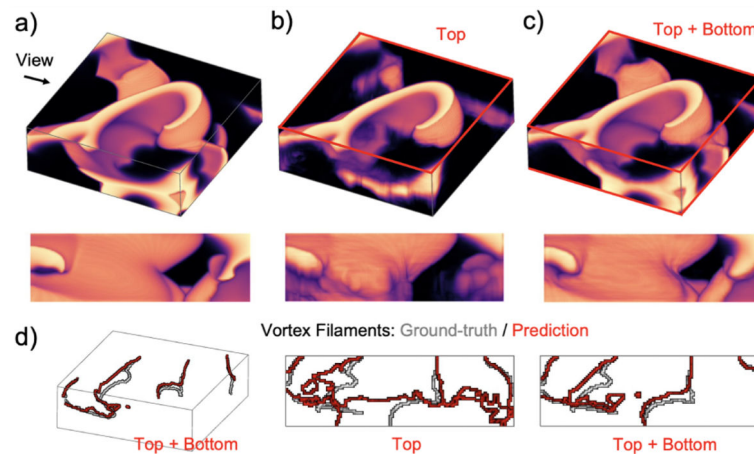


FIG. 4.

Predictions of three-dimensional ‘laminar’ scroll wave dynamics from two-dimensional observations using deep convolutional encoding-decoding neural network (U-Net) in an anisotropic excitable medium ($128 \times 128 \times 24$ voxels). a) Ground-truth scroll wave dynamics (5 random representative snapshots). The simulations exhibit scroll waves with meandering and curved vortex cores, see also Fig. 5d), wavebreak, and dissociation between the top and bottom surface dynamics, see Fig. 6. b) Predictions from two-dimensional wave pattern visible only on the top surface of the bulk when the medium is completely opaque. The reconstruction accuracy decreases slightly with increasing depth (wave pattern becomes fuzzy towards the bottom). c) Predictions from two-dimensional projection of the whole three-dimensional dynamics along the z -axis in a transparent medium. The prediction accuracy is slightly better than in b), particularly towards the bottom layers of the bulk. Overall, the predictions and the ground-truth are visually difficult to distinguish from each other. The data was not seen by the network during training. See Supplemental Videos 1–4 for corresponding scroll wave dynamics [34].

**FIG. 5.**

Predictions of three-dimensional scroll wave dynamics (‘laminar’ parameter set) and their vortex filaments from either single- or dual-surface observations in a thick, opaque, anisotropic excitable medium ($128 \times 128 \times 40$ voxels), see also corresponding Supplemental Video 2 [34]. a) Ground-truth scroll wave dynamics. b) Prediction from the top surface only. c) Prediction from both the top and bottom surfaces. In dual-surface mode, the network is able to recover the dynamics sufficiently well. Arrow indicates direction of cross-sectional view. d) Ground-truth vortex filaments (gray) and vortex filaments calculated from predicted scroll wave dynamics (red). In single-surface mode, predictions become unreliable (wave pattern becomes fuzzy / vortex filaments do not match) towards the bottom of the bulk. Reconstructions performed with U-Net.




FIG. 6. Different electrical wave patterns as seen in the a) top layer, b) bottom layer and c) projection of all layers of a three-dimensional bulk. Bulk sizes are $128 \times 128 \times 24$ voxels (left and right) and $128 \times 128 \times 40$ voxels (center), respectively. With increasing bulk thickness or smaller scroll waves, the top and bottom layers are dissociated because the dynamics become increasingly three-dimensional. The projection is calculated for a given (x, y) -coordinate by averaging the u -values ($u \in [0, 1]$) along the depth (z -direction) of the bulk. Data from left to right shown in Figs. 4, 5 and 9, respectively. The left and center snapshots are from the ‘laminar’ parameter regime, while the ones on the right are from the ‘turbulent’ chaotic parameter regime.

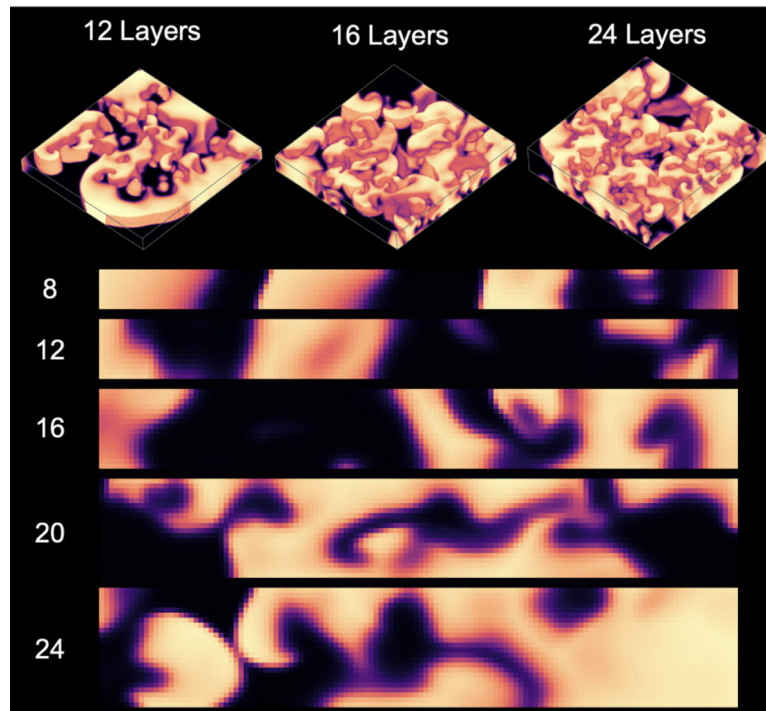


FIG. 7. Bulk thickness and transmurality of scroll wave dynamics. ‘Turbulent’ scroll wave chaos with different bulk thicknesses $d_z = \{8, 12, 16, 20, 24\}$, also shown in corresponding cross-sections. The dynamics are quasi two-dimensional with $d_z = 8$. Dissociation between top and bottom layers starts to emerge at $d_z = 12$ as the dynamics become increasingly three-dimensional. At $d_z > 12$ the dynamics are fully three-dimensional.

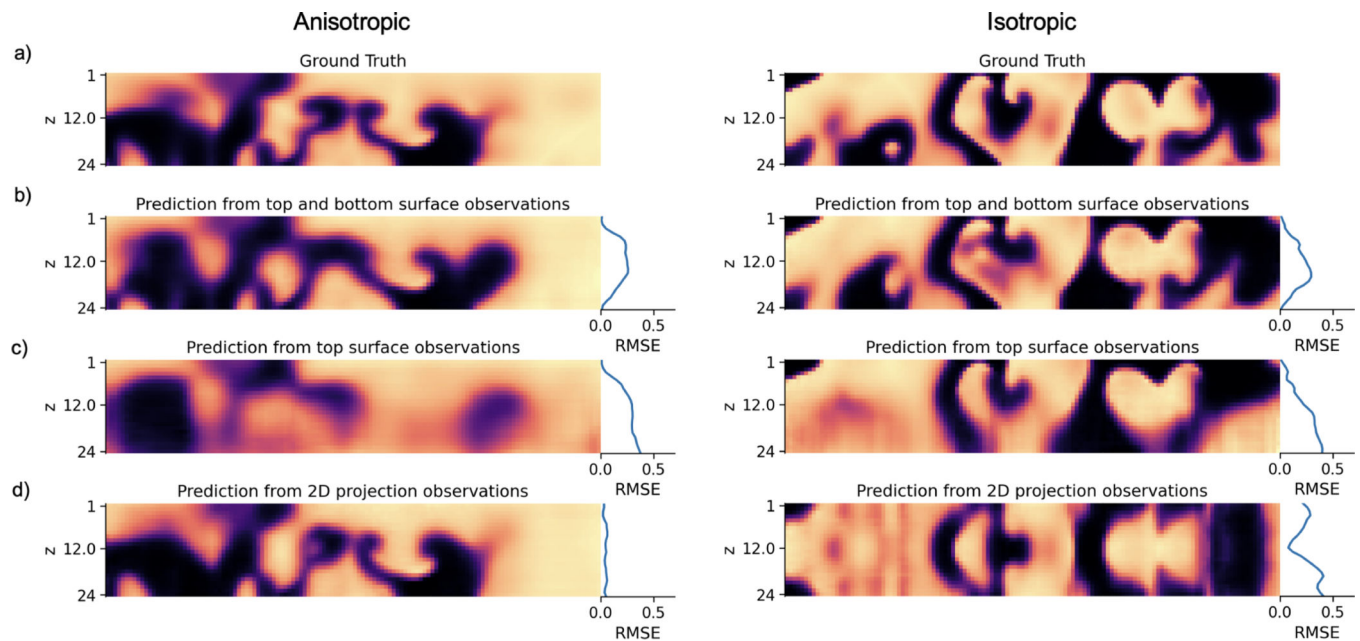
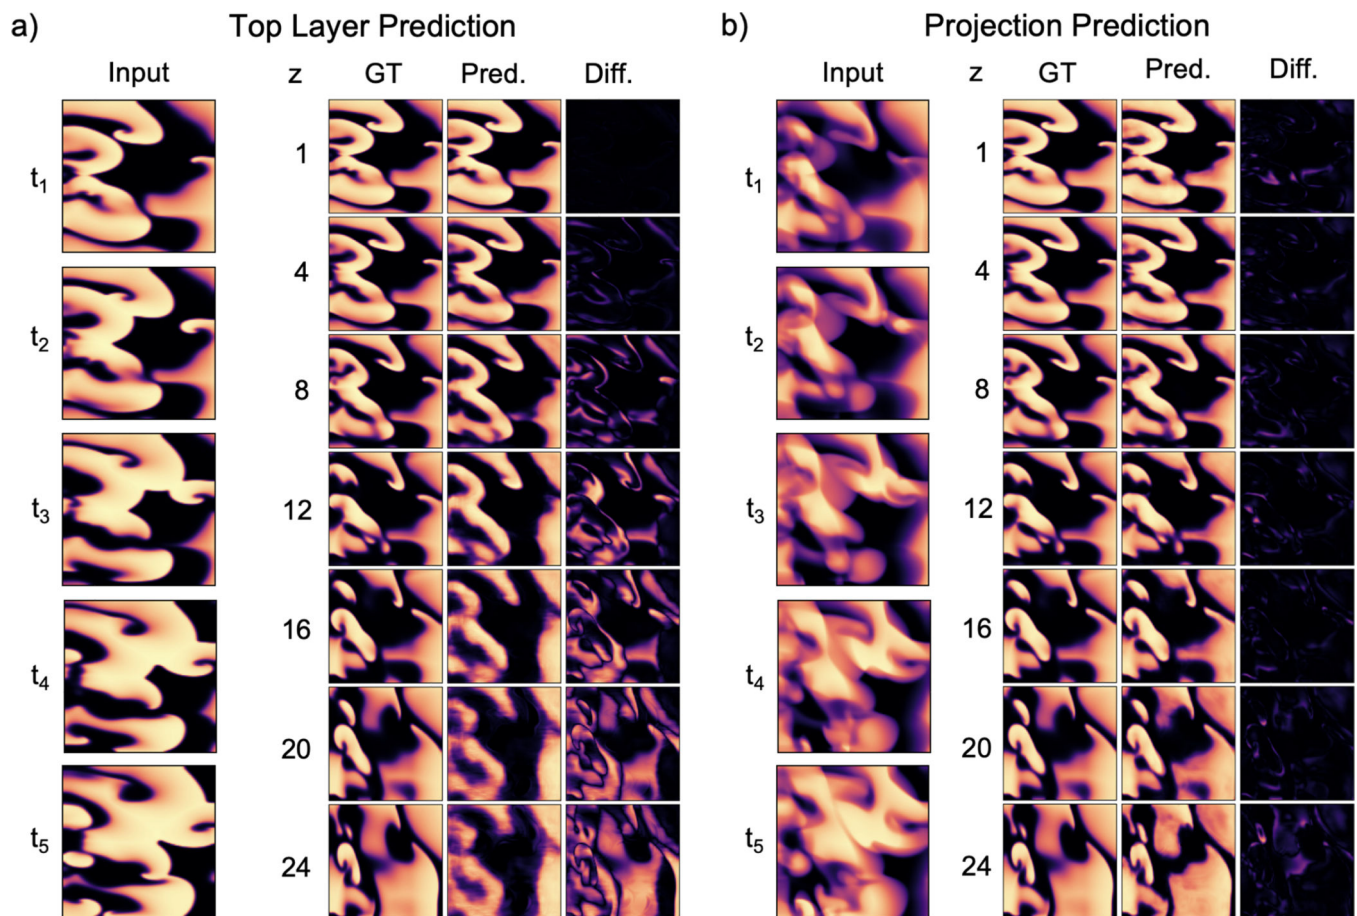
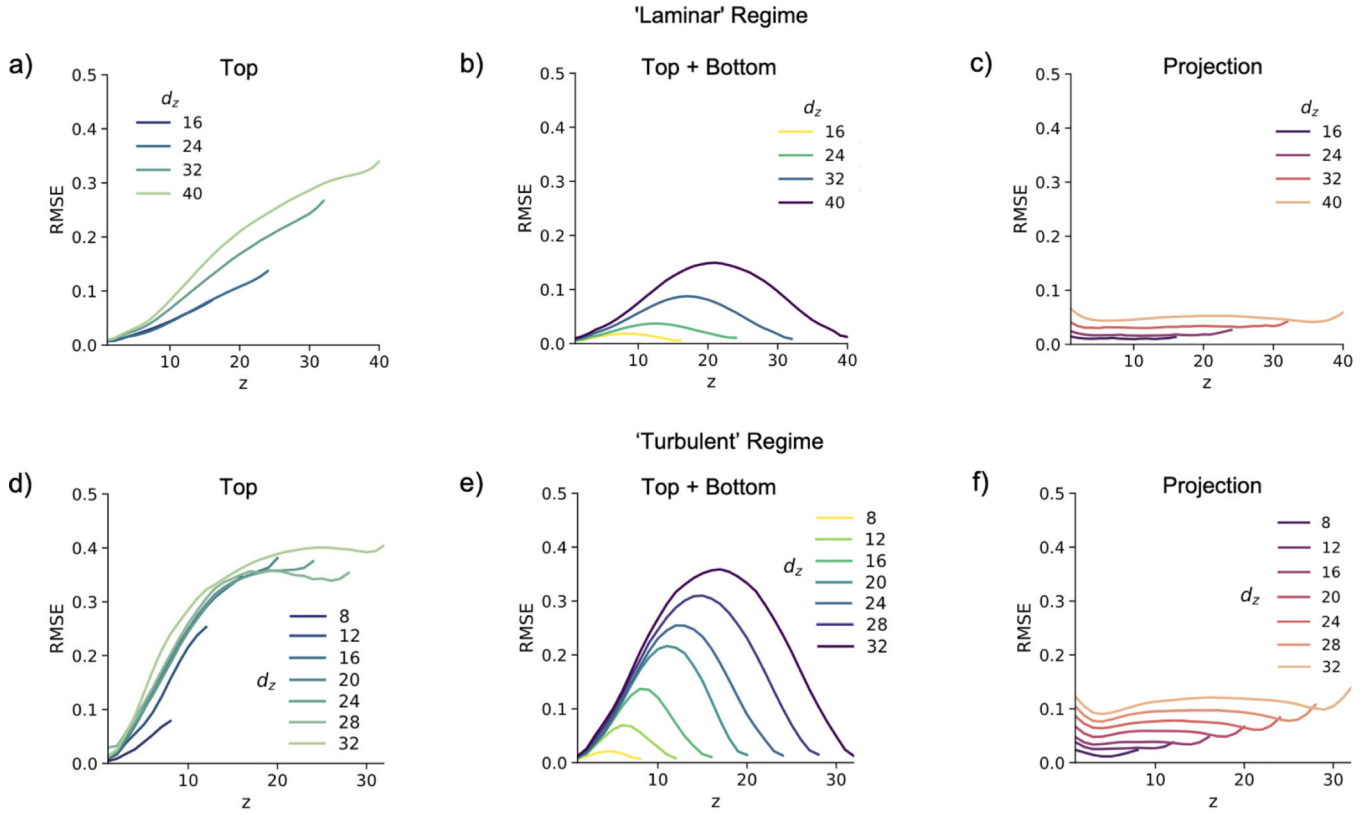


FIG. 8.

Predictions of ('turbulent') electrical scroll waves in subsurface layers of anisotropic (left) and isotropic (right) bulk tissue with dimension $128 \times 128 \times 24$ voxels. a) Ground-truth scroll wave dynamics (representative snapshots) b) Prediction in dual-surface mode analyzing the top and bottom layers of an opaque bulk tissue. c) Prediction in single-surface mode analyzing the top surface layer of an opaque bulk. d) Prediction analyzing the z -projection of the dynamics along its depth (or z -axis) in a transparent bulk. The depth-profile of the prediction error (RMSE: root mean squared error) along the z -axis is shown to the right of each prediction. The reconstruction is successful in anisotropic transparent media, but fails in isotropic transparent media, as shown in d). In opaque media, the reconstruction performs sufficiently well in dual-surface mode with larger errors emerging at midwall of the bulk, as shown in b). In transparent isotropic media, the subsurface prediction fails because the network is unable to infer the depth of the layers. Cross-sections intersect the bulk at its center. All reconstruction with U-Net.

**FIG. 9.**

Predictions of electrical scroll waves within subsurface layers of a bulk-shaped ‘turbulent’ anisotropic excitable medium from observing either a) the top layer of the bulk or b) the projection of the three-dimensional wave pattern in a transparent bulk along its depth (or z -axis). The bulk’s dimensions are $128 \times 128 \times 24$ voxels and predictions were performed with the U-Net architecture. Predictions are shown for the 24 layers along the z -axis of the bulk, where the first layer is the top layer and the 24th layer is the bottom layer. First column: The five two-dimensional frames (t_1, \dots, t_5) which are the input for the neural network prediction. Second column: Ground truth (GT) electrical excitation wave pattern within cross-sectional layers (1–24), of which layers 2–24 cannot be observed. Third column: Prediction of the current cross-sectional layer (1–24) by the neural network. Fourth column: Absolute difference per voxel between prediction and ground-truth.

**FIG. 10.**

Average reconstruction error over bulk depth (RMSE: root mean squared error along z -axis) in a,b,d,e) opaque or c,f) transparent excitable media with anisotropy (with varying bulk depths of $d_z \in \{8, 12, \dots, 40\}$). All reconstructions were performed with U-Net. a-c) ‘Laminar’ scroll wave dynamics as shown in Figs. 4 and 5: a) Single-surface mode, see also Figs. 4b) and 5b). b) Dual-surface mode. c) Projection, see also Figs. 4c) and 5c). ‘Turbulent’ scroll wave chaos as shown in Fig. 7: d) Single-surface mode, see also Figs. 8c) and 9a). e) Dual-surface mode, see also Figs. 2b) and 8b). f) Projection, see also Figs. 8d) and 9b). An error of 0.1 (RMSE) corresponds to a mean absolute error (MAE) of about 5%. In opaque excitable media, the reconstruction error increases approximately linearly with depth. In transparent media with anisotropy the error remains flat and below 0.1. We trained separate U-Net neural networks for each combination. See also Fig. S1 for a comparison with isotropic excitable media and Fig. S2 for a comparison of the deep learning-based reconstruction with a naive reconstruction in which the top layer is simply repeated in each following layer.

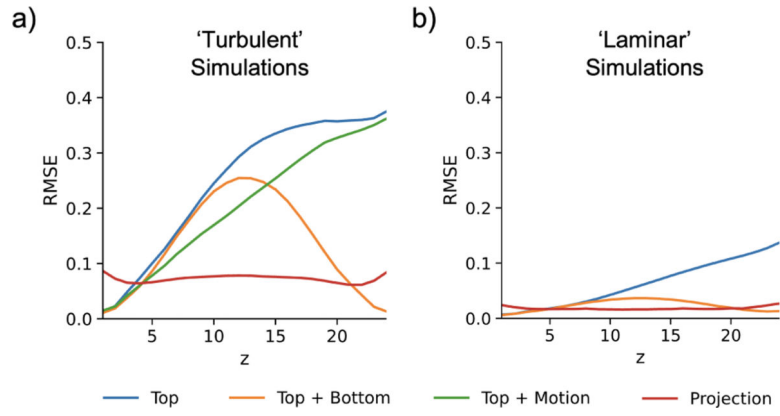


FIG. 11. Comparison of reconstruction errors obtained with different imaging configurations in opaque anisotropic excitable medium (with thickness $d_z = 24$ layers) with a) ‘turbulent’ scroll wave chaos and b) ‘laminar’ scroll wave dynamics. In the different configurations, the network (U-Net) analyzes i) in single-surface mode the electrics on the top layer (blue), ii) in single-surface mode the electrics and motion on the top layer (green), iii) in dual-surface mode the electrics on both top and bottom layers (orange), and iv) the z -projection of the three-dimensional electrics (red). All reconstruction errors were calculated per depth as root mean squared error (RMSE).




FIG. 12. Average reconstruction error per layer (depth) over time in an anisotropic excitable medium with thickness $d_z = 24$ shown for 5 different depths. Left: In opaque media and single-surface mode, the prediction error increases and fluctuates more with increasing depth. Right: In transparent media, the prediction error stays small throughout the different depths. Both plots derived for the ‘laminar’ scroll wave dynamics as shown in Fig. 4.

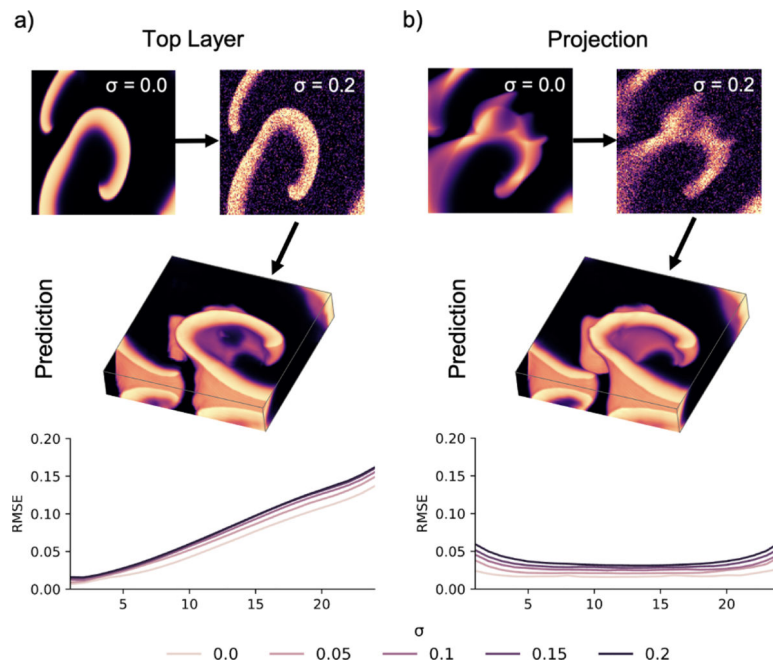


FIG. 13. Noise does not pose a limitation for the deep learning-based reconstructions (performed with U-Net on ‘laminar’ scroll wave dynamics). Reconstructions succeed in the presence of various noise levels ($\sigma = 0.05, \dots, 0.2$) in both a) opaque and b) transparent anisotropic excitable media. Top: Gaussian noise with standard deviation σ was added onto the input images. Bottom: Reconstruction error profiles rise slightly with higher noise (light pink curve: $\sigma = 0.0$, black curve: $\sigma = 0.2$).

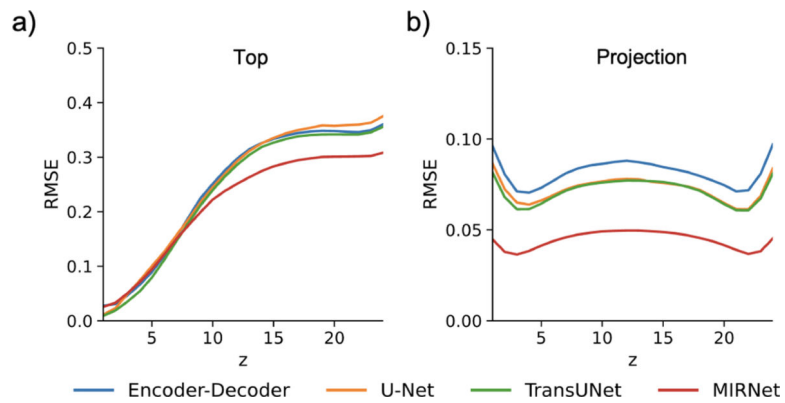


FIG. 14. Reconstruction errors obtained with different neural network architectures for ‘turbulent’ scroll wave chaos in a bulk with thickness of $d_z = 24$. a) Steep increase of reconstruction error with all networks (Encoder-Decoder, U-Net, TransUNet, MIRNet) in opaque excitable media from a single surface (top). b) Low and relatively flat reconstruction errors (below 10%) with all networks in transparent anisotropic excitable media (projection). MIRNet performs slightly better than the other networks. All reconstruction errors stated as root mean squared error (RMSE).

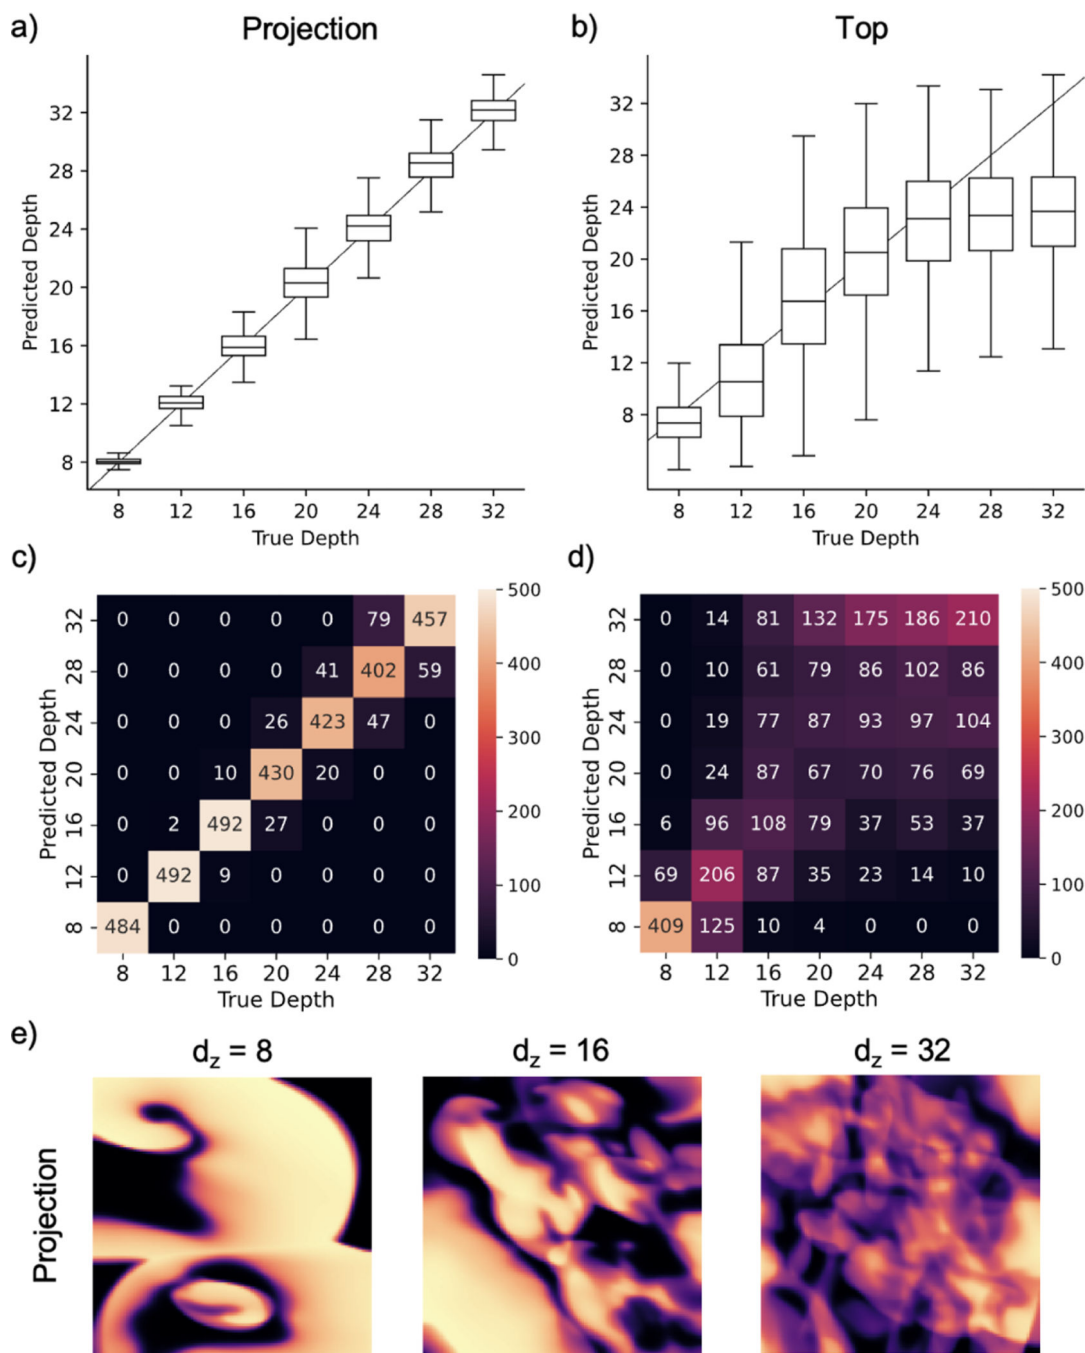


FIG. 15. Prediction of bulk thickness from surface observations of scroll wave chaos in a,c) transparent bulk medium with projection observations and b,d) opaque bulk medium with top surface observations. The bulk thickness was predicted using either a,b) a regression (black line shows ideal prediction $d_{\text{prediction}} = d_{\text{true}}$) or c,d) a classification neural network, respectively (all in anisotropic media). In transparent media, the thickness can be predicted from observations as shown in Fig. 6c), whereas in opaque media neither the regression nor classification neural networks predict the thickness correctly. e) Exemplary projection

images for bulk thicknesses or depths $d_z = 8$, $d_z = 16$, and $d_z = 32$. Due to the averaging, the contrast of the waves decreases with increasing depths.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE I.

Electrical (top) and mechanical (bottom) parameters used to simulate two different regimes of scroll wave dynamics: a ‘laminar’ regime, see Figs. 4, 5, and a fully ‘turbulent’ scroll wave chaos regime, see Figs. 6(right), 7 and 9. Electromechanical simulations were only performed with the ‘turbulent’ parameter set.

Parameter	‘Laminar’ Set	‘Turbulent’ Set
k	8	8
a	0.05	0.05
ϵ_0	0.002	0.002
μ_1	0.8	0.2
μ_2	0.3	0.3
k_T	—	3
k_{ij}	—	5
k_j	—	0.5
k_f	—	4
c_f	—	10

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE II.

Different neural network architectures used in this study and their respective number of trainable parameters and training times for 20 epochs. Training was performed on a single NVIDIA RTX A5000 GPU.

Model	Parameters	Training Time
Encoder-Decoder	6,829,309	44 min
U-Net	8,278,168	55 min
TransUNet	406,899,608	19 hours
MIRNet	145,358,026	17 hours
d_z Regression	426,593	5 min
d_z Classification	432,743	5 min

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript