Eurographics Workshop on Visual Computing for Biology and Medicine (2023)
T. Höllt and D. Jönsson (Editors)

Short Paper

Smoke Surfaces of 4D Biological Dynamical Systems

Marwin Schindler[1], Aleksandr Amirkhanov[1,2], and Renata Georgia Raidou[1]

[1]TU Wien, Austria [2]ZAnGeSa GmbH, Austria

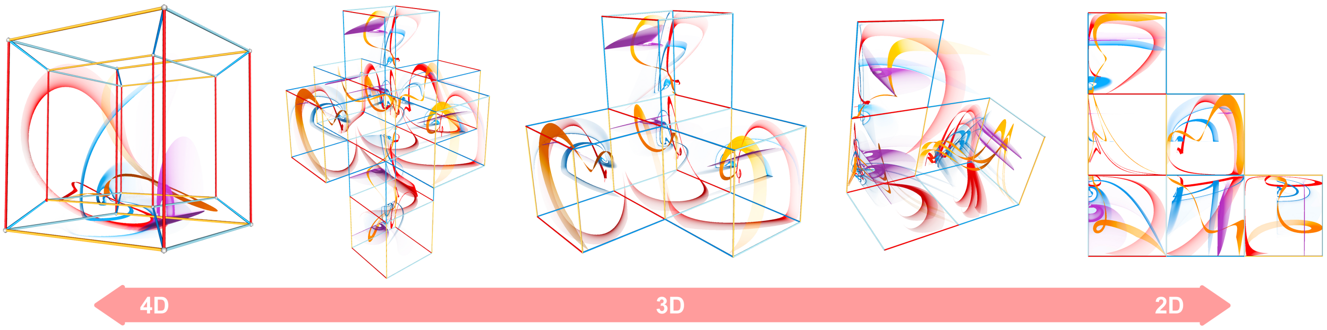

Figure 1: *Smoke surfaces of the Goldbeter model describing bipolar disorder. The tesseract on the left-hand side contains four smoke surfaces indicated by four distinct colors. Each smoke surface corresponds to one seeding line that denotes multiple initial conditions. The tesseract is consecutively unfolded to its lower-dimensional subspaces—drilling down to 3D and finally unfolding to 2D.*

Abstract

To study biological phenomena, mathematical biologists often employ modeling with ordinary differential equations. A system of ordinary differential equations that describes the state of a phenomenon as a moving point in space across time is known as a dynamical system*. This moving point emerges from the initial condition of the system and is referred to as a* trajectory *that "lives" in* phase space*, i.e., a space that defines all possible states of the system. In our previous work, we proposed* ManyLands [AKS*19]*—an approach to explore and analyze typical trajectories of 4D dynamical systems, using smooth, animated transitions to navigate through phase space. However, in* ManyLands *the comparison of multiple trajectories emerging from different initial conditions does not scale well, due to overdrawing that clutters the view. We extend* ManyLands *to support the comparative visualization of multiple trajectories of a 4D dynamical system, making use of smoke surfaces. In this way, the sensitivity of the dynamical system to its initialization can be investigated. The 4D smoke surfaces can be further projected onto lower-dimensional subspaces (3D and 2D) with seamless animated transitions. We showcase the capabilities of our approach using two 4D dynamical systems from biology* [Gol11, KJS06] *and a 4D dynamical system exhibiting chaotic behavior* [Bou15].

CCS Concepts

• ***Human-centered computing*** → ***Scientific visualization;*** • ***Applied computing*** → ***Life and medical sciences;***

1. Introduction

Mathematical biologists have established models to study various biological phenomena, such as the immune response to infection [KJS06] or mechanisms behind bipolar disorder [Gol11]. These models are defined by ordinary differential equations, i.e., dynamical systems that describe how the state of the biological phenomenon evolves from its initial condition across time—forming a *trajectory*. The dimensionality of the dynamical system depends on how many variables are necessary to describe the underlying phenomenon's state. The space that defines all possible states of the dynamical system is referred to as *phase space*.

In this work, we focus on 4D dynamical systems, as they are widely used to describe a multitude of biological phenomena. Mathematical biologists study the behavior of these systems to understand the underlying phenomena, but—being 4D—challenges for the visualization and analysis thereof arise. After a model is established, domain experts study it by investigating various characteristics of the system's behavior such as stability of equilibria, periodic oscillations, and switching phenomena [KS16]. To identify these characteristics, domain experts rely on 2D or 3D *phase portraits* of typical trajectories, which are biologically meaningful solutions of the system. Yet, an increasing number of variables in

DOI: 10.2312/vcbm.20231217

these dynamical systems leads to a multitude of 2D or 3D phase portraits, which makes the intuitive and holistic visualization of such systems challenging. Obtaining even qualitative insights on the sensitivity of a dynamical system to its initial conditions and comparing how changes in the initialization affect the resulting distribution of alternative trajectories is currently a tedious task.

In our previous work, we proposed *ManyLands* [AKS*19] to support domain experts in exploring 4D models of biological systems. *ManyLands* provides smooth, animated transitions for 4D trajectories across projections to lower-dimensional subspaces, i.e., from 4D down to 2D. However, *ManyLands* does not support the explicit comparison of trajectories emerging from different initial conditions. For two or more solutions shown simultaneously, the view in *ManyLands* is quickly affected by visual clutter that does not allow to easily compare multiple trajectories or to analyze the sensitivity of the model's behavior to initial conditions (Figure 2).

Our contribution is an extension of *ManyLands* through the integration of smoke surfaces to support the effective comparison of multiple trajectories of a 4D dynamical system. This is achieved by combining trajectories to form a transparent smoke surface. The transparent shading of the surface is optimized to decrease visual clutter and further enables the simultaneous rendering of several smoke surfaces. In this way, our method brings forward knowledge of the qualitative, global behavior of the system by providing functionality to compare distinct subspaces of initial conditions. Hence, our approach supports mathematical biologists to *comparatively represent, explore, and study multiple trajectories*, while also *investigating the sensitivity* of a dynamical system to its initialization.

2. Related Work

The exploration of dynamical systems has been extensively researched [Lor63, AS83, GWM*96, WLG97], focusing on the 3D representation of dynamical systems with additional encodings, such as color, to support more than three variables. Going beyond 3D, parallel coordinates have been employed in early work by Wegenkittl et al. [WLG97] and Grottel [GHWG14]. Only recently, Bartolovic et al. [BGG20] proposed a novel dimensionality reduction technique for high-dimensional dynamical systems that better preserves properties in the high-dimensional space.

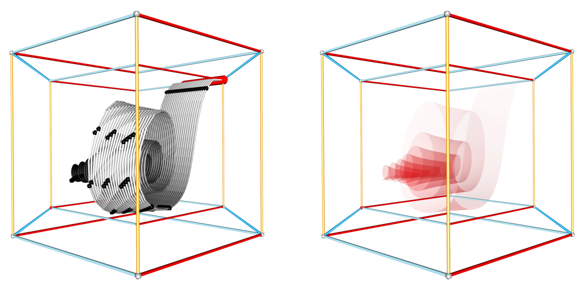

Figure 2: *(Left) Visual clutter in* ManyLands *from the visualization of twenty trajectories stemming from different initial conditions. Following the trajectories is not possible. (Right) Our proposed solution with a smoke surface encompassing the twenty trajectories.*

Our previous work on *ManyLands* enabled the exploration of 4D dynamical systems using interaction and smooth, animated navigation through phase space [AKS*19]. In *ManyLands*, we employ the metaphor of traveling across 4D *HyperLand* down to 3D *SpaceLand*, and finally 2D *FlatLand*. Additionally, we employ *TimeLines* to represent each variable of the dynamical system across time, and other visual representations to facilitate the analysis of segments of the trajectories. In the present work, we do not focus on the dimensionality of the dynamical system—instead on the missing ability to compare trajectories of multiple initial conditions in *ManyLands*.

We take inspiration from the work of Tian et al. [TG20], Hummel et al. [HGH*10], and von Funck et al. [vFWTS08] for the illustrative rendering of surfaces that serve in the comparison of initial conditions and their resulting trajectories in 4D dynamical systems. Our work also relates to previous studies of Günther et al. on how to compute optimal opacity values for streamlines and stream surfaces [GSE*14, GTG17]. By integrating the concept of smoke surfaces into *ManyLands*, we aim to provide increased understanding of the sensitivity of the 4D biological dynamical systems to their initializations, while at the same time supporting an easy comparison of trajectories stemming from different initial conditions. In this way, we provide comprehensive insights into the global behavior of the investigated 4D dynamical system.

3. Smoke Surfaces

In *ManyLands*, the analysis of multiple initial conditions in a 4D dynamical system requires visualizing one individual trajectory for each condition. Since the original implementation of a trajectory consists of cylindrical segments, the visualization of multiple trajectories quickly becomes cluttered and the behavior of the system is harder to understand (Figure 2). Smoke surfaces have been commonly used for the visualization of streak surfaces in time-dependent flows [vFWTS08]. Their advantage is that they can handle well non-regular triangles of the mesh, which is often an issue with rendering opaque surfaces. This leads to interactive rendering times and a representation that resembles the appearance of a smoke surface. Taking inspiration from this work, we employ smoke surfaces for the comparison of multiple trajectories emerging from different initial conditions of the 4D dynamical systems at hand. When using smoke surfaces, each time the user interacts with the application, the surface resulting from a set of initial conditions does not need to be recomputed to ensure a high-quality visual representation without visible large triangles of the mesh. This also enables smooth transitions across subspaces in real-time.

To generate the smoke surfaces we start by defining a 4D *seeding line* and a number of homogeneously distributed samples along this line to compute multiple initial conditions. For each initial condition, we solve the given system of differential equations using a fixed step size and integration duration to obtain a set of trajectories. For numerical integration, we employ a fourth-order Runge-Kutta method [Atk91]. This results in a trajectory with a fixed amount of nodes for each initial condition. Multiple seeding lines are employed to generate distinct sets of trajectories of interest—each rendered with a different color. We create a mesh for each set of trajectories by forming triangles between adjacent nodes.

The behavior of each trajectory over the integration time can

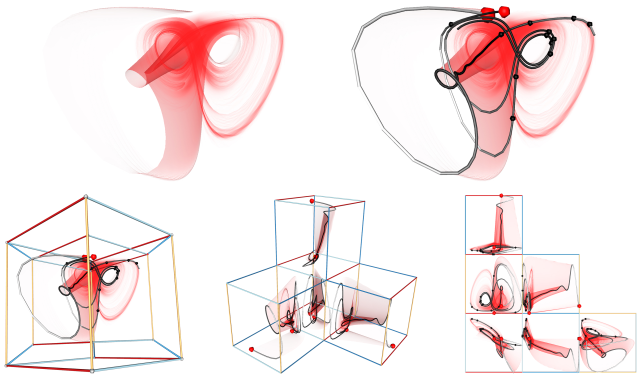

Figure 3: *The upper row depicts a smoke surface for the Bouali strange attractor without and with additional boundary trajectories of the seeding (black) line. The lower row shows the same smoke surface in the tesseract and its unfoldings to 3D/2D.*

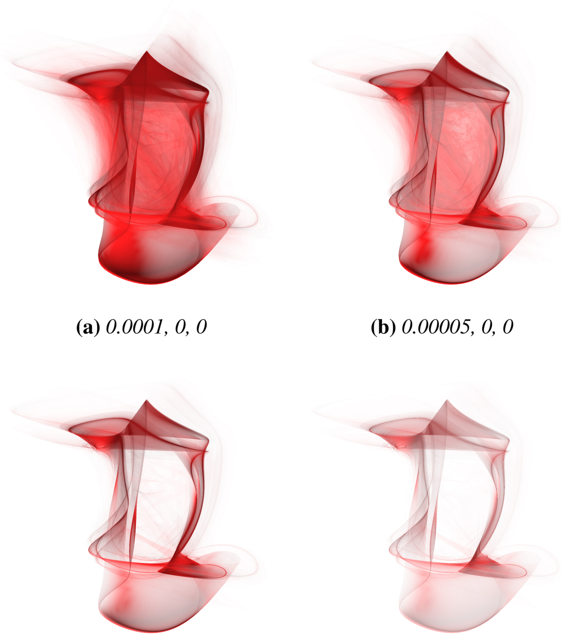

(a) *0.0001, 0, 0* **(b)** *0.00005, 0, 0* **(c)** *0.00005, 0, 5* **(d)** *0.00005, 0.1, 5*

Figure 4: *A smoke surface derived from the Bouali attractor for different parameters, which follow this order in the sub-captions: surface height, shape exponent, and curvature factor. In this case, the mesh is very dense, leading to a preference for low values for the surface height.*

strongly vary in speed and direction, even if the initial seeding positions are in close proximity. This happens when the system is *sensitive to its initialization*. The problem is even more apparent outside of the biological domain—when visualizing strange attractors. This results in smoke surfaces being heavily convoluted and suffering from self-occlusion, which deteriorates when the scene contains multiple smoke surfaces. To address this, we visualize each surface with varying transparency following the approach presented by von Funck et al. [vFWTS08]. Similarly to their approach, we consider each triangle as a prism with an additional user-given height and calculate its corresponding opacity using the formula:

$$\alpha = \alpha_{density}\alpha_{shape}\alpha_{curvature} \tag{1}$$

The terms α_{shape} and $\alpha_{curvature}$ encode information about the local speed and directional behavior of the trajectories (in relation to their neighboring trajectories) in the transparency of each vertex. This indicates that the speed of trajectories will influence the aspect ratio of the resulting triangles. Faster behavior will lead to a smaller α_{shape} value and less opacity for the triangle. As a region with large mean curvature in the surface cannot be represented adequately by a non-adaptive mesh, these areas are rendered less opaque using $\alpha_{curvature}$. We note that in the original work of von Funck et al. the α_{fade} term was used to decrease the opacity of older vertices since they release the vertices successively into the flow. We omit this term in Eq. (1) since our dynamical systems are steady and the smoke surface is rendered at a single pass.

For demonstration purposes, we employ the Bouali strange attractor (despite not being of biological interest) to showcase our smoke surfaces approach in Figure 3. In the upper row, we demonstrate our smoke surfaces without and with additional boundary trajectories of the seeding line (shown in black) for the 4D dynamical system. In the lower row, one can see how the smoke surface unfolds from the tesseract to its 3D and 2D subspaces. This example also demonstrates the use of varying transparency, which is able to bring forward the complex behavior of the attractor.

Vertex normals and the curvature at each vertex are computed on the GPU. In the interface, the user can adjust the visual properties with three parameters [vFWTS08] to achieve better visuals for the smoke surfaces: the *surface height* which affects the density of the smoke, the *shape exponent* which changes the influence of the triangle shape on the final transparency, and the *curvature influence* which changes the impact of the curvature at each vertex on the final transparency. In Figure 4, we show an example of the effect of the parameters on a smoke surface. To increase the visibility in dense areas of the smoke surface, one must decrease the *surface height*. Notice the effect of the *surface height* between the smoke surfaces in **(a)** and **(b)**: decreasing this parameter decreases the opacity of the mesh, depending on the area of the triangles and the viewing direction. To increase the visibility in areas of slow behavior, one must increase the *shape exponent*, as demonstrated between **(c)** and **(d)**. Finally, increasing the *curvature factor* decreases visibility in high-curvature areas, as shown comparatively between **(b)** and **(c)**. There is no parameter configuration suitable for all cases; we, therefore, recommend adjusting the parameterization depending on the phenomena or behaviors to be emphasized.

For multiple seeding lines, we obtain an equal number of smoke surfaces which all need to be rendered within and across the phase space. When we compute the final outcome and blend all computed surfaces, we do not require an explicit ordering of all transparent vertices. Instead of using *depth peeling* as proposed by von Funck et al., we employ another order-independent transparency technique called *weighted blended transparency* [MB13]. This technique has the benefit of better performance compared to *depth peeling* since results are achieved in a single geometry pass.

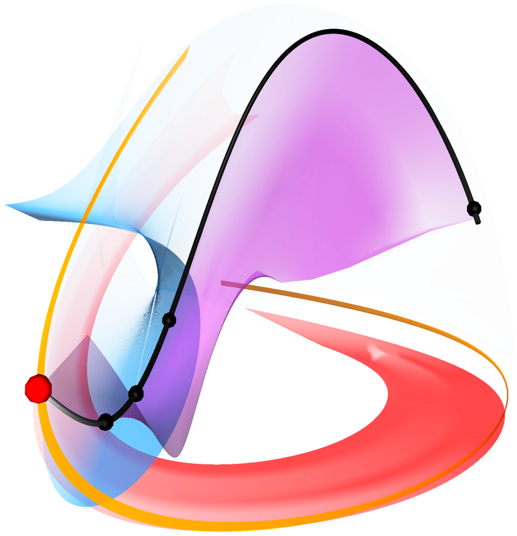

Figure 5: *Four smoke surfaces derived from the NF-κB Pathway model (indicated with the different colors) and depicted together with a single boundary trajectory (black line).*

The steps discussed above are applied to all subspaces. Drilling down to lower-dimensional subspaces remains similar to the initial implementation in *ManyLands*. However, when the smoke surfaces are unfolded to the subspaces, their computation is heavily parallelized. The users smoothly unfold the calculated smoke surfaces to their lower-dimensional subspaces (see Figures 1 and 3). This requires that smooth, animated transitions from 4D to 3D and subsequently from 3D to 2D, where in-between representations (e.g., the Dalí cross, as shown in the second sub-image of Figure 1) are employed and reduced to remove redundant symmetries before moving to the next dimension. For more details on the mechanisms of the unfoldings, we refer to the original paper [AKS*19]. The projections from one phase space to a lower or higher dimensional one are implemented on the GPU to make use of its parallel processing capabilities, reducing the computational time per frame. Users can also bring forward specific initial conditions by visualizing individual trajectories over the smoke surfaces (Figure 5, see black line). In this way, individual trajectories can be analyzed and compared to the underlying smoke surface.

Implementation: We employ OpenGL and compute shaders to calculate vertex normals, curvature, and projection for each vertex and per frame. Our implementation is a stand-alone application in C++ and is publicly available in our repository[†].

4. Results

We demonstrate the results of our proposed approach using three dynamical system cases. Two of these systems arise from biological phenomena (i.e., bipolar disorder [Gol11] and immune response to infections [KJS06]) and the third one is a strange attractor chosen for its intricate dynamics [Bou15]. These cases demonstrate how our smoke surfaces can be used and interpreted by experts and how insights can be derived from them.

Figure 6: *Five smoke surfaces derived from the Goldbeter model of bipolar disorder (left: 4D, right: 3D and 2D). Each color denotes a distinct initialization set of the dynamical system.*

The Goldbeter model of bipolar disorder [Gol11] represents alternations between manic and depressive states with respect to medications, formalized in the parameters of the 4D model. For certain initial conditions, the model exhibits oscillatory behavior between the two states. Corresponding results are shown in Figure 6 with the 4D system on the left, its 3D projections at the top right corner, and the 2D phase portraits at the bottom right corner for five sets of initial conditions (i.e., medication schemes), depicted with the five distinct colors. The oscillating behavior of the model is evident, while the strong impact of the five slightly different initializations is also depicted. Figure 1 also showcases this oscillating behavior in all subspaces, including also the in-between states as resulting from the smooth transitions across the phase space. This indicates that the system—and the underlying phenomena it represents—are sensitive to slight changes to their initialization, i.e., different therapeutic conditions, as demonstrated by the sizable differences across the distinct smoke surfaces.

The NF-κB Pathway [KJS06] plays a key role in regulating human immune responses to infections. This system also exhibits oscillatory behavior for certain initial conditions. Its smoke surfaces in conjunction with a single trajectory computed from this system can be seen in Figure 5 for a set of four initial conditions, indicated with the distinct colors. In particular, the purple smoke surface exhibits a fast–slow behavior (indicated by the black dots on the trajectory, which at first are close together and then spread out) and also sensitivity to the initial conditions in its slow part (indicated by the wider band towards the "end" of the smoke surface).

A Bouali attractor is a 4D strange attractor where the equation system is derived from the 2D Lotka-Volterra oscillator [Bou15]. It has no biological effect and it is used in this work purely for demonstration purposes, as it exhibits interesting chaotic behavior. The corresponding smoke surfaces are presented in Figure 7 for three initial conditions, depicted by the three distinct colors. The figure showcases the intricate dynamics that the chaotic system exhibits. In *ManyLands*, this would result in a view with heavy occlusions for an increasing number of initialization sets. Our approach, instead, can handle the occlusions that naturally occur due to this

† https://github.com/MarwinSc/ManyLands-SmokeSurfaces

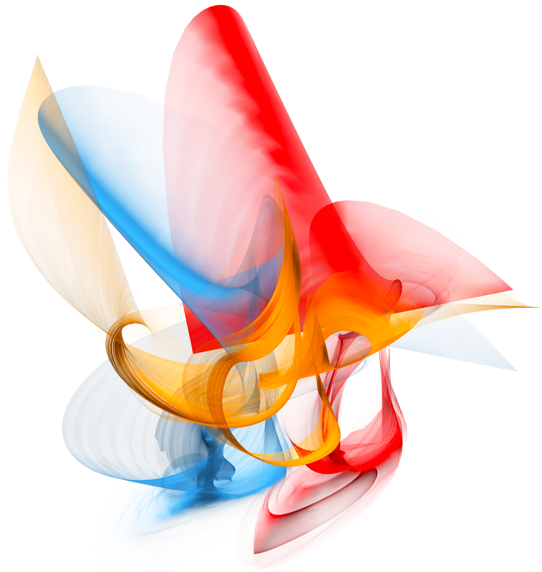

Figure 7: *Three smoke surfaces derived from the Bouali attractor, depicting the intricate dynamics that such models can exhibit when derived from systems with chaotic behavior.*

chaotic behavior. Adding the corresponding trajectory, as shown in Figure 3 (see the difference between the images in the first row), helps further with analyzing the behavior of such heavily convoluted surfaces. The accompanying 3D and 2D subspaces support the investigations of this chaotic behavior of the model in lower dimensions. Here, the parts of the trajectories that exhibit fast–slow behaviors are also visible (see the unequal spacing of the black dots on the trajectory in the top right representation in Figure 3).

Performance: The frame rate depends on the number of surfaces and trajectories in each surface, as well as the state of the application. The average for a single smoke surface rendered with the tesseract is 59 fps, and once unfolded to the eight 3D subspaces it averages at 33 fps. This frame rate was observed on a system with an NVIDIA RTX 3070 graphics card and a Ryzen 9 3900XT CPU.

5. Conclusion and Future Work

We extended *ManyLands* by enabling smoke surfaces for the representation of multiple stream surfaces that support the comparison of trajectories stemming from different initial conditions. Smoke surfaces support a holistic notion of the dynamics of the observed system and their sensitivity, as shown in our three examples. Comparing qualitatively multiple trajectories and understanding the sensitivity of the system to its initialization would not have been possible otherwise in a view without clutter, as shown in Figure 2.

We have also identified several directions for future work. While the slow–fast behavior is, to a degree, encoded in the surface transparency, labeling would further help to understand the system behavior. Extending the seeding approach can also enable the creation of trajectories along a curve in 4D, while investigating unsteady systems is also important. Other interesting directions for future work include extending or combining our approach with solutions for higher-dimensional systems (e.g., [BGG20]). Initial informal feedback from mathematical biologists indicates that our approach—except for being visually pleasing—provides them with qualitative insights into the sensitivity and behavior of their systems. Although a future user study to assess the usability of our approach and its usefulness to domain experts is pending, the integration of smoke surfaces enables mathematical biologists to explore and study the sensitivity of a dynamical system to its initialization.

References

[AKS*19] AMIRKHANOV A., KOSIUK I., SZMOLYAN P., AMIRKHANOV A., MISTELBAUER G., GRÖLLER M. E., RAIDOU R. G.: ManyLands: A journey across 4D phase space of trajectories. *Computer Graphics Forum 38*, 7 (2019), 191–202. 1, 2, 4

[AS83] ABRAHAM R. H., SHAW C. D.: *Dynamics—The geometry of behavior*. Aerial Press, 1983. 2

[Atk91] ATKINSON K.: *An introduction to numerical analysis*. John Wiley & Sons, 1991. 2

[BGG20] BARTOLOVIC N., GROSS M., GÜNTHER T.: Phase space projection of dynamical systems. *Computer Graphics Forum 39*, 3 (2020), 253–264. 2, 5

[Bou15] BOUALI S.: A new hyperchaotic attractor with complex patterns. *arXiv preprint arXiv:1503.08872* (2015). 1, 4

[GHWG14] GROTTEL S., HEINRICH J., WEISKOPF D., GUMHOLD S.: Visual analysis of trajectories in multi-dimensional state spaces. *Computer Graphics Forum 33*, 6 (2014), 310–321. 2

[Gol11] GOLDBETER A.: A model for the dynamics of bipolar disorders. *Progress in Biophysics and Molecular Biology 105*, 1 (2011), 119–127. BrainModes: The role of neuronal oscillations in health and disease. 1, 4

[GSE*14] GÜNTHER T., SCHULZE M., ESTURO J. M., RÖSSL C., THEISEL H.: Opacity optimization for surfaces. *Computer Graphics Forum 33*, 3 (2014), 11–20. 2

[GTG17] GÜNTHER T., THEISEL H., GROSS M.: Decoupled opacity optimization for points, lines and surfaces. *Computer Graphics Forum 36*, 2 (2017), 153–162. 2

[GWM*96] GRÖLLER E., WEGENKITTL R., MILIK A., PRSKAWETZ A., FEICHTINGER G., SANDERSON W. C.: The geometry of wonderland. *Chaos, Solitons & Fractals 7*, 12 (1996), 1989–2006. 2

[HGH*10] HUMMEL M., GARTH C., HAMANN B., HAGEN H., JOY K. I.: Iris: Illustrative rendering for integral surfaces. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (2010), 1319–1328. 2

[KJS06] KRISHNA S., JENSEN M. H., SNEPPEN K.: Minimal model of spiky oscillations in NF-kB signaling. *Proceedings of the National Academy of Sciences 103*, 29 (2006), 10840–10845. 1, 4

[KS16] KOSIUK I., SZMOLYAN P.: Geometric analysis of the goldbeter minimal model for the embryonic cell cycle. *Journal of Mathematical Biology 72* (2016), 1337–1368. 1

[Lor63] LORENZ E. N.: Deterministic nonperiodic flow. *Journal of Atmospheric Sciences 20*, 2 (1963), 130–141. 2

[MB13] MCGUIRE M., BAVOIL L.: Weighted blended order-independent transparency. *Journal of Computer Graphics Techniques (JCGT) 2*, 2 (December 2013), 122–141. 3

[TG20] TIAN X., GÜNTHER T.: Progressive rendering of transparent integral surfaces. In *Eurographics Conference on Visualization Short Papers* (Norrköping, Sweden, 2020), Eurographics, pp. 145–149. 2

[vFWTS08] VON FUNCK W., WEINKAUF T., THEISEL H., SEIDEL H.-P.: Smoke surfaces: An interactive flow visualization technique inspired by real-world flow experiments. *IEEE Transactions on Visualization and Computer Graphics 14*, 6 (2008), 1396–1403. 2, 3

[WLG97] WEGENKITTL R., LÖFFELMANN H., GRÖLLER E.: Visualizing the behavior of higher dimensional dynamical systems. In *Visualization* (1997), IEEE Computer Society Press, pp. 119–125. 2

Proceedings published by Eurographics - The European Association for Computer Graphics.