

In Situ Visualization for 3D Agent-Based Vocal Fold Inflammation and Repair Simulation

*Nuttiya Seekhao*¹, *Joseph JaJa*¹, *Luc Mongeau*², *Nicole Y.K. Li-Jessen*²

© The Authors 2017. This paper is published with open access at SuperFri.org

A fast and insightful visualization is essential in modeling biological system behaviors and understanding underlying inter-cellular mechanisms. High fidelity models produce billions of data points per time step, making *in situ* visualization techniques extremely desirable as they mitigate I/O bottlenecks and provide computational steering capability. In this work, we present a novel high-performance scheme to couple *in situ* visualization with the simulation of the vocal fold inflammation and repair using little to no extra cost in execution time or computing resources. The visualization component is first optimized with an adaptive sampling scheme to accelerate the rendering process while maintaining the precision of the displayed visual results. Our software employs VirtualGL to perform visualization *in situ*. The scheme overlaps visualization and simulation, resulting in the optimal utilization of computing resources. This results in an *in situ* system biology simulation suite capable of remote simulation of 17 million biological cells and 1.2 billion chemical data points, remote visualization of the results, and delivery of visualized frames with aggregated statistics to remote clients in real-time.

Keywords: *in situ, visualization, vocal fold, systems biology simulation, agent based modeling, tissue inflammation and repair, computational steering.*

Introduction

Agent-based modeling (ABM) is a powerful and widely used approach to simulate a system consisting of interacting components or *agents*. This form of modeling expresses a system at the microscale, and attempts to explain the emergence of higher order properties of the overall system [29]. As opposed to the analytical, equation-based approaches, the agent-based approach offers the ability to add complex behaviors to individual components or *agents*, making modeling of composite networks uncomplicated [8, 15]. In ABM, *agents* are used to represent a wide spectrum of entities such as animals in ecosystems [13, 18, 24, 25], consumers and markets in economic models [3, 10, 36–39], and cells and proteins in biological systems [9, 11, 12, 17, 20–23, 31, 32, 34, 41]. These entities interact among themselves and with their environment (ABM *world*) in discrete time steps following a set of stochastic and/or deterministic rules. The simulation area or ABM *world* is discretized into 3D cubes called *patches*. In ABM, *agents* can be mobile and move from *patch* to *patch*. Each *patch* maintains its states, which affect the action decision of the residing and neighboring *agents*.

In this paper, we use the ABM simulation approach to capture tissue injury and repair at the cellular level. More specifically, we focus on the vocal fold injuries. It is estimated that voice disorders afflict 1 in 13 adults [7], and nearly 1 in 12 children [1] in the the United State annually. During phonation, human vocal folds undergo continuous biomechanical stresses. Thus, voice overuse can lead to vocal fold mucosal tissue injury that triggers complex biological processes of inflammation and repair. Voice treatments are usually prescribed to patients with voice problems [16, 27]. However, the healing outcomes of the treatment depend on the patient's initial condition and biological profile [20], making the treatment-planning process restraining and difficult for physicians and therapists. Computational medicine is a promising approach to

¹University of Maryland, College Park, USA

²McGill University, Montreal, Québec, Canada

addressing this problem as it incorporates the patient's biological profile and initial conditions as parameters in determining an appropriate treatment [19, 26]. Thus, ABM offers a gateway to capturing inflammation and repair behavior to predict the outcomes of specific treatments for an individual.

The vocal fold (VF) ABM requires a high-resolution 3D grid to capture cell-cell and cell-substrate interactions with sufficient details in order to accurately predict the temporal tissue response to a provocation. Since the size of the *world* grid reflects the spatial resolution of the simulation, a high-resolution 3D VF ABM involves generation and analysis of large amounts of data. Hence, a fast and low I/O load visualization is highly desirable, making *in situ* visualization an ideal candidate for understanding the model output.

In this work, we extend our previous work on a fast GPU implementation of VF ABM simulation to incorporate *in situ* visualization at no extra cost to the overall simulation. The proposed scheme combines a data reduction technique to gain optimal visualization performance with a CPU-GPU scheduling technique to overlap simulation and visualization while hiding the execution costs of the visualization component. More specifically, our main contributions can be stated as follows.

- Reducing the amount of data analyzed, while maintaining a high fidelity visual resolution using an adaptive sampling scheme;
- An *in-situ* bio-simulation suite capable of processing 17 million biological cells and 1.2 billion chemical data points (a scale biologically representative of human vocal fold at the cellular level) as well as collecting aggregated statistics, which:
 - Completely bypasses I/O; and
 - Analyzes and renders results without causing an increase in the overall simulation time.

The rest of this article is structured as follows: section 1 introduces brief descriptions of concepts fundamental to this work including agent-based modeling (ABM), the use of ABM for inflammation and healing modeling, *in situ* ABM, and information on software and hardware platforms used. Section 2 discusses approaches taken to enhance visualization resolution, while keeping visualization performance optimized to enable a complete masking of the visualization cost, using our CPU-GPU task scheduling scheme [32, 33]. Performance and resolution enhancement evaluation is discussed in section 3.

1. Background and Related Work

1.1. Agent-Based Modeling (ABM)

The basic components of ABMs are:

- **Agents** - Autonomous objects that perform actions and interact with other agents and the environment;
- **Agent Rules** - Behaviors of each type of agents; and
- **World** - The environment which all agents belong to.

Multiple types of agents can be modeled in a single ABM. Each type of agents behaves according to a set of predefined rules, which can be deterministic or stochastic. For example, a simulation related to tissue inflammation may have various biological cell types, such as neutrophils, macrophages and fibroblasts, as *agents*. The predefined rules are determined using the best available knowledge in the literature of each component of the system. The autonomous

agents are mobile and make decisions based on their states and their environment. *Patch* size is uniform across the world, and thus the resolution of the simulation environment is inversely proportional to the *patch* size. The temporal dimension of ABMs is discrete and the simulation progresses in a sequence of synchronous iterations (sometimes referred to as *ticks*).

1.2. Modeling Vocal Fold Inflammation and Repair with ABM

A Vocal Fold (VF) ABM simulating inflammation and repair was developed and partially verified against empirical data [20]. The biological cells were implemented as mobile ABM *agents*. These *agents* perform their functionality and make action decisions based on the states of their surrounding. Briefly, at the time of injury, the damaged mucosal tissue triggers platelet degranulation [21, 22]. Platelets secrete chemicals by modifying the states of their residing *patches*. In the inflammatory phase, these chemical gradients stimulate vasodilation and attraction of inflammatory cells, namely, neutrophils and macrophages. These inflammatory cells then get activated once they reach the wound site. Activated cells then clean up cell debris and produce more chemicals to attract cells called fibroblasts, which are responsible for tissue structure maintenance. In the healing phase, fibroblasts, activated by tissue damage, synthesize and deposit extracellular matrix (ECM) proteins such as collagen, elastin, and hyaluronans to form building blocks in tissue repair [40]. These ECM proteins then form a scaffold for supporting fibroblasts and other cells migration and wound repair activities [5].

1.3. *In-Situ* Agent-Based Modeling

In this section, we summarize the most related works. To the best of our knowledge, there has not been much *in situ* visualization work involving agent-based modeling (ABM). A quadtree-based ABM is proposed in [17] to reduce the amount of irrelevant data analyzed in-situ, where the work in [35] attempts to accomplish the same goal with a bitmap-based approach. There are tools available on Paraview [14], a popular visualization framework, which can be used for ABMs. Paraview Catalyst [4, 6] was developed to process simulation output data in-situ according to the user's co-processing script. An image-based approach built on top of Paraview Catalyst was presented in [2] to efficiently manage rendered images created in-situ by Paraview Catalyst. As much as all these works [2, 4, 6, 17, 35] reduce I/O loads, none completely by-passes I/O or can be used to achieve anything close to the desired performance for our problem.

1.4. Platforms Used

The model was implemented using C++ as the main language for speed and portability. The tasks executed on CPUs and GPUs are parallelized using Open Multi-Processing (OpenMP) application programming interface (API) for shared-memory parallel programming and Compute Unified Device Architecture (CUDA) API [28], respectively. The visualization component was implemented using Open Graphics Library (OpenGL), a cross-platform API for 2D and 3D graphics rendering. The model outputs are visualized *in situ*, and the visualized frames are delivered to the remote user via VirtualGL, an open source software which allows any Unix or Linux remote display software to run OpenGL applications by using the server's powerful 3D accelerator to perform rendering calls and send only rendered images to the client [30].

The model was tested and benchmarked on a single compute node with 44-core Intel(R) Xeon(R) CPU E5-2699 v4 @ 2.20GHz host and two attached accelerators, NVIDIA Tesla M40.

The host has 128 GB of main memory and the accelerator has 24 GB of memory per device. The program occupies the compute node fully while running and does not allow for sharing with other processes as the program requires most of the memory available on the node.

2. A Novel Approach

2.1. Data Reduction Techniques

The processing of large amounts of data is inevitable in high-fidelity simulations. Expressive visualization enables the human eye to easily extract insightful information about the simulated system. Our model consists of approximately 17 million agents and produces 1.7 billion bio-marker data points in each iteration. The visualization component includes cell migration, chemical diffusion, and damage tracking. The most time consuming component is chemical diffusion, which needs to access 154 million points of data during each iteration. To optimize the visualization of such large amounts of data, we employ sampling techniques and study their effects on the simulation visual output and corresponding performance enhancements.

2.1.1. Constant Sampling

The first sampling approach is simply constant sampling. The world environment is divided into tiles of size $grid_x \times grid_y \times grid_z$, where the corresponding grid point represents the values of the tile centered at the point and within the radius of $grid_x/2$, $grid_y/2$ and $grid_z/2$ in the x-, y-, and z-dimension respectively. This naive approach was used to improve the visualization speed in our earlier work [33]. At the perspective of the entire simulated volume, the appearances of sampled simulation were indistinguishable for up to $(6 \times 6 \times 6)$ -segment sampling. However, the output became pixelated when the view is magnified and the camera focuses on a smaller area. In particular, the wound area was not rendered with reasonable fidelity.

2.1.2. Adaptive Sampling

Adaptive sampling is used to optimize the data access while enhancing the resolution of the visual output in important areas. This sampling scheme helps keep the execution time low, and yet the details in the output presented are not compromised.

For models aiming to capture injury undergoing inflammation and repair processes, the wound site is the most active area. Therefore, the highest importance index was assigned to the wound site volume. The margin around the wound site, and the rest of the tissue are then respectively assigned less and least importance. The adaptive sampling pipeline diagram in Fig. 1 illustrates data processing steps used to sample and send data to the visualization pipeline. The program, by default, divides the whole tissue volume into three sections by first inspecting the initial wound position and size, followed by adding a margin around the wound based on user inputs, and then labeling this volume as the most active (region 1). The program lets the user specify the volume ratio between medium- and low- activity area, and splits the rest of the volume into regions 2 and 3 accordingly. As our visualization intends to highlight wound activity, the re-sampling is performed only once, thus the memory footprint is not heavily affected by the re-sampling process.




Figure 1. Adaptive sampling pipeline. From the whole set of data (left), the data are read with resolution conforming to the importance index (middle). The sampled data are then processed and sent to the visualization pipeline to be rendered (right)

2.2. Simulation and Visualization Scheduling

Given the amounts of data updated at each time step, our model requires a carefully planned scheduling mechanism for optimal performance. Accelerators such as GPUs need a CPU host, each of which has a number of cores that can be exploited using parallel programming techniques. However, whenever accelerated performance is the goal, the focus usually shifts towards GPUs due to their exceptional data parallel computation capabilities. Often, this tendency results in idle CPU cores, since their only job is to transfer data and launch GPU tasks, while the GPUs perform all the computing work. We proposed a host-device computation overlap technique in our earlier work [32], which results in state-of-the-art performance for multi-scale 2D bio-simulation ABMs. Since the 3D model is substantially more computationally demanding, we significantly extend the methods described there to achieve extremely high speed 3D simulation with *in situ* visualization.

Model operations are divided into sub-tasks and categorized as coarse or fine [32]. Since coarse-grain tasks (inflammatory cell and ECM function) are complex but less data-intensive, they are deemed CPU-suitable, where simple data-intensive fine-grain (chemical diffusion) and rendering tasks run most effectively on GPUs. The blue boxes in Fig. 2 illustrate the coarse-grain tasks that are executed in parallel on the CPU using OpenMP, where the green boxes denote the fine-grain tasks executed on the GPUs. The coarse-grain tasks execute on all CPU cores with the exception of $N_{GPU} + 1$ cores, where N_{GPU} denotes the number of GPUs. The rest of the N_{GPU} CPU cores are then used to dispatch fine-grain tasks and manage data to and from the GPUs, while the last core is spared to issue rendering calls to execute visualization on GPU. In [33], constant sampling was used, which was fast at the cost of lower resolution. Thus, placing the visualization execution in the GPU idle period was simple. However, in some circumstances, the user may want to focus on a smaller sections of the world, which means a sampling technique to enhance the resolution while keeping the visualization execution time smaller than GPU idle

period is required. The adaptive sampling technique discussed in section 2.1.2 was introduced to solve this performance-resolution trade-offs problem.




Figure 2. Diagram illustrating the scheduling and coordination of CPU-GPU computation and visualization overlap. For simplicity, this diagram is depicting a system with a single GPU. For multi-GPU, diffusion kernels launched are simply divided up and dispatched to multiple GPUs

3. Results

The results discussed below were benchmarked with 32 threads on a single compute node with 44-core Intel(R) Xeon(R) CPU E5-2699 v4 @ 2.20GHz host and two attached accelerators, NVIDIA Tesla M40. The *in situ* simulation suite performance reported was measured with the remote user running the simulation off-site on a typical 50/50 Mbps WIFI connection. The size of the world grid is 110 x 1390 x 1006.

3.1. Visualization Component Performance

For the ABM simulation, CPU tasks (excluding updates) take about 4.7 seconds while GPU tasks only take 2.5 seconds for each iteration. Thus, there is an idle period on the GPUs waiting for the CPU to finish the coarse-grain tasks. Our goal is to make the visualization component fast enough so that the 2.2-second window gap would allow us to integrate visualization with computation on the GPUs without increasing the total execution time. As discussed earlier, our work in [33] used constant sampling, which was fast enough but did not achieve good enough resolution when zooming in areas of interest such as the wound site. It was observed that for visualization view of the entire simulated volume (Fig. 3), the appearance of sampled simulation was indistinguishable from $(1 \times 1 \times 1)$ - up to $(6 \times 6 \times 6)$ -segment sampling. Thus,

In Situ Visualization for 3D Agent-Based Vocal Fold Inflammation and Repair Simulation

for optimal performance, the least important areas in adaptive sampling are set to use the coarsest resolution of 6^3 sampling windows. The medium and most important areas are then sampled with finer 4^3 and 2^3 windows, respectively. With 2–4–6-sampling resolution, as shown in Fig. 4, the resolution of the visualization improved significantly, while the execution time of the visualization component only increased to 1.9 seconds, which is still below our visualization time budget, hence no increase in the overall execution time.




Figure 3. A screenshot of *in situ* visualization of the simulation captured at the client side. The 2D charts plot total concentration for each type of a chemical. The left most 3D volume displays the distribution of one of the eight chemicals specified by the user. The second and third volumes show macrophage (brown) and neutrophil (red) distributions respectively. The last volume on the right displays current damage (pink) and the distribution of fibroblasts (blue)

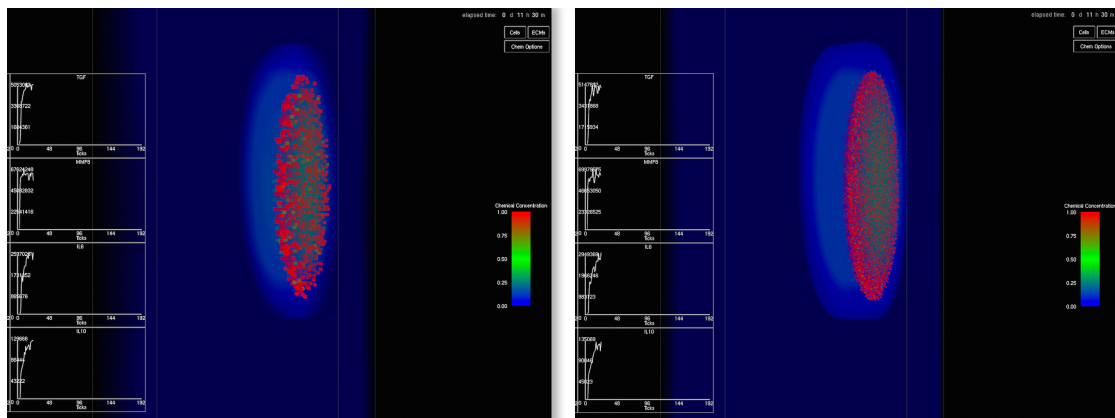


Figure 4. Screenshots comparison of 6^3 sampling windows (left) and 2–4–6-sampling resolution (right) when zoomed in to high-activity area

3.2. Coupled Simulation and Visualization Performance

The performance of the simulation suite is shown in Fig. 5. Without sampling, visualization took 23 seconds to complete. With adaptive sampling, visualization of chemical diffusion

decreased to 1.9 seconds. The visualization execution was performed during the idle period on one of the GPUs keeping the total execution time unchanged at 6.2 seconds per iteration on average. This results in the ability to run the simulation from the start of the iteration, remote computation, remote visualization to the moment the frame gets rendered on the client’s machine in under 7 seconds/frame. To the best of our knowledge, this is by far the fastest known complex ABM simulation and visualization of a problem of that physiological scale.

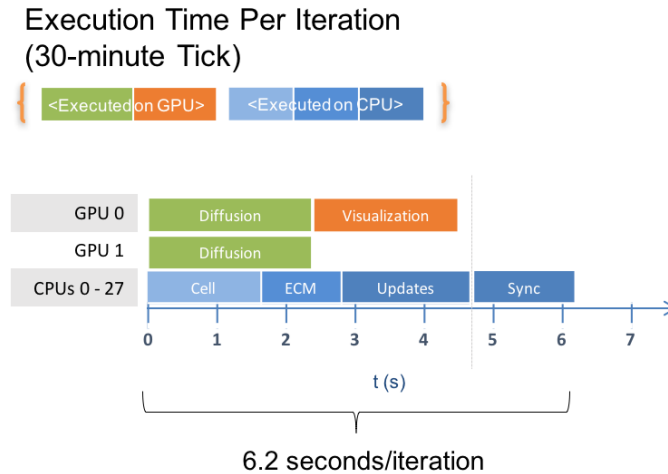


Figure 5. Simulation suite performance. Chart demonstrating overlapped visualization and computation executions on GPUs and CPUs

Conclusion

In this paper, we presented novel techniques to achieve *in situ* 3D ABM visualization at almost no cost to the overall simulation. As a result, we can simulate and render the VF inflammation and repair in real time. An effective task scheduling and management approach was used to orchestrate the execution of coarse-grain cellular functions, which are parallelized on the multi-core CPU, with the execution of fine-grain GPU tasks, including the overlapping with the in-situ visualization component. This results in optimal concurrent utilization of both multi-core CPU and GPU, including the fact that the execution time of the GPU visualization component is completely hidden behind the CPU tasks. We are able to simulate the total of 17 million inflammatory cells and 1.7 billion bio-marker data points, as well as analyze and render the same number of cells and 154 million bio-marker data points on the server and send result frames to the remote user in under 7 seconds per iteration. The model is currently being developed to incorporate more visualization functionality, which includes an integration of ECM proteins visualization and direct volume rendering into the simulation suite to give users the ability to extract meaningful information and explore the output data in different ways. The goal is to keep the overall performance of the simulation the same even with more comprehensive visualization components.

Acknowledgments

The authors would like to thank Caroline Shung, Yun (Yvonna) Li and Alireza Najafi Yazdi for their contributions to the development of the initial and base sequential model. Sujal Bista for guidance in developing the visualization component. And UMIACS staff for assistance with VirtualGL and TurboVNC configuration.

Research reported in this publication was supported by National Institute of Deafness and other Communication Disorders of the National Institutes of Health under award numbers R03DC012112 and R01DC005788. The authors gratefully acknowledge the support provided by NSF, Contract Number CNS-1429404 MRI project.

This paper is distributed under the terms of the Creative Commons Attribution-Non Commercial 3.0 License which permits non-commercial use, reproduction and distribution of the work without further permission provided the original work is properly cited.

References

1. Voice, speech, and language quick statistics. <http://www.nidcd.nih.gov/health/statistics/vsl/Pages/stats.aspx> (2012), accessed: 2016-10-3
2. Ahrens, J., Jourdain, S., OLeary, P., Patchett, J., Rogers, D.H., Petersen, M.: An image-based approach to extreme scale in situ visualization and analysis. In: Damkroger, T., Dongarra, J. (eds.) SC14: International Conference for High Performance Computing, Networking, Storage and Analysis. pp. 424–434. IEEE (nov 2014), DOI: 10.1109/sc.2014.40
3. Arthur, W.B.: Chapter 32 out-of-equilibrium economics and agent-based modeling. In: Handbook of Computational Economics, pp. 1551–1564. Elsevier (2006), DOI: 10.1016/s1574-0021(05)02032-0
4. Ayachit, U., Bauer, A., Geveci, B., O'Leary, P., Moreland, K., Fabian, N., Mauldin, J.: ParaView catalyst. In: Weber, G.H. (ed.) Proceedings of the First Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization - ISAV2015. pp. 25–29. ACM Press (2015), DOI: 10.1145/2828612.2828624
5. Bainbridge, P.: Wound healing and the role of fibroblasts. Journal of Wound Care 22(8), 407–412 (aug 2013), DOI: 10.12968/jowc.2013.22.8.407
6. Bauer, A.C., Geveci, B., Schroeder, W.: The paraview catalyst users guide (2013)
7. Bhattacharyya, N.: The prevalence of voice problems among adults in the united states. The Laryngoscope 124(10), 2359–2362 (may 2014), DOI: 10.1002/lary.24740
8. Bonabeau, E.: Agent-based modeling: Methods and techniques for simulating human systems. Proceedings of the National Academy of Sciences 99(Supplement 3), 7280–7287 (may 2002), DOI: 10.1073/pnas.082080899
9. Brown, B.N., Price, I.M., Toapanta, F.R., DeAlmeida, D.R., Wiley, C.A., Ross, T.M., Oury, T.D., Vodovotz, Y.: An agent-based model of inflammation and fibrosis following particulate exposure in the lung. Mathematical Biosciences 231(2), 186–196 (jun 2011), DOI: 10.1016/j.mbs.2011.03.005

10. Caiani, A., Russo, A., Palestrini, A., Gallegati, M.: *Economics with Heterogeneous Interacting Agents: A Practical Guide to Agent-Based Modeling*. Springer (2016)
11. Cilfone, N.A., Kirschner, D.E., Linderman, J.J.: Strategies for efficient numerical implementation of hybrid multi-scale agent-based models to describe biological systems. *Cellular and Molecular Bioengineering* 8(1), 119–136 (nov 2014), DOI: 10.1007/s12195-014-0363-6
12. D’Souza, R.M., Lysenko, M., Marino, S., Kirschner, D.: Data-parallel algorithms for agent-based model simulation of tuberculosis on graphics processing units. In: Wainer, G.A., Shaffer, C.A., McGraw, R.M., Chinni, M.J. (eds.) *Proceedings of the 2009 Spring Simulation Multiconference*. p. 21. Society for Computer Simulation International, SCS/ACM (2009)
13. Gras, R., Devaurs, D., Wozniak, A., Aspinall, A.: An individual-based evolving predator-prey ecosystem simulation using a fuzzy cognitive map as the behavior model. *Artificial Life* 15(4), 423–463 (oct 2009), DOI: 10.1162/artl.2009.gras.012
14. Henderson, A., Ahrens, J., Law, C., et al.: *The ParaView Guide*. Kitware Clifton Park, NY (2004)
15. Janssen, M.A.: Agent-based modelling. *Modelling in ecological economics* pp. 155–172 (2005)
16. Johns, M.M.: Update on the etiology, diagnosis, and treatment of vocal fold nodules, polyps, and cysts. *Current Opinion in Otolaryngology & Head and Neck Surgery* 11(6), 456–461 (dec 2003), DOI: 10.1097/00020840-200312000-00009
17. Krekhov, A., Grninger, J., Schlvoigt, R., Krger, J.: Towards in situ visualization of extreme-scale, agent-based, worldwide disease-spreading simulations. In: *SIGGRAPH Asia 2015 Visualization in High Performance Computing on - SA '15*. ACM Press (2015), DOI: 10.1145/2818517.2818543
18. Lampert, A., Hastings, A.: Stability and distribution of predator-prey systems: local and regional mechanisms and patterns. *Ecology Letters* 19(3), 279–288 (jan 2016), DOI: 10.1111/ele.12565
19. Li, N.Y.K., Abbott, K.V., Rosen, C., An, G., Hebda, P.A., Vodovotz, Y.: Translational systems biology and voice pathophysiology. *The Laryngoscope* 120(3), 511–515 (mar 2010), DOI: 10.1002/lary.20755
20. Li, N.Y.K., Verdolini, K., Clermont, G., Mi, Q., Rubinstein, E.N., Hebda, P.A., Vodovotz, Y.: A patient-specific in silico model of inflammation and healing tested in acute vocal fold injury. *PLoS ONE* 3(7), e2789 (jul 2008), DOI: 10.1371/journal.pone.0002789
21. Li, N.Y.K., Vodovotz, Y., Hebda, P.A., Abbott, K.V.: Biosimulation of inflammation and healing in surgically injured vocal folds. *Annals of Otology, Rhinology & Laryngology* 119(6), 412–423 (jun 2010), DOI: 10.1177/000348941011900609
22. Li, N.Y.K., Vodovotz, Y., Kim, K.H., Mi, Q., Hebda, P.A., Abbott, K.V.: Biosimulation of acute phonotrauma: An extended model. *The Laryngoscope* 121(11), 2418–2428 (oct 2011), DOI: 10.1002/lary.22226

23. Martin, K.S., Blemker, S.S., Peirce, S.M.: Agent-based computational model investigates muscle-specific responses to disuse-induced atrophy. *Journal of Applied Physiology* 118(10), 1299–1309 (feb 2015), DOI: 10.1152/jappphysiol.01150.2014
24. McLane, A.J., Semeniuk, C., McDermid, G.J., Marceau, D.J.: The role of agent-based models in wildlife ecology and management. *Ecological Modelling* 222(8), 1544–1556 (apr 2011), DOI: 10.1016/j.ecolmodel.2011.01.020
25. McLane, A.J., Semeniuk, C., McDermid, G.J., Tomback, D.F., Lorenz, T., Marceau, D.: Energetic behavioural-strategy prioritization of clark’s nutcrackers in whitebark pine communities: An agent-based modeling approach. *Ecological Modelling* 354, 123–139 (jun 2017), DOI: 10.1016/j.ecolmodel.2017.03.019
26. Mi, Q., Li, N.Y.K., Ziraldo, C., Ghuma, A., Mikheev, M., Squires, R., Okonkwo, D.O., Verdolini-Abbott, K., Constantine, G., An, G., Vodovotz, Y.: Translational systems biology of inflammation: potential applications to personalized medicine. *Personalized Medicine* 7(5), 549–559 (sep 2010), DOI: 10.2217/pme.10.45
27. Misono, S., Marmor, S., Roy, N., Mau, T., Cohen, S.M.: Multi-institutional study of voice disorders and voice therapy referral. *Otolaryngology-Head and Neck Surgery* 155(1), 33–41 (jul 2016), DOI: 10.1177/0194599816639244
28. Nvidia, C.: Compute unified device architecture programming guide (2007)
29. Page, S.E.: Agent based models. *The New Palgrave Dictionary of Economics*. Palgrave MacMillan, New York (2005), DOI: 10.1057/9780230226203.0016
30. Project, T.V.: VirtualGL background. <http://www.virtualgl.org/About/Background> (2015)
31. Richmond, P., Walker, D., Coakley, S., Romano, D.: High performance cellular level agent-based simulation with FLAME for the GPU. *Briefings in Bioinformatics* 11(3), 334–347 (feb 2010), DOI: 10.1093/bib/bbp073
32. Seekhao, N., Shung, C., Jaja, J., Mongeau, L., Li-Jessen, N.Y.K.: Real-time agent-based modeling simulation with in-situ visualization of complex biological systems: A case study on vocal fold inflammation and healing. In: 2016 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW). IEEE (may 2016), DOI: 10.1109/ipdpsw.2016.20
33. Seekhao, N., Shung, C., JaJa, J., Mongeau, L., Li-Jessen, N.Y.: High-resolution 3d vocal fold repair simulation using highly-parallelized agent-based modeling. submitted to *PloS one* (2017)
34. Shi, Z., Chapes, S.K., Ben-Arieh, D., Wu, C.H.: An agent-based model of a hepatic inflammatory response to salmonella: A computational study under a large set of experimental data. *PLOS ONE* 11(8), e0161131 (aug 2016), DOI: 10.1371/journal.pone.0161131
35. Su, Y., Wang, Y., Agrawal, G.: In-situ bitmaps generation and efficient data analysis based on bitmaps. In: Kielmann, T., Hildebrand, D., Taufer, M. (eds.) *Proceedings of the 24th International Symposium on High-Performance Parallel and Distributed Computing - HPDC '15*. pp. 61–72. ACM Press (2015), DOI: 10.1145/2749246.2749268

36. Tesfatsion, L.: Agent-based computational economics: modeling economies as complex adaptive systems. *Information Sciences* 149(4), 262–268 (feb 2003), DOI: 10.1016/s0020-0255(02)00280-3
37. Tesfatsion, L.: Agent-based computational economics: Growing economies from the bottom up. *Artificial Life* 8(1), 55–82 (jan 2002), DOI: 10.1162/106454602753694765
38. Tesfatsion, L.: Chapter 16 agent-based computational economics: A constructive approach to economic theory. In: *Handbook of Computational Economics*, pp. 831–880. Elsevier (2006), DOI: 10.1016/s1574-0021(05)02016-2
39. Tesfatsion, L., Judd, K.L.: *Handbook of computational economics: agent-based computational economics*, vol. 2. Elsevier (2006)
40. Velnar, T., Bailey, T., Smrkolj, V.: The wound healing process: An overview of the cellular and molecular mechanisms. *Journal of International Medical Research* 37(5), 1528–1542 (oct 2009), DOI: 10.1177/147323000903700531
41. Wang, Z., Butner, J.D., Kerketta, R., Cristini, V., Deisboeck, T.S.: Simulating cancer growth with multiscale agent-based modeling. *Seminars in Cancer Biology* 30, 70–78 (feb 2015), DOI: 10.1016/j.semcancer.2014.04.001