

About the Cover

Simulating the development of a brain tumor

This month's cover was suggested by Rick Durrett's article on cancer in this issue. The images are still frames taken from an animation simulating 180 days of tumor growth in a 1cm x 1cm section of brain tissue. The simulation was produced by Paul Macklin and John Lowengrub at the University of California at Irvine. Macklin (now at the University of Southern California) is also responsible for Figure 3 in Durrett's article.

Each frame advances the simulation by 10 days. The top image of each pair shows the developing tumor. Red marks the region of growing tumor tissue, blue tumor tissue behind this growing front is oxygen-starved (hypoxic), and brown regions behind these are dying (necrotic) tumor cells. The tumor plot is laid over the original brain tissue, in which the white region is the cranium, light gray is white matter, dark gray is gray matter, and black is cerebrospinal fluid. The tumor grows at noticeably different rates within each environment. Indeed, the principal feature of this simulation is its ability to deal with heterogeneous tissues.


The bottom image of each pair shows the distribution of pressure generated by the growing tumor, ranging from low pressure (blue) to high pressure (oranges and reds).

It is pressure that causes the first symptoms felt by the patient, such as headaches, dizziness, or vision impairment.

The full animation can be found at

<http://MathCancer.org/Multimedia.php>

It also tracks several other features of the tumor development. Here is the full frame for the last day:



Macklin commented on this simulation, "There is extensive nonlinear feedback between the extracellular matrix (ECM) density and its degradation by the tumor, the tissue's biomechanical properties, the resulting tumor growth profile, and the distribution of oxygen and hypoxia in and around the tumor. As the simulation progresses, the blue region in the ECM frame enlarges. This is due to the secretion of matrix degrading enzymes (MDEs) by the tumor (lower middle). Oxygen is released by the

pre-existing vasculature (blood circulation system) in non-degraded tissue, which perfuses through the tissue domain. Necrosis provides mechanical stress relief. The tissue's mechanical compliance is directly dependent upon the ECM density. Hypoxic regions of the tumor secrete tumor-angiogenic growth factors (TAFs). These stimulate the cancer's construction of its own blood circulation system, called angiogenesis.


"In the simulation, oxygen gradients emerge that drive nonuniform growth on the outer tumor edges toward better oxygenated regions. This growth outpaces cell adhesion, increasing the magnitude of the oxygen gradients and thus also the nonuniformity of the tumor growth. This leads to the morphological instabilities seen here: invasive fingering growth and tumor fragmentation in regions of cell death. This is another example of a nonlinear feedback in the tumor-microenvironment system. Both these types of shape instabilities can be observed in real tumors."

The principal reference for the cover simulation is the paper "A New Ghost Cell/Level Set Method for Moving Boundary Problems: Application to Tumor Growth", by Macklin and John Lowengrub, in volume 35 of the *Journal of Scientific Computing*. More references can be found at

<http://www.MathCancer.org/Publications.php>

<http://math.uci.edu/~lowengrub>

Macklin's laboratory projects at USC have become more ambitious in recent times. Here is a frame from a 3D animation:



You can see more of this type of graphic at

<http://MathCancer.org/AMS2013>

We were intrigued by the great complexity and potential impact of this work, and asked Macklin to tell us more.

• *What are the long term goals of your project?*

In the long term, we'd like to create modeling tools that help oncologists to plan surgeries more accurately and choose better, patient-optimized therapies. We'd also like to create *in-silico* tools that help biologists to understand and extrapolate their *in-vitro* experimental findings to

(Continued on page 370)

(About the cover continued from page 325)

in-vivo cell behavior. It is my hope that this process will also drive the development of new and novel mathematical modeling and computational methods.

• *How will you measure its success?*

This is a really great question. First, does our project influence the thinking of biologists and clinicians? Do they start to think beyond the molecular biology of single cells, towards interconnected systems of cells with physical constraints, such as those imposed by transport limitations of oxygen and therapeutic compounds? Second, can we make tools that are sufficiently descriptive and realistic to provide new insights on the underlying cancer biology, and possibly lead us to revisit and refine our operating biological hypotheses? Third, given a set of *in-vitro* or patient measurements, can we quantitatively predict something biologically or clinically important (e.g., predict growth rates and overall chemotherapy response, based solely upon imaging and pathology inputs)? Fourth, can we create validated models that are sufficiently mechanistic that we can not only predict behavior, but also apply controls to change behavior? That is, we would like to reach the point where we are so good at describing cancer progression and therapy response for individual patients that we can choose a desired outcome (e.g., stay within organ X with size under Y for Z years) and optimize treatment to attain that outcome.

Computational and mathematical oncology is having increasing success at #1 and #2. We're starting to make headway on #3, and even on #4.

• *How do you measure the accuracy of your simulations?*

First, we'd measure it qualitatively: does it match behavior as observed in the clinic, on reasonable space and time scales? If so, can we make a macroscopic prediction (e.g., on growth rates with and without therapy) that match within some relative error bound for most of our patients? What does it predict about the microstructure of the cancer? Do expected correlations pop out, such as between cell position and relative frequency of cell mitosis? Does it predict new correlations that we can verify in existing data?

• *What size are the meshes in the cover animation? What time interval? How long did the computation take?*

It used a 20-micron mesh. The time scale was set to dynamically satisfy a Courant-Friedrichs-Levy condition, typically 1 hour or less. These simulations were done single-threaded on desktop computers about six years ago, and required on the order of days to a week to complete. Assuming four 18-month Moore's law doublings in computer speed, that simulation should take several hours today, and it would be faster if parallelized.

• *What machines and OS are you using for parallelization in the recent 3D stuff?*

The recent 3D work is written in standards-compliant, object-oriented C++, with parallelization in OpenMP.

(OpenMP is a cross-platform/cross-architecture programming interface that supports shared memory parallelization in C/C++ and Fortran. In shared memory parallelization, all threads have access to a shared pool of memory, eliminating the need for message passing between threads.) We tend to run the bigger simulations on a machine with two 6-core CPUs with hyperthreading (up to 24 simultaneous execution threads) and 48 GB of memory, running Ubuntu Linux. Bigger simulations (around a half million cells in a 3-5 mm³ domain) take between several hours and a weekend, depending upon the complexity and the size of the computational domain. I run these on Windows and Mac OSX machines during development, to make sure the simulator works across platforms.

—Bill Casselman
Graphics Editor
(notices-covers@ams.org)