

PROSTATE CANCER: Crunching the Numbers

By Alexander Gelfand

The numbers tell the story: Prostate cancer is a killer. It's the second most common form of cancer and the second leading cause of cancer death among American men. More than 230,000 new cases are expected by the end of this year alone, and nearly 30,000 men are expected to die from it in the same period.

Early diagnosis is vital, but current methods are far from perfect; and once it metastasizes to the bones, the disease is incurable. Some researchers and clinicians are embracing computation as a means of improving both diagnosis and treatment: They're finding better ways of detecting and treating the cancer with robots and magnetic resonance imaging, figuring out how prostate cancer evolves, and homing in on the genes that regulate the disease.

Image Enhancement

Anant Madabhushi, PhD, wants to use computers to get a better picture of prostate cancer.

Doctors typically use manual ultrasound scans to guide their biopsy needles into patients' prostates or to implant radioactive seeds in the gland during a treatment called brachytherapy—an approach that can destroy tumors before they spread. But studies show that ultrasound-guided biopsy fails to detect prostate cancer in at least 20 percent of patients who have it.

MRI images, with their high resolution and excellent soft-tissue contrast, can do a better job of enabling both targeted biopsy and more narrowly focused treatment of prostate tumors. They could also fuel computer-aided detection and diagnostic algorithms, and be combined with other kinds of imaging data to improve computer-assisted surgical navigation and radiotherapy. But manually segmenting MRI images to identify a tumor's borders within the prostate is not currently standard practice, partially because it is difficult and time consuming. Segmentation algorithms would seem to offer a faster and better way of parsing MR data, but they aren't yet reliable enough for clinical use—especially since the scans themselves tend to be highly idiosyncratic,

with differences amongst scanners generating lots of variability in image appearance and quality.

Madabhushi, who is associate professor of biomedical engineering and director of the Center for Computational Imaging and Personalized Diagnostics at Case Western Reserve University, has therefore been trying to draw more algorithm developers into the fray. In 2012, he was the lead organizer for the Prostate MR Image Segmentation challenge (aka PROMISE12), which had 11 teams from industry and academia compete to see whose algorithms—some fully automated, some highly interactive—could best segment scores of MR images provided by imaging centers in the United

States and Europe.

States and Europe. In some particularly tricky cases, the interactive algorithms, which relied on human users to digitally paint large parts of the images, outperformed their fully automated counterparts. But much to Madabhushi's surprise, the two algorithms that scored best overall were completely automated. Both employed active appearance models, which use large data sets to construct models of the shape and appearance of the prostate; and on at least some measures, both managed to outdo even the inexperienced human annotator. (In a second challenge organized in 2013, Madabhushi's own team from Case Western won with a semi-automated algorithm.)


Madabhushi hopes that the algorithms will continue to improve, but he doesn't think they'll ever completely replace expert human annotators. "You have the autopilot, and you can go on cruise. But you still want the pilot there when you're taking off or landing," he says, pointing out that the inexperienced annotator still outshone most of the algorithms in the 2012 challenge.

States and Europe.


After tuning their algorithms on a training data set that included a reference standard comprised of manual segmentations by expert human annotators, the teams downloaded and segmented one test set that did not include such a benchmark, and were handed yet another at a live workshop in Nice, France. The organizers ranked the algorithms based on how closely they approached the reference standard, and on how well they did compared to an inexpe-

Madabhushi hopes that the algorithms will continue to improve, but he doesn't think they'll ever completely replace expert human annotators. "You have the autopilot, and you can go on cruise. But you still want the pilot there when you're taking off or landing," he says, pointing out that the inexperienced annotator still outshone most of the algorithms in the 2012 challenge. Which begs the question: what would happen if a whole bunch of different algorithms, all using different approaches, joined forced


Three sets of images representing three different cases from the PROMISE12 challenge. Different colors are used to illustrate prostate segmentations by different teams; on average, case 3 (images a, b, and c) had the best algorithm scores, case 10 (images d, e, and f) had reasonable scores, and case 25 (images g, h, and i) had the worst scores. The two algorithms with the best overall scores in the contest were fully automated; but in case 25, a large area of fat around the gland caused most of the algorithms to make large errors in prostate volume, and a more interactive algorithm did best. Reprinted from Litjens G, Evaluation of prostate segmentation algorithms for MRI: The PROMISE12 challenge, Medical Image Analysis 18:359-373 (2014), with permission from Elsevier.




(a)




(b)




(c)




(d)




(e)




(f)



(g)



(h)



(i)

with a whole bunch of inexperienced humans?

Madabhushi describes one possible scenario in which large numbers of non-experts (e.g., high school and college students) segment batches of MR images, and automated algorithms check their annotations for accuracy. The algorithms could identify the best of the inexperienced human annotators for future reference; segment the harder cases themselves; and send the ones that even they can't handle on to expert human annotators, who could then pass their own properly segmented images back down to the algorithms for training purposes. Whether such a system would adequately preserve patient privacy or gain FDA approval remains to be seen. But the ensuing virtuous circle of data-sharing and analysis amongst experts, amateurs, and algorithms could, says Madabhushi, yield something "more enriched, and perhaps more accurate, than any individual source of information."

Open-Source Revolution

If Madabhushi envisions a future where crowdsourcing and computer automation enhance prostate cancer diagnosis and treatment, **Gabor Fichtinger, PhD**, a professor in the School of Computing at Queen's University in Kingston, Ontario, and adjunct professor of computer science and radiology at Johns Hopkins University, sees one dominated by open-source software.

In collaboration with the National Alliance for Medical Imaging Computing (NA-MIC), Fichtinger led the development of Prostate Nav, a prostate-specific module within 3D Slicer (www.slicer.org), the Alliance's free, open-source platform for visualization and image analysis. Prostate Nav allows researchers and clinicians to use medical robots to perform biopsies and brachytherapy. It can create an interface between a robot and the rest of the equipment (scanners, navigation systems) in the medical suite; register the robot to the same coordinate system that 3D Slicer uses to pinpoint the location of any other tracked surgical instrument; issue commands to the device; and even cause an animated model of it to appear on the operator's screen. Fichtinger and his colleagues have used Prostate Nav to support an entire family of MR-compatible robots that can function inside the bore of an MRI scanner, guiding needles into patients with far greater accuracy than the standard manual ultrasound-guided method can achieve.

Now Fichtinger is exploiting open-source software to more quickly and effi-

ciently build systems that combine robots and tracked surgical tools with preoperative MRI, intraoperative ultrasound, and other imaging modalities. For example, he and his colleagues recently added color stereo optical imaging to their software platform to accommodate researchers who are interested in laparoscopic prostate surgery. And last year, Fichtinger's team developed a custom system for MRI- and ultrasound-guided prostate intervention research at Harvard's Brigham and Women's Hospital in just eight weeks—and it only took that long, Fichtinger says, because of some "funky requests" by the clinicians, such as simultaneous image acquisition from multiple ultrasound transducers. Within the next decade, Fichtinger predicts that the technology will have matured to the point where it will be possible to derive working, clinical-grade applications from open-source platforms such as 3D Slicer in a matter of days—though getting FDA approval for them "will still take a good bit of time."


An Ecological Approach

Identifying and treating tumors in the prostate is critical. But prostate cancer becomes truly lethal when it migrates elsewhere. **David Basanta, PhD**, and **Arturo Araujo, PhD**, in the Integrated Mathematical Oncology department at the Moffitt Cancer Center in Tampa, Florida, have therefore built a computational model that

combines agent-based techniques with conventional mathematical modeling methods to simulate how prostate cancer metastasizes to bone in order to better understand, and hopefully foil, the process.

Previously, Basanta used another hybrid model to investigate how the protein TGF-beta affects tumor growth. He has also employed evolutionary game theory to explain how interactions between prostate cancer cells, normal cells, and their shared microenvironment influence cancer progression, comparing tumor cells to invasive species that disrupt the ecosystem of healthy tissue. His latest work, carried out in conjunction with a group led by molecular biologist **Conor Lynch, PhD**, and reported in 2014 in the journal *Cancer Research*, builds on those earlier efforts, using a hybrid cellular automaton model to illustrate how metastatic prostate cancer cells are able to exploit elements of the bone ecosystem, including TGF-beta and another signaling molecule called RANKL, to their own advantage.

The agents in Basanta's model include not only prostate cancer cells, but also the osteoclasts and osteoblasts that break down and build up bone tissue during the course of normal bone maintenance. Partial differential equations, meanwhile, are used to mimic the production, diffusion, and decay of TGF-beta, RANKL, and other molecules that coordinate normal bone maintenance yet also facilitate the proliferation of cancer cells. In simulations



The ecosystem in a prostate to bone metastasis comprises several types of cancer cells interacting with other cellular populations such as osteoblasts, osteoclasts, osteocytes and stem cells. Tumor cells compete and cooperate for resources such as nutrients, space and growth factors. Reprinted with permission from Basanta D, Anderson A, Exploiting ecological principles to better understand cancer progression and treatment, Interface Focus, 3, 20130020 (2013).

that ran for 240 virtual days, Basanta's model demonstrated how prostate cancer cells manipulated levels of TGF-beta and RANKL to create a vicious cycle of aggressive tumor growth and abnormal bone formation and resorption. The model was also able to predict the efficacy of two types of drugs that are commonly used to slow the progress of bone metastasis, and offered some insight into how one of them—an anti-RANKL inhibitor—might be used more effectively in the clinic.

Basanta and Lynch are now testing the efficacy of TGF-beta inhibitors using both *in silico* and *in vivo* tools, and they have joined forces with several clinicians to develop a computationally and mathematically enhanced method of personalizing treatments for patients with metastatic prostate cancer. Basanta hopes to use models to predict how tumors with particular mutations might evolve and grow in response to different drugs, then use that information to optimize the sequence of treatments a patient receives “in order to reduce the tumor burden in the bone and, presumably, extend quality of life—and improve their chances of coming out of this alive.”

The Root of the Problem

Andrea Califano, PhD, professor of chemical systems biology and chair of the department of Systems Biology at Columbia University, is pursuing the same goals with a different set of computational tools. Ultimately, Califano wants to personalize cancer treatments by reconstructing the regulatory networks, or interactomes, that control different kinds of tumors. This approach would allow researchers to look beyond the bewildering array of genetic mutations that accompany the various tumor types and focus instead on the master regulators of the disease: those genes that are necessary for the survival of a given form of cancer. Because they rarely harbor genetic mutations, these master regulators cannot be identified through standard genetic sequencing. “But you can find them by analyzing these networks,” Califano says. And once found, they may be inhibited by existing drugs.

That was the case in a study that Califano and his colleagues, **Cory Abate-Shen, PhD**, and **Michael Shen, PhD**, recently published in *Cancer Cell*. They began by using an algorithm called ARACNe (Algorithm for the Reconstruction of Accurate Cellular Networks) to reconstruct two interactomes: one responsible for producing prostate cancer tumors in human beings,

and one responsible for producing them in mice. Reverse engineering each network required sifting through hundreds of thousands of possible interactions between thousands of transcription factors and their target genes. The team then ran a different algorithm to determine which cancer-related transcription factors controlled genetic programs that were conserved between mice and people, and were there-

By combining the genes discovered in the indolent tumor and master regulator studies, Califano and Abate-Shen hope to develop a comprehensive five-gene panel that can give prostate cancer sufferers a “complete report” on the aggressive potential of their tumors.

fore most likely to be significant.

Califano and his collaborators then used an algorithm called MARINa (Master Regulator Inference Algorithm) to identify the transcription factors that were most likely to induce the genetic signature observed in aggressive prostate tumors. The seven conserved master regulators that emerged were then computationally analyzed for poten-

tial synergistic interactions among themselves, and a single pair of synergistic master regulators—the genes FOXM1 and CENPF—were found to drive aggressive prostate cancer in both mice and humans. Silencing one gene slowed cancer growth in the mouse models; silencing both shut it down completely. And protein expression analysis of prostate tissue samples taken from more than 900 prostate cancer patients at Memorial Sloan-Kettering Cancer Center revealed that patients with elevated expression levels of both genes experienced by far the worst outcomes—including shortest time to metastasis, and death. Abate-Shen and Califano have already identified two drugs that can inhibit these master regulators.

In addition to identifying the master regulators that induce aggressive prostate cancer, Califano and his colleagues have found a cluster of genes that can be used to predict whether tumors that seem indolent, or slow-growing, are destined to stay that way. It's a crucial task, since overtreatment of prostate cancer is both costly and potentially risky, yet the only thing worse than unnecessarily treating a person with an indolent tumor is failing to treat one whose tumor only appears to be so.

Califano's indolent tumor work, published last year in *Science Translational Medicine*, began with a manually curated list of 377 genes associated with the tumor-inhibiting processes of cellular aging and senescence. He and colleagues used Gene Set Enrichment Analysis (GSEA), which ranks genes on a spectrum from most to least expressed, to identify 17 senescence genes that were over-expressed in indolent mouse tumors and under-expressed in aggressive human ones; then applied a decision-tree algorithm to prune them down to a trio of genes with the greatest predictive power. All three were validated in the lab by Abate-Shen and were found to be under-expressed at the protein level in biopsies taken from prostate cancer patients whose tumors initially appeared to be indolent, but nonetheless became aggressive.

By combining the genes discovered in the indolent tumor and master regulator studies, Califano and Abate-Shen hope to develop a comprehensive five-gene panel that can give prostate cancer sufferers a “complete report” on the aggressive potential of their tumors.

More accurate diagnoses, new biomarkers, and improved therapies: add it all up, and computation just might make prostate cancer's numbers look a little less menacing after all. □