


WINNING ESSAYS

ENCOURAGING COMMUNICATION THROUGH AN ANNUAL WRITING CONTEST

WINNER



by Kenley Pelzer

CAN PEELING AN ONION CURE CANCER?

AWARDING COMMUNICATION

The DOE CSGF launched an annual essay contest in 2005 to give current and former fellows an opportunity to write about their work with a broader, non-technical audience in mind. The competition encourages better communication of computational science and engineering and its value to society to non-expert audiences.

In addition to recognition and a cash prize, the winners receive the opportunity to work with a professional science writer to critique and copy-edit their essays. The latest winning essays are published here.

For more information on the essay contest, visit www.krellinst.org/csgf.

WITH NO OFFENSE to the quantum theorist in the room,” the professor said, casting me a quick look and smiling gently, “you should never, ever deal with quantum effects unless you have to.” As an eager young graduate student in theoretical quantum chemistry, I was enthralled by the wonderful and bizarre laws Einstein and other great scientists had discovered. Yet I knew the professor had a point: If you unnecessarily involve these strange principles in your approach to a scientific problem, it will be harder — or perhaps impossible — to solve it.

In the world of quantum mechanics, objects behave in very peculiar ways. An object can pass through walls. It can be in two places at the same time. No matter how hard you try, you can never be sure where it is or where it is going. Fortunately for scientists, these quantum effects are only important for the smallest particles and can be ignored in many situations.

However, it is absolutely crucial to consider quantum theory in some cases, because its messy, complicated rules can help us understand diseases — and how to cure them.


The human body contains a multitude of protein molecules that play a central role in whether we stay healthy or get sick. Each protein contains thousands of electrons — tiny, negatively charged particles that are attracted to positive charges. These attractions influence whether a protein “sticks” to another molecule.

This biological stickiness is especially important in treating cancer. In some types of cancer cells, there are protein molecules with a region called an “active site.” When certain molecules attach to the active site, they trigger a cascade of events that leads to disease. But if a drug can target a particular protein and stick to its active site, it can block other molecules from binding and prevent this dangerous spiral. If a drug can’t stick to the protein — if it falls off and circulates in the bloodstream — it may just float around doing nothing or, worse yet, cause toxic side effects.

To predict which drugs will attach firmly to an active site, scientists must understand how protein molecules and their electrons behave. But because electrons are so small, they obey all of the complicated laws of quantum mechanics.



This image shows a molecular orbital calculated using a quantum mechanics program. By calculating the shapes and orientations of the molecular orbitals, which may contain electrons, scientists obtain detailed information on the distribution of the electrons in space.



So how can anyone possibly study a protein with thousands of these slippery, mysterious little particles?


Since electrons are too small to see through a microscope, chemists predict their behavior with sophisticated computer simulations. Not only do these simulations have the power to study electrons, they also have economic advantages: Once the software has been developed, it costs virtually nothing to use it to study large numbers of potential drugs. Given the massive costs involved in pharmaceutical research, this advantage has important implications in the search for effective medicines.

The bad news is that — thanks to the influence of quantum properties — these simulations may take weeks to generate information about just a few electrons. Since a protein has thousands, accurately predicting the distribution of electrons could take decades. Not very helpful when cancer patients are hoping for a new drug *now*.

This is where the idea of peeling an onion comes to the rescue. With the aptly named ONIOM method (“Our own N-layered Integrated molecular Orbital molecular Mechanics”) the protein molecule is divided into a chosen number (N) of layers. The program treats each with a different simulation technique, carefully selected based on what information is needed. When highly accurate information on the electronic structure is needed, the program uses a “molecular orbital” method that rigorously incorporates quantum theory. The details of quantum theory are so important for these molecular orbital calculations that



Pictured here is dihydrofolate reductase, a protein whose function is crucial for cell proliferation. Drugs that bind to the active site of this protein and inhibit its activity are used to stop the rapid proliferation of cancer cells. The yellow shaded region surrounds the active site.



they are usually referred to as “quantum mechanics” calculations.

When a lower level of accuracy is acceptable for a particular layer, the program may use a less advanced (and less time-consuming) quantum mechanics method. In cases where even more approximate information is sufficient, the computer performs a “molecular mechanics” calculation. Because molecular mechanics methods are less precise and incorporate less quantum mechanical theory, they’re fast — really fast. So by applying molecular mechanics and working with information that’s a little less accurate, the simulations gain a lot of speed.

To understand how the ONIOM approach can help us answer important questions, consider the example of a drug that needs to bind to a particular protein. At the protein’s active site — the onion’s center — the computer uses an advanced quantum mechanics method to calculate the behavior of electrons and predict whether a drug will “stick.”

Then the computer must simulate the next layer of the onion — or rather, the next layer of the protein — which wraps around the active site but isn’t part of the active site. Since this layer doesn’t actually stick to the drug of interest, we don’t need to worry quite as much about its electrons. On the other hand, we can’t totally ignore this layer, because changes in its shape or

electron distribution might affect the active site’s behavior. So scientists compromise and use a less precise method: either a quantum mechanics method that is a little less accurate (and hence faster), or a molecular mechanics method.

The farther we are from the active site, the less we need detailed information about electrons and their quantum mechanical behavior. So we use a less precise simulation technique with each successive layer. By using faster techniques to treat the outer layers of the “onion,” it’s possible to predict how a large protein will interact with a particular drug. And fortunately for patients waiting for the next new medicine, the answers can be obtained in days or weeks — not decades.

By peeling apart a protein as though it were an onion, scientists follow my professor’s advice to never, ever deal with quantum effects unless you have to. The beauty of this approach is that for many physiological proteins, neither quantum mechanics nor molecular mechanics alone could effectively answer our questions. Fortunately, collaboration between scientists who specialize in each method has led to ingenious hybrid programs that can contribute great insights to drug design. And then with the press of a key, a computer can guide us in the urgent quest to develop life-saving medicines.

HONORABLE MENTION



by Hayes Stripling

ON THE QUANTIFICATION OF “MAYBE”: A NICHE FOR COMPUTATION

CONSIDER PERHAPS THE MOST

pressing question facing the global community in this and coming decades: “Is human activity adversely affecting the environment and global climate?” A brief skimming of newspaper headlines, scholarly journals and political debates reveals that credible responses to this “yes/no” question all have a flavor of “probably,” “maybe” or “probably not.” In fact, most scientifically based stances concede we cannot provide a definitive “yes/no” answer without many more years’ worth of data, investigation and discovery. So the correct answer today is just “maybe,” with some evidence and proponents on either side.

In response to scientific and political pressure for a more definitive answer, researchers have focused on this reformulation of the question: “To what extent does human activity affect the environment and global climate, and how do we manage the risks of this activity as the scientific investigation continues?” This is not posed as a “yes/no” question; instead, it calls for a quantification of “maybe” — an estimation of uncertainty in our response. Answering the question also requires a multidisciplinary effort — involving experts in science, engineering and policy development — that will evolve as we gain understanding.

These kinds of questions have led to an exploding demand for knowledge. In reaction, researchers are expanding the scientific method, which is founded on the long-lived pillars of theory and experiment, to explore new domains and strategies to support decisions and conclusions. These include computation, an invaluable tool many now accept as the third pillar of

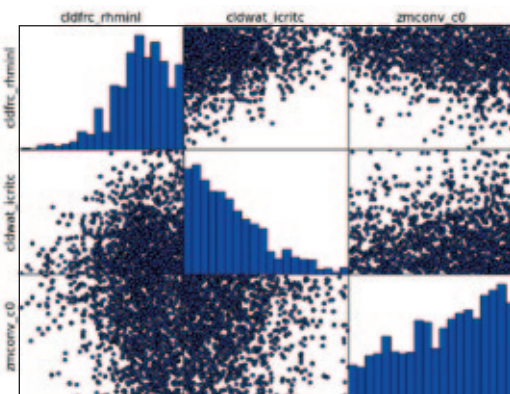
science. Using computational power to simulate a system or experiment is relatively inexpensive and has increasing potential to provide accurate, thorough results. Most importantly, computation also has surfaced as the most valuable tool for exploring uncertainty in “yes/no” questions.

The specific branch of computation that addresses this issue is uncertainty quantification (UQ). UQ — or the quantification of maybe — aims to comprehend how uncertainty in nature affects a particular quantity of interest (QOI). It then uses this understanding to make predictions or informed decisions about the quantity. Achieving this goal, however, is not so simple, for the QOIs we seek are complex results of multidisciplinary systems that are confounded by uncertain inputs and less-than-perfectly-understood physics.

To illustrate the utility of computation in uncertain systems, let’s consider a simple example: determining the forces a passenger experiences in a car crash. Imagine we can identify the five most important factors (such as speed or seat-belt use) contributing to our QOI but that we can’t know exactly what the settings or statuses of these inputs will be (that is, our inputs are uncertain). If we restrict each input to one of two settings (for example, speed is either fast or slow) and wish to test each possible combination of inputs, we would require the design, construction, crashing and analysis of 2^5 (32) test cars. This may be an acceptable number of experiments, depending on economic, administrative and/or political constraints. But five inputs probably are too few. A more realistic approach may consider 15 inputs — that’s

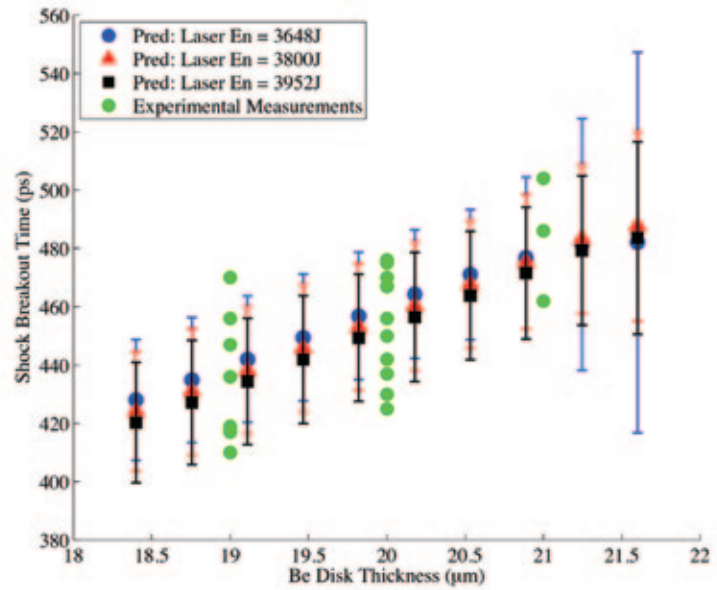


Each column of the figure illustrates a calibrated or posterior distribution of a single uncertain input to a massive climate model. The bar charts (histograms) represent the single-variable distribution and the scatter plots illustrate how the calibration of one variable correlates to the other two variables. For example, these results suggest larger values of *cldfrc_rhminl* and smaller values of *cldwat_icritc* will improve the accuracy and reduce uncertainty in future model runs.





The error bars represent predictions (with uncertainty) of the time required for laser energy to ablate a beryllium disk. We use existing experimental data to “tune” our uncertain parameters and reduce the magnitude of the predictive uncertainty. The analysis shows that between experimental results, the error bars are comparable to the experimental variability. Predictions at extrapolated (larger) disk thicknesses, however, are less informed by the experimental results and have larger uncertainty.



2^{15} , or 32,768 experiments — certainly a number too expensive and time-consuming to consider.

Computation has proven extremely valuable in such cases, in which experimentation is economically or politically infeasible. What if we can afford only five crash tests? We may approach the problem computationally instead. The computer model would solve equations that govern (or simulate) the physics that take place during a crash. We could design the model to accept any number of adjustable inputs that contribute to the final QOI (in this case, the forces a passenger experiences). Inputs could span a range instead of adhering to a binary choice, like fast or slow speed. It would be common, given sufficient computer resources, to run hundreds or thousands of “computer experiments” corresponding to hundreds or thousands of input combinations. The final result is a distribution of our QOI, allowing us to make an informed statement of the form, “We are 99 percent certain that the passenger G-force will be below the injury threshold in 99.9 percent of vehicle collisions.” Such a statement is much more useful to a policy-making board or

consumer than results from only a handful of experiments.

Cutting-edge simulations running on some of the largest computer architectures in the world use this kind of approach to address society’s toughest challenges. But in some cases, the problems are becoming too large and complex. For example, models designed to describe the long-term global climate predict thousands of quantities of interest from thousands of uncertain inputs. To help handle data of this magnitude we employ sensitivity analysis, a sub-discipline of uncertainty quantification. By sampling results from previous computer experiments we can determine which outputs are most sensitive to which inputs. For example, if we find that the computed prediction of global surface temperature isn’t sensitive to lunar cycles but is highly sensitive to cloud coverage, we can adapt our sampling strategies to explore cloud coverage more thoroughly and hold off on varying the lunar effects. This will give us a more precise understanding of global surface temperature behavior and use fewer evaluations of the (potentially time-consuming) computer model. It’s also

something we could never determine using physical experiments alone.

Of course, it’s impossible for computation to ever make experimentation a moot practice. We must compare experimental measurements with computational results to ensure our models are valid. Further, experimentation’s maturity as a technique and the ability to witness its physical results with your own eyes makes it more credible to humans than computational number crunching. For example, let’s return to our car crash case: Would you be more inclined to believe a vehicle’s safety report based on five physical experiments using real crash-test dummies or 1,000 simulated smash-ups in which you can’t see, hear or feel the impact? Would you be more inclined to believe the computer if it exactly predicted the results of the five physical crashes we could afford? We can never abandon experimentation — but we can leverage validated computer models to develop and explore scientific concepts that will guide our society’s policies in the face of the great challenges ahead.