


Dogs, Doses AND Devices:

The FDA's Ambitious Plans for Computational Modeling

By Alexander Gelfand





What role does computational modeling play at the United States Food and Drug Administration (FDA)?

If you ask **Paul Watkins, MD**, director of the Hamner—University of North Carolina Institute for Drug Safety Sciences, you'll hear a story about a dog.

Once upon a time, Watkins' son's dog, Bailey, swallowed four tablets of Motrin. While that dose of ibuprofen would almost surely have been harmless to a person, it landed Bailey in an animal hospital with acute liver and kidney failure. (He survived.) The moral of this story? "Had Motrin been tested in dogs," Watkins says, "it probably never would have gone into man."

It's a personal anecdote with a larger point: computational modeling can help fill gaps in how we develop and review new drugs and devices.

Watkins himself is using computer models to fill one of those gaps—species-specific variation in drug sensitivity, the very problem that affected his son's dog—as part of a collaborative project with the FDA. "By going from one compound to another that has different mechanisms of toxicity in different species, we're building out a model that we believe will be applicable to virtually any compound you bring in, even without giving the actual new drug to a living animal or man," Watkins says.

It's an ambitious vision, but one that is entirely consistent with the FDA's growing use of computational modeling as part of research and review. The ultimate objective is to provide greater consistency and predictability in the development and review of drugs and devices, thereby saving time and money while improving safety and efficacy.

Jogarao Gobburu, PhD, director of the division of pharmacometrics at the FDA's Center for Drug Evaluation Research (CDER), which regulates

over-the-counter and prescription drugs, says that the number of new drug applications that employ modeling and simulations has increased six fold in the past 10 years. And researchers at the Center for Devices and Radiological Health (CDRH), which regulates medical devices, initiated the FDA's purchase of high-performance computing facilities in order to validate the many complex simulations submitted by device manufacturers, according to **Brian Fitzgerald**, deputy director of the division of electrical and software engineering at the CDRH.


Both centers use modeling to further the basic research they do to inform regulatory decisions, and both are confronted with the need to verify the simulations that industry includes in its applications for the approval of new drugs and devices. But the CDER is already pushing the envelope farther: It is using models to extend clinical findings, guide additional clinical testing, determine appropriate doses, decide wording on drug labels and, yes, solve the problem of Watkins' dog: species-specific dosing.

DEVICES

At CDRH, modeling does not play as prominent or pervasive a role yet as it does at the CDER. For CDRH, modeling remains firmly in the Office of Science and Engineering Laboratories (OSEL), where it is used to further basic research. The modeling results inform the work of the Office of Device Evaluation (ODE), which clears medical devices for clinical trials and marketing. However, the ODE does not actually run simulations for

evaluating a specific device.

Instead, reviewers like **Tina Morrison, PhD**, co-principal investigator on a project to promote computational modeling in the design and evaluation of cardiovascular devices, analyze the models used by the device manufacturers.

 They also do independent research to further their understanding of the modeling process, and gather boundary condition data from *in vitro* testing and clinical imaging that can be used to refine and improve upon the models used by industry and at the agency.

Validating Submitted Simulations

Modeling is ubiquitous in the medical device industry. Manufacturers rely heavily on techniques such as finite element analysis and computational fluid dynamics to model the behavior of hip implants, heart pumps and other devices under extreme virtual conditions, identifying scenarios for testing in clinical trials while saving considerable amounts of time and money. In a situation where a single stent can cost upwards of \$3000, going virtual has its advantages. "Imagine," Morrison says, "if you wanted to take two dozen of those just to determine on the benchtop what load would cause it to deform the most."

Yet before they can trust the simulations that developers include with their applications for device approval, Morrison and her colleagues must validate and verify their models, ensuring that the companies are using the right equations and running them correctly. That means having sponsors perform validation studies in which they justify the parameters they have set and compare the results of their simulations with data from bench tests or clinical studies. "We don't expect them to line up exactly," Morrison says, "but we do expect them to be within an acceptable range of error."


In-house Computational Modeling

Understanding why different manufacturers generate different simulated results for similar devices is also important. For example, Morrison estimates that some 40 companies currently manufacture stents, devices used to open up diseased, narrowed blood vessels. All of

them use computational modeling to simulate the ways in which these devices deform *in vivo* when blood is pulsing through them, and their simulations can differ markedly. "Are the outcomes different because the designs of the stents are different, or because the models they use are different?" Morrison asks.

To answer that question, the agency

"We would like to be able to use computational modeling more as a tool for evaluation and not just have the sponsor use it as a tool for development," Morrison says.



is sponsoring a round-robin investigation of the computational flow models that are used to predict potential blood damage from cardiovascular devices. (Areas of low blood flow can cause clots, while areas of high flow can damage blood cells.) By comparing multiple simulations of two devices—one a simple nozzle, the other a more complex ventricular assist device—to one another and to physical bench tests, the agency hopes to glean information that will help it standardize the models used by industry. Thus far, the answers have been illuminating.

“What we learned is that people interpret problems differently,” Morrison says. “When setting up the computational models, people make different choices regarding materials, forces, and boundary conditions. So their answers may be very different.” And those answers don’t just differ from simulation to simulation. “Even expert modelers got very different results than what we saw in the physical experiments.” The study has helped reviewers understand some of the basic challenges involved in modeling and

armed them with fresh questions for device sponsors.

At the same time, Morrison and her colleagues are collecting boundary condition data that they hope will lead to a set of “gold standard” models that can be distributed to industry in open-source fashion. With such models in hand, a device manufacturer could virtually implant an artificial hip joint in a simulated *in vivo* environment. If the device failed to behave as expected within the FDA’s own trusted model, the regulator could ask the manufacturer to do further bench testing or even redesign the device. “We would like to be able to use computational modeling more as a tool for evaluation and not just have the sponsor use it as a tool for development,” Morrison says.

It’s an ambitious goal, and one that will take time to achieve. Even within the realm of cardiovascular devices, where researchers rely on modeling and simulations to advance both device design and basic science, much work remains to be done.

For example, **Richard Gray, PhD,**


a biomedical engineer at OSEL who studies electrical arrhythmias that lead to sudden cardiac death—a topic of significance for the evaluation of defibrillators—says that the heart “is the most advanced, mathematically understood organ in the body,” and the one for which the most sophisticated models have been constructed to date.

Yet while researchers have a clear picture of the 3-D geometry and structure of the heart, they still lack a thorough understanding of the cellular dynamics that drive fibrillation and defibrillation. “We’re getting very close to being able to model the effect of an electrical field on the heart,” Gray says, but predicting whether an electrical shock administered during fibrillation will cause the organ to defibrillate is “a whole other ballgame.”

Nonetheless, progress is being made. For the past three years, the FDA has co-sponsored an annual workshop in conjunction with the National Science Foundation and the National Heart, Lung, and Blood Institute to pool cardiovascular modeling expertise

“We’re getting very close to being able to model the effect of an electrical field on the heart,” Gray says, but predicting whether an electrical shock administered during fibrillation will cause the organ to defibrillate is “a whole other ballgame.”

Image courtesy of Richard Gray, FDA.



from industry, academia, and government. And Morrison recently made a presentation on computational modeling to the Office of Science and Technology Policy in the Executive Office of the President that she anticipates will lead to funding specifically for the use of modeling in device evaluation.

DRUGS

Modeling to Aid Decision-Making

But to see the future of computational modeling at the FDA, you need look no further than the Center for Drug Evaluation Research.

“Our focus is not on the models,” says Gobburu. “Our focus is on deci-

the relationships between clinical trial characteristics such as exclusion and inclusion criteria. These efforts fall under the broad heading of pharmacometrics, or quantitative pharmacology, which seeks to interpret and analyze pharmacology in a quantitative fashion by integrating data pulled from such diverse sources as clinical trials, drug chemistry, and biology. Pharmacometric models are already able to simulate the relationships between drug exposure (or pharmacokinetics); drug response (pharmacodynamics); and individual patient characteristics, teasing out the factors that contribute to both desired and undesired effects. And as the models continue to improve, the agency is able to use them in increasingly sophisticated ways.

During the H1N1 influenza pandemic of 2009, for example, the FDA used

tease inhibitors boceprevir and telaprevir—for patient populations that were not specifically covered in the clinical trials conducted by their manufacturer. “We approved dosing for patients who were not directly studied in the registration trials based on analysis across different trials and biological reasoning,” Gobburu says. (The agency is currently developing an Antiviral Information Management System (AIMS) comprising an automated modeling and trial simulation tool linked to a database of hepatitis C and HIV trials.)

And in July 2011, the FDA used a similar pharmacometric “bridging” approach to approve the anti-epilepsy drug Topamax (topiramate) for monotherapy among patients aged 2-10, even though the drug was never tested for such use on that age group. In monotherapy, a single

“Our focus is not on the models,” says Gobburu. “Our focus is on decisions—either regulatory decisions, or drug development decisions.”

sions—either regulatory decisions, or drug development decisions.”

Those decisions, which have to do with approval, labeling, and the design of clinical trials, depend more and more on computational modeling. The center has long performed a variety of in-house modeling in support of the regulatory review process, simulating diseases, modeling drug characteristics, and analyzing

pharmacometric modeling to project and approve a safe pediatric dose of intravenous peramivir, an experimental antiviral drug, despite the fact that the drug’s manufacturer had never actually tested it on children. (Later trial results proved “very close” to the agency’s simulated outcomes.)

More recently, the agency approved two new hepatitis C drugs—the pro-

drug is used to treat a condition. However, because monotherapy trials use placebos as controls, they cannot be conducted ethically on young children who suffer from serious conditions such as epilepsy, given the risk of injury or even death. To circumvent this problem, researchers modeled data derived from adults and pediatric patients who were given the drug as adjunctive therapy, and for patients aged up to 16 years as monotherapy, to predict that Topamax would also be effective as the single mode of treatment in much younger children. They then ran simulations based on exposure ranges among older patients to determine safe dosing guidelines for the 2-10 age range.

In each of these cases, computational modeling safely curtailed the need for several additional years of clinical trials—time that translates into patients’ lives and billions of dollars.

The CDER has also benefited greatly from the advent of physiology-based pharmacokinetic (PBPK) models that take into account not only traditional inputs such as age, sex, and disease state, but also what **Shiew-Mei Huang, PhD**, acting director of the office of clinical pharmacology (OCP), calls “micro factors” like turnover rates for drug-metabolizing enzymes and transporters. In a recent paper in *Clinical Pharmacology & Therapeutics*, a journal in the Nature Publishing Group, Huang and her col-





leagues described four cases between the years 2008 and 2010 in which the FDA used PBPK modeling to make decisions regarding the design of clinical trials and the language in drug labels.

In one case, the agency's simulations of the prostate cancer drug cabazitaxel confirmed the need for further *in vivo* studies and also predicted increased sensitivity among individuals suffering from hepatic impairment, allowing the agency to guide the design of a study involving patients with impaired liver function—in effect, telling the drug sponsor to perform a clinical study for a particular population and recommending appropriate dosing levels. On the other hand, simulation results suggested that certain drug interaction studies would be unnecessary. As Huang points out, modeling exercises like these can help the FDA evaluate drug applications critically and make recommendations concerning the design of clinical studies.

Collaborative Research to Develop and Verify Models

As at the CDRH, simulating tests and trials that have been performed in the real world is crucial to verifying the models. For example, the FDA is currently engaged in a collaborative research agreement with Archimedes, a Bay Area firm whose software tracks what happens to virtual patients as they make their way through a simulated healthcare system. The project is developing a computer model to reproduce the results of the Sibutramine Cardiovascular Outcome Trial (SCOUT), which showed that patients who received the weight-loss drug sibutramine were 16 percent more likely to suffer serious cardiac events such as heart attack and stroke than those who received a placebo. If Archimedes is successful, its model will also be used to virtually extend SCOUT 10 years into the future.

“This is really a pilot project to see the utility of their approach in terms of thinking about how we model our clinical trials, and how we make inference from findings in our clinical trials,” says **Darrell Abernethy, MD, PhD**, associate director for drug safety in the OCP. Although “quite pleased” with the results so far, Abernethy stresses the need to vet every simulation performed by an outside party. “We need to be able to reproduce the findings based on the inputs that the particular company used,” he says. “If we can't, that's a big problem.”

Which brings us back to dogs (and other animals). As Watkins' family discovered, different species can react very differently to the same drugs. Enter the Computer Models for Human DILI (drug-induced liver injury) Project, or DILI-sim, a collaborative venture between the FDA, the Hamner Institutes, and the pharmaceutical industry that seeks to predict species-specific variations in liver toxicity in order to improve dosing decisions. Watkins, who also sits on a subcommittee of the FDA Scientific Advisory Board that reviews pharmacovigilance procedures at the CDER, notes that many drugs that are found to be safe in animals cause serious liver toxicity (termed “hepatotoxicity”) as soon as they go into people during clinical trials. Conversely, some potentially useful drugs induce liver damage in animals and are removed from development before they are ever tested in human subjects—a matter of serious concern to regulators and pharmaceutical companies alike.

“It's said that when drugs first go into people, they fail for safety and efficacy equally, but that's misleading. Often, companies don't try high enough doses because they didn't have adequate safety margins in animals to get approval for higher doses in man,” Watkins says, adding that “hepatotoxicity is the organ toxicity that's now most problematic, and where there's been almost no progress for 30 years. This is an area that's really ripe for major advancements.”

The DILI-sim project aims to understand how species differences affect dosing. The results will help guide initial tests in humans.

So far, Watkins and his colleagues have accurately predicted species differences for dosing acetaminophen and the antihistamine methapyrilene, and they are currently working on furosemide, a powerful diuretic. The model they use incorporates the most up-to-date understanding of pharmacokinetics, reactive metabolite formation, and the like, along with *in vitro* data for mice, rats, dogs, and humans. Within a year,

Watkins expects that drug developers and the FDA will be able to use DILI-sim to determine dosing for human subjects. Eventually, he anticipates that the model will be capable of predicting even the rarest toxicities, and will play a role not only in Phase I trials—when the drug is first tested in humans and assessed for safety—but also during the initial approval process.

But the DILI-sim project, which was



During the H1N1 influenza pandemic of 2009, for example, the FDA used pharmacometric modeling to project and approve a safe pediatric dose of intravenous peramivir, an experimental antiviral drug, despite the fact that the drug's manufacturer had never actually tested it on children.

initiated through a collaborative research agreement with Entelos, a Boston-based company that develops platforms for *in silico* biomedical testing, has hit a few speed bumps. There were problems with Entelos' approach, Abernethy says, and their model ran on a proprietary platform. At Hamner, DILI-sim is being transferred to the widely used MATLAB format and will

be entirely in the public domain, thereby making it far more accessible and useful to the agency. (Entelos continues

Yet even if that day arrives, computational modeling may never be entirely unproblematic for the agency. As

significant the interaction would be. As a result, the label was left more vague than some outside experts and industry representatives might have liked.

“The comment we got was, ‘You haven’t gone deep enough,’” Huang says. But in the absence of an actual clinical trial by the sponsor, that was as deep as the regulator could go. “The key point is, they did not do a study. So the discussion point is, how comfortable are we with those numbers?” In this case, the answer was clear: comfortable enough to make a general statement, but not enough to go any further.

That balance between the utility of computational modeling and its inherent limitations lies at the heart of the nuanced attitude that many at the FDA hold toward this powerful tool: an attitude that blends caution over the assumptions and data that power the models with a growing enthusiasm for the promise that the models hold, and the benefits that have already accrued from their use.

“We know our limitations,” says Gobburu. “But we also know the opportunities.” □

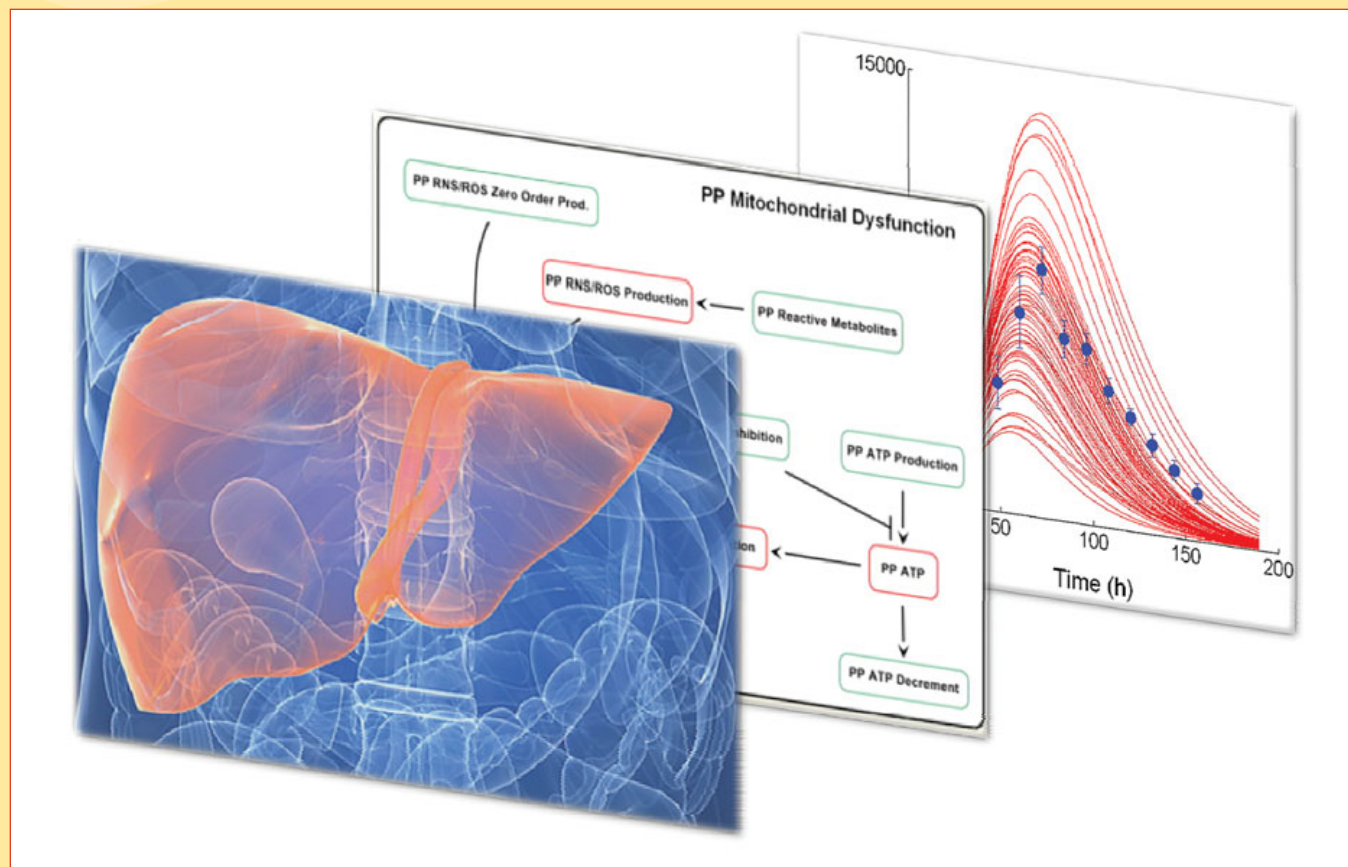
Within a year, Watkins expects that drug developers and the FDA will be able to use DILI-sim to determine dosing for human subjects.

to work with the FDA on a separate project involving cardiovascular drugs.)

A Nuanced Balance of Modeling’s Utility and Limitations

Abernethy and his colleagues over at the CDRH all dream of a day when modeling will become more transparent; a day when government and academic scientists will enjoy unfettered access to both data and code, and in which the verification and validation of simulated phenomena, from drug-drug interactions to arterial stents, will be (pardon the pun) virtually hassle-free.

Abernethy points out, models can fail due to inadequate inputs or incorrect assumptions; and as long as biomedical science remains a work in progress, both of those limitations will apply. In one of the cases described in Huang’s paper, the agency was able to determine that the pulmonary arterial hypertension drug sildenafil would interact less strongly with the anti-HIV drug ritonavir when administered intravenously rather than orally. But due to gaps in the *in vitro* data that were available and in some of the basic assumptions underlying the model itself, the agency could not specify exactly how *much* less signif-



The DILI-sim project models species-specific differences in drug induced liver injury. The graph at the right shows (in blue) combined data of patients’ alanine transaminase (ALT) levels in response to acetaminophen (Tylenol), a sign of liver toxicity; and (in red) several simulations using DILI-sim’s methodology for creating heterogeneous populations. The blue data comes from Schiodt, FV, et al., *Temporal Profile of Total, Bound, and Free Gc-Globulin After Acetaminophen Overdose, Liver Transpl* 7(8):732-8 (2001).