AN AUDIENCE WITH...

Karen Akinsanya


Credit: Schrödinger

Computational chemistry is already embedded in the drug discovery process. Schrödinger — a company that was founded more than 30 years ago to develop chemical simulation software for biopharmaceutical partners — believes that it should be more foundational still. Having co-founded several biotechs in the past decade, including Nimbus Therapeutics and Morphic Therapeutics, Schrödinger launched its own drug discovery pipeline in 2018 to expand this model. Heading up that effort is Schrödinger Chief Biomedical Scientist Karen Akinsanya. A pharmacologist by training, Akinsanya has more than 20 years industry experience working at the bench, the bedside and then in the boardroom. She now goes back to her research roots, leading the screening of hundreds of billions of compounds against targets of interest. She spoke with **Asher Mullard** about Schrödinger's physics-based approach to computational chemistry, the bottlenecks in this approach, and the new opportunities it can open up.

Why is Schrödinger, historically a software company, expanding into drug discovery?
The company is 30 years old, and has spent a lot of that time creating a software platform that is based on physics-based approaches to drug discovery. Our approach leverages the power of quantum mechanics and molecular dynamics to analyse molecular interactions at an atomistic level and predict properties such as solubility and binding affinity with a high degree of accuracy using computation. Importantly, it does this without the need for a training set, as is required with AI approaches.

About 10 years ago our CEO Ramy Farid realized that unless we were essentially involved in drug discovery, we couldn't really know what problems we were trying to solve with this software. We couldn't create this software solution in a vacuum; we had to be in the trenches, figuring out some of these difficult drug design challenges. And that is where the idea for Nimbus Therapeutics [a company co-founded by Schrödinger in 2009] came from, as a means of exploring how we could use these methods throughout the course of drug discovery. That hasn't changed, and our software is continually being improved.

But, if we're going to go out and perfect this software, we also want to show that it's not just useful in a follow-on mode — finding new drugs for validated targets — but also for working on novel targets. So this is all a natural extension of the work that we started with Nimbus. We're working with leading lights who have unique structural insights and running these drug discovery programmes in collaboration with those folks, but also increasingly building internally a capability to go after some challenging targets that we think have high potential.

At the R&D interface, we can work on some of the most challenging projects from a structural biology point of view. And what we're finding with our first five programmes is that there are breakthroughs when we try to solve something like selectivity. We've figured out how to use the software, for example, to ideate not just chemical space, but also what happens when you change an amino acid in the protein using the physics-based methods. And, if you change atoms in both the protein and the ligand to see what happens to binding, you can learn about what governs selectivity.

These breakthroughs are so important that we feel having our own drug discovery effort is critical to continually improving the software, and making sure that that improvement gets out to the rest of the community.

Computational chemistry, and your software, is already used throughout the industry. Do you plan to use it differently?
It's hard to give exact numbers, but let's say, for example, most pharmaceutical companies have about a 10:1 ratio of medicinal chemists to computational chemists. And most people use computational chemistry during the course of a programme that's already started in a more traditional way.

On our teams, we have more like a 1:1 or 2:1 ratio. We also use our software at its full scale and capacity. That means that we have unlimited access to these calculations. We use them from the very first step — which is target analysis, to think about how we are going to interrogate this protein. We don't usually start with high-throughput screens or medicinal chemistry, we start with an understanding of the structure, function, and how these physics-based methods might allow us to validate a computational assay. And our medicinal chemists and the computational chemists are also working closely together not just to ideate through chemical space, but to prioritize compounds before we even go into the wet testing and synthesis.

That's quite different, I think, than how most people use computational chemistry software.

What do you make of debates about the relative importance of medicinal chemistry versus computational chemistry?
I think we should use all of the methods available to us in our armamentarium. As far as I'm concerned, computational chemists along with traditional chemistry, working hand in hand, offer complementary capabilities that we should be leveraging to the greatest extent that we can to design the best molecules we can. I think the debate between the power of each of these methods is not quite the point. The point is how can we use this incredible power that we have access to, to shorten the length of time it takes to come up with the best-looking molecules?

The 10^8 compounds described in the *Chemical Abstracts Service* represent just a fraction of drug-like chemical space, which is estimated to be 10^{60} potential compounds. Why should we be limited by traditional approaches? As drug hunters we should

seek to explore as widely as we can to find molecules with ideal therapeutic properties. That's what our platform enables.

How much chemical space have you explored?
So far this year, we've ideated through something like 275 billion compounds to figure out how they interact with our targets, and if they have good physicochemical properties.

Molecules that score well in computational assays can lie in novel chemical space, needing long chemistry campaigns before wet testing. How do you handle this?
It's true. When you run these large chemical space enumerations, some of the compounds you end up looking at are quite unique. They don't sit in an existing library that you can pluck from a vial on a shelf to go ahead and do pharmacology testing with. Some of these compounds take quite some time to synthesize. And it takes a little bit of trust, and a leap of faith, that that molecule you ideated that is taking some time to synthesize is actually going to score in the wet lab the same way it scored on the computer. But we've learned to be patient.

In one of the very first programmes I worked on here, the structures looked quite unique. And the medicinal chemists looked at a structure and said "Yeah, that's not something I necessarily would have drawn. Let's see how it performs". And it came back from testing, and it was picomolar. And that's not something you usually get to start a programme with. We've found over and over again, across our own programmes and those of our collaborators, that when you get these predictions based on the physics of a compound binding at high affinity, those are usually spot on.

Obviously, this depends on the crystal structure, and how accurate that is. But if you have a good, accurate crystal structure, those scores tend to give you molecules that then perform very, very well.

Beyond the need for crystal structures, what other rate-limiting steps do you face?
I think that one of the things that biologists and pharmacologists have realized is that yes, it takes time to synthesize those initial molecules, but the design–make–test cycle can move extremely quickly once you get started. And so ultimately biology — not having access to the right biochemical, biophysical and cellular assays early on — can slow down a programme more than the access to the molecules. So we need to move much faster, and we've learned how to do that and how to make sure that our assays are up and running as soon as possible to validate the findings of the enumerations.

Another bottleneck is that although our methods are powerful, they don't necessarily predict all the properties of a molecule. One of the things that we're all very concerned about, obviously, is unexpected off-target activity. Traditional safety assessments and toxicology studies have timelines that we can't change.

You've disclosed five targets to date: CDC7, WEE1, MALT1, HIF2α and SOS1–KRAS. Why these?
We wanted to start in oncology, because there is an enormous sense of urgency there to accelerate drug discovery for the benefit of patients who often don't have very long to live. And we found a number of targets there that spanned different design challenges.

Our software works really, really well on kinases, so we are working on a few of these.

CDC7 inhibitors, for instance, have traditionally been pretty hard to design with the level of potency that we believe is required to interact in the cell cycle the way that you need these to work. Multiple companies went after CDC7 inhibitors, and if they found potency it was very difficult to also get either selectivity or great drug-like properties. It's called the whack-a-mole problem by a lot of people at Schrödinger. Being able to get molecules that had exquisite potency at the picomolar level, and that had great drug-like properties, that was a key design challenge.

In the case of WEE1, there are a lot of WEE1 inhibitors out there, but it has been challenging to dial out activity against other kinases. And that was an area where we believed that coming up with a very selective WEE1 inhibitor that had again, great drug-like properties, has always been challenging.

Other programmes are more protein–protein interaction (PPI) projects. KRAS has been a Holy Grail for drug developers, for example, and after looking at the KRAS–SOS interaction, we thought that this PPI was probably a good place for us to work.

Because industry has worked extensively on several of these targets, you get to benefit from the structures, ligands and insights they previously generated. How does your approach fare against novel targets?
It's a great observation. Our first cohort of five programmes does include those well-trodden targets that people have gone after, where we have lots of ligands and structures that we can use to build our models. However, we believe that there are also targets that aren't necessarily precedented with large numbers of structures and large numbers of compounds. If you look at the ACC programme that Nimbus worked on early on, for example, the whole industry had gone in one direction to try to find these selective small molecules. And Nimbus's breakthrough actually came from a natural peptide that bound to an allosteric site.

> This year, we've ideated through something like 275 billion compounds

Also, early on in these programmes you sometimes don't have the perfect structures. We now have the ability to go out and get those structures. With our HIF2α programme, we're doing something that people might find surprising, which is going out and doing a screen against a novel crystal structure for a mutant form of HIF2α to find a starting point that will allow us to move into that very innovative space.

I would add that behind the publicly disclosed pipeline, there are more examples of targets where an allosteric site or a natural ligand has allowed us to essentially ignite these projects with a very interesting and novel starting point. Going forward, the work we're doing to obtain more and more structures we think is going to be an exciting future for our pipeline.

The ultimate test, of course, is the clinic. What's your timeline for when your compounds will enter clinical trials?
We started our own programmes in the second half of 2018. We're on track to put our first molecules into safety studies in the first half of 2021. The typical timeline from the start of those studies to first-in-human trials is 9–12 months.

But these five programmes don't necessarily represent our first crop of drugs. There are molecules we've designed with other companies that have already entered the clinic. There are multiple Schrödinger software-designed molecules that are as far as phase IIb already. And work that we did with Agios has led to two drugs that are already on the market.