DDT Vol. 8, No. 1 January 2003

Chris Lipinski discusses life and chemistry after the Rule of Five

Interviewed by Joanna Owens

Christopher A. Lipinski
Pfizer Global R&D

Christopher A. Lipinski is a Senior Research Fellow at Pfizer in Groton, Connecticut, USA. He obtained a PhD from the University of California, Berkeley, and did his postdoctoral training at Caltech, supported by the National Institute of General Medical Sciences. After joining Pfizer in 1970, Chris supervised medicinal chemistry laboratories, discovering multiple gastrointestinal and diabetic clinical candidates. In 1990, he established a laboratory combining computations and experimental physical-property measurements. He is the author of 'The Rule of Five', a widely used filter for drug-like properties. He is a member of the American Chemical Society and the American Association of Pharmaceutical Scientists, an adjunct faculty member at Connecticut College, New London, Connecticut, and most recently joined the Scientific Advisory Committee of the Global Alliance for TB Drug Development.

Can you tell us a bit about yourself and your career, to date?
The most significant thing is that I retired on June 1st [2002], but I still come in every day and continue my scientific activities. I've got access to all the scientific information that I had when I was still working and I'm really enjoying working here. The only major difference is I don't have to worry about the small amount of annoying bureaucracy anymore! And also my pocket's a little lighter because I'm not getting paid!

So are you doing that on a consultancy basis?
I do a little bit of consultancy work but actually I'm on several scientific advisory boards and I have a very active schedule of invited talks – by the end of this year I will have given 26 invited lectures at various places! So I'm on the road a lot.

What is your research currently focused on?
I have several research themes. One is the whole issue of drug-like properties (and the thing I'm probably most famous for is the 'Rule of Five'), and how drug-likeness fits into combinatorial library design. Then there are a couple of additional recent topics: the whole issue of compound distribution and compound management systems, and the questions related to compound precipitation and chemical instability when materials are stored in DMSO solutions. Another recent theme is how to efficiently capture chemical data from the literature, especially in terms of constructing quantitative structure–activity relationship (QSAR) datasets.

What main areas do you think computational methods are being applied to within the drug discovery process?
One long-term objective is the gradual work on trying to replace experimental assays relating to drug-likeness with computational assays. So, this includes computational approaches to ADME/Tox properties. The ADME properties will be easier than the toxicity properties, but this objective as a whole will be pretty difficult to achieve. Another issue, which I think is relatively poorly addressed, is how information technology can help with some of the 'people' issues in research [for example, training between different disciplines]. For example, the first piece of information about a compound that the medicinal chemist has is its *in vitro* activity; but that's not good enough, you need other information content about drug-like properties. Typically, there is a delay between obtaining *in vitro* activity and other information, and it is very hard to get information concurrently, even in a very efficient organization. Even a delay of a few days or a week can be enough so that the human brain focuses on the thing it sees first, rather than the whole picture. This is an issue that's not purely computational, and I don't know how to solve it, but I think it's an important issue.

What computational tools for ADME are actually reliable?
The greatest degree of reliability comes from those tools that are used for combinatorial library design. I think tools to predict permeability and solubility are definitely useful at this stage, but I wouldn't particularly trust them with small numbers of compounds or if I needed a really accurate result. Pfizer and other companies have filtering and other computational tools that will bin compounds, good to bad, or good/intermediate to bad, and these tools might be extended to look at drug–drug interactions and metabolic stability. They are useful in that they will perhaps enable us to identify the really 'miserable' compounds, and I anticipate these tools will improve in time. But there are some tools that are really lacking. For example, we are still at the early stages with tools that can predict drug–drug interactions. There are workable tools for CYP2D6, but the tools are pretty shaky for CYP3A4, although with recent progress I anticipate things improving a lot. One thing I would love to have right now would be a computational program that will reliably calculate whether a compound can form an intramolecular hydrogen bond. This is very important because a single intramolecular hydrogen bond can easily give you a tenfold improvement in permeability, and this knowledge is missing from current permeability calculations. Also, computational tools for biological transporters are at an extremely primitive stage; transporter assays are only now becoming amenable to medium-throughput format and there's just not enough experimental data at this time to have reliable models.

How do you think computational chemists interact with other areas of drug discovery?
It depends on the company and different companies are structured in different ways.

1359-6446/02/$ – see front matter PII: S1359-6446(02)02556-4

I can tell you the way I think is most effective, and that is to have computational chemists working in teams with medicinal chemists – that works best.

Do you feel that computational approaches have a strong enough role within pharmaceutical research?
No, I think everybody recognizes the importance of such approaches, but the problem is finding people with the necessary expertise. There is a tremendous shortage of computational chemists, they are extremely hard to hire, there's just not enough people being trained, and those that are good probably have multiple, multiple job offers! There are only a limited number of really good universities that produce computational chemists, and they are tremendously in demand. Because almost every single small biotech start-up, almost by default, has to have one computational person, so there's just this huge, huge demand.

How are current computational tools for library design allowing the incorporation of medicinal chemistry knowledge into library planning?
One of the most difficult things to do is to capture metadata, that is, the reason why decisions were made; this is poorly dealt with. Where computational chemistry tools really shine is capturing the types of chemistry functionality you want to avoid and they're also pretty good at filtering by property, but they do a poor job of taking into account chemical feasibility.

'Computational chemistry implementation has become a real bottleneck.'

So, using these tools to design for diversity – that is the oldest, established method and it's in pretty good shape; designing for drug-like properties is a work-in-progress and is coming along well. However, the ability to computationally predict success in chemistry is quite poor, and there are reasons for that. Currently, people set up sophisticated protocol databases where they track their success in chemistry and then they come up with new chemistry; they see how similar it is to something they've done before and they call up that information. But the ability to computationally predict whether a reaction will work, particularly a reaction with different substrates, is pretty poor. This is because the technology for that died out in the early 1990s. There was technology along that route, such as Corey's LHASA (Logic and Heuristics Applied to Synthetic Analysis) program, but in the early 1990s reaction databases, such as those from MDL and Beilstein, came along and they were set up in a very easy-to-use format so chemists could do the searching themselves, without having to go to an IT professional or computational chemist. So those database-searching programs replaced the computational chemistry calculation programs. I think computational chemistry implementation has become a real bottleneck.

How are computational chemists approaching the blood–brain barrier?
The standard way is to use one of two methods. The first is to use calculations of the solvent accessible polar surface area and set a cut-off value. So, for example, anything with a polar surface area below 90 $Å^2$ has a decent chance, and this will take care of compound gross physicochemical properties. The second thing you have to do is determine whether the compound is going to be a P-glycoprotein substrate, which is a major barrier to compounds getting into the CNS. That's a slow process because you have to get a good experimental database at a relevant concentration for the CNS. Setting up the database or program based on higher concentration assays that were meant to measure, for example, intestinal permeability, probably won't work, you need to go to a lower concentration.

Can you point to any tangible evidence that the trend toward higher molecular weight and lipophilicity in clinical candidates has been reversed after the publication of your 'Rule of Five'?
Well, at Pfizer we basically do not buy, and combinatorially do not synthesize, compounds that break two parameters in the Rule of Five. Also if you look at companies that are selling compounds they usually quote Rule of Five compliance rates and typically what you find is that the parameter that is most difficult to control combinatorially is lipophilicity. So nowadays you seldom find compounds with two parameters out of range and if you do, then one of these is usually lipophilicity. A good library might have 20% of compounds above the Rule of Five lipophilicity cut-off, but if you're not careful and have a sloppier library, or in some of the older works, the drug-likeness can look really bad. I will say this though, you read in some articles that combinatorial chemistry has not lived up to its promise, and I think to some degree that is true. I really think that the first five years of combinatorial chemistry was a total disaster, because people tapped into pre-existing technology, which was Merrifield solid-phase peptide synthesis, where you have lots of compounds with peptidic backbones and we now know this is a real mistake. Second, the thinking was dominated by the idea of maximal chemical diversity and so the idea would be to stick in as much interesting functionality and space into a library as possible to try to pick up *in vitro* activity. This worked, but in the process compounds were made that were very large, and this essentially defeated our ability to get orally active compounds.

'The first five years of combinatorial chemistry was a total disaster.'

Do you think we will see a new blockbuster or class of drugs coming directly from advances in combinatorial chemistry?
Well without going into details, I know within my organization there are some dynamite compounds that came out of combinatorial chemistry. But let me put it this way, there is no choice, because if you believe that a rate-limiting step in getting orally active compounds is chemistry, then you have to be able to make the compounds for initial screening in an efficient manner, and that means you have to use combinatorial chemistry.

How do you think the 'Rule of Five' and other similar approaches have impacted on the reality of the 'fail fast, fail cheap' ethos in drug discovery?
The 'Rule of Five' deals with physicochemical properties and then the computational approaches deal with the metabolism and elimination parts of ADME, so they are both part of trying to filter things early. But I don't think these fail early approaches are going to get the 200% increase in productivity that senior executives in big pharma are talking about.

It's going to take something more than that, and the problem is that fail early approaches only deal with the things that we know about. I would say that if we implement things sufficiently and screen early then we might get a 25–35% increase in productivity, because it's the things you don't know about that kill you in drug development, toxicity in particular.

'I don't think fail early approaches are going to get the 200% increase in productivity that senior executives in big pharma are talking about.'

What are the limitations of the druggability rules you have developed?
The Rule of Five was intended as a very crude filter, so it can be improved on. First, it cuts across all therapeutic areas, so you would want to modify those rules for CNS compounds, for example, and try to obtain reduced molecular weight and a reduced number of hydrogen bond donors and acceptors. Second, the Rule of Five obviously doesn't work if there's a biological transporter system involved, so if you want an orally active compound that's out of the Rule of Five limits and you want decent absorption, then you need to either determine this in a screening assay, or design the compound as a substrate for absorptive transport. Developing assays for absorptive transport is a huge area starting up right now, the reagents are beginning to become available and people are setting up medium-throughput transporter assays, it's like HTS but delayed by ten years.

How can you quantify the effect of your rules on internal discovery at Pfizer?
Well, here's a specific example, if we take a retrospective look at the quality of our initial HTS hits, in about the first 20% of our HTS history, we could see that there were problems with those compounds. But if you look at the past 80% there were no problems at all, that's because we got smart and we put in computational filters and experimental filters. The other way you can quantify it is to look at the pattern of compounds made in lead optimization in our medicinal chemistry laboratories. We actually implemented the Rule of Five at Pfizer in 1995 before it came out in the literature in 1997, and about one year after we implemented it – while at the same time using medium capacity solubility assays – we saw a big improvement in the quality and physicochemical properties of the compounds the chemists were making, and those improvements have continued to this day. If you have the correct computational filters and good assays in place, you can get the advantages of HTS while minimizing the downside of poor physicochemical properties.

As with any set of rules, obviously there are occasional exceptions. Is there any way to generalize the exceptions to the Rule of Five?
Yes, in certain therapeutic areas orally active compounds don't comply with the Rule of Five. For example, the infectious disease area includes many orally active drugs – such as the cephalosporins – that are substrates for absorptive transporters. Another area where you find many exceptions is natural products, especially natural products from plant origin, and I think this is because during evolution, mammals, including humans, have developed biological systems for effluxing, to prevent xenobiotics being absorbed. So, we get lots of natural product features that don't comply. I think that is the case for natural products of plant origin, but it might not be the case for marine natural products.

What kinds of molecules tend to surprise you?
I think if I was to pick one important antibacterial class I would say macrolides and cyclic peptides. These are huge compounds with molecular weights of ~750 Da, and they might have two to three basic groups, yet some of them are very well absorbed, very orally active. So they break the Rule of Five, and although some of them are highly soluble it's not exactly clear why they are so well absorbed. These types of compounds are exactly those for which it is extremely difficult to come up with a computational prediction, because they have multiple charges and are conformationally flexible, and the calculations that we have don't work terribly well on compounds like that.

Why do you think compounds such as cyclosporin and FK506 are orally bioavailable?
I don't think anybody knows! Cyclosporin is actually only moderately permeable, so I don't know the answer to that. What I will say though is that although cyclosporin and FK506 are orally active it took a huge amount of chemistry and pharmaceutical sciences effort to take those compounds to where they are now. You have to have a very important therapeutic area with a complete lack of any useful medicine and a huge medical need to justify that sort of effort. Also, cyclosporin is not the sort of compound you are going to find through combinatorial chemistry, you're never going to make anything like that because it's too complex.

How early are aqueous solubility and permeability data provided in the discovery process so that later oral absorption problems can be avoided?
The first time that such data are used is during combinatorial library design. Calculations are good enough, especially in the binning sense [a prediction that categorizes a response], to enable you to actually design in a bias towards good aqueous solubility. Permeability in a combinatorial library is generally not a problem, because you have to work at the chemistry to make impermeable compounds, it is much easier to make insoluble compounds than it is to make impermeable compounds. The next step – that we tried at Pfizer but it didn't work – is to intercept compounds in early protocol development, that is, at the first stage of carrying out a combinatorial library synthesis, and improve the quality of the library by measuring solubility.

'You hear people saying, ''We're being swamped with testing for drug-like properties because we're getting hundreds of compounds from chemists'',…something is wrong there.'

This turned out to be unsuccessful for a purely 'people' reason: protocol development is typically a rate-determining step, and it was extremely difficult to get real compounds at a very early stage. What we got were compounds for experimental testing when about 75% of the work had been done. When somebody has already done 75% of the work on the rate-determining step,

and you then come back to them because the compounds are not very soluble, it is too late to change anything; too much effort has already sunk into it and the library just goes ahead into production anyway.
Where I think solubility data really works well, however, is at the hit-to-lead stage: you do your HTS and identify multiple chemotypes, you confirm that they really are active and then it's the best time to do the profiling so that you can pick the best few chemotypes to work on; that's where screening really helps. Subsequently, it is crucial that during lead optimization, when you look at *in vitro* activity and so on, you also provide information on drug-like properties. At conferences you hear people saying, 'We're being swamped with testing for drug-like properties because we're getting hundreds of compounds from chemists', and something is wrong there. I think this happens because the initial screening libraries are poorly designed. If you have poorly designed libraries you get lots of very active hits, which is just miserable, but if you have good screening libraries upfront, you don't get such a large number of compounds for drug-like property profiling. I think, parathetically, this ties in with what I said originally about the failure of the first half of combinatorial chemistry. If from 1991 through to 1996 your company made lots of non-drug-like compounds, and those compounds are now going into the screens, they will come out active but they will swamp the system. The trick is to have quality in the original screening library, I think that is really the key, and it's an area where there is tremendous leverage.

It seems like compounds are failing now at Pfizer and other companies more because of toxicity problems than bioavailability. How can we address this?
I think the only way we can address it is to have multiple chemical series, especially if you are going into a new therapeutic area. But even in existing targets, at the very beginning you never know which chemical series is the one that is going to succeed. If you look at drugs that are marketed, it is usually one chemotype, or one or sometimes two chemical series, for a particular target, but you don't normally find four or five series. I think having breadth in chemistry is the way to go.

What do you think is the major key to slashing the statistics on failed drug candidates?
I think there are two extreme viewpoints in the industry on this. One says that the solution is in genomics, and the other says the solution is in chemistry! I happen to think, with my bias as a chemist and because I work at Pfizer, that the solution is more likely to be in quality chemistry. I would actually refer people to some work by my colleagues at Pfizer on the druggable genome [1]. It presents a chemistry-centred view of the world, and says that a lot of the problems will be solvable in chemistry and in the early explorative clinical stage.

'I'll bet you that in three to five years, computational docking and scoring on selected targets will be competing with actual experimental testing.'

How can consideration of drug-like properties be extended to identify effective drug targets?
I think this is a very important area. This is where the genomics and chemistry viewpoints of the future differ. From a purely genomics perspective there are targets that, if druggable, would do something good for medicine. But if you're trying to get an orally active medicine, there are some target classes for which it is probably impossible to get a small, orally active, rule-of-five compliant ligand. That is one reason why, of the proportion of drugs that has been approved for marketing, the number of biologicals has been steadily climbing. So there are good targets from a genomics perspective that are probably nearly hopeless from a traditional small-molecule drug discovery viewpoint.

What advantages does 'property-based design' have over structure-based drug design?
They are actually complementary. Obviously, with structure-based design, you are limited to those target classes where you can get an X-ray crystal structure, which effectively means that you can't use structure-based drug design for 50% of drug targets because you can't get any X-ray structures for GPCRs. The hottest area right now is the field of kinases. Property-based design can obviously help with selectivity and drug-like properties because you can identify a part of your ligand or binding site that is solvent accessible, and add on a functionality that might help with solubility. Here at Pfizer we use both, one doesn't exclude the other.

What do you think computational chemistry will have achieved for the industry in 20 years time?
First of all, I think for drug-like properties that are biologically based, we will undoubtedly have very good computational models. The issue is not a lack of good computational science, but a matter of not having the underlying data. So that will be a really big win. I also expect that the computational efforts for structure-based drug design will continue to expand. One reason is that, on the experimental side, there has been great progress with high throughput crystallization, which is getting to be almost a factory operation. Docking on X-ray structures is already pretty decent, but the scoring could be greatly improved. We're already beginning to see it but I'll bet you that in three to five years, computational docking and scoring on selected targets will be competing with actual experimental testing. This is the case for certain target classes, but for many target classes we do not have the X-ray structure and certain targets can be difficult to handle. Kinases are probably the best example where computational docking and scoring will play out in the future.

In five years from now what* in silico *process will we be talking about?
I think that five years from now, we will have really good programs to handle the whole big issue of drug–drug interactions, because you can almost see it coming, the database is getting better. For factors such as P-glycoprotein efflux, a few companies have already accumulated large experimental databases on this, so we should start to get decent tools for studying that; they might not be accessible to everybody but I think some companies will have them. One reason for many drugs being pulled off the market has been cardiotoxicity, specifically QT prolongation. But I think this will be under very good control within five years.

What* in silico *approach will have the next major impact on drug design?
It's in very early stages, but I think it will be the biological systems modelling that is just beginning to come out. The analogy that has been talked about in conferences that I really like is the Boeing 747 design: one of the most recent models was built entirely based on *in silico* models and it flew. So, as things get better and we work out more and more pathways I think we will see *in silico* modelling being applied to biological systems much more. Right now, *in silico* work has very much a chemistry flavour, but in the long-term, I think the greatest impact will be in the modelling of cellular processes.

What professional achievements are you most proud of?
Probably the Rule of Five; I really feel rather good about it because it has helped my company, Pfizer, and it has helped many other people, so that would be number one. Number two would be how active I am on the speaking circuit and that I can answer questions and help people out. I feel very fortunate that Pfizer has been very supportive to me in that effort. Of course, it has benefited both the company and myself. I feel that we are facing problems that are two extremes: technical problems and 'people' problems. With technical problems, it is perhaps a good idea to be a little cautious on disclosure to avoid giving away any enabling technologies or structures of compounds or so on.

'Contributing some kind of technology or idea that can help everybody, that makes me feel very good.'

But when there are 'people' problems, it is sometimes a good idea to present them on the outside because, in the process of presenting externally, people within your own company hear and it really makes an impact. That's one of the things about the Rule of Five. It was presented internally at Pfizer and it had an effect internally, but it didn't really hit its maximum effect until people heard it on the outside. I think there are many other issues that the industry is facing like that. For example, the issue with compound distribution in DMSO that I mentioned earlier, there are several 'people' issues related to that, but by presenting information and sharing information at meetings and so on, it helps everybody.

What would you like to have achieved by the end of your career?
I would like to have made a positive impact on the ability of the industry in general to discover medicines that are useful for the treatment of human disease. The most straightforward way [to treat human disease] if you're a chemist is to make a compound that goes all the way and becomes a drug. But if that doesn't happen then contributing some kind of technology or idea that can help everybody, that makes me feel very good.

Reference

1 Hopkins, A.L. and Groom, C.R. (2002) The druggable genome. *Nat. Rev. Drug Discov.* 1, 727–730

Christopher A. Lipinski
Adjunct Senior Research Fellow
Pfizer Global R&D
Groton Labs
Eastern Point Road
MS 8200-36
Groton, CT, USA