Innovations

Piecing together drugs
Abbott Laboratories

Any drug company would covet a machine that identifies the building blocks for the perfect drug and then indicates how the building blocks should be joined together to create that drug. Perhaps, say researchers at Abbott Laboratories, that machine has been present in most pharmaceutical companies all along. A new method dubbed SAR by NMR (structure–activity relationships by nuclear magnetic resonance spectroscopy) uses an NMR machine to screen through thousands of building blocks, identifying those that bind to the site of interest in a protein. These weak binders are then linked together to create powerful inhibitors that bind at nanomolar concentrations.

The strong from the weak

Researchers in both academia and industry are excited by the method, published by Stephen Fesik's group in *Science* late last year. "Apparently every pharmaceutical company in the universe with an NMR department is jumping on this," says Michael Rosen of Sloan Kettering Cancer Research Center.

"It combines screening with an aspect of thermodynamics that is very elegant," continues Rosen. First, when the two weak binders are linked, their free energies of binding are additive, so the new binding affinity is the product of the two old binding affinities. And second, explains Fesik, "you can get a bigger boost even than that" because the linking removes one of the negative entropy terms. That boost that can convert micromolar binders to nanomolar binders.

Fesik used one of the strengths of NMR, its ability to identify weak binding interactions, to find the micromolar binders. The Abbott team used an ^{15}N-labeled target (in this case FK506-binding protein (FKBP)), eliminating the background that usually arises from non-specific binding of labeled ligands. The ligands that bound the target altered the electronic environment around the protein amides. This was detected using a standard two-dimensional (2D) NMR method called heteronuclear single-quantum correlation spectroscopy (HSQC; see box).

The testing process had a decimal theme: ten days to test 10,000 ligands in batches of ten. When a positive pool was identified, each member of that pool was tested individually. Once the team had identified a molecule that bound FKBP, they repeated the process to find another that bound nearby. This latter test was conducted in the presence of the first binder, thus ensuring that the two final molecules would not overlap unfavorably.

To determine how to link the two molecules, Fesik used computer modeling of the NMR-derived structure of FKBP with the untethered ligands. The best of the five final products had a binding constant (K_d) of 19 nM, very much better than the K_ds of the starting compounds (2 μM and 100 μM).

The genesis of an idea

The conceptual parent of SAR by NMR is combinatorial chemistry, which involves the reaction of large numbers of building blocks with each other in all possible combinations, either in mixtures or massively parallel syntheses. The new method also relies on the combination of small molecule building blocks, but the building blocks are selected out by the NMR experiment.

Thus, even though a huge 'virtual' library is theoretically sampled, the number of synthetic reactions is reduced from thousands to a handful. This means faster development times, even as more complex chemistries are used. The diversity of the building blocks is limited only by the need for millimolar solubility.

The idea of using NMR to probe small molecule interactions with proteins is not new. A number of groups have added single, simple organic chemicals to proteins and used NMR to determine where the chemical interacts (usually weakly) with the protein. Such information can give clues as to what chemical shapes fit well into particular protein pockets. What Fesik has done, says Ad Bax of the National Institutes of Health, is "put two and two together and come up with a realistic screening procedure."

Experiment or compute

If one extreme of drug discovery is purely experimental combinatorial chemistry, its polar opposite is computational drug design. Using a computer to predict which small molecules will bind a protein has proven difficult for a number of reasons. It is difficult to predict both the extent of entropy effects (e.g., the energetic cost when the mobility of a ligand decreases upon binding, or the energetic gain when water molecules are liberated by ligand binding), and whether ligands or proteins will undergo conformational changes upon binding. And to explore enough molecules, the forces involved in binding have to be simplified. But the most troubling variable is the treatment of water and ions.

"The others can all be addressed in one manner or another," says Irwin Kuntz of the University of California at San Francisco. A lack of experimental data means that this starting point is lacking for water placement. "It's unclear whether you include the water as part of the protein and try and come up with ligands for it," says Fesik, "or do you take it out and make a ligand that would displace it."

Innovations is a new monthly column that surveys the science done at biotechnology companies.

"The NMR experiment gives you an experimental result whereas the computational approach gives you a hypothesis," acknowledges Kuntz. The advantages are not all with the NMR method, however. The speed of computational screening, which requires only fractions of a second per sample, cannot be beaten.

Changing drug discovery

Although the NMR divisions of drug companies were adept at generating three-dimensional structures of proteins, they needed new ways to use these structures. In the past, NMR was only brought in at the later stages of drug development, to see exactly how a lead was binding and so suggest modifications. In the early stages, biological testing selected leads from amongst the compounds made by medicinal chemists. "Now," says Rosen, "the spectroscopists can generate their own leads."

In academia, Gerhard Wagner of Harvard Medical School sees the potential for a lot of collaborations. For each experiment the NMR structure of the target must be solved and an appropriate library of building blocks must be made or be otherwise available. "For a single laboratory that's pretty unrealistic," says Wagner. But in a collaboration the academic laboratory could contribute one element, and finish up with a selective inhibitor to use in further basic research.

Kuntz is looking forward to collaborations in which computational techniques are used to improve on the existing NMR method. The speed of computational methods means that they can be used to screen through huge libraries and the results can be used in selecting building blocks for the NMR experiment. "These are potentially complementary approaches," he says.

William A. Wells, Biotext Ltd,
211 Hugo Street, San Francisco,
CA 94122-2603, USA; wells@biotext.com.

A Beginner's Guide to NMR

NMR works because some nuclei act like tiny magnets. The basis for this magnetism is a pair of properties that protons and neutrons share with electrons: they have intrinsic spin and they are distributed in discrete orbitals. Protons and neutrons both pair in orbitals, and the two members of the pair have opposite spin and cancel each other out. But if either protons or neutrons are unpaired in an orbital (as happens in ^{1}H and ^{13}C, but not ^{12}C), the nucleus as a whole has spin. As these spinning nuclei are charged, they generate a magnetic field.

When a magnetic field is externally applied to a molecule, more of the nuclei in the molecule align with the magnetic field than against it, generating a bulk magnetization in the direction of the external field. As with a gyroscope, which is spinning but affected by a gravitational field, the combination of spin and the magnetic field causes the nuclei to rotate (or, more correctly, precess) around the direction of the external field.

Describing individual atoms: 1D spectra

An NMR experiment starts with a radiofrequency (rf) pulse that generates a magnetic field perpendicular to the external magnetic field. This rotates the nuclei and therefore their bulk magnetization, and the change is detected by a receiver coil. The signal oscillates at the rate at which the nucleus is spinning.

The signal is useful because the frequency of oscillation varies for different nuclei (e.g., spinning hydrogen and nitrogen atoms have different angular momenta). This frequency is modified depending on how well the nucleus is shielded (by electrons, for example) from the applied magnetic field; this value is termed the chemical shift. For example, an electron-withdrawing group nearby will change the chemical shift.

In its original form, NMR involved a continuous scan over different radiofrequencies to perturb different nuclei individually. This time-consuming approach has been replaced by pulse NMR: one short, strong pulse, which excites all nuclei simultaneously. Pulse NMR is not unlike hitting a bell instead of playing a scale to find which note the bell resonates to.

The single pulse gives a complicated output of many overlapping waves. Fortunately, a Fourier transform, in a flurry of almost incomprehensible mathematics, can be used to do two important things. It converts the time dimension into frequency, and it separates the overlapping waves into individual signals, one per nucleus. The result is a one-dimensional (1D) spectrum of frequency versus intensity.

Connecting atoms: 2D spectra

The simple, 1D spectra described above tell you a lot about the environment each atom finds itself in. But for structure determination, spectroscopists need two-dimensional (2D) spectra to give them information about relationships between atoms. The atoms are grouped, or correlated, by transferring magnetization between atoms.

The magnetization transfer is a complicated process involving quantum properties of the intervening bond, but two simple models are helpful. First, taking each nucleus as a magnet, we know that two magnets that are close to each other will affect each others behavior.

Second, the two nuclei can be modeled as two pendulums, and the bond between them represented by a stiff rod. When an initial pulse sets one of the pendulums in motion, the energy in this pendulum is gradually transferred to the other, which eventually takes over all the motion.

In the method used by Fesik's group, the first pulse is directed at the protons in the (ligand-free) protein. For those protons that are in amide groups, the change in magnetization can transfer along the bonds to the amide nitrogens. A series of pulses is used to transfer magnetization back to the proton, and the proton signal is recorded.

These transfers are useful because the signal does not returned unchanged. The chemical shift of the nitrogen is now superimposed on that of the proton. As the length of time between magnetization transfers is varied over multiple experiments, the nitrogen chemical shift oscillates on top of the proton chemical shift. The separation of these two values gives the coordinates of a unique point on a 2D plot, where the two axes represent the chemical shifts of the proton and nitrogen, respectively.

Determining structure and ligand binding

The amide peaks are assigned to specific residues in the protein by continuing the transfer of magnetization along the protein backbone, one bond at a time. Modification of the signal from the originating nucleus by neighboring nuclei allow the neighbors to be correlated. The direction of transfer is defined by pulses whose frequency targets only the originating and destination nuclei; the same pulses effectively cancel out any transfer in the other direction along the protein chain.

When a ligand binds to the protein target, it alters the electronic environment around nearby residues. This shows up as differences in the spectrum only at those few residues, thus indicating both the occurrence of binding and its location.