# The SHAPES strategy: an NMR-based approach for lead generation in drug discovery

Jasna Fejzo, Christopher A Lepre, Jeffrey W Peng, Guy W Bemis, Ajay, Mark A Murcko and Jonathan M Moore

Background: Recently, it has been shown that nuclear magnetic resonance (NMR) may be used to identify ligands that bind to low molecular weight protein drug targets. Recognizing the utility of NMR as a very sensitive method for detecting binding, we have focused on developing alternative approaches that are applicable to larger molecular weight drug targets and do not require isotopic labeling.

Results: A new method for lead generation (SHAPES) is described that uses NMR to detect binding of a limited but diverse library of small molecules to a potential drug target. The compound scaffolds are derived from shapes most commonly found in known therapeutic agents. NMR detection of low (μM–mM) affinity binding is achieved using either differential line broadening or transferred NOE (nuclear Overhauser effect) NMR techniques.

Conclusions: The SHAPES method for lead generation by NMR is useful for identifying potential lead classes of drugs early in a drug design program, and is easily integrated with other discovery tools such as virtual screening, high-throughput screening and combinatorial chemistry.

Address: Vertex Pharmaceuticals Incorporated, 130 Waverly Street, Cambridge, MA 02139-4242, USA.

Correspondence: Jonathan M Moore
E-mail: moore@vpharm.com

Key words: drug design, library, NMR, screening, SHAPES

Received: 7 April 1999
Revisions requested: 24 May 1999
Revisions received: 5 July 1999
Accepted: 14 July 1999

Published: 16 September 1999

Chemistry & Biology October 1999, 6:755–769
http://biomednet.com/elecref/1074552100600755

1074-5521/99/$ – see front matter


## Introduction

In a target-directed drug discovery program, many different strategies may be applied in the search for a clinical candidate. Although these approaches to drug discovery might follow significantly different pathways of optimization to a highly potent and bioavailable drug molecule, they all share a common origin: they must begin with a lead compound. Because many properties of the final compound (e.g. target affinity, solubility, ease of synthesis, toxicity and bioavailability) are highly dependent on those of the initial lead, the methods by which leads are identified in the early stages could significantly impact the success of the project in the latter stages.

How, then, are leads chosen? A lead molecule may typically be a known drug molecule, or analog of a known drug. Alternatively, if the target is an enzyme, the lead may be a substrate or substrate analog. In many cases, a lead is discovered simply by random screening of commercially available or proprietary compound libraries. However, many potential problems may arise using the above strategies. For example, starting with a known drug may not lead to sufficient diversity in the final class of compounds to describe them as unique intellectual property. Starting with a bioactive natural product may necessitate the design of high molecular weight analogs, with poor synthetic accessibility and difficult scale-up problems. Alternatively, using random screening to generate leads might result in a novel class of compounds, but the class might not contain a practical range of synthetically accessible compounds such that problems related to solubility, bioavailability or toxicity may be overcome without overly compromising potency.

In cases where random screening is used to generate leads early in a drug design program, several approaches may be taken. A brute force approach is to screen very large (>100,000) numbers of compounds, and hope to identify a potent binder or inhibitor of the drug target, then optimize. A more rational approach is to use information-driven methods for virtual screening of databases to select a smaller subset of compounds for screening [1,2]. For example, a typical protocol might employ similarity-searching or docking methods (in a structure-based program) to search available databases for good screening candidates. A subset of representative compounds is then assayed for binding or inhibition. Weak binders or inhibitors become leads for iterative structure-based drug design or SAR (structure–activity relationship)-based optimization, depending on whether structural information regarding the target is available. However, when one uses subset selection methods to attempt to represent several hundred thousand compounds with several thousand, several questions arise. Have a sufficient number of compounds been chosen? Have the correct subset of compounds been chosen to represent a much larger group in a compound space representing all commercially available compounds? And is sufficient diversity

represented in these compounds? If one now considers the possibility of filtering through a virtual library containing over $10^9$ compounds, the problem becomes even more challenging. Clearly, at this stage, more information to guide the selection of compounds for screening would be extremely beneficial.

We have developed a novel approach to assist in the process of early lead generation. Using simple nuclear magnetic resonance (NMR) techniques to assess binding, we screen a small but diverse library of low molecular weight, soluble compounds — the 'SHAPES library' — against the target of interest. The library is composed largely of molecular shapes that represent frameworks most commonly found in known drug molecules [3]. Weakly binding ($K_d$ ~μM–mM) hits, many of which would be missed in a standard enzymological assay, are then used to bias the filtering of large chemical databases or virtual libraries, and select a larger group of compounds for purchase or synthesis followed by high-throughput screening.

To demonstrate the feasibility of such an approach, we have screened the SHAPES library against a number of drug targets. We were able to rapidly differentiate between binding and nonbinding compounds in small noninteracting mixtures, demonstrating the utility of the method with several enzyme targets, including the p38 MAP kinase, a 42 kDa protein, and inosine-5′-monophosphate dehydrogenase (IMPDH), a multimeric enzyme target of 224 kDa.

## Results

### Library design

Our goal was to design a library of small molecules that optimized a large number of factors at once (cost, synthetic accessibility, solubility, separation of NMR peaks and diversity). The strength of such an approach is that a set of simple drug-like molecules provides an array of binding information for almost any target. These molecules and their derivatives translate into hits whose chemistry tends to be well understood and whose physical properties tend to be desirable for lead compounds.

### The SHAPES concept

Design of the SHAPES library is based to a large extent on previously published work, in which commercially available therapeutics were examined for common chemical features that made these compounds 'druglike' [3]. In this study, all drug molecules were broken down into systems consisting of rings, linkers and sidechains. The union of cyclic arrays of atoms in a molecule ('rings'), and those atoms that directly connect them ('linkers') were classified as 'frameworks'. An analysis of the comprehensive medicinal chemistry (CMC) database (MDL Information Systems, San Leandro, CA, v.94.1) indicated that 1179 simple frameworks described the 5120 relevant entries, and that surprisingly, only 32 different frameworks, or 'shapes', described ~50% of all known drugs. When atom type and bond order were incorporated into the analysis, 2506 complex frameworks described the 5120 entries, and 41 frameworks described 24% of all drugs. For details of the analyses, see the original work of Bemis and Murcko [3]. In this paper, the terms 'complex framework' and 'scaffold' are used interchangeably.

One may extend this type of analysis further. Drug sidechains are the linear arrays of atoms connected to molecular frameworks. A similar study of the CMC database revealed a set of preferred sidechains found in known therapeutics (details to be presented elsewhere). The details of the juxtaposition of sidechains and frameworks are likely to be important in many cases for drug binding. The results of the drug framework and sidechain analyses may therefore be used together as a general method to select a set of commercially available compounds with drug-like character.

### Compound selection from CMC-based frameworks

All library molecules were selected on the basis of frameworks in the CMC database. First, a semi-automated approach was used that combined the most frequently occurring drug frameworks of the 41 mentioned above with the 30 most common drug sidechains (see the Supplementary material section). A substructure search was done by searching the available chemicals directory (ACD) database (MDL Information Systems, San Leandro, CA; ACD98.2) for each molecular framework, leaving all sidechain valences open. The resulting hit list was then filtered to select examples that contained only sidechains specified by our analysis. These examples were biased towards molecules that in our estimation would have high aqueous solubility. We assumed that solubility is correlated with the number of heteroatoms, so all of the chosen molecules contained at least one nitrogen or oxygen. Approximately one third of the 41 frameworks were not used because of either their inherent synthetic complexity or the absence of sufficiently soluble analogs in the ACD. In our current library, approximately half were selected as described above. The remainder were chosen to represent drug classes that occur frequently in the CMC (MDL Information Systems, San Leandro, CA; CMC3D98.1) but are not among the 41 most common frameworks.

The current set of frameworks from which the 'SHAPES' library is derived is shown in Figure 1. All library compounds were commercially available, soluble and nonaggregated in water at 1 mM, chemically and isomerically pure, and nonreactive. Molecular weights of library compounds are in the range 68–341 Da with an average of 194 Da, contain 6–22 heavy atoms, and have a calculated logP of –2.2 to 5.5. For NMR screening analysis they must also yield a simple, well-resolved $^1$H NMR spectrum, and contain at least two protons within 5 Å of one another.

**Figure 1**

Molecular frameworks used for library selection. Attachment points for sidechains are indicated by single electrons or lone pairs. 'X' represents a C, N, O or S atom. Frameworks in the top three rows were found in the original analysis [3]. Additional frameworks are shown in the bottom two rows, as well as in the third row, depending on the identity of atom X.

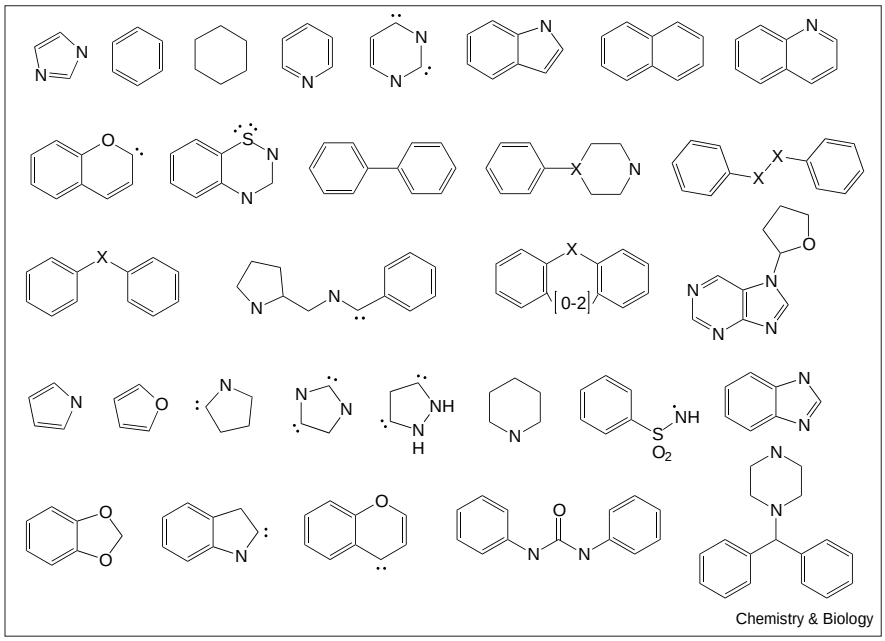


Presently, the SHAPES library consists of a pool of 132 compounds, from which a subset is screened depending on solubility considerations at a given pH. For screens in standard buffers (e.g. phosphate at pH 7.0 and Tris at PH 8.0), approximately 120 compounds are used.

### Detection of binding by NMR – theory of exchanging systems

The values of certain NMR observables, such as transverse and longitudinal relaxation rates ($R_1$ and $R_2$), the nuclear Overhauser effect (NOE), and the diffusion coefficient, are highly dependent on molecular size and shape. The NMR parameters for a small molecule are therefore exquisitely sensitive to interactions with larger molecules. Changes in line widths, relaxation rates and NOE values may be used to characterize and quantitate the binding of a small molecule to a larger target molecule such as a protein [4].

In a typical protein–ligand interaction there are three species in equilibrium: the protein, E, the ligand, L and the molecular complex EL:

$$E + L \underset{k_{off}}{\overset{k_{on}}{\Leftrightarrow}} EL \qquad (1)$$

The above equilibrium is described by the dissociation constant, $K_d = k_{off}/k_{on}$, where $k_{off}$ is the off rate, or rate of dissociation of ligand from the protein–ligand complex, and $k_{on}$ is the on rate, or rate constant for association of the ligand with the protein. The diffusion-limited on rate is often estimated at $10^8\ M^{-1}\ s^{-1}$. Ligands that bind tightly (e.g. $K_d \sim 10^{-9}$ M; $k_{off} \sim 10^{-1}\ s^{-1}$) are considered to be in 'slow exchange' on the NMR time scale, whereas ligands that bind very weakly (e.g. $k_{off} > 10^3\ s^{-1}$) are considered to be in 'fast exchange'.

Using the notation of Ni [4], and describing any NMR observable as P, the magnitude of the effect observed for a small-molecule ligand resonance ($P_{obs}$) exchanging between the free and bound states with a receptor molecule is a superposition of the parameters due to both the free ($P_{free}$) and bound states ($P_{bound}$). $P_{obs}$ is a function of protein [E] and ligand [L] concentration, the ratio of ligand to target, and $K_d$:

$$P_{obs} = (1 - p_b)P_{free} + p_b(P_{bound}) \qquad (2)$$

where $p_b$ is the fraction of ligand in the bound state. For example, if $[L] << [E]$, ligands in extremely slow exchange on the NMR time scale remain bound during the entire process of nuclear spin excitation and relaxation. In this case, $p_b \sim 1$, and $P_{obs} \sim P_{bound}$. The small molecule will then have $R_1$, $R_2$ and NOE values characteristic of the entire protein–ligand complex. Ligands in fast exchange, however, come off the protein tens or hundreds of thousands of times before the signals relax. In this case $p_b$ will be very small, $P_{obs} \sim P_{free}$, and the resonances will display parameters closer to those of the free ligand. Note that we use a protein here as an example of a typical target, but targets may be any molecule that interacts with the ligand, such as DNA, RNA or another biomolecule.

**Figure 2**

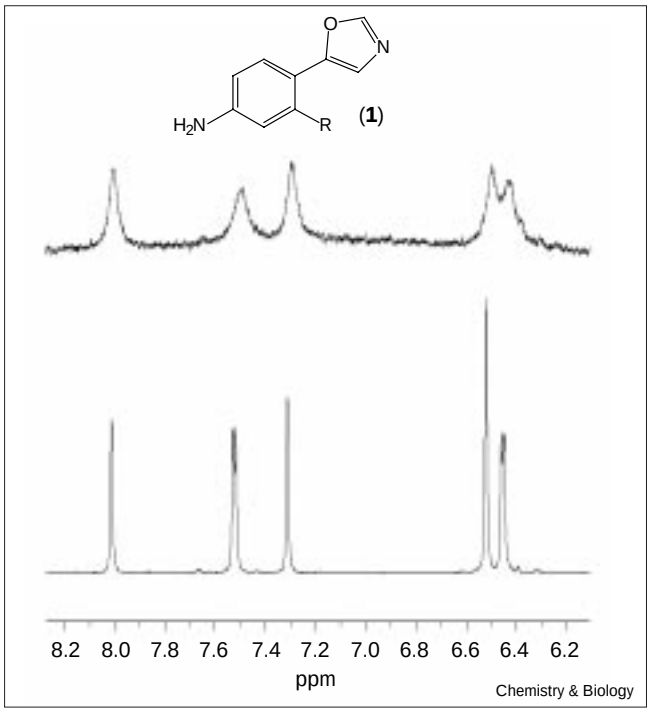


1D-$^1$H spectra of free ligand (1, lower spectrum) and 1 in the presence of IMPDH (upper spectrum). The significant line broadening observed in the presence of the enzyme indicates binding of the small molecule. Line widths at half height for the furthest downfield component of 1 (leftmost peak) are 3 Hz (lower trace) and 30 Hz (upper trace). The sample contained 1 mM ligand + 100 μM IMPDH in 25 mM d-TRIS pD* = 8.4, 300 mM KCl, 5% d-glycerol, 5 mM DTT. 1D spectra were collected at 277 K with 1 s low power presaturation of the residual HDO signal, using a Bruker DRX-800 spectrometer.

**Figure 3**

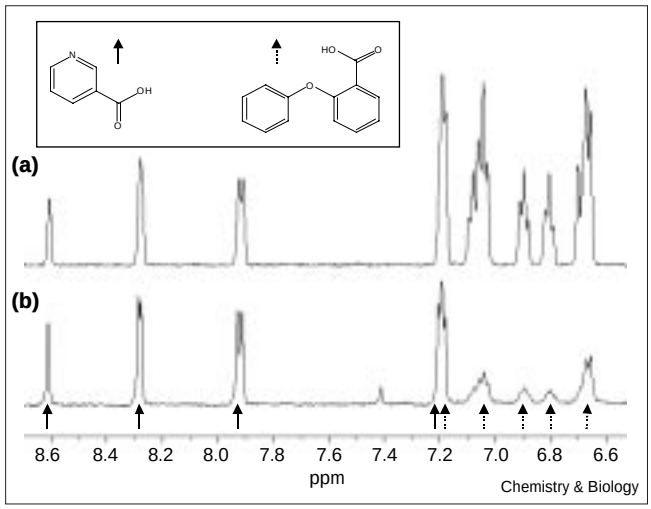


1D-$^1$H spectra of (a) a mixture of two ligands and (b) the mixture of the ligands in the presence of p38 MAP kinase. Resonances from nicotinic acid (left structure) and 2-phenoxy benzoic acid (right structure) are marked with solid and dashed arrows, respectively. The peak at 7.2 ppm consists of overlapping resonances from both compounds. Line broadening, suppression of fine structure, and attenuation of ligand resonance peak height due to relaxation filter in the bottom spectrum indicate 2-phenoxy benzoic acid binds to p38, whereas nicotinic acid does not. The sample contained 1 mM ligands, 0.2 mM p38 MAP kinase, 25 mM deutero-Tris, 10% deutero-glycerol, 20 mM deutero-DTT at pD* = 8.4. Experiments were carried out at 278 K. 1D NOESY spectra were collected with 16 K data points, 128 transients and a relaxation delay of 3 s. A relaxation filter was used after the preparatory delay to attenuate broad resonances arising from the protein.

### Large versus small proteins

Binding studies using NMR are often performed by measuring ligand line broadening in the $^1$H spectrum of a ligand–protein mixture. Line broadening for ligand resonances in an exchanging complex arises from $T_{2obs}$ relaxation ($T_{2obs} = 1/R_{2obs}$) associated with molecular tumbling, as well as additional contributions resulting from differences in the chemical shift between the free and bound ligand. Qualitatively, the most significant line broadening is observed for a ligand:protein ratio of 1:1 or less, where a higher percentage of the ligand is bound to protein. In practice, when monitoring changes in ligand parameters to assess binding, it is useful to have ligand concentrations >0.2 mM, which ensures adequate signal to noise in the NMR spectra over a short experiment time. In most cases, however, the target protein will be the limiting reagent, so a tradeoff exists between increased protein consumption with greater sensitivity, and low protein consumption with reduced sensitivity.

Of course, the sensitivity of simple differential line broadening (DLB) measurements is inherently dependent on the molecular weight of the target (Figures 2,3). Figure 2 shows the best-case scenario for line-broadening measurements. The target protein IMPDH is a 224 kDa multimeric enzyme essential for *de novo* guanine nucleotide biosynthesis. The furthest downfield component exhibits a line width at half height ($LW_{½}$) of 3 Hz. The ligand at the same concentration in the presence of 100 μM IMPDH ($K_i$ ~60 μM) shows significant line broadening ($LW_{½}$ = 30 Hz). For smaller target systems (e.g. 20–60 kDa), observed line-broadening effects are not always as pronounced. The one-dimensional (1D) spectra of two small-molecule compounds in the presence (ligand:protein = 5:1) and absence of p38 are shown in Figure 3. Noticeable line-broadening effects and attenuation of fine structure are observed for one of the compounds, 2-phenoxybenzoic acid ($K_d$ ~70 μM). For compounds that bind with weaker affinity ($K_d$ ~1–10 mM), however, it is difficult to obtain reproducible and reliable line-broadening data with smaller targets, even under the most favorable ligand:protein ratios.

Despite the difficulties in obtaining reliable line-broadening data, it is often possible to determine qualitatively whether a molecule is binding the target by looking for

**Figure 4**

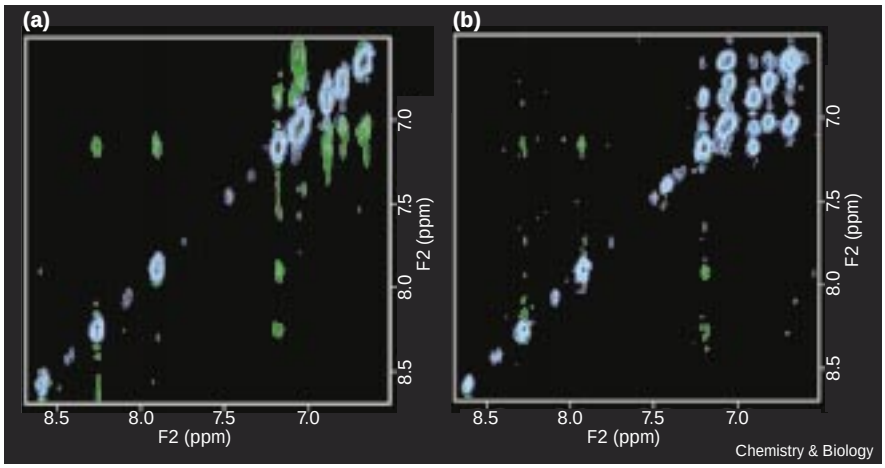


2D NOESY spectra of the same mixture of ligands shown in Figure 3. Positive contours are cyan and negative contours are green. (a) Both ligands in the mixture without protein present have weak negative cross peaks and positive diagonal peaks. (b) In the presence of the protein cross peaks remain opposite in phase to the diagonal peaks for the downfield resonances corresponding to nicotinic acid, indicating this compound does not bind. However, the sign of the upfield cross peaks of 2-phenoxy benzoic acid is the same as that of the diagonal peaks, indicating this compound binds to the protein. The peak at 7.2 ppm consists of overlapping resonances from both compounds (see Figure 3 with 1D spectra) and thus both positive and negative cross peaks are found at this frequency. 2D NOESY spectra were collected with 400 $t_1$ increments and 2 K complex points in $t_2$, with mixing times of 50 and 200 ms, relaxation delay of 2 s and 16 transients per $t_1$ increment. A spin echo sequence after the first proton pulse was used in the NOESY experiments as a relaxation filter.

changes in intensities of ligand peaks in the presence of the target versus the free state in concentration-matched relaxation-filtered spectra. Relaxation-filtering elements are routinely used in NMR pulse sequences to suppress resonances from the protein (see the Materials and methods section). Reductions in peak heights are regularly observed for compounds that bind in the several hundred micromolar to millimolar range. These spectral changes are apparent from simple inspection of the data, or can be quantitated using NMR difference methods, as described by others [5]. It is difficult to assess accurately relative affinities of binding using these methods, however, because line broadening arises because of both differences in $T_2$ and chemical shift between the free and bound states of the ligand. Line broadening can therefore occur to differing magnitudes for two ligands binding to the same target with the same affinity, as well as for different proton sites within the same ligand molecule.

To avoid these problems, we have relied on the two-dimensional (2D) transferred NOE experiment (tNOE or tNOESY) — a qualitative, but extremely reliable and sensitive diagnostic for binding. The tNOE experiment is a well-known NMR technique commonly used for determining the bound conformations of small molecules undergoing rapid exchange with a protein target. Here, however, its utility is in determining whether one of a mixture of compounds binds to a target. The 2D tNOE spectrum of small molecules in a mixture without protein exhibit very weak NOE cross peaks, with sign opposite to those of the diagonal peaks (Figure 4). In the presence of the protein, cross peaks of ligands that do not bind the target maintain phase opposite to that of the diagonal peaks. Ligands that bind to the protein, however, will transiently adopt the same rotational correlation time (or tumbling time) as the protein during their bound lifetimes. As the exchange process occurs many times before the NMR signal dies, strong NOEs built up by the ligand in the bound state are transferred to the easily detectable $^1$H NMR signals of excess free ligand, and the sign of the cross peaks change such that they are the same as the sign of diagonal peaks (Figure 4). These spectral changes allow the ligands that bind and those that do not to be unambiguously distinguished. Another group [6] has reported the use of tNOE methods in a similar approach to examine binding of a library of oligosaccharides to a protein target.

It is clear from the discussion above that large proteins (MW >60 kDa) represent the best case scenario for the automated screening of compound libraries using NMR, because 1D spectra at high ligand:protein ratio may be collected in minutes using a robotic sample changer. Reliable data may also be collected on smaller targets, although in a longer time period, as long as 2D tNOE methods are used.

Although the examples and discussion in the present work pertain to protein targets, the discussion is equally applicable to nucleic acids, protein–nucleic-acid complexes, proteins in detergents or micelles, targets tethered to beads, or even large subcellular structures.

## Screening the SHAPES library

### *The basic strategy*

The basic strategy of SHAPES screening is to use the SHAPES library to represent drug-like scaffolds, then assess binding of this limited library to a drug target using

simple NMR methods. In a typical NMR SHAPES screen, a series of experiments are carried out, some before, and others after, the actual determination of binding. Preliminary experiments include optimization of protein stability, collection of reference spectra for the free ligands, and design of mixtures. Following the SHAPES screen, experiments may be performed to determine binding specificity and affinity. Details of the experiments are described in the following sections.

*Prescreening protein stability*

The NMR screening process exposes the target molecule to what most biologists would consider extreme conditions. For example, a typical protocol for NMR screening involves dissolving a protein in 500 μl of a suitable aqueous buffer at a concentration of 50–100 μM. Small aliquots (e.g. 1 μl) of a concentrated stock of ligand in deuterated dimethyl sulfoxide (DMSO) solution are added to bring the final concentration of each ligand to ~1 mM. Depending on how many ligands are used in a mixture, the final DMSO concentration will be 0.2–1% or more. Furthermore, adding acidic or basic ligand at these concentrations can significantly alter the pH of the screening sample. Finally, the NMR samples are placed at room temperature for significant time periods, both in the sample changer rack and in the magnet. It is therefore important to examine target protein stability and activity as a function of organic solvent concentration, pH (buffer type and buffer concentration), temperature and time before a screen is carried out, so that valuable protein is not unnecessarily lost, and to ensure that the protein under these experimental conditions represents a viable target for lead generation.

For most targets, we routinely employ the microdrop protein screening method [7], which uses sub-milligram quantities of protein to optimize sample conditions before NMR screens are carried out. To assess the behavior of the target protein in the presence of DMSO, microdrops containing varying amounts of DMSO, with and without common stabilizers present, are monitored visually. Protein samples that appear stable under screening conditions are assayed to ensure the target is still catalytically active or otherwise competent. It is also important to assess whether sample pH is affected by the presence of ligand mixtures by measuring the pH of ligand mixtures in aqueous buffer in the absence of the target. We have found empirically that use of buffer concentrations much higher than needed to maintain the pH of the protein (e.g. 100 mM) are required to maintain the pH of the solution in the presence of a mixture of 1 mM organic acids or bases.

*Collecting reference data and mixture design*

When confronted with the task of screening a library of compounds using NMR, the question arises as to whether one should screen single compounds or mixtures of compounds. Several problems could arise in screening mixtures of compounds. For example, any single compound with a ten- or hundred-fold higher binding affinity than others in the mixture would compete successfully for most protein binding sites, particularly if ligands are present in excess, and therefore one or more potential binders might be missed. This problem is important because simple modifications to a weak binding scaffold might result in enormous increases in binding affinity at later design stages, so missing potential scaffolds early in the screening process might eliminate an entire compound class from further development. Ultimately, a tradeoff exists between resources (protein and spectrometer time), and ease of data analysis and reliability. In our experience, the competitive binding phenomenon described above has never been problematic (as confirmed by enzymological screening), probably because of the relatively small size and intentionally weak affinities of the compounds in the SHAPES library.

Another problem with using mixtures is spectral overlap. If overlap is severe, analysis of 1D spectra might not be possible, and analysis of 2D spectra would be difficult. Using automated methods to search all possible combinations of components, however, one can select mixtures in which overlap is minimized, allowing a reliable line broadening or tNOE analysis. We currently use mixtures of 1–4 compounds, although when using 2D tNOE spectroscopy to identify binders, we have successfully used mixtures containing up to 12 compounds (see the Materials and methods section for details).

*Screening the library and identifying target binders*

The next step, screening the mixtures of SHAPES compounds against the target, involves collecting either 1D or both 1D and 2D tNOE NMR spectra for each mixture in a separate sample with the target present. Ligand peaks in the 1D and 2D spectra of the mixtures (with and without the protein) are deconvoluted and assigned using the 1D spectra of the individual components. As described above, binding of components of the mixtures to the target is revealed by differential line broadening or inversion of the NOE cross peaks. Neither differential line broadening nor tNOE methods can distinguish specific binding from nonspecific binding, so all 'hits' obtained using NMR screening should be followed up with more rigorous studies such as competition experiments to characterize the specificity of the observed interaction.

*Throughput*

For a high-molecular-weight target (>60 kDa), the entire SHAPES library may be screened in several hours, as only 1D spectra are necessary. For smaller targets (10–60 kDa), which require both 1D line broadening and 2D NOE spectra, the process requires several days. As the SHAPES library is relatively small, the data collection time is not a

limiting factor. In fact, in a typical screen it is useful to collect tNOE spectra at several mixing times. Observing differences in cross peak sign and intensity as mixing time is increased allows a rough assessment of the relative affinities of SHAPES hits for a given target.

### *Protein requirements*

The quantity of protein required is significant when compared with that of high-throughput enzymological screening, but not necessarily when compared with a structure-based drug-design program employing X-ray crystallographic or NMR structural methods. Protein requirements depend on the molecular weight of the target. In initial studies, screens were carried out at ligand:protein ratio of 10:1, using 100 μM protein and 1 mM ligand, but more recently we have found that similar data quality and sensitivity may be obtained using a ligand:protein ratio of 20:1, and all screens are now carried out with protein concentrations of 50 μM. For 32 mixtures averaging four compounds each, each sample containing 0.5 ml 50 μM protein, screening the SHAPES library would require: IMPDH (57 kDa subunits): (32 samples) × (0.5 ml) × (2.85 mg/ml) = 45.6 mg; p38 (42 kDa): (32 samples) × (0.5 ml) × (2.1 mg/ml) = 33.6 mg. If no deterioration of the protein is observed during the course of the NMR measurements, the protein may be recycled and used again for additional screens. Also, with recent developments in NMR hardware, it is likely that protein consumption could be reduced significantly by using probes designed for smaller diameter tubes (2.5 mm), using liquid handling systems for direct injection of samples into a flow cell, or 'tubeless' NMR probe [8], or using a high-sensitivity cryogenically cooled probe [9].

### *Determination of binding specificity*

As mentioned previously, identification of ligand binding in a SHAPES screen implies only that the compound binds to the target. The observed binding could be specific, acting at either an active or a remote site, or nonspecific, with binding at many distinct sites. In the case of multisubstrate enzymes as targets, the 'active site' will be composed of several substrate binding sites, so defining a compound as binding specifically to the active site could also refer to any substrate subsite. When a compound is present in a 10–20-fold excess over the target, it is not surprising that the compound will bind to several sites on the protein. For example, even a compound with nanomolar affinity for a target will show a very high affinity for the tight binding site, and once that site is saturated, weaker (e.g. mM) binding might occur at one or more different sites. We therefore routinely carry out simple competition experiments for each hit in a SHAPES screen versus a known high-affinity active-site-directed inhibitor. In these studies, we define 'specific binding' as binding to a particular substrate site, for which the observed experimental line broadening or tNOE effects can be significantly attenuated or reversed in the presence of a known compound or inhibitor with a high affinity for that site. In our experience, most SHAPES hits are specific binders, but show detectable contributions from nonspecific or nonactive site binding components as well (data not shown). If no site-specific inhibitor is known, one can attempt to examine the dilution dependence of signal intensity or percent NOE as described by Murali *et al.* [10].

### *Determination of binding affinity using NMR*

Although in many cases relative binding affinities can be assessed on the basis of observed line broadening and tNOE data, it is often useful to determine binding constants for SHAPES screening hits under NMR conditions. Determining binding constants using NMR provides a more quantitative analysis of relative binding affinities of the SHAPES hits and establishes a basis for comparison with $K_i$ values determined by enzyme inhibition studies or $K_d$ values obtained with other biophysical methods, such as fluorescence measurements. Traditional methods of determining a $K_d$ value by NMR using $T_1$, $T_2$ or NOE measurements (reviewed in [4]) are very time consuming, requiring data measurements at a number of different ligand concentrations. These measurements actually involve two titrations: the addition of ligand, and the incrementation of a relaxation delay to determine, for example, $R_1$ or $R_2$ for each concentration of ligand. The time required to collect data with a reliable signal to noise ratio for longer relaxation delays makes a single $K_d$ titration prohibitively long (several days). If one now considers measuring a $K_d$ value for a series of screening hits, relaxation-based measurements become completely impractical. To address this problem, we have employed methods for $K_d$ determination based on NMR diffusion measurements [11].

It is well known that pulsed-field-gradient NMR experiments can measure the diffusion coefficients of molecules in solution [12,13]. The diffusion coefficients of a small molecule and of a protein can differ by an order of magnitude or more. If the diffusion coefficient for a small molecule is measured in the presence of a protein to which it binds, its diffusion coefficient will be a weighted average of the coefficients of the free and bound states of the ligand, and will appear less than that measured for the free ligand. The magnitude of this reduction is directly related to the bound lifetime of the ligand on the much more slowly diffusing protein, and may be used to calculate the $K_d$ value.

Figure 5 shows experimental pulsed-field-gradient NMR diffusion data for 4-(3-pyridin-4-yl-1H-pyrazol-4-yl) pyrimidine (VRT-34,306) binding to p38 MAP kinase. The decay curves show peak intensity versus $K^2$ — a parameter proportional to the square of the gradient amplitude — for protein, free ligand and ligand in the presence of protein. The dependence of peak height on gradient amplitude

**Figure 5**

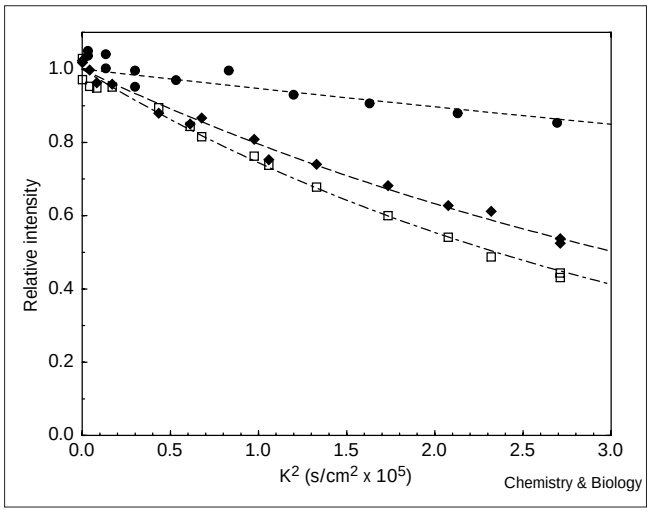


Diffusion curves showing peak integrals versus $K^2$ ($s/cm^2 \times 10^5$). Peak integrals for free 4-(3-pyridin-4-yl-1H-pyrazol-4-yl) pyrimidine, 4-(3-pyridin-4-yl-1H-pyrazol-4-yl) pyrimidine in the presence of p38, and p38 protein aromatic peaks are shown as open squares, filled diamonds, and filled circles, respectively. $K^2 = \gamma^2\delta^2G_z^2(\Delta\text{-}\delta/3)$ and $4\gamma^2\delta^2G_z^2(\Delta + (5\delta/3))$ in the case of the p38 aromatic peaks and ligand peaks, respectively.

appears greatest for the free ligand, but is attenuated in the presence of p38. Fitting of the decay curves yields values for diffusion coefficients of the protein ($D_{bound}$), the free ligand ($D_{free}$), and the ligand in rapid exchange between the free and bound states ($D_{obs}$). One can solve for percentage of ligand bound ($p_b$), and subsequently $K_d$ (described in the Materials and methods section). The decay curve shown for VRT-34,306 gives a $K_d$ value of $250 \pm 90$ μM for the p38 MAP kinase.

The advantages of determining binding constants using a diffusion-based approach are clear. First, the $K_d$ is determined under the same conditions as the NMR screening (e.g. the same NMR concentrations and buffers), so the NMR parameters observed by NMR screening may be correlated quantitatively with binding affinity. More importantly, the diffusion-based $K_d$ determination requires no titration. Titrations often demand the continued addition of ligands dissolved in solvents hostile to the protein. As ligand concentrations increase, so do the concentrations of organic solvent, making the results of such titrations difficult to interpret. Eliminating the need for titrations also reduces the time necessary for $K_d$ measurements from several days to several hours, making a series of measurements on a number of different SHAPES hits feasible. One must exercise caution, however, in interpreting these data, as small errors in measurement of diffusion coefficients may lead to large errors in the calculated dissociation constant. For example, statistical errors of 5% in measured diffusion coefficient can lead to errors of over 100% for certain ranges of binding constant (see the Materials and methods section for details). For this reason, $K_d$ values generated using diffusion methods should be considered approximate. Despite the uncertainties involved, the relative speed and ease of carrying out these measurements make them very useful for quickly discriminating binding affinities of small molecules for large proteins and for rank ordering binders to a common target.

## Follow-up strategies

### *Selection of compounds for high-throughput screening*

Effective follow-up of initial SHAPES hits is crucial if information from weak NMR binders is ultimately to lead to the discovery of potent drug leads (Figure 6). Our goal is to use the initial SHAPES hits to guide subsequent screening efforts and provide a starting point for structure-based drug design. On the basis of SHAPES hits, follow-up compounds may be chosen from in-house libraries or purchased from a database of commercially available chemicals, and screened using non-NMR-based enzymological or binding assays. In parallel, the data from SHAPES screens may be used in conjunction with computational techniques such as virtual screening to purchase additional compounds or direct the synthesis of combinatorial libraries. Other strategies, such as secondary NMR screens to assess binding affinities of follow-up compounds, soaking into crystals for characterization by X-ray crystallography, or even completely NMR-based methods similar to the 'SAR by NMR' technique [14–16], are possible as well.

Several approaches may be employed to select compounds for follow-up screening. For example, in a hypothesis-based strategy the relative binding affinities of similar SHAPES compounds are compared and used to construct a hypothesis that correlates molecular features (such as substitution pattern, location of H-bond donors and acceptors, or hydrophobic moieties) with binding. This 'structure–affinity relationship' is then used to construct a library of compounds in which the original scaffold is maintained, but the type and location of substituents is systematically varied to test the putative pharmacophore. Such a small and focused library will probably yield a high percentage of hits, but is unlikely to dramatically improve potency unless the original hit was already close to a minimum in compound space. In a framework-based strategy the original SHAPES scaffold is extended into larger frameworks by adding common rings and linkers, while the original scaffold core is allowed to vary within defined similarity limits. Because the frameworks are larger than the original compounds, they offer more potential sites of interaction with the target and hence affinity may be increased by several orders of magnitude. The overall hit rate will probably be lower than for a narrowly focused library, but the greater diversity allows better searching of compound space.

**Figure 6**

Flowchart indicating how information derived from a SHAPES screen may be integrated into the drug design process.

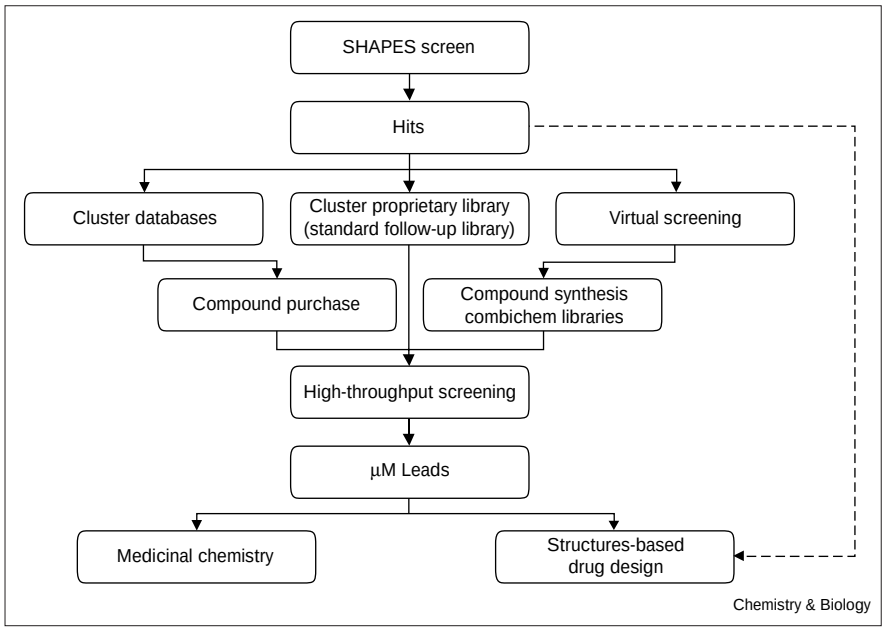


Typically, both strategies are used in combination. As some of the follow-up compounds will be relatively small (particularly those most closely resembling the original SHAPES hits), they are likely to be weak inhibitors. To detect inhibition for these compounds, follow-up assays may need to involve higher ligand concentrations than those used for general screening so it is important to remove from the follow-up libraries those compounds most likely to give false positives in enzymatic assays. We eliminate compounds that violate the 'rule of five' [17], contain inherently reactive or unstable functional groups [18], or are dyes, stains, radioactive, and so on. Only about 17% of the compounds in the ACD survive this filtering process.

For SHAPES screens carried out thus far with six proprietary targets, we have observed hit rates of 10–20% with libraries of 64–80 compounds. Of the initial SHAPES binders, 30–40% were found to be inhibitors, with $IC_{50}$ values between 150 μM and 5 mM. The NMR hits were followed up using libraries of 100–300 compounds, screened at 30 μM in enzymatic assays. The hit rates for SHAPES-derived follow-up compounds were 5–6%, which is 3–5-fold higher than for non-SHAPES compounds in the same assays (hit defined as ≥30% inhibition). The more potent hits (1–13 μM) were soaked into enzyme crystals or co-crystallized with target protein and the structures solved crystallographically.

To compare hit rates for SHAPES with those for non-SHAPES compounds using larger libraries, an independent analysis was done using 'virtual follow-up screens'. HTS data for thousands of general screening compounds were retrospectively analyzed to determine whether compounds that contain SHAPES scaffolds with known target affinity occurred more frequently as hits than compounds containing either no SHAPES scaffold or SHAPES scaffolds that did not bind the target. Compounds containing scaffolds from SHAPES binders gave a 2.5–4-fold greater frequency of hits at the 30% inhibition level, and an 8–10-fold higher frequency of hits at 50% inhibition.

### *Ligand design with SHAPES*

The following example illustrates another approach through which SHAPES screening may be employed to identify a scaffold class, and how potent target directed inhibitors might be designed on the basis of that scaffold class. Figure 7 shows binding data from a SHAPES screen of the p38 MAP kinase, along with data for selected follow-up compounds. In an initial screen, the simple imidazole core did not appear to bind p38. Several tethered bicyclic compounds containing an imidazole (or close derivative) and an aryl moiety (pyridyl, phenyl or benzoic acid) showed weak binding. As imidazole by itself does not bind, this group may form a core on which to fuse two of the tethered bicyclics or their derivatives, creating tricyclic molecules with aryl derivatives as sidechains. Two such tricyclic fusions of weak binding fragments showed much tighter binding than the bicyclic compounds, on the order of 200–300 micromolar versus 1–7 millimolar, respectively (Figure 7). Further substitution of the imidazole results in compounds similar to the trisubstituted imidazole shown in

**Figure 7**

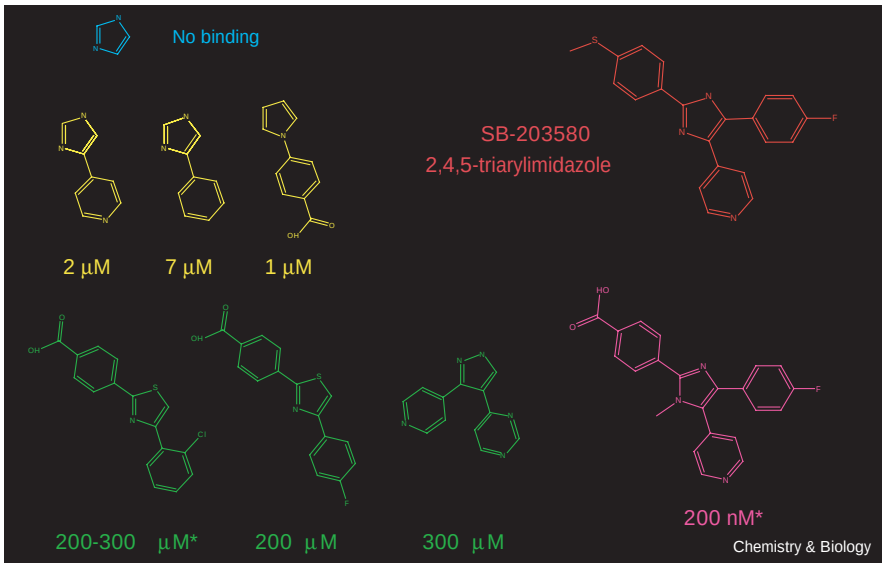


An example showing how fusion of fragments with a common scaffold can lead to a potent inhibitor. Dissociation constants of the compounds for p38 as determined by NMR diffusion measurements are given below each compound. Dissociation constants for compounds marked with asterisks could not be determined using diffusion methods. The values given for the compounds shown in green at the lower left are estimated from line broadening and transferred NOE data, and the value shown for the compound in magenta at the lower right was determined enzymatically [19].

Figure 7, a compound that binds too tightly to determine a binding constant using NMR, but has a $K_i$ value of ~200 nm in a p38 MAP enzyme assay [19]. This trisubstituted imidazole is in the same class of molecules as the SmithKline Beecham 2,4,5-triarlyimidazoles, including SB-203580, a potent and well-characterized p38 inhibitor [20].

Although the fragments used in this example were hand picked from simple, commercially available databases, it is also straightforward to employ a combinatorial chemistry strategy once a central scaffold is identified. Such an approach, coupled with HTS, takes advantage of higher throughput technologies to identify leads, rather than relying on slower, standard medicinal chemistry approaches, or attempt ligand design entirely using NMR.

Ligand design using NMR as a primary tool, such as in the SAR by NMR method [14–16], is also possible with the SHAPES library. Figure 8 shows the identification of two ligands that bind simultaneously to different sites on a target protein using simple homonuclear NMR methods. The transferred NOE spectrum shows positive cross and diagonal peaks for two compounds, VRT-19,962 and VRT-13,578, indicating both compounds bind to this 60 kDa enzyme drug target. In addition, the two compounds show intermolecular NOEs, indicating the two fragments are bound in close proximity (<5 Å) on the target. Subsequent experiments indicated unambiguously that both molecules were binding simultaneously and that VRT-19,962 binding was required before VRT-13,578 binding could occur. Such information is very useful in the process of lead generation, but, in general, observation of these events is rather fortuitous unless all possible combinations of SHAPES hits are examined.

A hybrid approach, using ideas from both SAR by NMR and SHAPES, may be used in an effort to optimize an existing lead by replacing only a fragment of the lead. For example, a binding core of a known inhibitor can be used to saturate a subsite on the target, and the SHAPES library can be screened to identify molecules that bind at an adjacent subsite. Transferred NOEs between protons on the different fragments should be observable if they are in reasonable proximity and if the fragments are in the appropriate exchange regime. This method may be useful in overcoming undesirable physical chemical properties of a lead compound resulting in poor bioavailability, pharmacokinetic or toxicological profiles.

## Discussion

We have described a general method for screening a library of ligands for target affinity using NMR spectroscopy. The NMR experiments are based on the well-known techniques of differential line broadening and tNOE, and may be applied to targets with no limitation on molecular weight and no requirement for isotope labeling. NMR SHAPES screening offers a complementary approach to standard enzymological or ligand binding HTS. Although enzymological techniques are superior when screening for inhibitors that bind in the nanomolar to micromolar range, the affinity of more weakly binding inhibitors will be below the level at which inhibition can be reliably detected. In addition, small-molecule cores that bind, but do not inhibit,

the enzymatic activity will also go undetected, eliminating the possibility of designing molecules that take advantage of binding to a remote site on the target.

Although we have thus far examined only soluble protein targets, one can envision these techniques being applied to membrane-bound proteins in micelles or detergents, nucleic acids, large subcellular structures, or even to targets that have been chemically cross-linked to a solid support. The methods described are extremely useful in the early stages of a drug-design program, or alternatively in a well-established program in search of second or third generation compounds.

#### Comparison with other NMR-based screening techniques

It is useful to compare our SHAPES strategy with the 'SAR by NMR' method, recently described by the Abbott group [14]. SAR by NMR may best be described as a method for primarily NMR-driven ligand design. Although this method is capable of providing detailed information regarding ligand binding and localization, there are significant barriers to its implementation in a drug discovery setting.

Because SAR by NMR requires the iterative screening of large compound libraries, hundreds of milligrams of $^{15}N$-labeled protein must be produced, which is problematic for drug targets that cannot be expressed to high levels in bacterial cells. Moreover, the method may be used only for smaller drug targets (up to 20–30 kDa), and requires that the NMR assignment and structure determination for the protein be completed before screening commences. Such requirements are highly limiting, as the NMR structure of a 30 kDa protein is well beyond the reach of many protein NMR labs, and, even in labs that are capable of such a structure determination, by the time it is completed the information might be of little value.

In many regards, the SHAPES strategy is of significantly greater utility than SAR by NMR. For example, SHAPES screening does not require isotopically labeled protein, may be applied to proteins with no molecular weight limit, and does not require that the structure of the target be solved. The SHAPES strategy does not involve iterative screens of large libraries, but rather a single screen of a much smaller library, so requires significantly less protein, which may be derived from either recombinant or natural sources. A disadvantage of the SHAPES approach is that although one can easily assess which small-molecule compounds bind the target, one cannot tell where these compounds bind. Follow-up experiments are needed to determine whether binding occurs at an active site, and whether the binding is specific.

Of the above considerations, perhaps the most important from a practical standpoint is target size. The SHAPES strategy evolved as method for lead generation aimed at larger molecular weight targets primarily because our in-house drug targets were all very large by NMR standards. Although deciding what makes a good drug target is a highly subjective process, in our opinion there have been only a handful of good drug targets of less than 20–30 kDa, making an approach like SAR by NMR, despite its advantages, of limited utility. To address this issue, other groups have also developed and used NMR techniques that are capable of examining larger molecular weight systems. For example, Meyer *et al.* [6] have used transferred NOE techniques to examine binding of a library of oligosaccharides to a protein target. In addition, Hajduk *et al.* [5] have described methods using relaxation and diffusion edited NMR difference methods to examine mixtures of unlabeled ligands with unlabeled target proteins.

Perhaps the greatest differences between SHAPES and SAR by NMR are in their basic philosophies and objectives. SAR by NMR is a very elegant but resource-intensive method of designing very tight binding (nanomolar

**Figure 8**

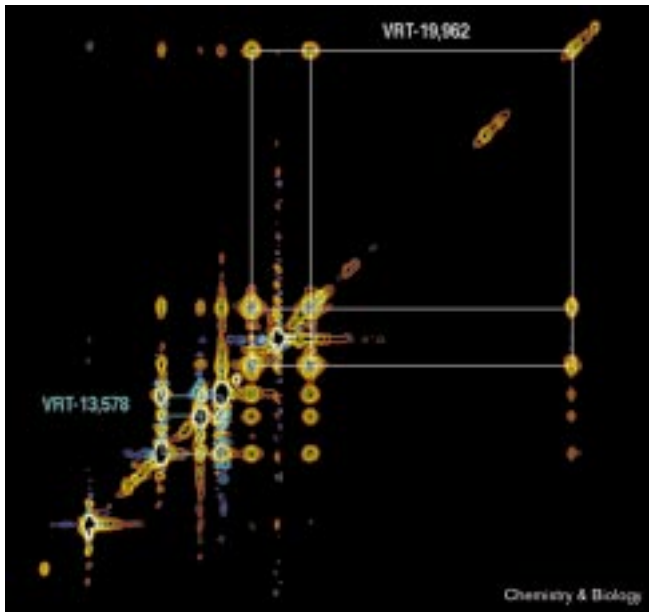


2D NOESY spectrum (τm = 80 ms) of a mixture of compounds with an undisclosed target. Two SHAPES ligands in this mixture bind the target. Intraligand cross peaks for VRT-13,578 are highlighted with blue boxes in the downfield region of the spectrum, whereas intraligand NOEs for VRT-19,962 are shown in white in the upfield region. Cross peaks falling outside the boxes are interligand NOEs, implying that these ligands bind in close proximity to one another. The screening sample contained a 10:1 ratio of ligands to protein; 1 mM ligands + 100 μM protein. Data were collected at 295 K with 1 s low power presaturation of the residual HDO signal, using a Bruker DMX-500 spectrometer. 16 Scans were collected for each of 200 $t_1$ increments. 1024 complex points were collected in the $t_2$ dimension. The $t_1$ dimension was expanded to 256 complex data points using linear prediction, then zero filled to 512 points before Fourier transformation.

**Figure 9**

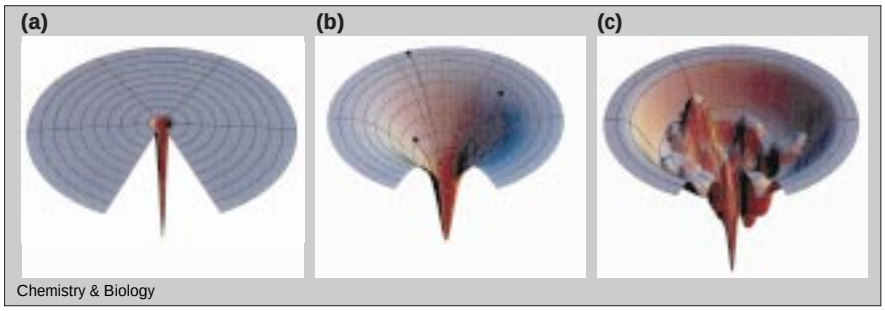


Inhibitory landscapes illustrating the conceptual approach to maximum inhibitory potential in compound space using (a) high-throughput screening and (b) SHAPES screening. A more realistic picture, containing peaks and valleys indicating local minima, is shown in (c). These landscapes were originally created by Dill and Chan [29] and used to describe protein folding energy landscapes. Images are used with permission.

affinity) inhibitors using NMR as the principal technology. The goal of SHAPES screening and the subsequent follow-up experiments is to establish one or more lead classes of micromolar affinity compounds that may be optimized using more efficient, higher throughput methods for compound synthesis, screening and structure-based drug design. We believe the latter methods are much better suited for exploring the larger regions of chemical diversity space necessary to maximize *in vitro* potency.

### Implementation of the SHAPES strategy and follow-up protocols

Given sufficient protein (or other target), SHAPES screening is simple to implement and uses hardware already found in a typical NMR lab. However, it is the follow up to the SHAPES hits that makes the technique versatile, and one can easily imagine different laboratories using SHAPES data to complement existing in-house technologies or expertise. For example, for a low molecular weight target that is amenable to SAR by NMR screening, SHAPES may be used as an inexpensive way to identify binders initially. SHAPES hits or follow-ups from HTS could then be examined using chemical shift perturbation experiments with $^{15}$N-labeled protein to localize the binding site and suggest modifications, or to carry out optimization of additional subsites.

In our research setting, SHAPES screening is a valuable component of a discovery effort based on numerous biophysical and computational approaches. Experimental data from SHAPES hits are integrated with computational methods such as virtual screening, where databases of molecules, either real or virtual, are assessed according to their synthetic accessibility, shape, flexibility, strain energy, similarity to known leads, pharmacophores, docking scores or other properties relevant for *in vivo* activity such as LogP or metabolic stability [2]. NMR-based screening methods may therefore be considered as an additional 'filter' in the process of virtual screening.

What are the advantages of an approach to lead generation combining NMR screening and computational protocols? The virtual screening process becomes significantly more information based; NMR screening has the potential for identifying screening compounds predisposed towards higher target affinity; starting the lead optimization process with small, soluble binding cores could increase the probability of a final drug candidate with an attractive solubility profile; and starting with a diverse set of NMR screening hits could result in a greater diversity of potential scaffolds. Having the flexibility of pursuing several compound classes at once to identify active drug molecules provides a significant advantage in a preclinical discovery program, where efforts often become trapped in a local minimum in compound space. In this situation, one typically has a compound with good *in vitro* potency but poor pharmacokinetic or toxicological properties. Although optimization of the pharmacokinetic properties might be possible, it cannot be achieved without sacrificing *in vitro* potency. Having several lead classes of micromolar potency would be very useful in such a situation, and would also be an advantage from the perspective of intellectual property protection.

Figure 9 shows a comparison of the inhibitory landscape of the SHAPES strategy with that of an approach based solely on HTS of compound libraries. Figure 9a shows a 'golf course', or flat landscape with a very narrow window into a maximum inhibitory potential. The flat landscape arises because of the insensitivity of a typical HTS assay, in which only inhibitors stronger than a certain threshold in the low micromolar or high nanomolar range may be detected. Alternatively, techniques like SHAPES are capable of detecting weak binding inhibitors (Figure 9b), and through subsequent activities such as screening of follow-up libraries, virtual screening, or design of combinatorial libraries based on SHAPES hits, will follow a more gradual descent into a hypothetical maximum inhibitory potential. Of course, this is a radically oversimplified view of two different approaches to drug discovery, and a more realistic landscape is shown in Figure 9c. This landscape has several peaks and valleys indicating local minima in compound space (i.e. trajectories from which a global maximum with regard to inhibitory potential cannot

be approached without significantly changing chemical composition). One should also note that compound space is rarely regarded as a three-dimensional space, but rather an N-dimensional space where one must optimize other parameters aside from *in vitro* potency such as pharmacokinetics, toxicology and solubility factors related to formulation. Understanding compound space in such higher dimensionality to guide compound selection remains a major focus of research efforts for computational chemists in drug discovery.

## Significance

A new method for lead generation (SHAPES) is described that uses simple NMR techniques to detect binding of a limited but diverse library of low molecular weight, soluble compounds to a potential drug target. The NMR experiments, based on the well-known techniques of differential line broadening and transferred nuclear Overhauser enhancement (tNOE), do not have a limitation on molecular weight and do not require isotope labeling of the targets. SHAPES library compounds are derived largely from molecular frameworks most commonly found in known therapeutic agents [3].

The SHAPES strategy is useful for identifying potential lead classes of drugs early in a drug design program. SHAPES screening may be easily integrated into a discovery effort based on numerous biophysical and computational approaches. Following screening, weak binding ($K_d$ ~μM–mM) hits, most of which would be missed in a standard enzymological assay, may be used to guide virtual screening of in-house or available compound databases, guide synthesis of combinatorial libraries, and bias the first compounds that undergo high-throughput screening. Data derived from several in-house drug discovery programs indicate that high-throughput screening hit rates for follow-up compounds chosen on the basis of an initial SHAPES screen are up to tenfold higher than hit rates for compounds chosen randomly. Binding data from SHAPES hits may also be used advantageously in computational efforts such as virtual screening, in which databases of molecules, either real or virtual, are assessed according to their synthetic accessibility, shape, pharmacophores, strain energy or other properties.

To demonstrate the feasibility of such an approach, we have screened the SHAPES library against a number of drug targets. We were able to rapidly differentiate between binding and nonbinding compounds in small noninteracting mixtures, demonstrating the utility of the method with several enzyme targets, including the p38 MAP kinase, a 42 kDa protein, and inosine-5′-monophosphate dehydrogenase, a multimeric enzyme target of 224 kDa.

## Materials and methods

### *NMR screening*

*Samples.* Stock solutions of SHAPES compounds were made at concentrations of 100 mM or 500 mM in *d*-DMSO and stored at –20°C. Reference NMR samples of individual compounds were made for each target in the appropriate aqueous buffer system by adding stock solution such that final compound concentrations were 1 mM in 0.5 ml. Mixture samples for p38 MAP kinase screening contained 1 mM ligands, 0.2 mM p38 MAP kinase, 25 mM deutero-Tris, 10% deutero-glycerol, 20 mM deutero-DTT at pD* = 8.4. Mixture samples for IMPDH screening contained: 1 mM compound (1), 0.1 mM IMPDH, 25 mM d-TRIS pD* = 8.4, 300 mM KCl, 5% d-glycerol, 5 mM DTT.

*Design of mixtures.* Experimental and computational protocols for mixture design may be found the Supplementary material section.

*NMR measurements.* NMR screening for all targets were carried out on a Bruker DMX 500 MHz spectrometer equipped with a sample changer and automation control software from the vendor (IconNMR, Bruker Instruments, Billerica, MA). 1D and 2D spectra were acquired at 277K (IMPDH) and 278K (p38 MAP kinase) using standard pulse sequences and phase cycling. A 1D NOESY pulse sequence was used for collection of 1D spectra due to better suppression of residual HDO. Recycle delays were chosen to ensure sufficient relaxation for the small molecule mixtures. 2D NOESY parameters were set to minimize total experiment time yet contain adequate resolution in both spectral dimensions. Details for individual experiments are given in the Figure legends. Acquisition data was processed using vendor software (XWinNMR 1.2, Bruker Instruments, Billerica, MA).

*Reducing spectral contributions from the target protein.* Relaxation filtering elements [21,22] in the pulse sequences can be used to attenuate the protein signals. Thickening agents may also be used to slow the molecular tumbling of the protein. Refer to the Supplementary material section for details.

### *$K_d$ determination by NMR diffusion measurements*

The apparent diffusion coefficient of the SHAPES ligand in the presence of protein is a weighted average, determined by the relative amounts of free and bound ligand. Specifically, if $p_b$ represents the fraction of ligands that are bound to protein, then we can express the observed diffusion coefficient as a specific case of the observable P in equation 2:

$$D_{app} = (1\text{-}p_b)D_{free} + p_b(D_{bound}) \quad (3)$$

where $D_{app}$ is the apparent diffusion coefficient of the SHAPES ligand in the presence of protein, $D_{free}$ is diffusion coefficient of SHAPES ligand by itself, and $D_{bound}$ is the diffusion coefficient of the bound ligand. This value is the same as that of the protein, and is readily measured. Using equation 3, we solve for the bound fraction, $p_b$. Then, from $p_b$, we get the $K_d$ using the well-known expression:

$$K_d = (L_{tot}/p_b)\{p_b^2\text{-}p_b(1 + P_{tot}/L_{tot}) + (P_{tot}/L_{tot})\} \quad (4)$$

where $L_{tot}$ is the total ligand concentration, and $P_{tot}$ is the total protein concentration.

The pulse sequences for measuring the diffusion included the water-sLED pulse sequence [23], and a 1D PFG-TOCSY sequence [24] employing bipolar gradients [25]. In both pulse sequences, molecular diffusion occurring during the time between the phase-encoding and phase-decoding gradients (or corresponding bipolar pairs in the TOCSY sequence) attenuates the peak heights of the resulting spectra. This attenuation is exacerbated with increasing gradient amplitude. Denoting a given peak integral by I, the attenuation is described by the decay:

$$I = A \exp(-DK^2) \quad (5)$$

In the above expression, D is the desired diffusion coefficient, A is the peak integral in the absence of the two phase encoding/decoding gradients, and $K^2$ is a factor proportional to the square of the gradient amplitude. In the case of the water-sLED sequence $K^2 = \gamma^2\delta^2G_z^2(\Delta-\delta/3)$ and for the 1D PFG-TOCSY, $K^2 = 4\gamma^2\delta^2G_z^2(\Delta + (5\delta/3))$. In both cases, $\gamma = {}^1H$ gyromagnetic ratio, $\delta$ = gradient pulse length, G = gradient amplitude and $\Delta$ = time delay between gradients (diffusion time). Fitting a data file consisting of peak integrals versus $K^2$ to equation 4 gives the diffusion coefficient, D, and the prefactor, A. The spin lock between the phase-encoding and decoding bipolar gradient pairs in the 1D-TOCSY pulse sequence provides additional relaxation filtering of the protein signals, and is helpful for a better estimation of ligand diffusion coefficients. Alternatively, the water-sLED experiment permits longer diffusion waiting times, and is preferable for measuring smaller diffusion coefficients, such as those of protein targets.

To maximize the percent of ligand bound to protein ($p_b$), diffusion measurements should be carried out with concentrations of protein and ligand at close to equimolar ratios. For the example shown in Figure 5, two NMR samples were prepared. One sample contained only 4-(3-pyridin-4-yl-1H pyrazol-4-yl) pyrimidine at 0.5 mM ,and the other contained both p38 and 4-(3-pyridin-4-yl-1H-pyrazol-4-yl) pyrimidine at 0.2 mM and 0.4 mM, respectively. As the diffusion experiments used (1D PFG–TOCSY) involve Hahn spin-echo segments, in which protein signals decay as $R_2$, this provided a serendipitous relaxation filter that suppressed most of the protein signals, with the exception of some resolved aromatic resonances. Use of the TOCSY spinlock further decreased protein resonance contributions. Appearance of protein signals in the spectrum of the p38/4-(3-pyridin-4-yl-1H-pyrazol-4-yl) pyrimidine sample, conveniently allowed measurement of diffusion coefficients for both the protein ($D_b$), and ligand ($D_f$) from the same sample.

Errors in $K_d$ values calculated from the diffusion coefficients may arise because of random errors or systematic errors arising from intermolecular cross relaxation [26]. Small errors in D arising from either source, can result in large errors in $K_d$. Errors due to chemical exchange effects can be eliminated by use of bipolar gradients [27]. Errors due to cross relaxation [26] can be minimized by working with short diffusion times. Random errors cannot be eliminated and the effects of these errors should be considered. For example, using elementary propagation of errors and assuming that errors in the measured diffusion coefficients are statistically independent and ~5%, we find in terms of fractional error: 5% error in $D_a \rightarrow$ 98% error in $p_b \rightarrow$ 125% error in $K_d$, for $K_d$ = 1 mM; 5% error in $D_a \rightarrow$ 20% error in $p_b \rightarrow$ 57% error in $K_d$, for $K_d$ = 100 μM; 5% error in $D_a \rightarrow$ 10% error in $p_b \rightarrow$ 127% error in $K_d$, for $K_d$ = 10 μM.

Clearly, the fractional precision suffers at both low- and high-affinity ends, but the absolute precision gets worse only for weaker binders. Uncertainties in $K_d$ will decrease or increase for lower and higher error estimates, respectively, for $D_a$, $D_b$, and $D_f$.

NMR diffusion experiments were carried out at 295 K on a Bruker DRX-800 MHz spectrometer. All gradient pulses are rectangular shaped and applied along the z-axis. The diffusion time, $\Delta$, was 30 ms and 148.5 ms for ligand and protein measurements, respectively. For the free ligand and ligand–protein samples, 16 data sets were recorded corresponding to increasing strengths of the phase encoding/decoding gradients. Data sets for both samples were Fourier-transformed using vendor software (XWinNMR 1.2, Bruker Instruments, Billerica, MA). The resulting peaks were integrated and diffusion coefficients were determined by using the Levenburg–Marquardt algorithm [28] to fit the integrals to equation 5. The 4-(3-pyridin-4-yl-1H-pyrazol-4-yl) pyrimidine sample yielded $D_{free} = 0.294 \pm 0.012 \times 10^{-5}$ cm$^2$/s. The resolved ligand resonances of the p38/4-(3-pyridin-4-yl-1H-pyrazol-4-yl) pyrimidine sample yielded $D_{app} = 0.229 \pm 0.002 \times 10^{-5}$ cm$^2$/s, whereas the resolved p38 aromatic resonances of the same sample yielded $D_{bound} = 0.053 \pm 0.002 \times 10^{-5}$ cm$^2$/s.

### *Selection of follow-up compounds*

*Pharmacophore-based search.* A pharmacophore model was constructed using three structurally similar compounds for which only NMR screening data were available. Models for the three-dimensional structures of the hits were constructed and minimized in Insight (version 950, Molecular Simulations, Inc.), and then overlapped using molecular field similarity optimization. The models were first pre-aligned using their moments of inertia, then superimposed using an equal combination of steric shape and electrostatic potential factors. A pharmacophore hypothesis was constructed, comprised of approximate distances between a common aromatic centroid, aliphatic carbon, and two hydrogen bond donors, from the superposition giving the highest similarity score. Using this model, a three-dimensional substructure search was then executed using Isis (v. 2.1.3, MDL Information Systems, Inc.) against minimized structures from the ACD.

*Structure-based searches.* ACD compounds resembling the screening hits were identified by performing substructure searches (100% identity, permitting heteroatoms at selected sites) and similarity searches (substructure search at 90% similarity, superstructure search at 40–70% similarity) using Isis. Some screening hits possessed a large number of commercially available analogs. In such cases, the number of candidates was reduced by giving preference to compounds containing sidechains commonly found in known drugs (see Supplementary material section), particularly those conferring aqueous solubility: carboxy, methoxy, hydroxy, N-methyl, amino, sulfo, carbonyl, and $CH_2$-OH .

*Selection filters.* The raw list of candidate compounds was subjected to several stages of filtering. To remove compounds likely to interfere with enzymatic assays (~20% of the total), highly colored compounds (e.g. biological stains, industrial dyes), radioactives, and solutions (e.g. forensic HPLC standards) were removed. Purchasing delays were minimized by purchasing only from a list of ten preferred vendors. To remove compounds with non-drug-like properties, a REOS filter was applied (REOS = rapid elimination of swill) [2], that incorporates rule of five [17], MW, logP and functional group criteria [18]. Approximately 83% of the compounds from the ACD are eliminated by the filtering process, most of them in the latter step.

*Virtual follow-up screening analysis.* In this analysis, HTS assay data were used to model the results of hypothetical follow-up assays for several enzyme targets. First, SHAPES screening was carried out, and core scaffolds were identified for hits that have inhibitory activity (typically yielding 5–6 classes of scaffolds). Previously collected 30 μM assay data (for between 1700 and 26,500 compounds) were edited to remove compounds that appeared to interfere with the assay (e.g. absorbed at the readout frequency of a colorimetric assay) or gave more than 30% relative standard deviation between replicate data points. In addition, the analysis was restricted to compounds that passed the REOS filter, because only those compounds would actually be selected for a typical follow-up library. Separate analyses were carried out by defining 'hits' as compounds with either 30% or 50% inhibition at 30 μM. An Isis substructure search was then used to identify the subset of compounds screened by HTS that contained scaffolds from SHAPES inhibitors; these subsets comprised the virtual follow-up libraries. The hit rates for the virtual follow-up libraries were defined as the percentage of compounds that were hits in the enzymatic assay. These hit rates were compared to those obtained for the pooled set of all other compounds in order to compute the difference in the frequency of hits between SHAPES and non-SHAPES derived compounds.

### *Supplementary material*

Supplementary material including the frequence of occurrence for sidechains most commonly found in drugs, details about the experimental and computational protocols for mixture design and details about reducing spectral contributions from the target protein is available at http://current-biology.com/supmat/supmatin.htm.

## Acknowledgements

We would like to thank Pat Walters for use and help with software tools for chemical database analysis. We also acknowledge Steve Chambers, Joyce Coll, Maureen Dwyer, Matt Fitzgibbon, Ted Fox, John Fulghum, Kurt Morgenstern and Elmar Nimmesgern for expression, purification and stability analysis of the numerous proteins used in these studies, and Cameron Stuver for enzymological data.

## References

1. Agrafiotis, D.K., Myslik, J.C. & Salemme, F.R. (1998). Advances in diversity profiling and combinatorial series design. *Mol. Divers.* **4**, 1-22.
2. Walters, W.P., Stahl, M.T. & Murcko, M.A. (1998). Virtual screening – an overview. *Drug Discov. Today* **3**, 160-178.
3. Bemis, G.W. & Murcko, M.A. (1996). The properties of known drugs. 1. Molecular frameworks. *J. Med. Chem.* **39**, 2887-2893.
4. Ni, F. (1994). Recent developments in transferred NOE methods. *Prog. Nucl. Magn. Reson. Spectrosc.* **26**, 517-606.
5. Hajduk, P.J., Olejniczak, E.T. & Fesik, S.W. (1997). One-dimensional relaxation – and diffusion-edited NMR methods for screening compounds that bind to macromolecules. *J. Am. Chem. Soc.* **119**, 12257-12261.
6. Meyer, B., Weimar, T. & Peters, T. (1997 ). Screening mixtures for biological activity by NMR. *Eur. J. Biochem.* **246**, 705-709.
7. Lepre, C.A. & Moore, J.M. (1998). Microdrop screening: a rapid method to optimize solvent conditions for NMR spectroscopy of proteins. *J. Biomol. NMR* **12**, 493-499.
8. Keifer, P.A. (1997). High-resolution NMR techniques for solid-phase synthesis and combinatorial chemistry. *Drug Discov. Today* **2**, 468-478.
9. Busse-Grawitz, M.E. & Roeck, W. (1998). NMR probe head with cryogenically cooled preamplifiers. United States Patent 5, 814,992. USA, Spectrospin AG, Fallenden, Switzerland.
10. Murali, N., Jarori, G.K., Landy, S.B. & Rao, B.D. (1993). Two-dimensional transferred nuclear Overhauser effect spectroscopy (TRNOESY) studies of nucleotide conformations in creatine kinase complexes: effects due to weak nonspecific binding. *Biochemistry* **32**, 12941-12948.
11. Lennon, A.J., Scott, N.R., Chapman, B.E. & Kuchel, P.W. (1994). Hemoglobin affinity for 2,3-bisphosphoglycerate in solutions and intact erythrocytes: studies using pulsed-field gradient nuclear magnetic resonance and Monte Carlo simulations. *Biophys. J.* **67**, 2096-2109.
12. Stejskal, E.O. & Tanner, J.E. (1961). Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *J. Chem. Phys.* **42**, 288-292.
13. Stilbs, P. (1987). Fourier transform pulsed-gradient spin-echo studies of molecular diffusion. *Prog. Nucl. Magn. Reson. Spectrosc.* **19**, 1-45.
14. Shuker, S.B., Hajduk, P.J., Meadows, R.P. & Fesik, S.W. (1996). Discovering high-affinity ligands for proteins: SAR by NMR. *Science* **274**, 1531-1534.
15. Hajduk, P.J., *et al.*, & Fesik, S.W. (1997). NMR-based discovery of lead inhibitors that block DNA binding of the human papillomavirus E2 protein. *J. Med. Chem.* **40**, 3144-3150.
16. Hajduk, P.J., *et al*., & Fesik, S.W. (1997). Discovery of potent non-peptide inhibitors of Stromelysin using SAR by NMR. *J. Am. Chem. Soc.* **119**, 5818-5827.
17. Lipinski, C.A., Lombardo, F. & Dominy, B. (1997). Experimental and computation approaches to estimate solubility and permeability in drug discovery and development settings. *Adv. Drug Deliv. Rev.* **23**, 3-25.
18. Rishton, G.M. (1997). Reactive compounds and *in vitro* false positives in HTS. *Drug Discov. Today* **2**, 382-384.
19. Wilson, K.P., *et al*., & Su, M.S. (1997). The structural basis for the specificity of pyridinylimidazole inhibitors of p38 MAP kinase. *Chem. Biol.* **4**, 423-431.
20. Cuenda, A., *et al*., & Lee, J.C. (1995). SB 203580 is a specific inhibitor of a MAP kinase homologue which is stimulated by cellular stresses and interleukin-1. *FEBS Lett.* **364**, 229-233.
21. Campbell, I.D., Dobson, C.M., Williams, R.J. & Wright, P.E. (1975). Pulse methods for the simplification of protein NMR spectra. *FEBS Lett.* **57**, 96-99.
22. Glaudemans, C.P., Lerner, L., Daves, G.D., Jr., Kovac, P., Venable, R. & Bax, A. (1990). Significant conformational changes in an antigenic carbohydrate epitope upon binding to a monoclonal antibody. *Biochemistry* **29**, 10906-10911.
23. Altieri, A.S., Hinton, D.P. & Byrd, R.A. (1995). Association of biomolecular systems via pulsed field gradient NMR self-diffusion measurements. *J. Am. Chem. Soc.* **117**, 7566-7567.
24. Birlirakis, N. & Guittet, E. (1996). A new approach in the use of gradients for size-resolved 2D-NMR experiments. *J. Am. Chem. Soc.* **118**, 13083-13084.
25. Wider, G., Dötsch, V. & Wüthrich, K. (1994). Self-compensating pulsed magnetic-field gradients for short recovery times. *J. Magn. Reson. Ser. A* **108**, 255-258.
26. Chen, A. & Shapiro, M.J. (1999). Nuclear Overhauser effect on diffusion measurements. *J. Am. Chem. Soc.* **121**, 5338-5339.
27. Chen, A., Johnson, C.S., Lin, M. & Shapiro, M.J. (1998). Chemical exchange in diffusion NMR experiments. *J. Am. Chem. Soc.* **120**, 9094-9095.
28. Press, W.H., Teukolsky, S.A., Vetterling, W.T. & Flannery, B.P. (1996). *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge, Cambridge University Press.
29. Dill, K.A. & Chan, H.S. (1997). From Levinthal to pathways to funnels. *Nat. Struct. Biol.* **4**, 10-19.

**Because *Chemistry & Biology* operates a 'Continuous Publication System' for Research Papers, this paper has been published via the internet before being printed. The paper can be accessed from http://biomednet.com/cbiology/cmb – for further information, see the explanation on the contents pages.**