MedChemComm

Cite this: *Med. Chem. Commun.*, 2012, **3**, 174

www.rsc.org/medchemcomm

CONCISE ARTICLE

SAR mining and its application to the design of TRPA1 antagonists

James E. J. Mills,* Alan D. Brown, Thomas Ryckmans, Duncan C. Miller, Sarah E. Skerratt, Carolyn M. Barker and Mark E. Bunnage

Received 18th August 2011, Accepted 17th September 2011
DOI: 10.1039/c1md00213a

Given the large amounts of screening data now available, empirical methods derived from matched-molecular pairs are being used as a means for suggesting bioisosteric replacements to the medicinal chemist. The pairwise analysis of compounds has been extended to the pairwise analysis of series to bring further context to these suggestions. A validation dataset derived from recent literature has been used to demonstrate that, given a series of active compounds, this approach would be expected to predict a more potent compound, if it exists, in around 46% of cases. The approach has been successfully applied to a series of TRPA1 antagonists.

Introduction

Bioisosteres are substituents or groups with similar physical or chemical properties that produce, when introduced into a given molecule, broadly similar biological properties. In the optimisation of a lead compound, substituents or functional groups are varied with the aim of not only improving potency against the biological target of interest, but also parallel optimisation of properties such as metabolic stability and selectivity over secondary targets. The concept of bioisosterism is often utilised whilst carrying out such modifications. To assist with this process, a range of empirical and property-based methods has been developed to select potential bioisosteres of given functional groups that may retain or improve potency.[1,2]

Empirical methods for bioisostere selection have evolved from tapping into the personal experience of the individual medicinal chemist into more sophisticated *in silico* approaches that incorporate the collective experiences of many medicinal chemists, for example the rule-based software Drug Guru.[3] Beyond this, access to large databases of published and corporate structure activity data has led to large-scale mining of bioisosteric pairs in which both structure and biological activity are compared. Methods such as Matched Molecular Pairs[4] (MMPs) and WizePairZ[5] have been developed to mine the vast amount of information available in these datasets. Further analysis of MMPs has led to approaches for R-group modification that are likely to give rise to more potent compounds either in the general case,[6] a family-specific context,[7] or compounds with an alternative activity profile.[8]

Similarly, property-based methods have evolved from simple decision-tree based methods such as the Topliss tree[9] to large-scale similarity searches in which the properties of a given group are matched against the calculated properties of a set of potential isosteric fragments;[10] these methods have been implemented in software packages such as Brood,[11] QID[12] and Novartis's "Substituent Bioisosteric Search".[10]

In this communication, we wish to present our own empirical approach to series optimisation that combines pairwise comparisons and automated structure–activity relationships (SAR) analysis. This approach takes the bioisostere concept a step further, developing the pairwise analysis of compounds into a pairwise analysis of series that allows us to utilise historical SAR more fully in series optimisation.

A novel approach to utilising prior SAR

In the traditional approach to isostere identification, the medicinal chemist designer asks the question "When group A has been used in other series against other targets, which other groups have given rise to similar biological activity? Can I apply these learnings to my series?". The designer may take this information from their own personal (or colleague) experience, knowledge of functional group replacements from literature sources, or from *in silico* analyses of databases of structure–activity data. However, if there are already further data available from the same series (*e.g.* group B gave rise to a 10-fold more potent compound than group A), it should be possible to refine the suggestions based on the activity of these analogues to suggest more context-specific alternative groups. In other words, the inclusion of more data should implicitly build up a more accurate picture of the binding pocket and therefore enable more useful predictions.

A preliminary analysis based on Pfizer in-house data shows that there are differences in pocket properties across different SAR types (Fig. 1). Series (defined in this paper as a group of compounds in which all variation occurs at one attachment point) were classified on the basis of the relationship between potency and clogP, and pockets were classified on the basis of their degree of burial.[13] At the simplest level, a series in which potency is apparently driven by clogP is more likely to occur

Pfizer Global Research and Development, Ramsgate Road, Sandwich, Kent, CT13 9NJ. E-mail: james@bespokecompchem.co.uk

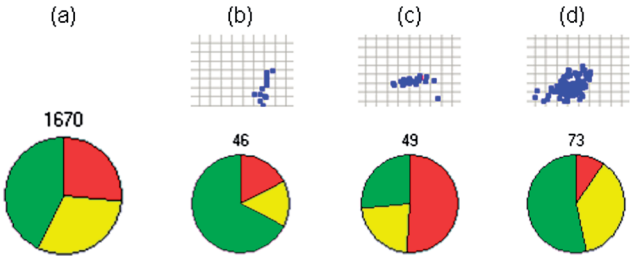

Fig. 1 Exploring the relationship between SAR type and the properties of the sub-pocket known, from ligand-protein co-crystal information, to contain the region of the molecule that varies. (a) 1670 binding sub-pockets were assigned as high (green), medium (yellow) or low (red) burial. 168 of the SARs for series were manually classified as (b) spiky, (c) flat or (d) clogP-driven, according to the plotted relationship between pIC50 and clogP (example shown above each pie chart). Spiky SAR (large changes in potency that are independent of clogP) tend to occur more often in buried pockets, flat SAR tend to occur more often when a group is solvent–exposed, and clogP-driven SAR rarely occur when the group is solvent-exposed.

when the point of variability is located in a sub-pocket of high burial. Therefore, if a functional group is known to be well-buried, it is quite likely that a compound of higher clogP could increase potency further. It should be noted however that these are merely trends, and that there are examples of series with clogP-driven SAR that have solvent-exposed sites of variability.

The approach described in this communication builds on the pairwise analysis of compounds to incorporate more than two compounds from each series in the data-mining step. Thus, once the initial members of a series (the target series) of compounds have been synthesised and biologically evaluated, the designer can examine a structure–activity database across multiple targets and series and determine whether a similar trend in activity has been observed previously. If such corresponding series are identified, the designer can extract further information from the SAR in these series (the donor series) to identify functional groups that resulted in equivalent or superior activity. Groups that gave a potency advantage in the donor series and that have yet to be incorporated into the target series can then be considered for incorporation into the target series (Fig. 2). We envisaged that, by implementing this SAR-mining approach, we would be able to aid the designer in selecting functional groups with an increased probability of achieving improved potency. This should contribute to an improvement in the efficiency of the lead optimization process through a reduction in the number of compound syntheses required to achieve target compound profiles.

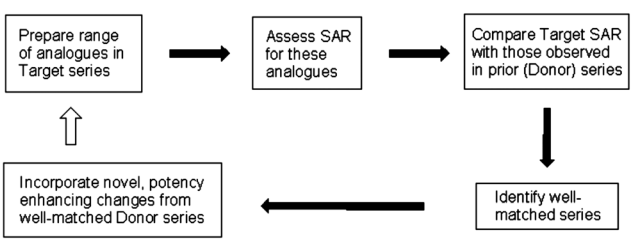

Fig. 2 Schematic illustrating the process of SAR mining and its application to compound design.

To illustrate the potential utility of this methodology, we initially examined SAR described in the recent medicinal chemistry literature. We then successfully applied the approach to a series in an active ion channel project.

Methods

Pairwise analysis

Two data sources for bioisostere mining were used in this work. The Pfizer internal database of IC_{50} and EC_{50} values contains endpoints covering around 2 million compounds and 400 targets that could be unambiguously assigned. An alternative dataset of some 117,000 compounds covering 730 targets has also been assembled from a number of publicly available sources.[14] This dataset, although useful, tends to contain fewer compounds per series, since organisations tend not to release all SAR generated into the public domain. Indeed key data on inactive compounds are often not included at all in publications. The database from publications therefore provides fewer examples of matched molecular pairs and therefore will be used in this publication solely for method illustration.

There are a number of recently published methods that can be used to calculate matched molecular pairs, all of which could be applied equally well to our datasets.[15,16] In this work, a simple substructure search based method (named SWAP) was used to probe the databases and identify pairs of molecules differing only in the replacement of a user-defined input functional group. The algorithm is illustrated in Fig. 3, illustrating how it identified a pair of Factor Xa inhibitors from the literature.[17]

A set of 400 functional groups was selected from the most common functional groups used in Pfizer drug discovery programs. The SWAP methodology was applied to these groups to generate a secondary database (the SWAP database) containing matched pairs of compounds differing structurally by a single functional group change from one of these groups. In each case, both members of the pair had been evaluated for activity at the same biological target. For every pair, each compound's structure, unique ID, potency, assay identifier and

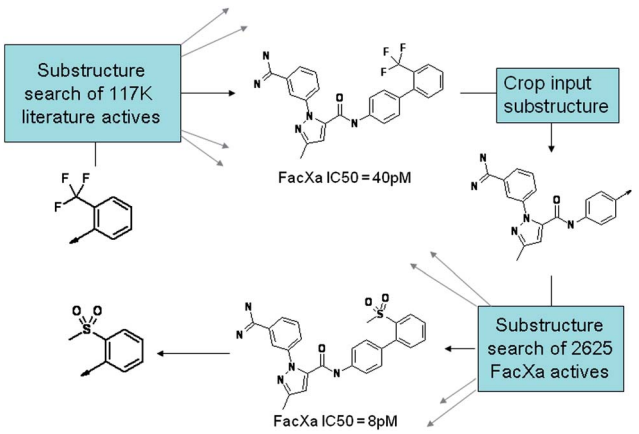

Fig. 3 Schematic workflow for SWAP algorithm, illustrating the search for isosteres of 2-CF_3-Ph. In this case, the 2-methylsulphonyl-phenyl group is identified as an isostere that provides a 5-fold improvement in the potency of a Factor Xa inhibitor.

molecular target were included. In order for comparisons to be rigorous, pairs of compounds should only be considered if they have been screened in the same assay. This for example avoids the mixing of functional and binding data, which can introduce unnecessary errors into analyses of this nature.

Identification of donor series with well-matched SARs

It would not be expected that a given isosteric replacement improves potency across all targets. Pockets that bind a given functional group are likely to have significant differences in the nature and spatial orientation of the interacting amino acid side-chains. It was therefore important to determine a process to triage the large number of potential isosteres identified from the SWAP methodology to allow the medicinal chemistry designer to focus resources on those with the highest probability of success in the series and protein target of interest. One process to triage the number of potential isosteric replacements was to widen the approach from simple compound pairwise analysis to series pairwise analysis. Comparison of these historical data with known SAR data from the target series could indicate to what extent the donor series could inform the design of the test series. In other words, we needed to determine how well matched a historic series was before we utilised its SAR. For example, if a replacement for a phenol was being sought and in the test series it was known that a pyridine was more potent than the phenol, it would seem logical to examine SAR from other series in which replacing a phenol with a pyridine had also improved potency by a similar magnitude. This is a purely empirical approach, in that the precise origins of the SAR similarity are not explored. The hypothesis is that by bringing in further series SAR, the dataset of matched pairs can be enriched more towards those that are more relevant. This could be because the binding pockets are similar, especially if the SARs come from targets from the same family, or that bind the same endogenous ligand. However, there could equally well be alternative explanations, for example the presence of similar imposed steric and conformational constraints in the two series.

Members of the target series were selected with a single template bearing R-group modifications at the same template attachment point, for which biological data were available for all analogues in the set. The SWAP database was then interrogated to find all other series for which there were at least 4 compounds containing the same combination of R groups as the target series. Any such series was analysed further to assess the correlation with the target series. In a plot of donor series activity *vs.* target series activity for matched compounds, if a donor series were to perfectly match the target series, a straight line with a gradient of 1 would be expected.

The initial approach used was to assess these series correlations by eye, since a number of factors needed to be considered. Confidence in the relevance of a correlation was more likely to be increased by:

- a high degree of correlation of the activity values, taking into account that fact that each IC_{50} value has an associated experimental error
- a large range of potency values covered by the correlation, in part to mitigate the risk of experimental error giving rise to false positive correlations between series with small ranges of potency
- an even spread of potency values within each series, to avoid the issue created by one outlier dominating the relationship
- the two series being active at similar protein targets, or targets with the same endogenous ligand, to increase the confidence that the donor series is more relevant to the target series.

However, in our experience, a genuinely useful correlation does not necessarily require all of these criteria to be met. Attempts were made to automate the recognition of similar SARs but assigning relative weightings to each of the criteria listed above did not yield rules that could be applied consistently with success (data not shown).

Results and discussion

Pairwise analysis

A typical set of results obtained using the SWAP methodology is illustrated using the example of replacing a 2-trifluoromethylphenyl group. Table 1 shows the performance of the most commonly observed potential isosteres, in terms of the likelihood of retaining potency within 3-fold of the starting group. Although there are broad trends relating increase of potency with increase in clogP (in that chloro substitution increases potency more often and methoxy substitution decreases potency more often), there are enough examples of the opposite behaviour to warrant further explanations. This analysis identified the 2-chlorophenyl group as an isostere with a statistically significant increased (chi-squared test $P < 0.02$) probability of retaining and increasing potency.

However, analysis by gene family (Table 2) shows that this isosteric replacement is more successful in some target families than in others. For example there are 262 pairs of compounds in the database with this transformation in which the potency difference has been measured against kinase family targets, with 42% instances of the replacement resulting in an improvement in potency and 53% instances of retaining potency. Only in 8% of cases did this change result in a reduction in potency against kinases. However, the same transformation is more likely to result in similar potency values when applied to ion channel targets (79%) or to decrease potency against GPCRs (20%). Clearly introducing further context, in this case target type, can add value to predictions arising from pairwise analysis. It should be noted here that although these differences are statistically

Table 1 Most common replacements for 2-trifluoromethyl phenyl as observed in the Pfizer dataset, and their propensities (expressed as a decimal 0–1) to increase, retain or decrease potency, using a 3-fold window to define an increase or decrease in potency. The average change in clogP observed for the molecular pairs is also reported

Group	Npairs	Increase	Retain	Decrease	ΔclogP
2-OMe-Ph	1212	0.13	0.70	0.17	−0.53
Ph	1009	0.24	0.56	0.21	−0.49
2-Cl-Ph	882	0.24	0.61	0.15	+0.26
2-Me-Ph	781	0.22	0.60	0.18	−0.01
2-F-Ph	763	0.22	0.58	0.20	−0.31
3-OMe-Ph	740	0.18	0.60	0.22	−0.52
4-OMe-Ph	716	0.21	0.50	0.28	−0.52
4-F-Ph	652	0.24	0.53	0.23	−0.32
3-Cl-Ph	601	0.25	0.56	0.18	+0.26
4-Cl-Ph	585	0.27	0.48	0.25	+0.25

Table 2 Propensity (expressed as a decimal 0–1) of the 2-trifluoromethylphenyl to increase, maintain or decrease potency relative to the equivalent 2-chlorophenyl compound, assessed for different target classes

Target Class	Npairs	Increase	Retain	Decrease
GPCR	433	0.21	0.60	0.19
Ion channel	234	0.10	0.79	0.11
Kinase	262	0.39	0.53	0.08
Protease	83	0.29	0.68	0.03
Transporter	115	0.39	0.58	0.03

significant (chi-squared test, $P < 0.01$) when taken at face value, no normalisation was carried out to prevent a small number of targets dominating this analysis, a problem prevalent in the analysis of corporate data. However, isostere performance has been shown to vary across target families by previous studies.[6,7]

An alternative way to add further context to this pairwise analysis was to introduce another compound from the same series. Table 3 illustrates the results when the relative potency of the unsubstituted phenyl group was known. In this particular case (and others, not reported here), local SAR was a more useful guide to prediction than the nature of the target. If the phenyl group gave rise to a more potent compound than the 2-trifluoromethyl-substituted phenyl group, the equivalent 2-chlorophenyl compound was more potent 73% of the time. If the phenyl group gave rise to a less potent compound, this proportion reduced statistically significantly (chi-squared test, $P < 0.01$) to 6%.

Identification of similar SARs

The process for identification of similar SARs is illustrated in Fig. 4 for a test series taken from Zablocki *et al.*,[18] in which the SAR of adenosine-2B (A2B) antagonists was presented. A number of series containing the same combination of R groups were identified in the collated publicly available activity dataset, including the four in Fig. 4. Of these series, only the 5-HT_4 series[19] was defined by eye to be similar. Further mining within this donor series indicated that the 2-substituted pyrazine was more potent against the 5-HT_4 receptor than any of the other compounds in this series. This would suggest that this might be a productive group to introduce into the A2B test series to generate a more active compound. However, as this was a literature example, it was not possible to assess whether this replacement would have been successful. A more generic

Table 3 Propensity (expressed as decimal 0–1) of 2-chloromethylphenyl to increase, retain or decrease potency relative to 2-trifluoromethylphenyl in the cases in which the equivalent unsubstituted phenyl has an increased, retained or decreased potency relative to the 2-trifluoromethylphenyl compound

		2-chloromethyl potency is		
Phenyl potency is	Ntriplets	Increased	Retained	Decreased
Increased	101	0.73	0.26	0.01
Retained	287	0.26	0.67	0.09
Decreased	108	0.06	0.56	0.39

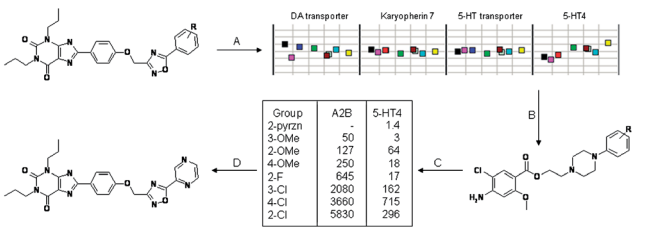

Fig. 4 Example of SAR mining. A. Equivalent series were sought for an A2B antagonist series, plotting pIC50 of donor series compounds against pIC50 for the equivalent A2B compounds. B. Triage of the donor series reduced the four illustrated series down to the 5-HT4 series. C. Further mining of the 5-HT4 series identified 2-pyrazine as a substituent generating a further improvement in potency. D. Incorporation of this moiety into the A2B series was predicted to improve potency.

validation experiment was therefore carried out making use of a larger sample of literature series.

Prediction of compounds with improved potency

To validate the applicability of this SAR-mining method and its expected success rate in the general case, the single-point of variation SAR was extracted from around 50 recent publications to serve as a test set. For each series, the most potent compound was removed and the ability of the methodology to predict the success of this compound was tested. Donor series in the Pfizer dataset were sought that contained the same set of R groups. These were split into series that did or did not correlate with the input SAR (in this case on the basis of visual inspection of plots akin to that shown in Fig. 4). A donor series was defined to be predictive if the predicted potency of the removed compound from the test series was, at worst, within 3-fold of the next best compound in the test series. In the 210 cases in which the series did correlate, there was a prediction success rate of 46%. This fell to 23% for the remaining 21,053 series, illustrating a 2-fold enrichment.

At the outset of this work it was unclear whether any enrichment would be achievable through the use of this SAR-mining approach because similarity of SAR between series against different targets could be coincidental, or merely clogP-driven. However, the enrichment seen suggests that this methodology can increase the probability of identifying R groups to incorporate that may give rise to analogues with improved potency. It is unsurprising that the predictivity is no higher than the 46% obtained, because no two series are going to be driven by identical issues *e.g.* the binding pockets are unlikely to be identical, so perfect SAR correlations would not be expected. However, there does appear to be value added by narrowing down suggestions to those predicted by well-matched donor series.

Application to TRPA1 compound design

The method has been applied to compound design on the TRPA1 project at Pfizer. TRPA1 is a member of the Transient Receptor Potential (TRP) ion channel family and is involved in the perception of nociceptive and inflammatory pain triggered by mechanical stimuli, endogenous mediators such as bradykinin and 4-hydroxynonenal, peroxides, prostaglandins and

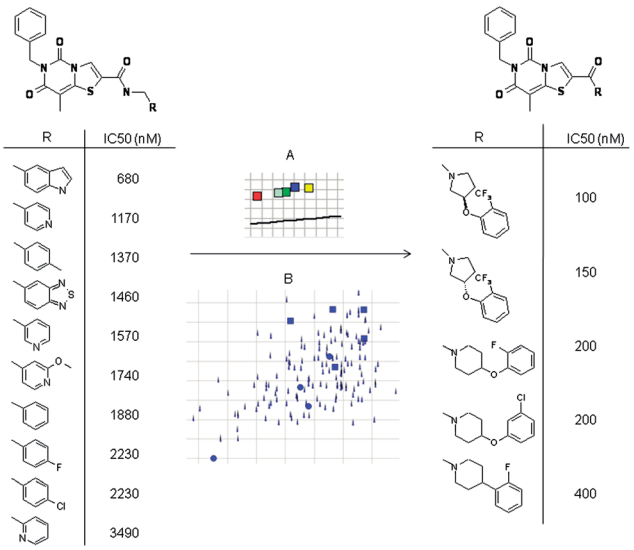

Fig. 5 Application of SAR mining to a TRPA1 series. A. SAR mining identified a 5-HT6 series with similar SAR, in that the plot of pIC50 for the 5-HT6 compounds against pIC50 for the matched TrpA1 compounds followed the same gradient as the line of equivalence (shown in bold). B. The plot of pIC50 *vs.* clogP for the matched compounds (circles) and the remainder of the 5-HT6 series showed around 20 compounds with predicted further increases in potency. Five of these compounds (squares) were selected and when the R groups were introduced into the TRPA1 series, this yielded compounds with increased potency relative to the original series.

exogenous chemical substances including pungent natural compounds and environmental irritants.[20,21] In an effort to assess TRPA1 as a therapeutic target for sensory disorders, around 50 compounds had been made in a lead series (Fig. 5), with the project unable to achieve potency below a threshold of 500 nM despite preparation of compounds across a wide range of lipophilicity and molecular weight chemical space. Using the SAR mining approach on the R group structure–activity data in Fig. 5 yielded a well matched 5-HT_6 series with a similar trend in activity. This series consisted of further compounds with potency increased relative to the best R group present in the TRPA1 series. Five of these R-groups were introduced into the TRPA1 series, with all five analogues synthesised proving more potent against TRPA1 than had previously been achieved in this series. The R groups suggested by this methodology were not of common structural classes that the medicinal chemist designer would have been likely to incorporate without the guidance of the SAR-mining methodology. They were also not obvious in terms of structural similarity to the starting structures: in other words this methodology is capable of suggesting novel structural motifs for inclusion into series.

Conclusions

This work has extended bioisostere selection from the concept of pairwise compound analysis to pairwise series analysis, adding further sophistication to the search for potential bioisosteres. It has been demonstrated to provide around two-fold enrichment in identification of isosteres with improved potency over cases where SARs were not seen to correlate. Thus this methodology can be used by the medicinal chemistry designer as a tool for idea generation to increase the probability of incorporating R groups that will lead to improved potency: the selected R groups are supported by data from other targets. However, as with all *in silico* design tools, this methodology should not be used as a "black box": the ideas should be analysed and interpreted by an experienced medicinal chemist to prioritise the emerging ideas such that they also satisfy other needs of the project *e.g.* optimisation of ADMET properties or selectivity over related targets.

This work can also provide a powerful resource for increasing understanding of what drives SAR when it is combined with the increasing wealth of structural data now available to us. In this case, the expected relationship between degree of solvent accessibility and nature of the SAR was confirmed, though, perhaps more interestingly, exceptions were identified. Further work will be required to understand if there are common explanations for these exceptions, and to expand the definitions of SAR types using some of the more elegant metrics that have recently been published.[22]

Notes and references

1 S. R. Langdon, P. Ertl and N. Brown, *Mol. Inf.*, 2010, **29**, 366–385.
2 M. Devereux and P. L. A. Popelier, *Curr. Top. Med. Chem.*, 2010, **10**, 657–668.
3 K. D. Stewart, M. Shiroda and C. A. James, *Bioorg. Med. Chem.*, 2006, **14**, 7011–7022.
4 G. Papadatos, M. Alkarouri, V. J. Gillet, P. Willett, V. Kadirkamanathan, C. N. Luscombe, G. Bravi, N. J. Richmond, S. D. Pickett, J. Hussain, J. M. Pritchard, A. W. J. Cooper and S. J. F. MacDonald, *J. Chem. Inf. Model.*, 2010, **50**, 1872–1886.
5 D. J. Warner, E. J. Griffen and S. A. St-Gallay, *J. Chem. Inf. Model.*, 2010, **50**, 1350–1357.
6 P. Hajduk and D. R. Sauer, *J. Med. Chem.*, 2008, **51**, 553–564.
7 A. M. Wassermann and J. Bajorath, *Med. Chem. Commun.*, 2011, **2**, 601–606.
8 Y. Hu and J. Bajorath, *ACS Med. Chem. Lett.*, 2011, **5**, 523–527.
9 J. G. Topliss, *J. Med. Chem.*, 1972, **15**, 1006–1011.
10 P. Ertl, *Curr. Opin. Drug Discovery Dev.*, 2007, **10**, 281–288.
11 J. J. Corkery, K. E. Schmidt, B. Kelley and A. G. Skillman, *Abstracts of Papers, 239th ACS National Meeting*, San Francisco, CA, United States, March 21–25, 2010, CINF-71.
12 M. Devereux, P. L. A. Popelier and I. M. McLay, *J. Chem. Inf. Model.*, 2009, **49**, 1497–1513.
13 L. H. Jones, N. W. Summerhill, N. A. Swain and J. E. Mills, *Med. Chem. Commun.*, 2010, **1**, 309–318.
14 L. Harland and A. Gaulton, *Expert Opin. Drug Discovery*, 2009, **4**, 857–872.
15 D. Y. Haubertin and P. A. Bruneau, *J. Chem. Inf. Model.*, 2007, **47**, 1294–1302.
16 J. Hussain and C. Rea, *J. Chem. Inf. Model.*, 2010, **50**, 339–348.
17 J. Sun, H. F. Chen, H. R. Xia, J. H. Yao and B. T. Fan, *QSAR Comb. Sci.*, 2006, **25**, 25–45.
18 J. Zablocki, R. Kalla, T. Perry, V. Palle, V. Varkhedkar, D. Xiao, A. Piscopio, T. Maa, A. Gimbel, J. Hao, N. Chu, K. Leung and D. Zeng, *Bioorg. Med. Chem. Lett.*, 2005, **15**, 609–612.
19 S. Curtet, J.-L. Soulier, I. Zahradnik, M. Giner, I. Berque-Bestel, J. Mialet, F. Lezoualc'h, P. Donzeau-Gouge, S. Sicsic, R. Fischmeister and M. Langlois, *J. Med. Chem.*, 2000, **43**, 3761–3769.
20 C. M. Fanger, D. del Camino and M. M. Moran, *Open Drug Discovery J.*, 2010, **2**, 64–70.
21 P. G. Baraldi, D. Preti, S. Materazzi and P. Geppetti, *J. Med. Chem.*, 2010, **53**, 5085–5107.
22 J. Bajorath, L. Peltason, M. Wawer, R. Guha, M. S. Lajiness and J. H. Van Drie, *Drug Discovery Today*, 2009, **14**, 698–705.

