

Quantifying biogenic bias in screening libraries

Jérôme Hert, John J Irwin, Christian Laggner, Michael J Keiser & Brian K Shoichet

In lead discovery, libraries of 10^6 molecules are screened for biological activity. Given the over 10^{60} drug-like molecules thought possible, such screens might never succeed. The fact that they do, even occasionally, implies a biased selection of library molecules. We have developed a method to quantify the bias in screening libraries toward biogenic molecules. With this approach, we consider what is missing from screening libraries and how they can be optimized.

High-throughput screening (HTS) is the dominant method of lead discovery in pharmaceutical research and chemical biology. A plurality of the new chemical entities in clinical trials may have their origins in this technique, as do at least two drugs¹. Whereas these screens have been productive against traditional drug targets, such as G protein-coupled receptors (GPCRs), ligand-gated ion channels and kinases, screening libraries of synthetic molecules has been problematic for others, such as antimicrobial targets and those identified from genomic studies. The reasons for these successes and failures have been widely debated^{2–5}. From a theoretical perspective, however, one might wonder not why screens of 10^6 molecules sometimes fail, but rather how they ever succeed.

Chemical space—that is, all possible molecules—is estimated to be greater than 10^{60} molecules with 30 or fewer heavy atoms⁶; 10 μg of each would exceed the mass of the observable universe. This figure decreases if criteria for synthetic accessibility and drug likeness are taken into account and increases steeply if up to 35 heavy atoms (about 500 Da) are allowed. Posing even a modest specificity of proteins for their ligand, the odds of a hit in a random selection of 10^6 molecules from this space seem negligible.

HTS nevertheless does return active molecules for many targets; how does it overcome the odds stacked against it? One might hazard two hypotheses. First, molecules that are formally chemically different can be degenerate to a target, and many derivatives of a chemotype may have little effect on affinity. This behavior, and the polypharmacology of small molecules^{7–9}, undoubtedly contributes to screening hit rates. Such chemical degeneracy seems unlikely, however, to overcome the long odds against screening. A second explanation is that screening libraries are far from random selections, but rather are biased toward molecules likely to be recognized by biological targets. This second hypothesis seems more plausible, as many accessible molecules are likely to resemble or derive from metabolites and natural products. Some of these will have been synthesized to resemble such biogenic molecules, while others will have used biogenic molecules as a starting material. The role of bias in screening has been mooted before^{10–13}, and indeed methods to measure metabolite or natural product likeness have been reported, thus permitting the design of these

features into screening libraries^{14,15}. How such bias might be quantified relative to what one would expect for an unbiased collection (and thus its extent and impact on screening and discovery) has remained unexplored.

Quantifying library bias requires three sets of molecules: one that represents all of chemical space, one that represents molecules that proteins are intrinsically likely to recognize (defining the optimal bias) and one that represents screening libraries. The set representing chemical space previously seemed inaccessible. Recently, however, researchers have calculated all of the synthetically accessible molecules with 11 or fewer nonhydrogen (heavy) atoms composed of first-row elements (C, N, O and F); there are over 26 million of these, not allowing for stereochemistry¹⁶. Whereas these molecules are small compared to most biologically interesting compounds, this Generated Database (GDB) is comprehensive, giving us the full unbiased set within its boundary criteria. For the molecules that proteins are likely to bind—defining the bias—several sets are possible, such as those molecules that have become drugs. Indeed, several investigators have adopted this approach when asking “what is drug likeness and how can libraries be biased towards it?”^{17,18}. Here, however, we ask why one should expect to find any hits from screening, and so we need a reference set that captures protein recognition in general. For this purpose, drugs are imperfect (they reflect many other criteria, such as bioavailability) and backward-looking (they capture information only on a small number of targets). We therefore chose metabolites and natural products from the Kyoto Encyclopedia of Genes and Genomes (KEGG; 2,018 molecules) and the Dictionary of Natural Products (141,985 molecules) databases, respectively. These molecules are recognized by at least one protein in the biosphere, often many, and are out-group molecules, meaning they are uninfluenced by human invention. For the set of molecules representing screening libraries, we used those molecules that are commercially available, reasoning that most HTS libraries, even in the pharmaceutical industry, are largely composed of molecules that have been purchased from commercial vendors, or that closely resemble them (for the Molecular Libraries Small Molecule Repository (MLSMR), the US national screening collection, almost all of the $\sim 300,000$ molecules are commercially

Department of Pharmaceutical Chemistry, University of California, San Francisco, San Francisco, California, USA. Correspondence should be addressed to B.K.S. (shoichet@cgl.ucsf.edu).

Received 30 January; accepted 10 April; published online 31 May 2009; doi:10.1038/nchembio.180




Figure 1 Overlap between commercially available molecules and the GDB gives the purchasable GDB.

sourced). To compare the commercially available molecules to those of the GDB, we restricted the former by the same criteria: only purchasable molecules with 11 or fewer heavy atoms composed of first-row elements were considered. There are 25,810 such molecules in the ZINC database of commercially available molecules (<http://zinc.docking.org/>); we refer to these as the purchasable GDB (**Fig. 1**).

As we will show, when metabolites are compared to both the purchasable GDB and the full GDB, the purchasable subset is almost 1,000-fold more similar to metabolites than the overall GDB, our proxy to full chemical space. The same bias is observed when the two sets are compared to natural products. The bias grows markedly with molecular size, which suggests that this bias will be greater still among larger “lead-like” or “drug-like” molecules in screening. This is consistent with the idea that these libraries are massively and productively biased toward biogenic molecules. We leverage this observation to determine what scaffolds occur among biogenic molecules but are absent from commercially available molecules. Almost 1,300 ring scaffolds are found among natural products that are missing from commercial libraries; these scaffolds provide criteria that could be used to further increase the bias in screening libraries toward those molecules that proteins have evolved to recognize.

RESULTS

We first compared the 26 million GDB and the 25,810 purchasable GDB molecules to the metabolites and the natural products. With the widely used ECFP₄ topological fingerprints¹⁹, we calculated the similarity between each biogenic molecule and each GDB and purchasable GDB molecule. We plotted the percentage of GDB and purchasable GDB molecules that had at least one metabolite or (separately) one natural product within a certain similarity, expressed as Tanimoto score (where a similarity score of 1 indicates identity between the pair of molecules and a similarity score of 0.2 indicates a similarity so low as to be essentially meaningless). As expected, almost all GDB and purchasable GDB molecules had a similarity coefficient of 0.2 or greater to at least one metabolite or natural product (**Fig. 2a,b** and **Supplementary Figs. 1a,b** and **2a,b** online). After this plateau, the percentages decreased rapidly as the similarity threshold became more

stringent. Critically, it diminished orders of magnitude faster for the GDB than it did for the purchasable subset of the GDB. Even by a similarity coefficient of 0.32 there were tenfold more molecules from the purchasable GDB, expressed as a percentage, than GDB molecules that resembled at least one metabolite. By a similarity coefficient of 0.53 this ratio was 100-fold, and by the time full identity was reached, considering only exact biogenic molecules in the purchasable subset of the GDB, 0.83% (215 compounds) and 10.5% (2,703 compounds) of the 25,810 purchasable GDB molecules were identical to metabolites and natural products, respectively—an enrichment of 922-fold and 618-fold, respectively, compared to the full GDB. Since there are 1,000-fold more GDB than purchasable GDB molecules, this ratio implies that most of the natural products or metabolites represented in the 26 million compounds of the GDB are captured by the 25,810-molecule subset that has actually been synthesized and may be purchased. This bias was unaffected by the size of the molecules used in the reference databases, and was the same whether we used the full set of metabolites and natural products or only those compliant with the GDB rules (**Supplementary Fig. 3** online). A bias toward metabolites and natural products of similar magnitude was also found among commercially available compound sets specifically designed for screening (**Fig. 2c,d**), whereas for the MLSMR library, the molecules subject to the GDB restrictions (MLSMR GDB) were even more biased toward metabolites and natural products (**Fig. 2c,d**). In all libraries this bias increases exponentially with molecular size (**Fig. 3**), and so we expect that most screening collections, which are made up of larger

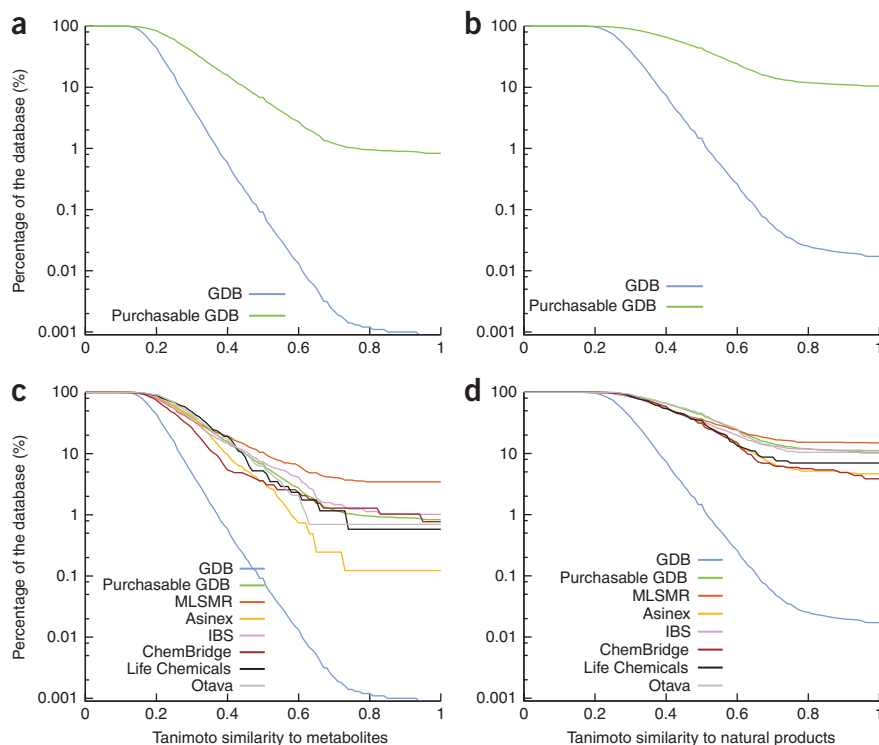


Figure 2 Compounds in screening libraries are biased toward biogenic molecules. (**a,b**) Percentage of the GDB and purchasable GDB databases as a function of the Tanimoto similarity to their nearest neighbor in the KEGG (**a**) and the DNP (**b**) databases. (**c,d**) Percentage of the GDB, the purchasable GDB, Asinex (360,042 compounds – 815 GDB compliant compounds), Chembridge (473,745 compounds – 389 GDB compliant compounds), IBS (424,806 compounds – 884 GDB compliant compounds), Life Chemicals (285,581 compounds – 172 GDB compliant compounds) and Otava (121,657 compounds – 287 GDB compliant compounds) databases as a function of the Tanimoto similarity to their nearest neighbor to the KEGG (**c**) and the DNP (**d**) databases.

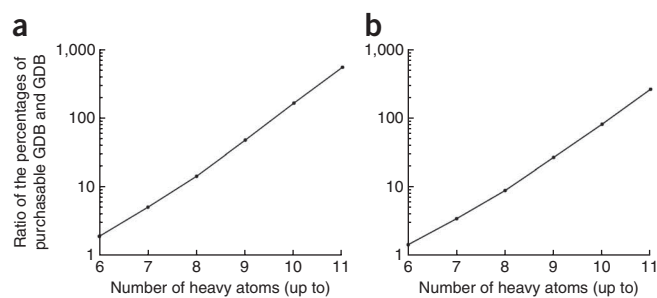


Figure 3 Biogenic bias increases with molecular size. **(a,b)** Ratio of the percentage of compounds in the purchasable GDB and GDB databases that had a similarity ≥ 0.75 to their nearest neighbor in the KEGG **(a)** and the DNP **(b)** databases versus the number of heavy atoms up to which the database compounds (in purchasable GDB and GDB) are considered.

molecules than considered here, will be even orders of magnitude more biased toward biogenic molecules.

We investigated whether there were particular chemical features that were most responsible for this bias. Whereas we do not pretend to undertake a comprehensive analysis of particular chemotypes here, a preliminary study may interest some investigators. We identified 871 bits from unfolded fingerprints that occurred in more than 1% of the metabolites or the natural products (there were 174,823 different bits among these molecules before folding). Twenty-eight of those bits (functional groups) also occurred in more than 0.5% of the purchasable GDB database and in less than 0.05% of the GDB database, which indicates at least a tenfold bias of the purchasable GDB relative to the entire GDB. The substructures corresponding to these bits illustrate some of the fragments contributing to the observed bias (**Supplementary Table 1** online). For instance, meta-substituted aryl alcohols such as 3-methylphenol occur among 1.3% of metabolites, 2.2% of natural products and 0.6% of the purchasable GDB, but only in 0.006% of the entire GDB. As an aside, notwithstanding the much greater similarity of the purchasable GDB to metabolites and natural products than the overall GDB, some functional groups were actually under-represented among the purchasable GDB. Thus, 38 bits occurred in more than 0.5% of the GDB and more than 1% of metabolites or natural products but in less than 0.2% of the purchasable GDB (**Supplementary Table 2** online). For instance, 1,2-trisubstituted epoxides occur in 2.6% of natural products and 1.7% of the GDB but in only 0.01% of the purchasable GDB. Were such substructures included in future commercial compounds, they would increase the already substantial bias still further.

If it is true that bias toward biogenic molecules contributes to the success of screening, it seems useful to ask whether this bias might be increased, productively, by adding scaffolds present among biogenic molecules that are currently unexplored in our libraries. We again turned to the metabolites, natural products and commercially available molecules, this time considering each without limit to molecular size or composition. To find scaffolds present among the two biogenic sets but absent from the commercially available molecules, we represented the molecules by their core rings (**Supplementary Fig. 4** online)²⁰. Each core ring scaffold among the biogenic molecules was matched to its counterpart among commercially available molecules, when available. There were 173, 15,637 and 29,496 unique rings among the 2,018 metabolites, 141,985 natural products and 9,131,254 commercially available molecules, respectively. 34 (20%) of the ring scaffolds among the metabolites and 12,977 (83%) of the ring scaffolds among the natural products were unrepresented among the commercial molecules. Even if

one restricts this set to molecular weight ≤ 350 Da and two or fewer stereocenters, there remain 1,891 ring scaffolds represented among natural products that have no counterpart among commercially available molecules, and by extension screening libraries.

DISCUSSION

Returning to our motivating question, a major reason why the screening of synthetic compounds ever finds notable hits is that our libraries are biased toward the sort of molecules that proteins have evolved to recognize. Thus, there are almost as many metabolites and natural products among the 25,810 purchasable GDB molecules as there are among the 26 million GDB molecules overall. This bias increases rapidly as molecules grow in size (**Fig. 3**), and the bias among larger lead-like and drug-like molecules is expected to be many orders of magnitude more still than that measured for the very small molecules explored here, where full enumeration¹⁶ allowed us to compare to a complete chemical space.

From this observation, two opposite inferences might be drawn. Since our libraries are already biased, then perhaps we should look for new screening molecules that are dissimilar to metabolites and natural products. Whereas this will certainly explore new chemotypes and ensure new scaffolds, we do not draw this conclusion. Chemical space is so large that, unless proteins are highly promiscuous, the likelihood of finding anything biologically interesting is remote. Instead, we suggest that screening libraries may be improved by increasing the bias toward biogenic molecules further still, by adding to libraries molecules resembling biogenic scaffolds that are now absent from them. After all, the bias in our current libraries is largely unintentional, the product of what organic chemists have synthesized since the birth of the field with urea in 1828 (though see refs. 21–24). This leaves room for intentional optimization. Indeed, 83% of the core ring scaffolds present among natural products are simply absent among commercially available molecules, and by extension screening libraries.

It is useful to compare these missing scaffolds with those from earlier studies that sought rings most common among drug-like molecules^{20,25–27}. An example are the six rings highlighted as characteristic of drug-like molecules in a previous study²⁵ (**Fig. 4**). Comparing these to those scaffolds found among natural products but unavailable commercially reveals molecules that are so similar to the drug-like rings that their absence from screening libraries is surprising (a few examples are given in **Table 1**). Earlier studies have suggested that scaffolds characteristic of drug-like molecules be sought when purchasing new molecules for screening; here we suggest that molecules containing scaffolds present in natural products but absent from commercial collections are places to begin expanding the biogenic chemistry available for screening. Biasing future screening


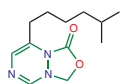
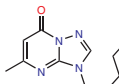
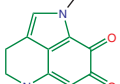
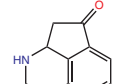
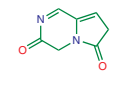
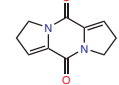
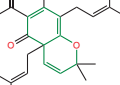
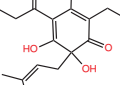
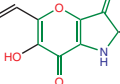
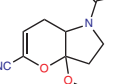
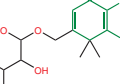
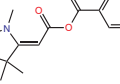
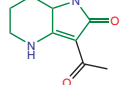
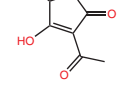
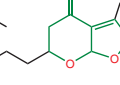
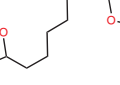
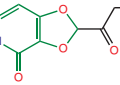
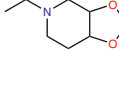
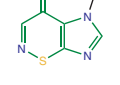
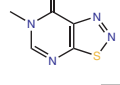
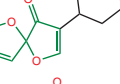
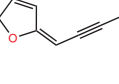
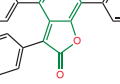
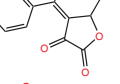
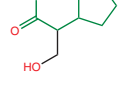
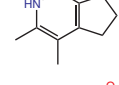
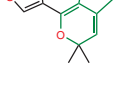
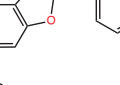
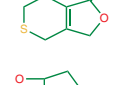
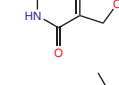
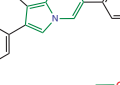
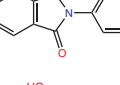
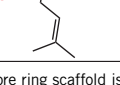
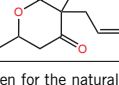
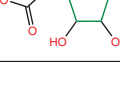
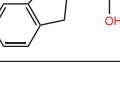


Figure 4 Core ring structures common among drugs and related molecules. From ref. 25.

Table 1 Some natural product scaffolds missing from commercial molecules

Natural product	Purchasable nearest neighbor	Natural product	Purchasable nearest neighbor
 7	 8	 25	 26
 9	 10	 27	 28
 11	 12	 29	 30
 13	 14	 31	 32
 15	 16	 33	 34
 17	 18	 35	 36
 19	 20	 37	 38
 21	 22	 39	 40
 23	 24	 41	 42

The core ring scaffold is shown in green for the natural products.

libraries to fill these systematic absences in our current collections will help address the new genomic targets with which we are increasingly confronted, and against which screening has had such mixed success.

METHODS

Chemical space. We used the Generated Database (GDB) as a proxy for chemical space^{16,28}. The GDB was obtained by exhaustively enumerating all the possible topologies for molecules composed of only first-row elements (C, N, O and F), up to 11 nonhydrogen atoms, taking into account the stability, synthetic accessibility and drug likeness of the resulting molecules. The GDB contained 26,429,328 unique compounds with no consideration of stereochemistry.

Biogenic space. We used two databases to approximate the space of molecules that occur in natural organisms: the KEGG database and the Dictionary of Natural Products (DNP). The KEGG (ligand) database contains 11,434 unique compounds²⁹, but many of these were xenobiotics such as hexachlorohexane. A xenobiotic-free subset of the KEGG database was generated by only considering primary metabolic pathways (**Supplementary Table 3** online). This subset

contained 2,018 unique compounds and was used as the reference to metabolic space. The CRC DNP³⁰ (DNP version 16.2) contained 141,985 unique structures and was used as the source of natural products.

Screening library space. ZINC³¹ contains 9,131,254 unique, commercially available compounds after adding the latest vendor catalogs and discarding some of the physical property filters normally used for docking purposes, such as molecular weight. ZINC was used as an approximation to a general-purpose screening library. To compare commercial compounds to the GDB, we filtered ZINC for molecules that conformed to the same rules as those used to generate the GDB. ZINC contains 25,810 compounds with 11 or less C, N, O or F atoms, the purchasable GDB.

The US National Institutes of Health (NIH) MLSMR contains 298,794 unique compounds (downloaded from <http://pubchem.ncbi.nlm.nih.gov/> on 18 June 2008). The MLSMR contains 866 compounds with 11 or less C, N, O and F heavy atoms—that is, it is subject to the same restrictions as the GDB. This subset was referred to as the MLSMR GDB.

Comparison of the chemical and screening spaces to the biogenic space.

Each molecule of the GDB and purchasable GDB databases was in turn compared to each molecule in the KEGG and DNP databases. Compounds were represented by their Scitegic Extended Connectivity FingerPrints (ECFP₄; using Scitegic Pipeline Pilot from Accelrys Software, Inc.), which encode the presence or absence of topological fragments (with no stereochemistry consideration) in the form of an extended connectivity string centered on a specific atom and calculated using a modification of the Morgan algorithm³². The initial code assigned to each molecule's atom is based on the number of connections, the element type and the mass. This code is hashed to produce the next order code, which is mapped into an address space of size 2³² (using Scitegic Pipeline Pilot), and the process iterates twice to describe features up to four bonds in diameter³³. The resulting fingerprint was further folded into a bit string of length 1,024 bits. The similarity between two molecules was measured by

comparing their respective ECFP₄ bit string with the Tanimoto coefficient. If a denotes the number of bits set to 'on' for molecule A, b denotes the number of bits set to 'on' for molecule B, and c denotes the number of bits set to 'on' in both molecules, the Tanimoto similarity between these two molecules is $T_c = c/(a + b - c)$.

Analysis of the bit frequencies. Each bit's exact substructure was exported as a SMARTS pattern (before folding) using the built-in function of the standard molecular fingerprint component of Scitegic Pipeline Pilot (Accelrys Software, Inc.). There were cases where a single bit corresponded to more than one substructure; often these substructures were related. For the frequency calculations or even for illustration purposes (see **Supplementary Tables 1 and 2**), one of these substructure was chosen arbitrarily.

Generation of the core ring scaffolds. Several approaches to extract the scaffold of a molecule are available^{20,26,34}; here we use the approach described in reference 20. Each molecule in the KEGG, DNP and ZINC databases was decomposed into its core ring scaffold using Pipeline Pilot. Core ring scaffolds consist of individual contiguous ring systems keeping atom types, bond orders,

aromaticity information and exocyclic double bonds but discarding stereochemistry and charges (see **Supplementary Fig. 4**). A canonical SMILES string was generated for each resulting core ring structure. The presence (or absence) of a particular KEGG or DNP scaffold was evaluated by matching its SMILES string to those obtained from the ZINC database.

Note: Supplementary information and chemical compound information is available on the Nature Chemical Biology website.

ACKNOWLEDGMENTS

This work was supported by US National Institutes of Health grant GM59957 to B.K.S. J.H. was supported by a Marie Curie fellowship from the 6th Framework Program of the European Commission; M.J.K. was supported by a US National Science Foundation graduate fellowship; C.L. was supported by a fellowship from the Max Kade Foundation.

AUTHOR CONTRIBUTIONS

The project was conceived of by J.H. and B.K.S. J.H. undertook most of the calculations, with molecular proof checking by J.J.L. and C.L. and algorithmic assistance from M.J.K. J.H. and B.K.S. wrote the manuscript, which was read and commented on by the other authors.

Published online at <http://www.nature.com/naturechemicalbiology/>
Reprints and permissions information is available online at <http://npg.nature.com/reprintsandpermissions/>

- Wilhelm, S. *et al.* Discovery and development of sorafenib: a multikinase inhibitor for treating cancer. *Nat. Rev. Drug Discov.* **5**, 835–844 (2006).
- Spencer, R.W. High-throughput screening of historic collections: observations on file size, biological targets, and file diversity. *Biotechnol. Bioeng.* **61**, 61–67 (1998).
- Fox, S., Farr-Jones, S., Sopchak, L., Boggs, A. & Comley, J. High-throughput screening: searching for higher productivity. *J. Biomol. Screen.* **9**, 354–358 (2004).
- Macarron, R. Critical review of the role of HTS in drug discovery. *Drug Discov. Today* **11**, 277–279 (2006).
- Pereira, D.A. & Williams, J.A. Origin and evolution of high throughput screening. *Br. J. Pharmacol.* **152**, 53–61 (2007).
- Bohacek, R., McMartin, C. & Guida, W. The art and practice of structure-based drug design: a molecular modeling perspective. *Med. Res. Rev.* **16**, 3–50 (1996).
- Roth, B., Sheffler, D. & Kroeze, W. Magic shotguns versus magic bullets: selectively non-selective drugs for mood disorders and schizophrenia. *Nat. Rev. Drug Discov.* **3**, 353–359 (2004).
- Paolini, G., Shapland, R., van Hoorn, W., Mason, J. & Hopkins, A. Global mapping of pharmacological space. *Nat. Biotechnol.* **24**, 805–815 (2006).
- Yildirim, M., Goh, K.-I., Cusick, M., Barabasi, A.-L. & Vidal, M. Drug–target network. *Nat. Biotechnol.* **25**, 1119–1126 (2007).
- Martin, Y.C. Diverse viewpoints on computational aspects of molecular diversity. *J. Comb. Chem.* **3**, 231–250 (2001).
- Breinbauer, R., Vetter, I.R. & Waldmann, H. From protein domains to drug candidates—natural products as guiding principles in the design and synthesis of compound libraries. *Angew. Chem. Int. Ed.* **41**, 2879–2890 (2002).
- Koehn, F. & Carter, G. The evolving role of natural products in drug discovery. *Nat. Rev. Drug Discov.* **4**, 206–220 (2005).
- Arve, L., Voigt, T. & Waldmann, H. Charting biological and chemical space: PSSC and SCONP as guiding principles for the development of compound collections based on natural product scaffolds. *QSAR Comb. Sci.* **25**, 449–456 (2006).
- Ertl, P., Roggo, S. & Schuffenhauer, A. Natural product-likeness score and its application for prioritization of compound libraries. *J. Chem. Inf. Model.* **48**, 68–74 (2008).
- Gupta, S., Aires-de-Sousa, J. Comparing the chemical spaces of metabolites and available chemicals: models of metabolite-likeness. *Mol. Divers.* **11**, 23–36 (2007).
- Fink, T. & Reymond, J.L. Virtual exploration of the chemical universe up to 11 atoms of C, N, O, F: assembly of 26.4 million structures (110.9 million stereoisomers) and analysis for new ring systems, stereochemistry, physicochemical properties, compound classes, and drug discovery. *J. Chem. Inf. Model.* **47**, 342–353 (2007).
- Sadowski, J. & Kubinyi, H. A scoring scheme for discriminating between drugs and nondrugs. *J. Med. Chem.* **41**, 3325–3329 (1998).
- Good, A.C. & Hermsmeider, M.A. Measuring CAMD technique performance. 2. How “druglike” are drugs? Implications of random test set selection exemplified using druglikeness classification models. *J. Chem. Inf. Model.* **47**, 110–114 (2007).
- Glen, R.C. *et al.* Circular fingerprints: flexible molecular descriptors with applications from physical chemistry to ADME. *IDrugs* **9**, 199–204 (2006).
- Bemis, G.W. & Murcko, M.A. The properties of known drugs. 1. Molecular frameworks. *J. Med. Chem.* **39**, 2887–2893 (1996).
- Schreiber, S. Target-oriented and diversity-oriented organic synthesis in drug discovery. *Science* **287**, 1964–1969 (2000).
- Haggarty, S., Clemons, P., Wong, J. & Schreiber, S. Mapping chemical space using molecular descriptors and chemical genetics: deacetylase inhibitors. *Comb. Chem. High Throughput Screen.* **7**, 669–676 (2004).
- Shang, S. & Tan, D.S. Advancing chemistry and biology through diversity-oriented synthesis of natural product-like libraries. *Curr. Opin. Chem. Biol.* **9**, 248–258 (2005).
- Gregori-Puigjané, E. & Mestres, J. Coverage and bias in chemical library design. *Curr. Opin. Chem. Biol.* **12**, 359–365 (2008).
- Ertl, P., Jelfs, S., Mühlbacher, J., Schuffenhauer, A. & Selzer, P. Quest for the rings. In silico exploration of ring universe to identify novel bioactive heteroaromatic scaffolds. *J. Med. Chem.* **49**, 4568–4573 (2006).
- Wester, M.J. *et al.* Scaffold topologies. 2. Analysis of chemical databases. *J. Chem. Inf. Model.* **48**, 1311–1324 (2008).
- Wetzel, S., Schuffenhauer, A., Roggo, S., Ertl, P. & Waldmann, H. Cheminformatic analysis of natural products and their chemical space. *Chimia* **61**, 355–360 (2007).
- Fink, T., Bruggesser, H. & Reymond, J.L. Virtual exploration of the small-molecule chemical universe below 160 Daltons. *Angew. Chem. Int. Ed.* **44**, 1504–1508 (2005).
- Kanehisa, M. & Goto, S. KEGG: Kyoto encyclopedia of genes and genomes. *Nucleic Acids Res.* **28**, 27–30 (2000).
- Buckingham, J. *Dictionary of Natural Products* (Chapman & Hall/CRC, United Kingdom, 2008).
- Irwin, J.J. & Shoichet, B.K. ZINC—a free database of commercially available compounds for virtual screening. *J. Chem. Inf. Model.* **45**, 177–182 (2005).
- Morgan, H.L. Generation of a unique description for chemical structures—a technique developed at Chemical Abstract Service. *J. Chem. Doc.* **5**, 107–113 (1965).
- Hert, J. *et al.* Comparison of topological descriptors for similarity-based virtual screening using multiple bioactive reference structures. *Org. Biomol. Chem.* **2**, 3256–3266 (2004).
- Koch, M. *et al.* Charting biologically relevant chemical space: a structural classification of natural products (SCONP). *Proc. Natl. Acad. Sci. USA* **102**, 17272–17277 (2005).