ELSEVIER

'Metabolite-likeness' as a criterion in the design and selection of pharmaceutical drug libraries

Paul D. Dobson, Yogendra Patel and Douglas B. Kell

School of Chemistry and The Manchester Interdisciplinary Biocentre, The University of Manchester, 131 Princess St, Manchester M1 7DN, UK

Present drug screening libraries are constrained by biophysical properties that predict desirable pharmacokinetics and structural descriptors of 'drug-likeness' or 'lead-likeness'. Recent surveys, however, indicate that to enter cells most drugs require solute carriers that normally transport the naturally occurring intermediary metabolites and many drugs are likely to interact similarly. The existence of increasingly comprehensive summaries of the human metabolome allows the assessment of the concept of 'metabolite-likeness'. We compare the similarity of known drugs and library compounds to naturally occurring metabolites (endogenites) using relevant cheminformatics molecular descriptor spaces in which known drugs are more akin to such endogenites than are most library compounds.

Introduction

The search for pharmaceutically active drugs with desirable properties and negligible side effects can be considered as a multi-objective optimisation problem over an enormous search space of 'possible' drugs [1,2]. It is usual to start the search by looking for hits and then leads [3], because, according to Oprea *et al.* [4], 'lead structures exhibit, on the average, less molecular complexity (less MW, less number of rings and rotatable bonds), are less hydrophobic (lower cLogP and LogD), and less druglike' than actual drugs (see also [5]). The process of optimising a lead into a drug with favourable ADMET properties [6,7] results in more complex structures [8] and system approaches [9–13] that consider not only a molecular target but also biochemical networks may be of value in understanding why.

In seeking to narrow the search space of chemically diverse candidate compounds, cheminformatic methods are used to constrain the compounds screened such that they tend to display 'lead-likeness' [4,14–16] or 'drug-likeness' [16–23] (and even 'CNS-likeness' [24]). The same concepts hold true for drugs with multiple intended targets (promiscuous drugs [25] or poly-pharmacology [2,26]).

The most common cheminformatic filter used to constrain pharmaceutical drug libraries is Lipinski and colleagues' celebrated 'rule of five' (Ro5) [27]. This states that poor absorption or permeation of a compound is more probable when there are more than five hydrogen-bond donors, the molecular mass is above 500 Da, the lipophilicity is high (*clogP* > 5) and when the sum of nitrogen and oxygen atoms is greater than 10. Other rules or filters consider generic and calculable properties such as the number of rotatable bonds and the polar surface area [28,29] or the ligand efficiency [30–33], and a 'rule of three' has been proposed [34] for fragment-based lead discovery (see e.g. [35,36]). It was recognised explicitly in the original review [27] that the Lipinski rules do not normally cover drugs that are derived from natural products [37,38], in which transporters are clearly involved in their disposition and it is, in fact, probable that this involvement of carrier molecules holds true for most other compounds too [39–41].

Descriptors such as those of Lipinski and colleagues [27] are essentially biophysical rather than structural in nature, and despite the widespread use of these measures it is not completely obvious how they should be understood mechanistically, given the enormous structural diversity of both drugs and libraries. Clearly, if drugs are mainly transported by carriers, this gives a ready explanation of why general descriptors are not normally going to be entirely effective in individual cases [41]; it also promotes the view that we need to understand the specificities for existing and candidate drugs of known drug transporters much better than we do now at a mechanistic level. Indeed, it is considered difficult to design for the use of active transporters because transporter selectivity is not well understood [42]. This said, at

Corresponding author: Dobson, P.D. (paul.dobson@manchester.ac.uk)

1359-6446/06/$ - see front matter © 2008 Elsevier Ltd. All rights reserved. doi:10.1016/j.drudis.2008.10.011

least for those uptake transporters that are known, it is reasonable to assume that they normally exist to transport common biochemical compounds involved in primary (intermediary) metabolism and the knowledge of these molecules may provide useful constraints for identifying other potential substrates by direct structural comparison or by SAR.

A drug may also interact with its target or targets in a manner that emulates native substrate binding (e.g. [43]). This will also constrain drugs towards metabolite space, given that the number of known protein folds is comparatively restricted and that once found they appear to be conserved in evolution [44,45]. Similarly, drugs may regulate activity allosterically at a native control site and this will also constrain such drugs towards metabolite-likeness. The guiding principle is that interacting with targets in any native-like manner imposes native-like constraints, and this will manifest itself in the drug discovery process by a tendency to develop towards the regions of chemical space occupied by metabolites, although other concerns in the drug development pipeline will also impinge upon the final drug structure(s) chosen.

Much of post-genomic drug discovery has concentrated on proteins and proteinaceous drug targets [19,46]. In recent years, however, we have witnessed the development of freely available, curated, reconstructed genome-scale metabolite networks (e.g. [47–51]) and of databases of human and other metabolites [52–59], where here the term 'metabolite' is used to refer to small molecule components of primary metabolism and not the products of the reaction of drugs with drug metabolising enzymes. To this end, we shall sometimes use the term 'endogenite' in this article to describe these endogenous, naturally occurring molecules. Nobeli and colleagues [60,61] have produced a very interesting summary of some of the properties of the known metabolome of *Escherichia coli* in particular (and we note that many microbially derived gut metabolites may also influence their human host (e.g. [62])). In a similar vein, the existence of databases of endogenite molecules allows us to ask the question as to whether existing drugs, that is those that have been successful in passing through the various phases of drug discovery to the marketplace, are more metabolite-like (i.e. endogenite-like) than are the typical contents of pharmaceutical screening libraries. To address this, known drugs and library compounds, representing the sorts of pre-drugs that might be screened in hit discovery, are compared to human metabolites in a variety of appropriate molecular descriptor spaces. We find that drugs are indeed considerably more similar to endogenous metabolites than are library compounds, and conclude that endogenite-likeness might be a useful filter in the design and analysis of pharmaceutical libraries for drug discovery.

Related comparisons between metabolites and other types of molecules have been considered previously. Gupta and Aires-de-Sousa [63] compared the distributions in chemical space of metabolites drawn from KEGG and compounds from the supplier library ZINC [64], concluding that discriminatory features include hydroxyl groups, aromatic systems and molecular weight when combined with other global descriptors.

In the major analysis of Karakoc *et al.* [65] relationships between drugs, drug-like compounds, antimicrobials, and human and bacterial metabolites were considered. One result finds that bacterial metabolites and antibiotics are highly similar and this mirrors the similarity between human metabolites and drugs we observe. There is also the suggestion, however, that human metabolites form a distinct class of molecules that are unlike bacterial metabolites, drugs or drug-like molecules, and they occupy a separate region of chemical space. This seemingly counter-intuitive result does not concur with that presented here. The set of 5333 metabolites used here is much larger (compared to 1104), giving far greater coverage of 'metabolite space' and so more fully represents the total diversity of human metabolites. Moreover, the redundancy measures of Karakoc *et al.* only removed exact duplicate molecules and this allows highly similar molecules to remain within the set. Inevitably this biases the set's properties towards the properties of over-represented molecules. Indeed their own analysis indicates the over-abundance of scaffolds drawn from sugar- and nucleotide-like molecules. Through the application of clustering to choose representative molecules for multiple represented 'types' of molecules the influence of redundancy within our sets is negated. We suggest that the differences observed between human metabolites and other classes actually reflect the construction of their human metabolite set and not a fundamental difference between the properties of human metabolites and other classes of molecule, and this is reinforced by our analysis. Also note that we do not claim that all human metabolites are similar to drugs; many clearly are not. If the human metabolite set of Karakoc *et al.* contains many of these (sugar scaffolds are prevalent among their metabolites but not their drugs) then this could also underpin the differences observed.

Finally, Ganesan [38] has very recently compared natural products and synthetic molecules released as drugs, in terms of their 'Lipinski-likeness', commenting that (only) 'half of the 24 natural products lie in what can be called the ''Lipinski universe''.'

In this article, the metabolites are drawn from human-specific databases and genome-scale metabolic reconstructions, and are greater in number than in previous studies, although many of the carbohydrates and especially lipids [66,67] that might usefully be considered metabolites in this context still remain undetermined. This said, it emerges that the types of drugs that exhibit virtually no metabolite-likeness in our analysis are very atypical and will probably remain so even when the missing metabolites are included.

Comparing drugs and library compounds to metabolites

To assess the relationship between drugs and metabolites we compare against a background of compounds of the kinds that typically make up screening collections for hit discovery, which we refer to as library compounds. These represent pre-drugs and can be considered as starting points for drug discovery and development. During these processes candidate drugs are selected and modified to enhance properties favourable to drug action and our hypothesis suggests that this optimisation process drives such starting molecules towards the regions of chemical space occupied by metabolites, because of the necessity to participate in native-like reactions (including those with transporter molecules).

In the analysis we therefore distinguish metabolites (endogenites), drugs and library compounds (Table 1). The molecules retrieved from source databases contained duplicate records and over-represented structural types. Thus, to avoid [68] biasing the

TABLE 1

Sources of drug, metabolite and pre-drug structures.

Class	Source	Compounds	Total
Metabolite	HMDB [58]	2835	5333 (5560)
Metabolite	Palsson [94]	806	
Metabolite	BioCyc [95]	772	
Metabolite	BiGG [96]	698	
Metabolite	Edinburgh [48,51]	2048	
Drug	DrugBank [97]	4152	7330 (8002)
Drug	KEGG Drug [98]	4435	
Pre-drug	Zinc [64]	62,390	

'Compounds' is the number of unique structures after washing and filtering. 'Total' is the total number of unique compounds for each class. The figures in brackets are before the semi-automatic correction of annotation errors in sources.

analysis towards common structural forms, the compounds were clustered and cluster centres used as representatives. Physicochemical distributions are shown in Fig. 1a–d. Figure 1a illustrates (i) that the distribution of the number of atoms among metabolites is markedly different from that of drug and library compounds, but (ii) that the similarity between drugs and library compounds suggests that screening sets do cover drug space in a sensible manner, at least with respect to atom number.

The distribution of *clogD* [69], a calculated value of lipophilicity able to account for charged species, is shown in Fig. 1b. Positive values indicate a preference for a hydrophobic environment and negative values a preference for a hydrophilic environment. The difference between metabolites and drugs is clear, and again the library compound distribution conforms to the drug distribution and suggests appropriate representation of existing drug space in the screening set. That drugs and libraries have similar distributions is expected because considerations of lipophilicity have played a major role in designing drugs and libraries with useful bioavailability, with the Ro5 being particularly influential, despite the fact that only approximately half of the marketed drugs obey it [70]. The fact, however, that metabolites are in general much more hydrophilic than are both drugs and library compounds is very

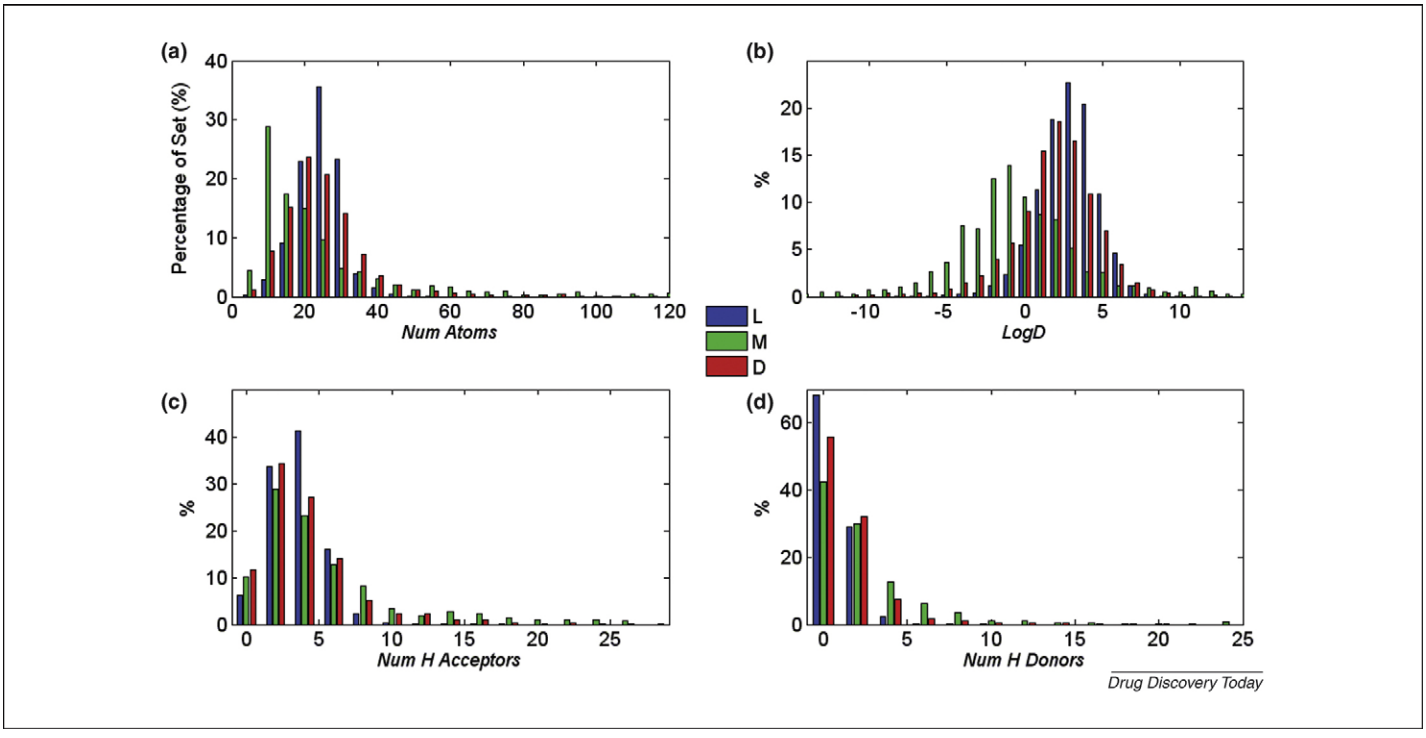

FIGURE 1

Histograms, normalised by class sizes, for different simple molecular properties. (a) Number of atoms (excluding hydrogens) over the drug, library and metabolite representative sets. Of note is the fact that above 40 atoms there are very few library compounds, although drugs and metabolites are still represented. (b) *clogD*, a lipophilicity calculation that takes charge into account. The distribution of library compounds is similar to that of drugs, but metabolites differ markedly. (c) Number of hydrogen bond acceptors. The number of hydrogen bond acceptors in Lipinski's rule of five is suggested to be not more than ten. The histogram illustrates that both known drugs and metabolites are still found above this number, although the library compounds are rare, suggesting that Lipinski and related measures of 'drug-likeness' have been used as a constraint in library design. (d) Number of hydrogen bond donors. The number of hydrogen bond donors in Lipinski's rule of five is suggested to be not more than five. As in (c), library compounds largely follow the rule, but there are many drugs and metabolites that do not.

noticeable, and is at least consistent with the requirement for specialised carriers to transfer them into and out of cells and between intracellular compartments.

Figure 1c and d shows the distributions of the numbers of hydrogen bond donors and acceptors. The Ro5 suggests that the number of hydrogen bond acceptors be not more than ten, and the number of hydrogen bond donors be not more than five. Both figures illustrate that whilst the library sets mostly follow these suggestions, there are numerous drugs (and metabolites) that do not. From this perspective, endogenites are considerably more like drugs than are library compounds.

Metabolite-likeness curves

Whilst physicochemical distributions provide a general overview of the relationships between types of molecules, it is their distributions in chemical space that are of most relevance. If the processes of drug discovery and development drive towards metabolite-likeness then this should manifest itself through considerable overlap between the distributions of drugs and (a subset of) metabolites, and one greater than seen between library compounds and metabolites. The notion of chemical space is abstract, but it can be represented and operated on by the techniques of cheminformatics that allow molecular similarity to be quantified.

The similarity of drug and library compounds to metabolites is here assessed by calculating the Tanimoto distance to the closest metabolite. A variety of molecular descriptors were computed, and similarities calculated using the Tanimoto coefficient [71]. The molecular descriptors used were connectivity fingerprints [72,73], paths [74,75], MDL Public Keys [76] and electrotopological state (E-state) keys [77,78].

Representative sets of drugs and library compounds were calculated in each space at thresholds that removed high-level redundancy, which will be described later. Redundancy within the metabolites was not addressed because it does not negatively affect the outcome because similarity is measured only with the closest metabolite.

Figure 2a–d shows the proportion of drugs and library compounds within a given distance to the closest metabolite, using the above four sets of molecular descriptors. For example, in Fig. 2a one can determine that 12% of drugs have a Tanimoto distance of 0.5 or less to their closest metabolite. By contrast, less than 2% of library compounds fall within the same threshold. Although the shapes of the curves vary for the different descriptors used, the drugs are consistently closer, often considerably so, to endogenous metabolites than are the contents of typical screening libraries. That the drug curves are consistently higher than the curves for library compounds, in a variety of descriptor spaces covering various ways of assessing molecular similarity, indicates that successful, marketed drugs are indeed much more like metabolites than are the typical library compounds.

Molecular similarity can be represented in different ways. Because of this, metabolite-likeness is calculated in several molecular descriptor spaces that capture different aspects of structure, and so illustrate that metabolite-likeness is not simply an artefact of a particular descriptor but a general phenomenon in each of the chemical spaces assessed. It appears to be generally true that drugs that are very close to metabolites are typically analogues of the native substrate of their targets. For example, Fig. 3 illustrates how the closest metabolite differs in the various spaces using the example query of atorvastatin (Lipitor). Different metabolites are retrieved in each space, and whilst there are features common to the query and each of the retrieved structures, the connectivity fingerprint-retrieved structure particularly recalls the native product structure (mevalonate) of the main atorvastatin target HMG-CoA reductase (although we note that statins can exhibit many pleiotropic effects, see e.g. [79–81]). Generally, the closest metabolite to a drug is quantifiably more similar than in the example shown in Fig. 3.

The valyl-ester prodrug of ganciclovir (valganciclovir), which is taken up by peptide transporters [82] of solute carrier family 15 [83], retrieves nucleoside-like metabolites that more closely resemble the active drug than do the valine modification that one might expect, although the relative contributions of the large drug and small valine probably bias molecular similarity measures towards the drug, and maximal common substructure methods may be of use. Another type of prodrug modification couples bile acids and drugs [84], including those designed to target the human apical sodium-dependent bile acid transporter (hASBT) [85], which transports bile acids including chenodeoxycholate, deoxycholate, cholate and ursodeoxycholate. By coupling bile acids via valine to acyclovir, enhanced uptake was observed *in vitro* and *in vivo*, most successfully for the prodrug acyclovir valylchenodeoxycholate, which lead to a twofold increase in acyclovir bioavailability in rats. Using acyclovir valylchenodeoxycholate as the drug query in the metabolite search all spaces retrieve bile acids (taurochenodeoxycholate) or intermediates in bile acid biosynthesis (choloyl-CoA). This emphasises that metabolite-likeness can be because of drugs mimicking metabolites in a pathway as opposed to those interacting with a specific target.

Dissimilar drugs

Whilst drugs are generally more similar to metabolites than to library compounds, certain drugs do not conform to this trend. The fraction that does not depends upon how one chooses to define the boundary between 'similar' and 'dissimilar'. In the connected fingerprint space a realistic choice for the limit of molecular similarity is a Tanimoto distance between 0.7 and 0.8, equating to 50–80% of drugs being metabolite-like. An illustrative selection of these 'remote' compounds is shown in Fig. 4, but clear trends towards particular types of drug or structural classes are not immediately discernable, although many remote compounds are heavily halogenated or sulphurated.

Discussion

That the processes of drug discovery and development lead largely to regions of chemical space already occupied by metabolites, although a novel discovery, is both expected from the arguments given in the introduction and observed experimentally in our analyses. This has major implications for future library design, which might beneficially take account of the structures and properties of endogenous metabolites now that usefully complete structural metabolomes are available. Of course, further efforts to elucidate the measured metabolome are ongoing, but it is of note that many metabolites observed experimentally have yet to be identified chemically [86–88], particularly lipids [66,89,90].

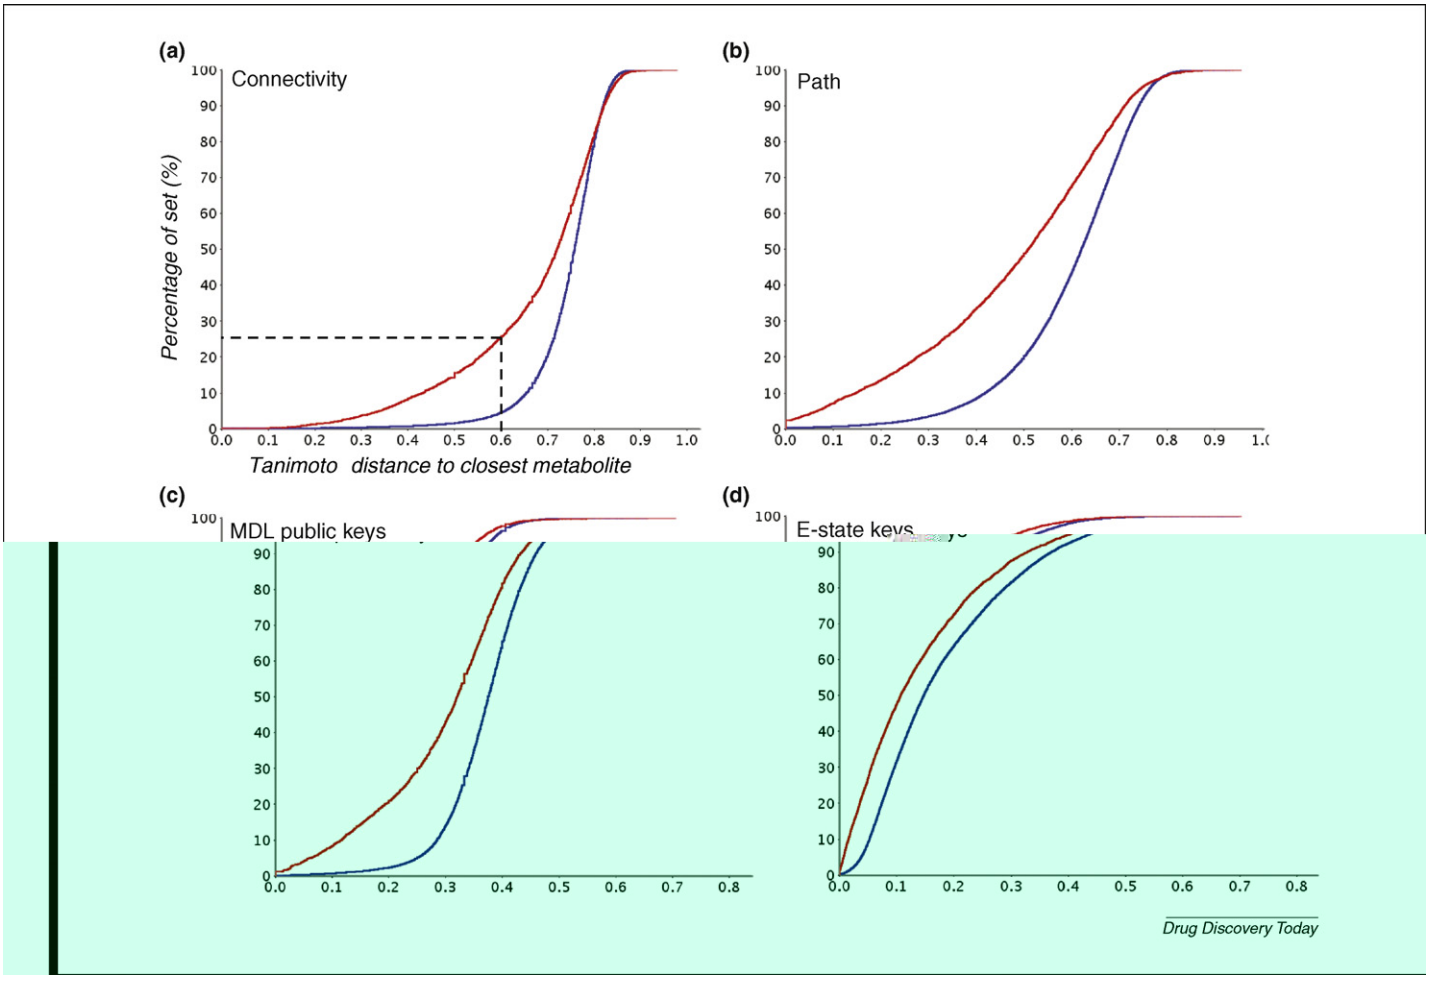

FIGURE 2

A comparison of drug and library distances to closest metabolites in various molecular descriptor spaces. (a) Connectivity fingerprint space. The proportion of the representative sets that lie within the specified Tanimoto distance can be seen for each class; the dashed line illustrates that 25% of drugs have a distance to their closest metabolite of less than or equal to 0.6, but fewer than 5% of library compounds are found within the same threshold closeness. (b) Path descriptor space. In this space some drugs and metabolites are equivalent, these being 2% of drugs with a Tanimoto distance to their closest metabolite of 0. Only 3.5% of library compounds are within a distance of 0.2 of a metabolite, compared with 22% of drugs. (c) MDL Public Keys space. The closeness of both kinds of compound to the nearest metabolite (endogenite) using these descriptors is numerically rather smaller than that using the extended connectivity fingerprint space (a) but again the distances are considerably smaller for endogenites than for library compounds. Only 2% of library compounds are within a Tanimoto distance of 0.2 of a metabolite, compared with 20% of drugs. (d) E-state keys space. The closeness of both kinds of compound to the nearest metabolite (endogenite) using these descriptors is numerically rather smaller than in (a) and (c), but the distances to closest endogenites are again considerably lower for drugs than for library compounds. Only 6% of library compounds are within a Tanimoto distance of 0.05 of a metabolite, compared with 23% of metabolites.

Understanding why drugs are more similar to metabolites than library compounds are is complex. High levels of similarity imply considerable structural equivalence between drug and metabolite molecules, whereas lower similarity levels might suggest that something more general has been derived that relates the molecules, such as similar physicochemical properties, but it is difficult to generalise any such interpretation over the set. The conclusion that drugs are more like metabolites than library molecules, as demonstrated using multiple molecular descriptor spaces, however, is clear and will, we believe, be of considerable value in shaping future drug discovery efforts. Unlike drugs, metabolites have not been through a human-guided development process that considers important factors such as developability, desirable PK/PD properties and other concerns that dictate how hits become leads and then drugs. This is, however, information that should be inductively learned by a well-constructed model of drug-likeness [19]. Conversely, such models are largely bounded by the limits of existing information and only describe the types of drugs already known. The search for new types of drug that exist outside the current drug space (see e.g. [70]) might usefully begin in the unexplored regions of metabolite space, particularly given that drugs probably do not cover the whole of metabolite space, our knowledge of which continues to grow. Furthermore, the spaces in which drugs exhibit enhanced metabolite-likeness will be of use in predicting drug–metabolite interactions.

For problems of molecular similarity it is important to consider the scope of descriptors and the extent to which they are able to capture useful relationships between molecules. As with any other kind of clustering, where utility is the most significant criterion [91], molecular similarity has no innate or absolute meaning and the appropriateness of one descriptor ahead of another is largely a subjective choice, and this problem of representation is well