Journal of Medicinal Chemistry

Perspective

pubs.acs.org/jmc

Exploring Activity Cliffs in Medicinal Chemistry

Miniperspective

Dagmar Stumpfe and Jürgen Bajorath*

Department of Life Science Informatics, B-IT, LIMES Program Unit Chemical Biology and Medicinal Chemistry, Rheinische Friedrich-Wilhelms-Universität, Dahlmannstrasse 2, D-53113 Bonn, Germany

INTRODUCTION

In recent years, structure−activity relationships (SARs) have increasingly been studied through mining of large compound data sets, leading to a renaissance of the *activity landscape* concept.[1−5] In general terms, an activity landscape is defined as any graphical representation that integrates similarity and potency relationships between compounds sharing a specific activity.[5] From activity landscape representations of different design, both global and local SAR features present in compound data sets can be extracted.[3,5] Large-scale SAR analysis complements classical quantitative structure−activity relationship (QSAR) modeling.[6] For example, activity landscape models make it possible to delineate regions of SAR continuity where gradual changes in compound structure lead to moderate changes in compound potency. The presence of SAR continuity provides a fundamental basis for QSAR analysis and resulting compound activity predictions.[6] Furthermore, activity landscapes also reveal regions of SAR discontinuity where small changes in chemical structure lead to large changes in compound potency, a scenario often encountered in lead optimization. However, the presence of SAR discontinuity falls outside the applicability domain of the QSAR paradigm.[7] However, regions of SAR discontinuity in activity landscapes are generally thought to provide much SAR information[5,8] because small structural changes of active compounds lead to large potency effects. The most prominent form of SAR discontinuity is provided by *activity cliffs*[5,7,8] that have been discussed in the medicinal chemistry relevant scientific literature since the early 1990s.[9] In general, an activity cliff is defined as a pair of structurally similar or analogous compounds having a large difference in potency.[5,7] As an extreme form of SAR discontinuity, activity cliffs represent the most prominent features of activity landscapes and are often the primary focal point of their analysis.

Modeling and rationalizing activity landscapes require the application of computational methods. As such, the activity landscape concept is of comparable relevance for chemoinformatics (having a strong focus on data mining and representation) and medicinal chemistry (given its immediate relevance for SAR analysis).

In a previous article,[5] different approaches to activity landscape design and analysis have been discussed in detail. This contribution provides a follow-up on this activity landscape perspective by specifically concentrating on the study of activity cliffs. We currently observe a trend to increasingly discuss activity cliffs in the computational and medicinal chemistry literature. There are probably several reasons for this. Without doubt, the activity cliff concept is intuitive from a medicinal chemistry point of view and therefore attractive to consider. Furthermore, the analysis of activity cliffs in compound data sets is also of interest for computational research, as it requires systematic compound similarity and potency comparisons. However, if one wanted to be a bit provocative, one could perhaps also argue that the term activity cliff has become a trendy buzz word in the SAR field that is often used without thinking too much about its scientific relevance. There are certainly a number of still open questions that should be addressed. For example, how can one describe activity cliffs in a formally consistent manner? What exactly makes cliffs interesting for medicinal chemistry? Are there general characteristics of activity cliffs or should one better consider them on an individual basis? As will be argued herein, the exploration of activity cliffs is a more complex task than often thought and requires exact definitions to be made. Moreover, the activity cliff concept is associated with substantial caveats, both from a theoretical and experimental point of view. Finally, whether or not activity cliff analysis provides information that is of immediate use for medicinal chemistry is often a matter of debate.

Herein, we describe in detail the multifaceted nature of activity cliffs, discuss underlying scientific concepts, and explain how their individual or systematic analysis might (or might not) provide useful information for medicinal chemistry programs.

ACTIVITY CLIFF DEFINITION

We begin the discussion by reiterating a generally accepted definition: activity cliffs are formed by pairs of structurally "similar" compounds with "large" differences in potency. Importantly, this definition highlights two aspects that substantially complicate the consistent assessment of activity cliffs in computational and medicinal chemistry. First, what does "similar" mean in this context? How do we assess structural similarity and when do we consider compounds to be similar? This criterion presents a major conundrum for activity cliff analysis (vide infra). Second, what are "large" potency differences, 10-fold, 100-fold, or more? Is a potential cliff formed between two compounds with 1 and 100 μM potency comparable to one formed by compounds having 1 and 100 nM potency? Should one only consider highly potent compounds as potential cliff partners? These are other relevant questions that are often not considered in the analysis of activity cliffs. For the exploration of activity cliffs, it is essential to clearly define these criteria and take potential ambiguities into account, as further discussed in the following.

Received: December 19, 2011
Published: January 11, 2012

dx.doi.org/10.1021/jm201706b | J. Med. Chem. 2012, 55, 2932−2942

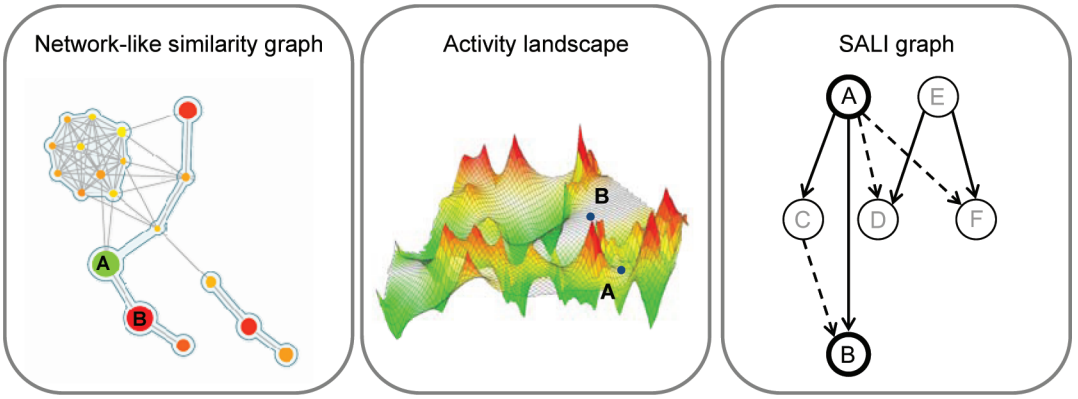

Figure 1. Activity cliff representations. Three alternative activity cliff representations are shown for a given compound pair A (weakly potent) and B (highly potent). These representations include a network-like similarity graph[10] (NSG, left), a 3D activity landscape model[12] (middle), and a SALI graphs[11] (right). In NSGs, nodes represent compounds and edges pairwise similarity relationships. In SALI graphs, nodes are also compounds and edges indicate activity cliffs of varying magnitude. The 3D landscape model results from a 2D projection of chemical reference space with an interpolated potency surface added as the third dimension. Here, compounds are not explicitly shown (their positions are defined in the underlying 2D projection). NSGs and 3D landscapes are colored by potency using a spectrum from green (low potency) to red (high). In NSGs, nodes are scaled in size according to the contribution of each compound to local SAR discontinuity (quantified via the discontinuity score[10]). A compound makes a large contribution to local SAR discontinuity if its potency significantly differs from the potency of its immediate structural neighbors. Hence, in NSGs, combinations of large red and green nodes indicate activity cliffs.

■ IDENTIFICATION OF ACTIVITY CLIFFS

Activity cliffs are the most prominent features of activity landscape representations. However, given the previous detailed review of activity landscape representations,[5] only one aspect of activity landscape modeling (focusing on activity cliffs) will be discussed here, i.e., how to best identify activity cliffs in large compound data sets. For this purpose, we adhere to the general definition of activity cliffs (vide supra). Following a conventional approach, compounds forming activity cliffs might, for example, be selected by inspection of standard R-group tables, although this is only possible for individual analogue series. Even for individual series this is quickly becoming an arduous task when they substantially grow in size. More comprehensive graphical access to the identification of activity cliffs in large and structurally heterogeneous compound data sets is provided by 2D or 3D activity landscape representations that emphasize compounds with high structural similarity and significantly different potency. This is illustrated in Figure 1 for alternative activity cliff representations.[10−12] From graph representations shown in this figure, activity cliffs can be readily selected on the basis of visual inspection.

A systematic account of activity cliffs beyond visual analysis is facilitated through the use of numerical SAR analysis functions that quantify similarity and potency relationships in a consistent manner.[2,3] Such functions have been introduced to analyze large compound data sets and globally characterize SAR continuity and discontinuity. Global SAR analysis functions include the structural similarity vs activity similarity formalism upon which the design of structure−activity similarity (SAS) maps is based[8] as well as the SAR index (SARI),[13] which yields a composite score of individual SAR continuity and discontinuity scoring functions. However, for the identification of activity cliffs, local scoring schemes focusing on SAR discontinuity are more relevant than global assessments. Such local SAR analysis functions include the compound discontinuity score,[10] a variant of the SARI formalism, which quantifies the individual contribution of compounds to local SAR discontinuity, and the structure−activity landscape index (SALI).[11] Both global and local SAR analysis functions are based on systematic pairwise compound comparisons.

The local discontinuity score is defined as

$$\mathrm{disc}(i) = \frac{\sum_{\{j|\mathrm{sim}(i,j)>0.65, i\neq j\}} \mathrm{potdiff}(i,j) \times \mathrm{sim}(i,j)}{|\{j|\mathrm{sim}(i,j) > 0.65, i \neq j\}|} \quad (1)$$

Here, "potdiff" stands for potency difference and "sim" for compound similarity. It is standardized with respect to all compounds within a set and normalized to the value range [0, 1].[10] According to this formalism, compounds obtain high discontinuity scores if their potency significantly differs from the potency of their immediate structural neighbors. Hence, pairs of structurally similar compounds with significantly different potency values obtain compound scores close to 1 and mark activity cliffs.

The SALI scoring scheme[11] is specifically designed to quantify activity cliffs:

$$\mathrm{SALI}(i,j) = \frac{P_i - P_j}{1 - \mathrm{sim}(i,j)} \quad (2)$$

In the SALI formula, "*P*" means potency and "sim" similarity. SALI also is a local pairwise score, but it is not normalized and has an infinite value range. It is applied to generate activity cliff-centric representations of activity landscapes. In SALI graph representations, nodes represent compounds and edges activity cliffs. Edges are depicted as arrows that are directed toward the more potent compound. Thus, two compounds are connected if their SALI score exceeds a given threshold value, for example, a score greater than 70% or 80% of all scores. This makes it possible to identify activity cliffs of increasing magnitude. The SALI graph represents series of pairwise connected activity cliffs.[11] By application of varying score threshold levels, a continuum of activity cliffs is monitored for a given compound data set.

■ DESCRIPTION OF ACTIVITY CLIFFS

Inherent in the general definition of activity cliffs are potential ambiguities associated with the assessment of molecular

dx.doi.org/10.1021/jm201706b | J. Med. Chem. 2012, 55, 2932−2942

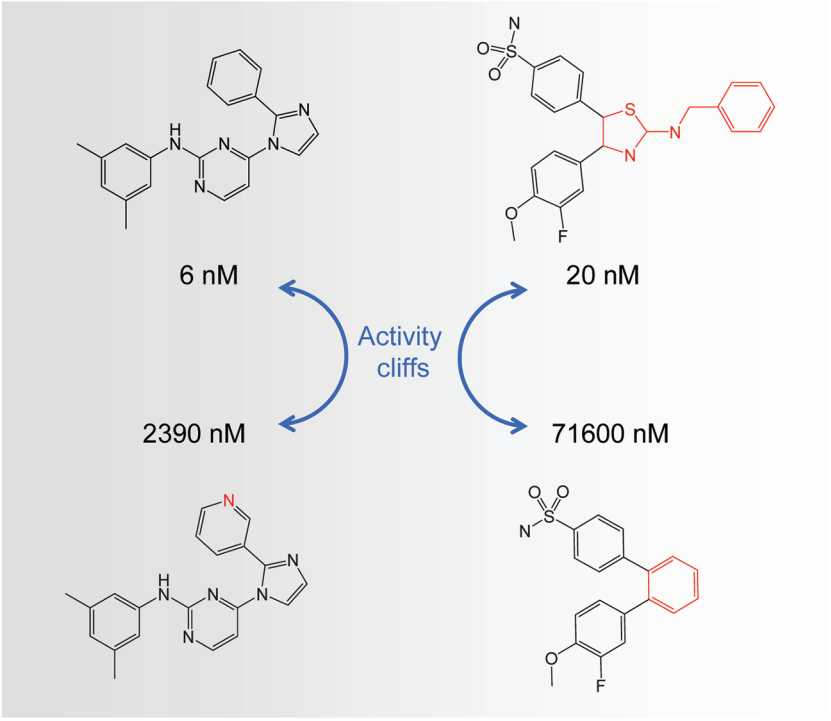

Figure 2. Exemplary activity cliffs. Two activity cliffs are shown where structural relationships between cliff partners notably differ. The activity cliff on the left is formed by two antagonists of vascular endothelial growth factor receptor-2. These two analogues are only distinguished by a minute chemical change (red), and the activity character of this similarity–potency relationship is undisputable. On the right, a pair of cyclooxygenase-1 inhibitors is shown that form another potential cliff. However, in this case, a central ring system is replaced (red), effectively producing two chemically distinct scaffolds. Thus, although these two inhibitors yield a MACCS Tanimoto similarity of 91%, the activity cliff character might be debatable in this case.

similarity and potency differences (vide supra). Currently, no generally accepted criteria for activity cliffs are available. Although the description of cliffs requires a quantitative readout, it is rather common to refer to activity cliffs in a qualitative manner. However, to describe activity cliffs consistently and render different analyses comparable, clear definitions of cliff parameters are required.

Continuous vs Discrete. The design of the SALI formalism points at an important distinction: should a continuous spectrum of activity cliffs be considered or only discrete (large-magnitude) cliffs? There is no general answer to this question; it depends on the application. An obvious advantage of considering a continuum of activity cliffs is that compound data sets of different composition can be scanned for interesting cliffs. This is helpful, for example, when searching bioassay data for activity cliffs of increasing magnitude.[14] On the other hand, a disadvantage of the continuum approach is that cliffs detected at a certain score threshold level might essentially be irrelevant (pseudo-cliffs) because they are only of small and chemically insignificant magnitude (potency difference). Furthermore, through scoring, the magnitude of cliffs is not determined; they are only compared on a relative scale. Regardless of whether a continuum of activity cliffs is considered or discrete states, it is important to note that numerical local SAR analysis functions will ultimately reveal the most significant activity cliffs that are present in a given compound data set, which is a major advantage of these approaches. However, identifying the most prominent cliffs through local SAR scoring is often not sufficient. For example, to determine activity cliff distributions via large-scale compound data mining and compare cliffs across different compound series or data sets, discrete definitions of activity cliffs are required. These definitions must include the applied similarity criterion, the potency measure and potency difference between cliff partners, and the potency range (interval) that is considered relevant for cliff formation. Concerning the potency criteria, for a number of analyses carried out in our laboratory it has been proven useful to consider activity cliffs only if one cliff partner has a potency in the nanomolar range and if there is an at least 100-fold difference in potency between two partners.[15] Depending on a specific application, these criteria might of course be modified. However, for any activity cliff analysis, they need to be clearly specified, which is often not the case.

Molecular Representation Caveat. The way in which compound similarity relationships are assessed is also critical for conclusions drawn from any activity cliff analysis. Figure 2 shows an example of an "undisputable" activity cliff, a pair of analogues with a small chemical modification that triggers a large potency change. In such a case, it is not required to calculate similarity values to represent cliffs because the high degree of similarity is immediately obvious. Such considerations essentially apply to all compound pairs that are a part of a single analogue series. However, for the analysis of large and heterogeneous data sets, similarity values must be calculated to systematically compare compounds and similarity threshold values for cliff formation must be defined. In another example in Figure 2, fingerprint Tanimoto similarity has been calculated to establish the compound similarity relationship. Here, the activity cliff character might well be a matter of debate, at least from a medicinal chemistry point of view, given the difference in core structures between these compounds, despite overall structural resemblance. This comparison illustrates a well-known conundrum in chemical space and activity landscape design: chosen molecular representations (descriptors) and, to a lesser extent, similarity measures substantially influence the

dx.doi.org/10.1021/jm201706b | J. Med. Chem. 2012, 55, 2932–2942

assessment of similarity relationships and, consequently, the formation and distribution of activity cliffs.[12,16,17] For example, activity cliffs that are present when using a particular molecular representation might be absent when using another one because pairwise similarity values change.[12,16] SAR analysis functions such as SARI or SALI are generally affected by the molecular representation dependence of similarity value distributions because they are based on calculated fingerprint Tanimoto similarities. In this context, an important point to consider is whether calculated similarities are chemically intuitive or not, as illustrated in Figure 2. As long as molecules share common frameworks and one can appreciate the similarity of cliff partners by eye, activity cliff representations are usually meaningful, regardless of how similarity values might be calculated. Some rules can be applied. For example, if we need to calculate fingerprint similarities for activity cliff analysis, we typically utilize a Tanimoto similarity value of 0.55 for the extended connectivity fingerprint with bond diameter 4 (ECFP4)[18] as a similarity threshold criterion[15] or, alternatively, a value of 0.85 for molecular access system (MACCS) structural keys.[19] For these standard fingerprints, the given Tanimoto similarity values approximately correspond to each other and typically identify visibly similar structures in pairwise comparisons. Nevertheless, these compound might contain different scaffolds that are chemically (synthetically) more or less similar to each other. In general, the more complex molecular representations are, the more weight is put on fine structural details and/or property differences and the more dissimilar compounds will be on the basis of calculated similarity values. Such representation or descriptor "artifacts" will rarely lead to false-positive activity cliff assignments; this would require classifying structurally distinct compounds as being similar, which is not very likely. Rather, the probability is higher that potential activity cliffs might be missed if descriptors are used to represent molecules that strongly abstract from chemical structure. However, if the chemical resolution of molecular representations is too low, false-positive cliff assignments are likely. A key issue in similarity evaluation is that any calculated similarity that cannot clearly be reconciled and understood on the basis of 2D molecular graphs is not suitable for large-scale SAR analysis. Hence, a meaningful assessment of similarity is a crucial aspect for activity cliff exploration.

INFLUENCE OF DATA VARIABILITY

Other major factors that influence the identification and description of activity cliffs include the type and intrinsic variability of experimental measurements.[20,21] In this context, it is important to understand how the use of alternative potency measurements such as IC_{50} and K_i values might affect activity cliff formation and distribution. Furthermore, if multiple potency values are available for active compounds, the question becomes which of these one should choose or how they should be combined. For activity cliff analysis, this is another important issue. To evaluate these factors, a systematic analysis of public domain compounds from BindingDB[22] has recently been carried out.[21] For this and other data mining investigations discussed below, compound pairs were considered to form activity cliffs if at least one of the compounds had a potency in the nanomolar range, if there was an at least 100-fold difference in potency, and if an ECFP4 Tanimoto similarity threshold value of at least 0.55 was reached. In our systematic analysis, it was found that activity cliff distributions notably changed when multiple available potency values were averaged and when minimum or maximum values were used. Over many different compound data sets, the selection of maximum potency values generally yielded the lowest number of activity cliffs. Only approximately half of the activity cliffs detected for alternative ways to represent multiple potency measurements were conserved. Figure 3 shows an example. For high-confidence data, i.e., when all available potency values agreed within an order of magnitude, fewer activity cliffs were consistently detected than for lower confidence (more variable) assay data. Moreover, the use of IC_{50} or K_i measurements significantly altered activity cliff

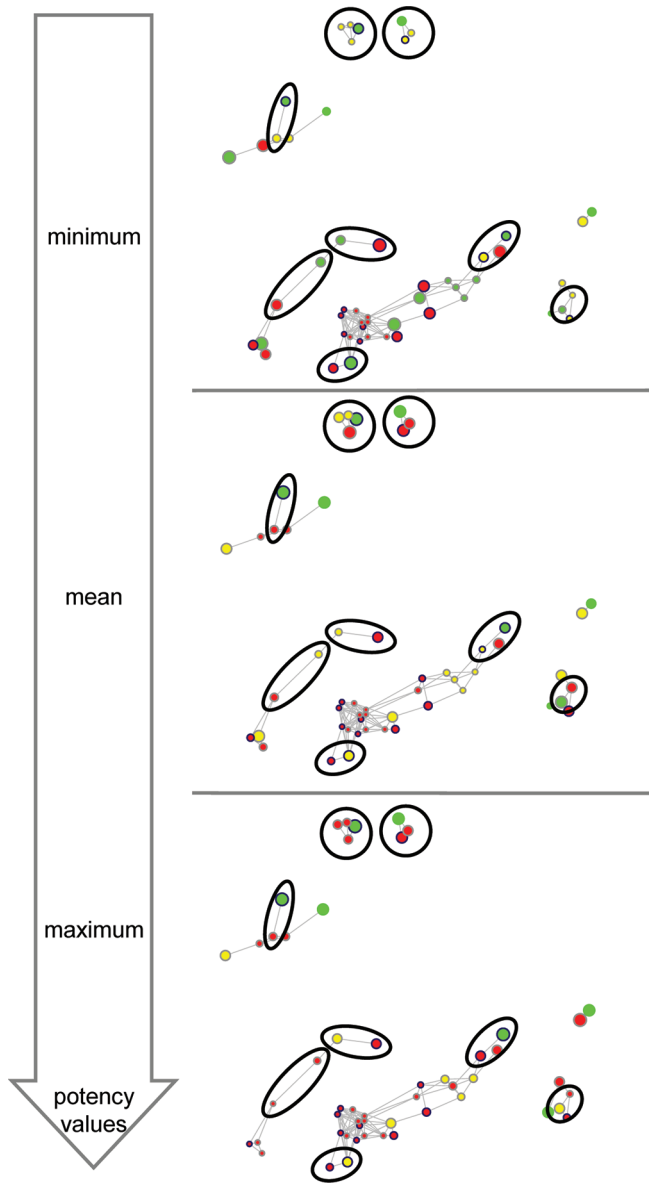

Figure 3. Activity cliff variability. Nonconserved activity cliffs found in a data set of dihydrofolate reductase inhibitors are shown in NSGs generated on the basis of alternative potency values (minimum, mean, or maximum). In order to focus the representation on activity cliff formation, nodes are color-coded according to different potency ranges: green, $pIC_{50} \leq 5$; yellow, $pIC_{50} > 5$ and $pIC_{50} \leq 7$; red, $pIC_{50} > 7$ (each connected pair of red and green nodes represents an activity cliff). Otherwise, the node representation is according to Figure 1. Exemplary activity cliffs that are not conserved for alternative potency measurements are encircled. The figure is adapted from ref 21.

dx.doi.org/10.1021/jm201706b | J. Med. Chem. 2012, 55, 2932−2942

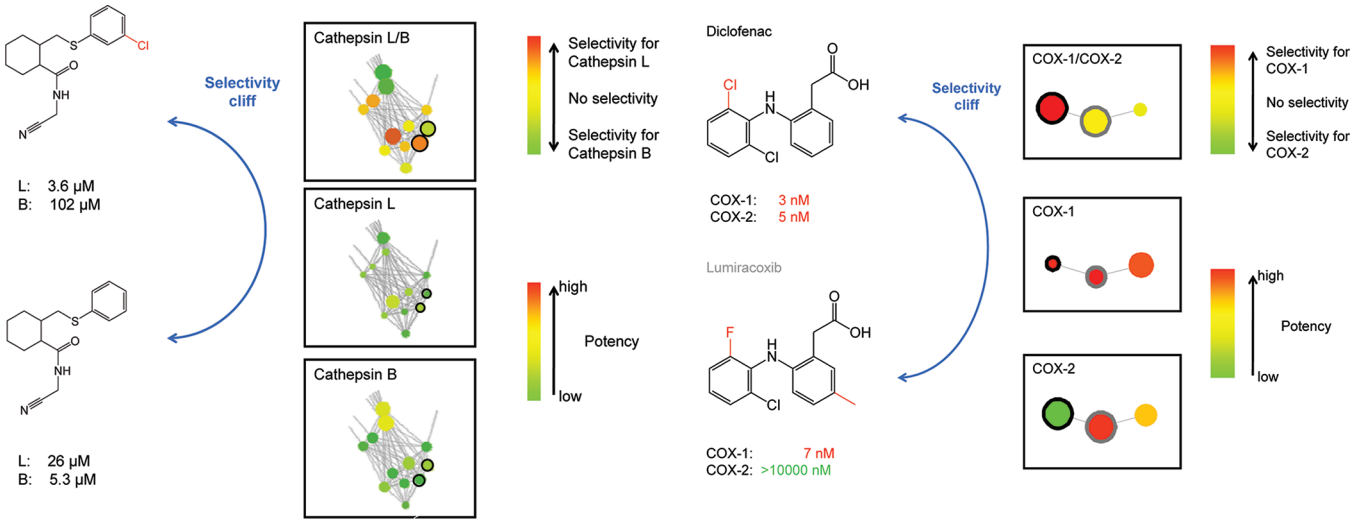

Figure 4. Selectivity cliffs. Shown are exemplary selectivity cliffs for two inhibitors of cathepsin L and B (left) and COX-1 and COX-2 (right), respectively. Structural differences between analogue pairs are shown in red, and potency values of the inhibitors are reported. NSG "windows" from representations of larger data sets are displayed according to Figure 1, and nodes corresponding to the pairs of inhibitors are encircled. For each compound pair, the NSG at the top captures the selectivity cliff while the other two NSG views illustrate the presence or absence of activity cliffs formed by each pair of inhibitors against the individual targets. The two capthepsin inhibitors do not form activity cliffs against cathepsin L or B but a notable selectivity cliff (within this data set of relatively weakly potent inhibitors). A different scenario is observed for the COX inhibitors. Diclofenac is highly potent against both COX-1 and COX-2, whereas lumiracoxib is highly potent against COX-1 but nearly inactive against COX-2. Hence, this compound pair forms a large-magnitude activity cliff against COX-2 but no activity cliff against COX-1. However, these two inhibitors form a prominent selectivity cliff, due to their dramatic difference in potency against COX-2.

distributions. K_i measurements, which represent equilibrium constants (different from IC_{50} values), yielded consistently fewer activity cliffs than IC_{50} measurements.[21] These findings clearly indicate that assay variability and approximate potency measurements generally lead to larger numbers of activity cliffs than the use of high-confidence data and equilibrium constants. Thus, limited accuracy of experimentally determined potency values causes a tendency of false-positive activity cliff assignments. Generally, most stable and reliable activity cliff distributions are produced by compounds for which multiple K_i values of comparable magnitude are available.[21] For such high-confidence data, the choice of alternative measurements has only little influence on cliff formation.

EXTENSIONS OF THE ACTIVITY CLIFF CONCEPT

Following the discussion of intrinsic limitations of activity cliff analysis, this section focuses on specific modifications of the activity cliff concept. Recently, the activity cliff definition has been further extended in different ways and several specialized cliff representations have been introduced.

Consensus Activity Cliffs. Given the influence of molecular representations on cliff formation (vide supra), attempts have been made to identify activity cliffs in compound sets that are consistently formed when alternative descriptors are used.[16] Accordingly, such cliffs, termed consensus activity cliffs,[16] are least affected by changes in molecular representations and similarity assessment. Thus, they should present interesting test cases for further analysis, both from a chemical reference space and medicinal chemistry perspective. For example, it might be interesting to study if invariant cliffs are preferentially found in specific compound classes.

R-Cliffs. As an extension of conventional R-group tables, series of analogues have been graphically organized with respect to multiple replacements at each individual substitution site.[23] This organization has also permitted the identification of activity cliffs formed at each site.[23] These site-specific activity cliffs are termed R-cliffs (in analogy to R-groups) and derived for single analogue series.

Selectivity Cliffs. Activity cliffs are typically determined for individual targets. However, it is straightforward to conceptualize selectivity cliffs for pairs of targets.[24] Following this idea, a selectivity cliff is formed by a pair of compounds having significantly different potencies against one or two targets of a pair. Exemplary selectivity cliffs are depicted in Figure 4. Importantly, compounds forming a selectivity cliff may or may not form activity cliffs against the individual targets, which results in a possible gain in information when considering the formation of selectivity cliffs,[24] as also illustrated in Figure 4.

Multitarget Activity Cliffs. A further extension of the activity cliff concept into bioactivity space is provided by the introduction of multitarget activity cliffs.[15,25] Going beyond the assessment of target pair selectivity, these types of cliffs are formed by compounds having different potency against series of targets, for example, members of a given protein family. For compounds with activity profiles against a specific number of targets, a hierarchy of all theoretically possible multitarget activity cliffs (of predefined magnitude) can be formally derived.[25] As illustrated in Figure 5, multitarget cliffs can be "directed" or "undirected". In the former case, one of the cliff partners has consistently high potency against its targets while the other has low potency; in the latter, the cliff partners have different high or low potency against at least one of several targets. These potency relationships determine the directionality of multitarget activity cliffs.

Mechanism Cliffs. In activity landscape representations, compound potency measurements have also been complemented with molecular mechanism-of-action information. This provides a basis for the exploration of structural changes that lead to "mechanism hopping".[26] The analysis of such effects is particularly relevant for receptor ligands (often of G protein

dx.doi.org/10.1021/jm201706b | J. Med. Chem. 2012, 55, 2932−2942

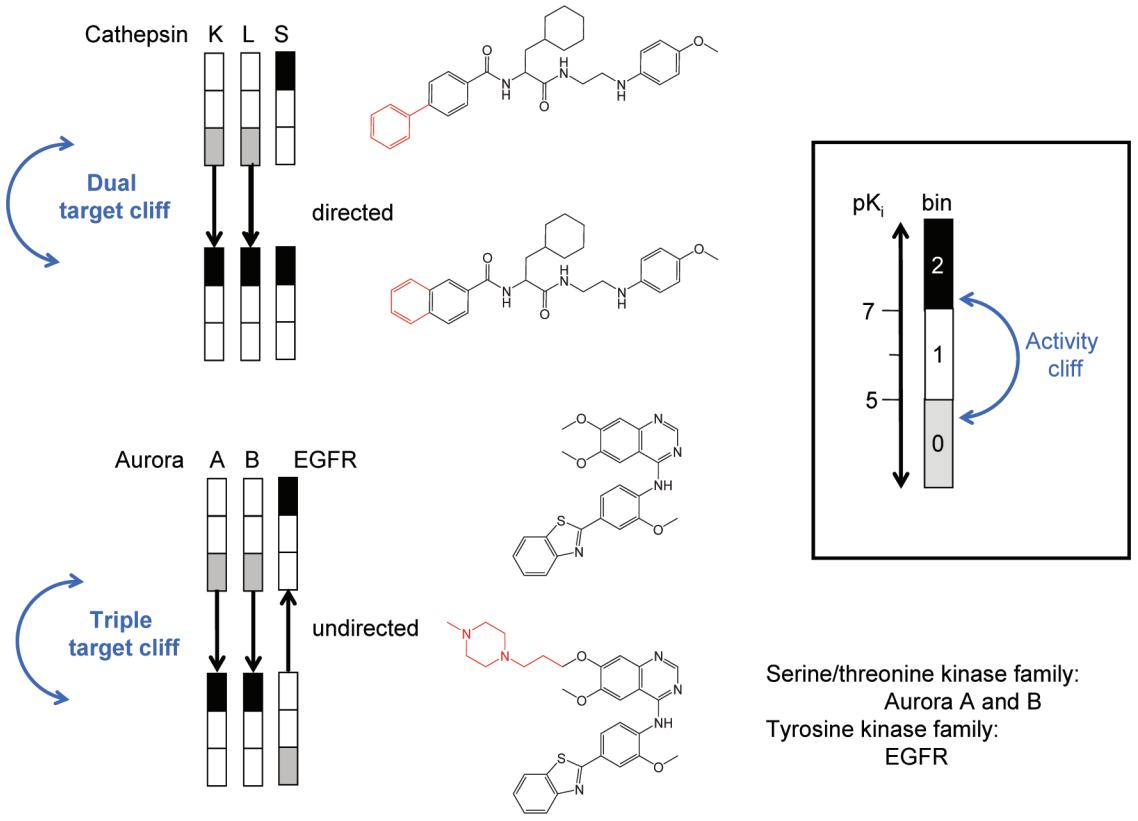

Figure 5. Multitarget activity cliffs. Exemplary directed and undirected dual- and triple-target activity cliffs are shown. For the schematic representation of these activity cliffs, potency values were binned as indicated on the right. A directed dual-target cliff is formed by two inhibitors of cathepsin K, L, and S (top). In addition, an undirected triple-target cliff is observed for two inhibitors of protein kinases belonging to different families: Aurora serine/threonine kinases A and B and the epidermal growth factor receptor (EGFR) tyrosine kinase. Structural differences between the compounds in each pair are indicated in red. The figure is adapted from ref 25.

coupled receptors) that might act by one of several alternative mechanisms including agonists, partial agonists, inverse agonists, or antagonists. Following this approach, mechanism cliffs are formed by pairs of compounds in series of analogues where small structural modifications induce a transition from one molecular mechanism to another. Exemplary mechanism cliffs are shown in Figure 6.

LIGAND-BASED VS STRUCTURE-BASED VIEWS

In medicinal chemistry, SAR information is typically extracted from ligand sets. However, SAR discontinuity is often (but not always) determined by specific receptor–ligand interactions. Hence, the formation of activity cliffs might be attributed to the presence or absence of one or more critical receptor–ligand contacts, for example, an important hydrogen bond, ionic interaction, or a complementary fit of an aromatic substituent into a hydrophobic pocket. The critical role of such interactions is ultimately also encoded as SAR information in sets of ligands. Given the direct link between ligand–target interactions and SAR characteristics, it is logical and attractive to also explore activity cliff formation at the target structure level, for example, with the aid of ligand–target complex structures. Thus far, only few attempts have been made to study activity cliff formation at the target structure level. For example, exploring prominent activity cliffs identified in compound data sets using structures of complexes has revealed that these cliffs can often, but clearly not always, be rationalized on the basis of short-range ligand–protein interactions formed in X-ray structures.[27] Of course, interactions seen in complex structures only provide an incomplete account of a ligand binding or inhibition process where entropic effects and/or desolvation penalties often play an important role. Nevertheless, rationalizing activity cliffs detected in ligand sets taking target structures into account is not only intellectually stimulating but can provide important information for structure-based drug design about the relevance of individual interactions seen in ligand–target complexes and reveal opportunities for further compound modifications.

Going beyond the analysis of selected activity cliffs on the basis of structural data, it has recently been attempted to systematically deduce activity cliffs from (experimentally observed or modeled) ligand–target interactions. Following this elegant approach termed identification of structure-based activity cliffs (ISAC),[28] computed protein–ligand interaction energies are utilized to derive an interaction fingerprint for each available active compound by scoring its contacts with individual target atoms. The interaction fingerprint then replaces conventional structural fingerprints for the calculation of Tanimoto similarity to derive SALI scores. For prominent activity cliffs formed by compounds with similar interaction fingerprints but different potency, individual interactions that distinguish the cliff partners from each other are extracted from their fingerprint representations. Target protein atoms involved in these discriminatory interactions are then considered "hot spots" for the formation of activity cliffs.[28] The cardinal feature of the ISAC approach that sets it apart from conventional activity cliff analysis is the replacement of structural ligand similarity with a measure of

dx.doi.org/10.1021/jm201706b | J. Med. Chem. 2012, 55, 2932–2942

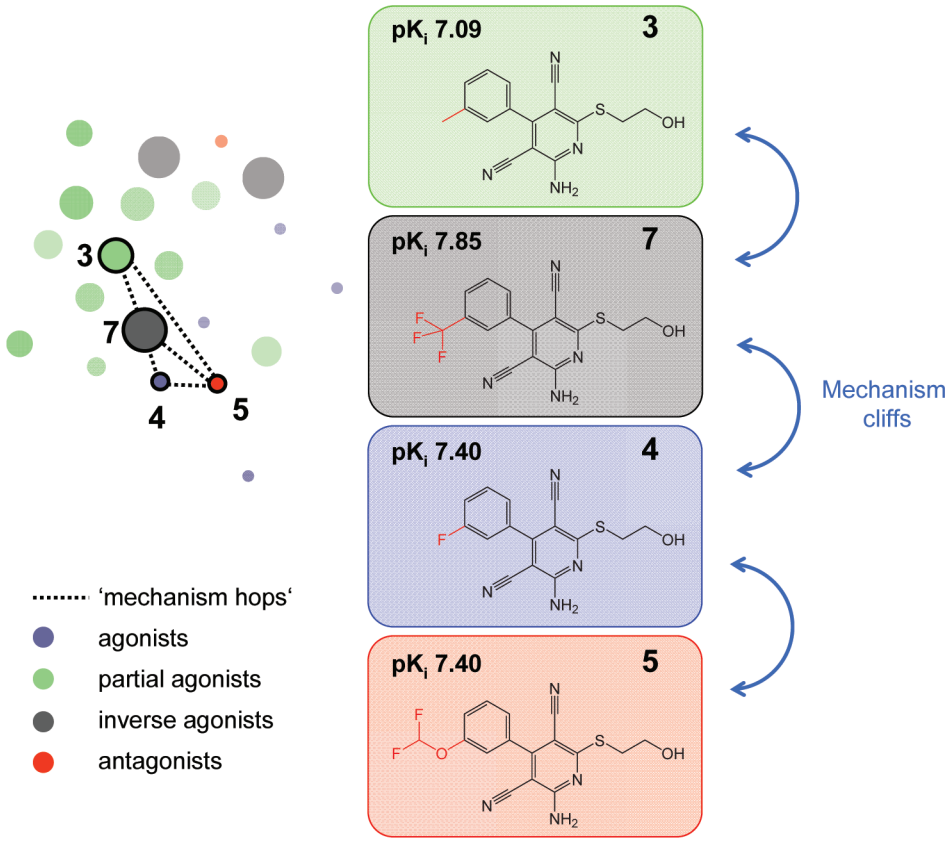

Figure 6. Mechanism cliffs. On the left, a subgraph of a mechanism-based NSG representation of adenosine A1 receptor ligands is shown where node colors represent different molecular mechanisms of action (and node transparency reflects potency). Nodes representing pairs of ligands that constitute mechanism hops are connected by dashed edges. Compounds and nodes are numbered correspondingly. Structural differences between analogues are highlighted in red and displayed on a background colored according to mechanism. The four analogues have comparable potency but constitute a series of mechanism cliffs. The figure is adapted from ref 26.

ligand−target interaction similarity in the context of the SALI formalism.

■ ACTIVITY CLIFF DISTRIBUTION

Key questions for activity cliff analysis and for judging the relevance of the activity cliff concept for medicinal chemistry include how frequently significant activity cliffs occur in compound data sets and how they might be distributed over different target families. To provide answers to these questions, a systematic search for activity cliffs has been carried out in BindingDB and ChEMBL[29] compound data sets.[15] It was found that approximately 12% of all bioactive compounds were involved in the formation of activity cliffs of at least 2 orders of magnitude (corresponding to 2% of all possible pairs of structurally similar compounds). However, only 4% of all activity cliffs were multitarget cliffs and nearly all of these cliffs were directed. Thus, it follows that activity cliff compounds with different selectivity for multiple targets are very rare, a finding with significant implications for the design of selective compounds. For example, if a compound is found to have high potency against a given target, it is likely that it will also be highly potent against related targets.[15] Thus, on the basis of currently available data, it might be difficult to modify the compound in such a way that it displays differential potency against series of related targets. For such considerations, multitarget activity cliff information is rather useful. It is of course well appreciated in medicinal chemistry that generating active compounds that are highly selective for one target over closely related ones is often (but not always) a very difficult task.

Furthermore, it was also found that cliffs were similarly distributed over different target families, without revealing a significant enrichment for individual families (as one might have expected). Thus, on the basis of currently available data, differences in the specifics of ligand−target interactions do not significantly alter activity cliff propensities.

From this systematic activity cliff survey, it can be concluded that more than 10% of all currently available active compounds are involved in the formation of at least one or two single-target activity cliffs of significant magnitude, with similar coverage over different target families. Hence, for medicinal chemistry applications, activity cliffs should indeed provide a substantial source of SAR information.

This view is further substantiated by activity cliff analyses at the level of molecular building blocks. For example, in a systematic analysis of chemical substitutions inducing activity cliffs, facilitated through the application of the matched molecular pair formalism,[30] approximately 200 R-group replacements have been identified that display a strong tendency to form activity cliffs in different compound classes across different target families.[31] Moreover, approximately 100 heteroatom-containing molecular scaffolds[32] of varying size and chemical complexity have been identified to preferentially occur in activity cliff forming compounds, also across different target families.[33] Thus, medicinal chemistry analysis of activity cliffs and compound design should not only focus on substitution

dx.doi.org/10.1021/jm201706b | J. Med. Chem. 2012, 55, 2932−2942

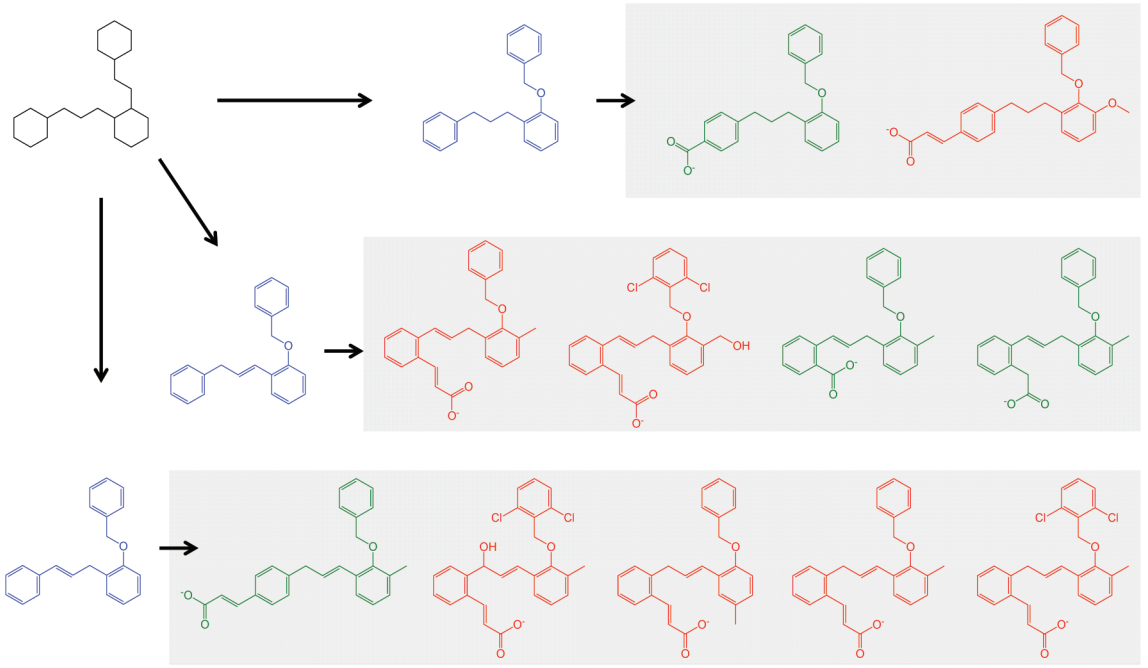

Figure 7. Activity ridge. Shown is a set of prostanoid EP3 receptor antagonists that form an exemplary activity ridge. A cyclic skeleton (black) covers three very similar molecular scaffolds (blue) each of which represents a varying number of highly (red) and weakly potent (green) analogues. Each pair of red and green compounds forms an activity cliff according to the criteria given in the text.

patterns in given analogue series but also include the exploration of alternative high-priority scaffolds. This might also provide a basis for a more detailed analysis of SAR transfer events than has been possible before. For example, cliff forming substitutions might be evaluated at corresponding sites in alternative scaffolds, thus providing opportunities for future research.

■ ACTIVITY RIDGES

An observation we have consistently made in compound data analysis is that activity cliffs often do not occur in isolation, i.e., in the absence of close structural neighbors, but involve groups of compounds forming multiple cliffs. Thus, the "classical" view of activity cliffs and compound pairs should be extended to include cliffs formed by multiple compounds. To assess this conjecture, a hypothetical data structure termed "activity ridge" was introduced.[34] This structure was proposed to consist of a "nanomolar layer" of at least five compounds (all with potency values within an order of magnitude) and another layer of at least five compounds with lower or higher potency, with an at least 100-fold potency difference between these two layers and the additional requirement that each compound in one layer must form pairwise activity cliffs with all compounds in the other. Hence, the activity ridge structure was envisioned to involve "combinatorial" activity cliff formation. When 242 compound data sets for this activity cliff-rich data structure were searched, a total of 125 activity ridges were identified in 71 compound activity classes involving up to 70 active compounds.[34] An exemplary activity ridge is shown in Figure 7. When activity ridge criteria are relaxed and fewer participating compounds are required, many more groups of molecules forming multiple activity cliffs are identified. Activity ridges are rich in SAR information and hence particularly interesting from a medicinal chemistry perspective.

Furthermore, as also illustrated in Figure 7, for the definition of activity ridges, hierarchical structural relationships between compounds, heteroatom-containing scaffolds,[32] and cyclic skeletons[35] (that further abstract from bond orders and heteroatom positions) were utilized as a structural similarity criterion to replace calculated Tanimoto similarity. Accordingly, scaffolds and the analogues they represent are considered similar if they correspond to the same cyclic skeleton (i.e., if they have the same topology). There are instances where the heteroatom content of topologically equivalent scaffolds substantially differs, which then questions the chemical similarity of theses scaffolds. However, for the majority of activity ridges that were detected and analyzed, this structural organization scheme identified analogue series around topologically equivalent scaffolds with comparable heteroatom arrangements. Hence, this intuitive similarity criterion is regarded as a viable alternative to calculated similarity values for activity cliff representation.

■ INTERPRETATION

In medicinal chemistry, SAR analysis should ultimately suggest new compounds to be made. Although activity cliffs reveal structural modifications that are of critical importance for biological activity, their identification does not automatically enable compound design, especially when cliffs are considered on a case-by-case basis. Clearly, activity cliffs need to be analyzed in light of general SAR trends and interpretable structural relationships. For calculated similarity values, this might not always be the case (vide supra). Accordingly, attempts have also been made recently to completely replace calculated similarity values in SAR networks with well-defined substructure relationships.[36] In order to maximize the information provided by activity cliff analysis, searching for multiple cliffs in structurally related series, as exemplified by activity ridges, is of high value. Insights obtained from a comparison of multiple cliffs help to transform activity cliff information into compound design suggestions.

dx.doi.org/10.1021/jm201706b | J. Med. Chem. 2012, 55, 2932−2942

Furthermore, it is worth noting that almost all compound data sets analyzed thus far contain activity cliffs of considerable magnitude, often with high frequency. Thus, for hit-to-lead or lead optimization projects focusing on popular targets, it is meaningful to carefully mine existing compound data for activity cliff information when new compound series are taken into consideration. It might then often be possible to focus on critical substitution patterns in different sets of compounds, explore corresponding substitutions on related scaffolds, and consider SAR transfer potential. Thus, while the identification of activity cliffs alone does not ensure interpretability of SAR information, the exploration of activity cliffs within the context of alternative compound series provides a promising basis for the identification of SAR determinants.

■ PRACTICAL CONSIDERATIONS

Although it is difficult to propose generally applicable procedures on how to best incorporate activity cliff analysis into practical medicinal chemistry projects (given the specific requirements of individual projects), following a few simple guidelines should make it possible to consistently take activity cliffs into consideration and avoid artificial bias in their analysis. Most compound sets contain activity cliffs, and it is hence useful to search for them in the context of any medicinal chemistry efforts on targets for which prior compound information is available. Graphical access to activity cliffs in compound data sets of any source is provided by freely available programs tools such as SALI graphs[11] or various tools implemented in the SARANEA software environment.[37] These graphical analysis programs are fairly easy to use and should be readily accessible to interested medicinal chemists, with some initial support by computational experts. These tools also provide a direct access to compound structures associated with cliffs, which further supports interactive analysis. Importantly, for a scientifically sound and practically relevant evaluation of cliffs, at least two further aspects should be considered. Whenever possible, equilibrium constants should be used as potency measurements to define activity cliffs, which helps to eliminate questionable cases. In addition, molecular similarity should be considered in a conservative manner. Hence, one should best limit cliff analysis to cases where chemical similarity is intuitive from a medicinal chemistry point of view. By contrast, one should certainly not rely on calculated similarity values to define cliffs in cases where similarity might be a matter of debate. Finally, activity cliffs should preferably not be considered as isolated occurrences. Rather, the structural neighborhood of cliffs should always be carefully inspected (which is easily done with the aid of graphical analysis tools). The environment of prominent activity cliffs might often reveal the formation of additional cliffs of varying magnitude or other interesting SAR information. For example, SAR continuity can frequently be observed in the neighborhood of activity cliffs, even within the same compound series.[38] Thus, for any practical purposes, activity cliff environments in compound data sets are a prime source of multilayered SAR information.

■ CONCLUDING DISCUSSION

Activity cliffs are a much discussed topic in chemoinformatics and medicinal chemistry, but are often not well-defined. Regardless of whether activity cliffs are considered as a continuum or as discrete states, any analysis must clearly define criteria for cliff formation. Furthermore, the study of activity cliffs is intrinsically biased by the use of alternative molecular representations and similarity measures (potentially leading to false-negatives) and by activity data integrity and variability (giving rise to false-positives). Hence, care must be taken to select appropriate molecular representations and pay attention to data quality. It also follows that chemical space design and the variance of chemical space representations will continue to influence activity cliff assessment. In this context, we promote the view that activity cliff analysis should best be limited to cases where structural relationships are obvious and that highly complex descriptor space representations should be avoided when attempting to study cliff formation in a systematic manner.

In general, activity cliffs are the most prominent features of activity landscapes and as such equally interesting for individual analysis or mining of large compound sets. Activity cliffs are currently mostly analyzed on the basis of ligand data but can also be studied in a complementary fashion using ligand−target complex structures. In the latter case, there is much less information available. However, exploring specifics of ligand−target interactions and binding modes that might favor or work against the introduction of cliffs is not only scientifically stimulating but also of immediate relevance for drug design.

Numerical local SAR analysis functions make it possible to identify the most prominent activity cliffs in any data set. However, whether those cliffs are significant and provide useful SAR information requires additional analysis. Moreover, by use of data mining methods, well-defined (discrete) activity cliffs can be systematically extracted from compound data sets. Although the large-scale exploration of activity cliffs (and of compound data structures containing multiple cliffs such as activity ridges) is still in its infancy, some conclusions can be drawn on the basis of currently available data. The picture emerges that large-magnitude activity cliffs frequently appear in sets of active compounds and are comparably distributed over different target classes. Furthermore, activity cliffs are often not formed in isolation but as parts of larger information-rich data structures. Hence, a significant body of activity cliff information is already available from which SAR information can be deduced. For medicinal chemistry applications, the interpretation of this information is a key aspect. Despite of their current popularity, the notion of activity cliffs alone does not lead to the generation of better compounds. Whether or not activity cliff information can be productively utilized will much depend on the SAR context, the observed compound potency distribution, and the interpretability of structural relationships.

The activity cliff concept has recently been extended in different ways, for example, by introducing consensus activity cliffs and selectivity or multitarget cliffs. It is anticipated that additional extensions will be considered for specific applications. Because the assessment of activity cliffs is a prime objective of activity landscape analysis, it is also expected that representations of activity cliffs will be further modified and refined as additional activity landscape models are introduced. These might include, among others, landscapes covering high-dimensional target space designed for the study of polypharmacology[39] or for chemogenomics[40] applications. In medicinal chemistry, it is expected that the activity cliff concept will continue to experience considerable interest not only in lead optimization but also in the context of large-scale data mining efforts. Despite the increasing notion of polypharmacological drug behavior and the utility of this concept in certain therapeutic areas (such as oncology), compound selectivity will continue to be a hallmark of medicinal chemistry efforts in many areas of drug discovery

(such as infectious or chronic diseases). Therefore, striking a balance between compound selectivity and promiscuity might often be of critical importance for future progress. From this point of view, applying the activity cliff concept and its extensions is expected to play a vital role in efforts focusing on compound selectivity against single targets as well as binding patterns in target families (for example, by exploring the formation of multitarget cliffs). Especially in the context of polypharmacology, a noteworthy feature of activity cliffs is that their presence provides a clear indication of specific interactions (causing SAR discontinuity). Hence, searching for activity cliffs should be very helpful to distinguish between "true" polypharmacological behavior and nonspecific binding events. In conclusion, it is hoped that an increasing number of case studies will become available in the near future that illustrate how the growing amount of activity cliff information (and the underlying scientific analysis concepts) might be transformed into practical compound design strategies.

AUTHOR INFORMATION

Corresponding Author

*Phone: +49-228-2699-306. Fax: +49-228-2699-341. E-mail: bajorath@bit.uni-bonn.de.

Notes

The authors declare no competing financial interest.

Biographies

Dagmar Stumpfe studied biology at the University of Bonn, Germany. In 2006, she joined the Department of Life Science Informatics at the University of Bonn headed by Prof. Jürgen Bajorath for her Ph.D. thesis, where she worked on methods for computer-aided chemical biology with a focus on the exploration of compound selectivity. Since 2009, Dagmar has been working as a postdoctoral fellow in the Department, and her current research interests include computational chemical biology and large-scale structure−activity relationship analysis.

Jürgen Bajorath is Professor and Chair of Life Science Informatics at the University of Bonn, Germany. He is also an Affiliate Professor in the Department of Biological Structure at the University of Washington, Seattle. His research interests include early phase drug discovery and the development of computational methods for chemical biology, molecular similarity analysis, and the systematic exploration and visualization of structure−activity relationships. For more details, see http://www.lifescienceinformatics.uni-bonn.de.

ABBREVIATIONS USED

ECFP4, extended connectivity fingerprint with bond diameter 4; ISAC, identification of structure-based activity cliffs; NSG, network-like similarity graph; QSAR, quantitative structure−activity relationship; SAR, structure−activity relationship; SAS, structure−activity similarity; 2D, two-dimensional; 3D, three-dimensional; SALI, structure−activity landscape index; SARI, structure−activity relationship index

REFERENCES

(1) Maggiora, G. M.; Shanmugasundaram, V.; Lajiness, M. S.; Doman, T. N.; Schulz, M. W. A Practical Strategy for Directed Compound Acquisition. In *Chemoinformatics in Drug Discovery*; Oprea, T. I., Ed.; Wiley-VCH: Weinheim, Germany, 2005; pp 317−332.

(2) Bajorath, J.; Peltason, L.; Wawer, M.; Guha, R.; Lajiness, M. S.; Van Drie, J. H. Navigating Structure−Activity Landscapes. *Drug Discovery Today* **2009**, *14*, 698−705.

(3) Peltason, L.; Bajorath, J. Systematic Computational Analysis of Structure−Activity Relationships: Concepts, Challenges and Recent Advances. *Future Med. Chem.* **2009**, *1*, 451−466.

(4) Bajorath, J., Maggiora, G., Lajiness, M., Organizers. The Emerging Concepts of Activity Landscapes and Activity Cliffs and Their Role in Drug Research. Section A of the Divisions of Chemical Information and Computers in Chemistry, 240th National Meeting of the American Chemical Society, Boston, MA, August 22−26, 2010.

(5) Wassermann, A. M.; Wawer, M.; Bajorath, J. Activity Landscape Representations for Structure−Activity Relationship Analysis. *J. Med. Chem.* **2010**, *53*, 8209−8223.

(6) Esposito, E. X.; Hopfinger, A. J.; Madura, J. D. Methods for Applying the Quantitative Structure−Activity Relationship Paradigm. *Methods Mol. Biol.* **2004**, *275*, 131−214.

(7) Maggiora, G. M. On Outliers and Activity Cliffs: Why QSAR Often Disappoints. *J. Chem. Inf. Model.* **2006**, *46*, 1535−1535.

(8) Shanmugasundaram, V.; Maggiora, G. M. Characterizing Property and Activity Landscapes Using an Information-Theoretic Approach. *Proceedings of 222nd National Meeting of the American Chemical Society*, Chicago, IL, August 26−30, 2001; American Chemical Society: Washington, DC, 2001; Abstract No. 77, Division of Chemical Information.

(9) Lajiness, M. Evaluation of the Performance of Dissimilarity Selection Methodology. In *QSAR: Rational Approaches to the Design of Bioactive Compounds*; Silipo, C., Vittoria, A., Eds.; Elsevier: Amsterdam, The Netherlands, 1991; pp 201−204.

(10) Wawer, M.; Peltason, L.; Weskamp, N.; Teckentrup, A.; Bajorath, J. Structure−Activity Relationship Anatomy by Network-like Similarity Graphs and Local Structure−Activity Relationship Indices. *J. Med. Chem.* **2008**, *51*, 6075−6084.

(11) Guha, R.; Van Drie, J. H. Structure−Activity Landscape Index: Identifying and Quantifying Activity Cliffs. *J. Chem. Inf. Model.* **2008**, *48*, 646−658.

(12) Peltason, L.; Iyer, P.; Bajorath, J. Rationalizing Three-Dimensional Activity Landscapes and the Influence of Molecular Representations on Landscape Topology and the Formation of Activity Cliffs. *J. Chem. Inf. Model.* **2010**, *50*, 1021−1033.

(13) Peltason, L.; Bajorath, J. SAR Index: Quantifying the Nature of Structure−Activity Relationships. *J. Med. Chem.* **2007**, *50*, 5571−5578.

(14) Wild, D. J.; Chen, B.; Zhu, Q. Exploring Activity Cliffs Using Large-Scale Semantic Analysis of PubChem. *Proceedings of the 240th National Meeting of the American Chemical Society*, Boston, MA, August 22−26, 2010; American Chemical Society: Washington, DC, 2010; Abstract No. 72, Division of Chemical Information.

(15) Wassermann, A. M.; Dimova, D.; Bajorath, J. Comprehensive Analysis of Single- and Multi-Target Activity Cliffs Formed by Currently Available Bioactive Compounds. *Chem. Biol. Drug Des.* **2011**, *78*, 224−228.

(16) Medina-Franco, J. L.; Martínez-Mayorga, K.; Bender, A.; Marín, R. M.; Giulianotti, M. A.; Pinilla, C.; Houghten, R. A. Characterization of Activity Landscapes using 2D and 3D Similarity Methods: Consensus Activity Cliffs. *J. Chem. Inf. Model.* **2009**, *49*, 477−491.

(17) Yongye, A. B.; Byler, K.; Santos, R.; Martínez-Mayorga, K.; Maggiora, G. M.; Medina-Franco, J. L. Consensus Models of Activity Landscapes with Multiple Chemical, Conformer, and Property Representations. *J. Chem. Inf. Model.* **2011**, *51*, 2427−2439.

(18) Rogers, D.; Hahn, M. Extended-Connectivity Fingerprints. *J. Chem. Inf. Model.* **2010**, *50*, 742−754.

(19) *MACCS Structural Keys*; Symyx Software: San Ramon, CA, 2005.

(20) Lajiness, M. Exploring and Exploiting the Potential of Structure−Activity Cliffs. *Proceedings of the 240th National Meeting of the American Chemical Society*, Boston, MA, August 22−26, 2010; American Chemical Society: Washington, DC, 2010; Abstract No. 61, Division of Chemical Information.

(21) Stumpfe, D.; Bajorath, J. Assessing the Confidence Level of Public Domain Compound Activity Data and the Impact of Alternative Potency Measurements on SAR Analysis. *J. Chem. Inf. Model.* **2011**, *51*, 3131−3137.

(22) Liu, T.; Lin, Y.; Wen, X.; Jorissen, R. N.; Gilson, M. K. BindingDB: a Web-Accessible Database of Experimentally Determined Protein−Ligand Binding Affinities. *Nucleic Acids Res.* **2007**, *35*, D198−D201.
(23) Agrafiotis, D. K.; Wiener, J. J. M.; Skalkin, A.; Kolpak, J. Single R-Group Polymorphisms (SRPs) and R-Cliffs: An Intuitive Framework for Analyzing and Visualizing Activity Cliffs in a Single Analog Series. *J. Chem. Inf. Model.* **2011**, *51*, 1122−1132.
(24) Peltason, L.; Hu, Y.; Bajorath, J. From Structure−Activity to Structure−Selectivity Relationships: Quantitative Assessment, Selectivity Cliffs, and Key Compounds. *ChemMedChem* **2009**, *4*, 1864−1873.
(25) Dimova, D.; Wawer, M.; Wassermann, A. M.; Bajorath, J. Design of Multi-Target Activity Landscapes That Capture Hierarchical Activity Cliff Distributions. *J. Chem. Inf. Model.* **2011**, *51*, 256−288.
(26) Iyer, P.; Stumpfe, D.; Bajorath, J. Molecular Mechanism-Based Network-like Similarity Graphs Reveal Relationships between Different Types of Receptor Ligands and Structural Changes That Determine Agonistic, Inverse-Agonistic, and Antagonistic Effects. *J. Chem. Inf. Model.* **2011**, *51*, 1281−1286.
(27) Sisay, M. T.; Peltason, L.; Bajorath, J. Structural Interpretation of Activity Cliffs Revealed by Systematic Analysis of Structure−Activity Relationships in Analog Series. *J. Chem. Inf. Model.* **2009**, *49*, 2179−2189.
(28) Seebeck, B.; Wagener, M.; Rarey, M. From Activity Cliffs to Target-Specific Scoring Models and Pharmacophore Hypotheses. *ChemMedChem* **2011**, *6*, 1630−1639.
(29) Gaulton, A.; Bellis, L. J.; Bento, A. P.; Chambers, J.; Davies, M.; Hersey, A.; Light, Y.; McGlinchey, S.; Michalovich, D.; Al-Lazikani, B.; Overington, J. P. ChEMBL: A Large-Scale Bioactivity Database for Drug Discovery. *Nucleic Acids Res.* **2012**, *40*, D1100−D1107.
(30) Kenny, P. W.; Sadowski, J. Structure Modification in Chemical Databases. In *Chemoinformatics in Drug Discovery*; Oprea, T. I., Ed.; Wiley-VCH: Weinheim, Germany, 2005; pp 271−285.
(31) Wassermann, A. M.; Bajorath, J. Chemical Substitutions That Introduce Activity Cliffs across Different Compound Classes and Biological Targets. *J. Chem. Inf. Model.* **2010**, *50*, 1248−1256.
(32) Bemis, G. W.; Murcko, M. A. The Properties of Known Drugs. 1. Molecular Frameworks. *J. Med. Chem.* **1996**, *39*, 2887−2893.
(33) Hu, Y.; Bajorath, J. Molecular Scaffolds with High Propensity to Form Multi-Target Activity Cliffs. *J. Chem. Inf. Model.* **2010**, *50*, 500−510.
(34) Vogt, M.; Huang, Y.; Bajorath, J. From Activity Cliffs to Activity Ridges: Informative Data Structures for SAR Analysis. *J. Chem. Inf. Model.* **2011**, *51*, 1848−1856.
(35) Xu, Y.-J.; Johnson, M. Algorithm for Naming Molecular Equivalence Classes Represented by Labeled Pseudographs. *J. Chem. Inf. Comput. Sci.* **2001**, *41*, 181−185.
(36) Wawer, M.; Bajorath, J. Local Structural Changes, Global Data Views: Graphical Substructure−Activity Relationship Trailing. *J. Med. Chem.* **2011**, *54*, 2944−2951.
(37) Lounkine, E.; Wawer, M.; Wassermann, A. M.; Bajorath, J. SARANEA: A Freely Available Program To Mine Structure−Activity and Structure−Selectivity Relationship Information in Compound Data Sets. *J. Chem. Inf. Model.* **2010**, *50*, 68−78.
(38) Namasivayam, V.; Iyer, P.; Bajorath, J. Exploring SAR Continuity in the Vicinity of Activity Cliffs. *Chem. Biol. Drug Des.* **2012**, *79*, 22−29.
(39) Hopkins, A. L. Network Pharmacology: The Next Paradigm in Drug Discovery. *Nat. Chem. Biol.* **2008**, *4*, 682−690.
(40) Jacoby, E.; Mozzarelli, A. Chemogenomic Strategies To Expand the Bioactive Chemical Space. *Curr. Med. Chem.* **2009**, *16*, 4374−4381.