Published in final edited form as:

J Chem Inf Model. 2008 April ; 48(4): 730–741. doi:10.1021/ci700369e.

Indirect Similarity Based Methods for Effective Scaffold-Hopping in Chemical Compounds

Nikil Wale^{*,†}, Ian A. Watson[‡], and George Karypis[†]

[†] *Department of Computer Science, University of Minnesota, Twin Cities, Minnesota*

[‡] *Eli Lilly and Company, Lilly Research Labs, Indianapolis, Indiana*

Abstract

Methods that can screen large databases to retrieve a structurally diverse set of compounds with desirable bioactivity properties are critical in the drug discovery and development process. This paper presents a set of such methods that are designed to find compounds that are structurally different to a certain query compound while retaining its bioactivity properties (scaffold hops). These methods utilize various indirect ways of measuring the similarity between the query and a compound that take into account additional information beyond their structure-based similarities. The set of techniques that are presented capture these indirect similarities using approaches based on analyzing the similarity network formed by the query and the database compounds. Experimental evaluation shows that most of these methods substantially outperform previously developed approaches both in terms of their ability to identify structurally diverse active compounds as well as active compounds in general.

1. Introduction

Discovery, design, and development of new drugs is an expensive and challenging process. Any new drug should not only produce the desired response to the disease but should do so with minimal side effects. One of the key steps in the drug design process is the identification of the chemical compounds (*hit* compounds or just *hits*) that display the desired and reproducible activity against the specific biomolecular target.¹ This represents a significant hurdle in the early stages of drug discovery.

A popular method for finding these hits is to use a structure-based ranked-retrieval approach. In this approach, an active compound is used as a *query* to identify and rank the compounds of a large database—based on how similar their structures is with that of the query. This approach relies on the structure–activity relationship (SAR)² principle, which states that that compounds sharing key structural features will most likely have similar activity against a biomolecular target. The structural similarity between the compounds is usually computed by first representing their molecular graph as a vector in a particular *descriptor-space* and then using a variety of vector-based methods to compute their similarity.^{2,3}

The task of identifying hit compounds is complicated by the fact that the query might have undesirable properties such as toxicity, bad ADME (absorption, distribution, metabolism, and excretion) properties, or may be promiscuous.^{4,5} These properties will also be shared by most

* To whom correspondence should be addressed. E-mail: nwale@cs.umn.edu.

Supporting Information Available: Data sets used in this paper. This material is available free of charge via the Internet at <http://pubs.acs.org>.

of the highest ranked compounds as they will correspond to very similar structures. In order to overcome this problem, it is important to identify (i.e., rank high) as many chemical compounds as possible that not only show the desired activity for the biomolecular target but also have different structures (come from diverse chemical classes or chemotypes). Finding novel chemotypes using the information of already known bioactive small molecules is termed as *scaffold-hopping*.^{4,6,7}

In this paper we address the problem of scaffold-hopping by developing a set of techniques that measure the similarity between the query and a compound by taking into account additional information beyond their structure-based similarities. This *indirect* way of measuring similarity enables the retrieval of compounds that are structurally different from the query but at the same time possess the desired bioactivity properties. The set of techniques we present derive the indirect similarities by analyzing the similarity network formed by the query and the database compounds. All of these techniques operate on the descriptor-space representation of the compounds and are independent of the selected descriptor-space.

We experimentally evaluate the performance of the proposed methods using three different descriptor-spaces and six different data sets. Our results show that most of these methods are quite effective in improving the scaffold-hopping performance over the standard retrieval schemes and that they substantially outperform schemes that were specifically developed for addressing the scaffold-hopping problem. Moreover, even though these methods were created to improve the scaffold-hopping performance, our results show that they are quite effective in improving the overall ranked-retrieval performance as well (i.e., rank the active compounds much higher than the inactive).

The rest of the paper is organized as follows. Section 2 describes the problems addressed in this paper. Section 3 introduces the definitions and notations used in this paper. Section 4 introduces the various descriptor-spaces for this problem. Section 5 describes the methods developed in this paper. Section 6 gives an overview of the related work in this field. Section 7 describes the materials used in our experimental methodology. Section 8 compares the results obtained, and finally, section 9 summarizes the results of this paper.

2. Problem Statement and Motivation

The ranked-retrieval and the scaffold-hopping problems that we consider in this paper are defined as follows:

Definition 1 (Ranked-Retrieval Problem). Given a query compound, rank the compounds in the database—based on how similar they are to the query in terms of their bioactivity.

Definition 2 (Scaffold-Hopping Problem). Given a query compound and a parameter k , retrieve the k compounds that are similar to the query in terms of their bioactivity but their structure is as dissimilar as possible to that of the query.

The solution to the ranked-retrieval problem relies on the well-known fact that the chemical structure of a compound relates to its activity (SAR).² As such, effective solutions can be devised that rank the compounds in the database—based on how structurally similar they are to the query.

However, for scaffold-hopping, the compounds retrieved must be structurally *sufficiently* similar to possess similar bioactivity but at the same time must be structurally *dissimilar* enough to be a novel chemotype. This is a much harder problem than simple ranked-retrieval as it has the additional constraint of maximizing dissimilarity that runs counter to SAR. Methods that have the ability to rank higher the compounds that are structurally different (different

chemotypes) have advantages over methods that do not. They improve the odds of being able to find a compound that is not only active for a biomolecular target but also has all the other desired properties (nontoxicity, good ADME properties, target specificity, etc.⁴) that the reference structure and compounds with similar structures might not possess. One of such compounds is then more likely to become a true drug candidate. Furthermore, scaffold-hopping is also important from the point of view of unpatented chemical space. Many important lead compounds and drug candidates have been already patented. In order to find new therapies and offer alternative treatments it is important for a pharmaceutical company to discover novel leads significantly different from the existing patented chemical space. Methods that perform scaffold-hopping can achieve those objectives.

3. Definitions and Notations

Throughout the paper we will use D to denote a database of chemical compounds, q to denote a query compound, and c to denote a chemical compound present in the database. Given two compounds c_i and c_j , represented in a certain descriptor-space, we will use $\text{dsim}(c_i, c_j)$ to denote their direct similarity, $\text{isim}(c_i, c_j)$ to denote their indirect similarity, and $\text{sim}(c_i, c_j)$ to denote either their direct or indirect similarity. The direct similarity between a pair of compounds is computed by a suitable similarity measure that takes into account their descriptor-space representation, whereas the indirect similarity is computed by taking additional information beyond their descriptor-space representation. Given a compound c_i and a set of compounds A , we will use $\text{sim}(c_i, c_j)$ to denote the average pairwise similarity between c_i and all the compounds in A . Given a query compound q , a database D , and a parameter k , we define to be the k compounds in D that are most similar to q . Given a compound c , a set of compounds A , and a similarity measure, its *k-nearest-neighbor list* contains the k compounds in A that are most similar to c . Finally, throughout the paper, we will refer to the task of retrieving active compounds as ranked-retrieval and the task of retrieving scaffold-hops as scaffold-hopping.

4. Descriptor Spaces

The similarity between chemical compounds is usually computed by first transforming them into a suitable descriptor-space representation.^{2,3} A number of different approaches have been developed to represent each compound by a set of descriptors. These descriptors can be based on physiochemical properties as well as topological and geometric substructures (fragments).^{8–15}

In this study we use three descriptor-spaces that have been shown to be very effective in the context of ranked-retrieval and/or scaffold-hopping. These descriptor-spaces are the graph fragments (GF),^{13,14} extended connectivity fingerprints (ECFP),^{11,12} and the extended reduced graph (ErG) descriptors.⁶

GF is a 2D topology-based descriptor-space^{13,14} that is based on all the graph fragments of a molecular graph up to a predefined size. The idea behind this descriptor-space is to capture all possible (arbitrary) fragment shapes up to a certain size in a given database in order to characterize this descriptor-space. We use seven as the upper limit on the size of the graph fragments used in this work. ECFP is also a 2D topological descriptor-space and many flavors of these descriptors have been described by several authors.^{11,12} The idea behind this descriptor-space is to capture the topology around each atom in the form of shells whose radius (number of bonds) ranges from 1 to l , where l is a user defined parameter. We use the ECZ3 variation of ECFP in which each atom is assigned a label corresponding to its atomic number and the maximum shell radius is set to three. Both extended connectivity fingerprints (ECFP) and GF have been shown to be highly effective for the ranked-retrieval of chemical compounds.^{11,13,14}

Extended reduced graph descriptors (ErG) is a pharmacophoric descriptor-space. A pharmacophore is defined as a critical 3D or 2D arrangement of molecular fragments forming a necessary but not sufficient condition for biological activity. The descriptors that rely only on 2D information are called 2D pharmacophoric descriptors whereas descriptors that utilize 3D information are called 3D pharmacophoric descriptors. ErG is a 2D pharmacophoric descriptor-space that combines the reduced graphs^{16,17} and binding property pairs¹⁸ to generate pharmacophoric descriptor-space. A detailed description on the generation of these pharmacophoric descriptors has been described in previous work.⁶

5. Methods

To improve the scaffold-hopping performance of ranked-retrieval approaches we developed techniques that are based on measuring the similarity between the query and a compound by taking into account additional information beyond their descriptor-space-based representation. These methods are motivated by the observation that if a query compound q is structurally similar to a database compound c_i and c_i is structurally similar to another database compound c_j , then q and c_j could be considered as being similar or related even though they may have zero or very low direct similarity. This indirect way of measuring similarity can enable the retrieval of compounds that are structurally different from the query but at the same time, due to associativity, possess the same bioactivity properties with the query.

5.1. Nearest-Neighbor Graph-Based Similarity

The set of techniques we develop to capture such indirect similarities were inspired by research in the fields of information retrieval and social network analysis. These techniques derive the indirect similarities by analyzing the network formed by a k -nearest-neighbor graph representation of the query and the database compounds.

The network linking the database compounds with each other *and* with the query is determined by using a k -nearest-neighbor (NG) and a k -mutual-nearest-neighbor (MG) graph. Both of these graphs contain a node for each of the compounds as well as a node for the query. However, they differ on the set of edges that they contain. In the k -nearest-neighbor graph there is an edge between a pair of nodes corresponding to compounds c_i and c_j , if c_i is in the k -nearest-neighbor list of c_j or vice-versa. In the k -mutual-nearest-neighbor graph, an edge exists only when c_i is in the k -nearest-neighbor list of c_j and c_j is in the k -nearest-neighbor list of c_i . These two graphs are illustrated in Figure 1. As a result of these definitions, each node in NG will be connected to at least k other nodes (assuming that each compound has a nonzero similarity to at least k other compounds), whereas in MG, each node will be connected to at most k other nodes.

Since the neighbors of each compound in these graphs correspond to some of its most structurally similar compounds and due to the relation between structure and activity, each pair of adjacent compounds will tend to have similar activity. Thus, these graphs can be considered as network structures for capturing bioactivity relations.

A number of different approaches have been developed for determining the similarity between nodes in social networks that take into account various topological characteristics of the underlying graphs.^{19,20} In our work, we determine the similarity between a pair of nodes as a function of the intersection of their adjacency lists, which takes into account all two-edge paths connecting these nodes. Specifically, the similarity between c_i and c_j with respect to graph G is given by

$$\text{isim}_G(c_i, c_j) = \frac{\text{adj}_G(c_i) \cap \text{adj}_G(c_j)}{\text{adj}_G(c_i) \cup \text{adj}_G(c_j)} \quad (1)$$

where $\text{adj}_G(c_i)$ and $\text{adj}_G(c_j)$ are the adjacency lists of c_i and c_j in G , respectively.

This measure assigns a high similarity value to a pair of compounds if both are very similar to a large set of common compounds. Thus, compounds that are part of reasonably tight clusters (i.e., a set of compounds whose structural similarity is high) will tend to have high indirect similarities as they will most likely have a large number of common neighbors. In such cases, the indirect similarity measure re-enforces the existing high direct similarities between compounds. However, the indirect similarity between a pair of compounds c_i and c_j can also be high even if their direct similarity is low. This can happen when the compounds in $\text{adj}_G(c_i) \cap \text{adj}_G(c_j)$ match different structural descriptors of c_i and c_j . In such cases, the indirect similarity measure is capable of identifying relatively weak structural similarities, making it possible to identify scaffold-hopping compounds.

Given the above graph-based indirect similarity measures, various strategies can be employed to retrieve compounds from the database. In our work, we used three such strategies. The first corresponds to that used by the standard ranked-retrieval method, whereas the other two are inspired by information retrieval methods used for automatic relevance feedback²¹ and are specifically designed to improve the scaffold-hopping performance.

5.1.1. Best-Sim Retrieval Strategy—This is the most widely used retrieval strategy and it simply returns the compounds that are the most similar to the query. Specifically, if A is the set of compounds that have been retrieved thus far, then the next compound c_{next} that is selected is given by

$$c_{\text{next}} = \arg \max_{c_i \in D-A} \{\text{isim}(c_i, q)\} \quad (2)$$

This compound is added to A , removed from the database, and the overall process is repeated until the desired number of compounds has been retrieved. Depending on whether the indirect similarities are computed on the NG or MG graph, we will refer to this retrieval strategy as BestSimNG or BestSimMG, respectively.

5.1.2. Best-Sum Retrieval Strategy—This retrieval strategy incorporates additional information from the set of compounds retrieved thus far (set A). Specifically, the compound selected, c_{next} , is the one that has the highest average similarity to the set $A \cup \{q\}$. That is,

$$c_{\text{next}} = \arg \max_{c_i \in D-A} \{\text{isim}(c_i, A \cup \{q\})\} \quad (3)$$

The motivation behind this approach is that due to SAR, the set A will contain a relatively large number of active compounds. Thus, by modifying the similarity between q and a compound c to also include how similar c is to the compounds in the set A , we obtain a similarity measure that is re-enforced by A 's active compounds. This enables the retrieval of active compounds that are similar to the compounds present in A even if their similarity to the query is not very high; thus, enabling scaffold-hopping. We will refer to this best-sum retrieval strategy as BestSumNG and BestSumMG, corresponding to the NG and MG graph-based similarity measures, respectively.

5.1.3. Best-Max Retrieval Strategy—A key characteristic of the retrieval strategy described in section 5.1.2 is that the final ranking of each compound is computed by taking into account *all* the similarities between the compound and the compounds in the set A . Since

the compounds in A will tend to be structurally similar to the query compound, this approach is rather conservative in its attempt to identify active compounds that are structurally different from the query (i.e., scaffold-hops).

To overcome this problem, we developed a retrieval strategy that is based on the best-sum approach but instead of selecting the next compound based on its average similarity to the set $A \cup \{q\}$, it selects the compound that is the most similar to *one* of the compounds in $A \cup \{q\}$. That is, the next compound is given by

$$c_{\text{next}} = \arg \max_{c_i \in D-A} \{ \max_{c_j \in A \cup \{q\}} \text{sim}(c_i, c_j) \} \quad (4)$$

In this approach, if a compound c_j other than q has the highest similarity to some compound c_i in the database, c_i is chosen as c_{next} and added to A irrespective of its similarity to q . Thus, the query-to-compound similarity is not necessarily included in every iteration as in the other schemes, allowing this strategy to identify compounds that are structurally different from the query. Depending on whether the indirect similarities are computed on the NG or MG graph, we will refer to this retrieval strategy as BestMaxNG or BestMaxMG, respectively.

5.1.4. Illustrative Example—We illustrate the above methods (Best-Sim, Best-Sum, and Best-Max) using a small example. The task is to retrieve five compounds from a database of seven compounds $\{c_1, \dots, c_7\}$ using the three retrieval strategies. Figure 2 shows (a) the indirect similarity matrix between the compounds in the database, (b) the indirect similarity between q and the database compounds, and (c) the five most similar compounds to the query (and their ranking) using different retrieval strategies.

Using the Best-Sim retrieval strategy (described in) for the query q , the first compound retrieved is c_6 as it has the highest similarity (0.53) to the query among all the compounds. This compound added to the set A and removed from the database. The next compound selected is c_7 as it now has the highest similarity to the query among all the compounds left in the database. It is added to the set A and removed from the database. The above procedure is repeated which will results in the retrieval of c_5 followed by c_3 , and finally c_1 .

Using the Best-Sum retrieval strategy (described in section 5.1.1) for the query q , the first compound retrieved is again c_6 . This compound is added to the set A and removed from the database. The next compound selected will have the highest sum of similarities to the query q and all the compounds in set A ($A = \{c_6\}$). This value is 0.7 ($0.32 + 0.38$ from Figure 2) for c_1 . Similarly, the sum of similarity values are 0.36 for c_2 , 0.37 for c_3 , 0.36 for c_4 , 0.69 for c_5 , and 0.68 for c_7 . Therefore, c_1 will be the next compound selected, added to A , and removed from the database. The above procedure is repeated which will results in the retrieval of c_5 followed by c_7 , and c_3 .

Using the Best-Max retrieval strategy (described in section 5.1.3) for the query q , the first compound retrieved will be c_6 . This compound is added to the set A and removed from the database. For each of the remaining compounds in the database, the method calculates their highest similarity value to the query or one of the compounds in the set A . This value is 0.38 (maximum between 0.32 and 0.38 from Figure 2) for c_1 . Similarly, the highest similarity values are 0.28 for c_2 , 0.34 for c_3 , 0.19 for c_4 , 0.35 for c_5 , and 0.38 for c_7 . The next compound selected will have the maximum among all such values which can either be c_1 or c_7 . The tie is broken by randomly selecting one of the two compounds. Let us assume that c_1 is selected. Following this procedure the next compound selected is c_7 followed by c_5 , and finally c_3 .

Therefore, each of the above method may result in a different sequence in which the compounds are retrieved.

5.2. Combining Parameters

One of the drawbacks of the NG and MG graph-based schemes is that given a problem, it is not clear what should be the right number of neighbors (k) used in constructing the NG and MG graphs. Moreover, our experience with this approach has shown that the best-performing value of k is to some extent data set dependent. In order to increase the robustness of the overall approach we introduce two extensions of the above methods that take as input multiple parameter values. In particular, if the number of input parameters is m , then the corresponding NG or MG graph is derived for each of the input parameter values. The indirect similarity values are calculated using for each pair of compounds c_i and c_j corresponding to each of the m graphs. A unique indirect similarity value is then derived by choosing either the sum of all the m similarities or taking the maximum of these values. Thus, for each of the six schemes BestSimNG, BestSumNG, BestMaxNG, BestSimMG, BestSumMG, and BestMaxMG, we have two extensions: one based on summation and the other based on choosing the maximum of all the similarity values. The two extensions of BestSimNG are referred to as BestSimNGsum and BestSimNGmax. The notation for the extensions of the other methods are derived in similar fashion.

5.3. Computational Complexity of Constructing Nearest Neighbor Graphs

The complexity of constructing the NG or MG graphs is dominated by the cost of finding the k -nearest neighbors for each compound. A simple all-against-all implementation of this step will require time that is quadratic on the size of the database. In the cases in which the dimensionality of the descriptor space is small (usually less than 30 dimensions^{24,37}), efficient tree-based algorithms like k - d trees, v - p trees, or quad-trees³⁷ can be used to significantly reduce that complexity. Unfortunately, since most of the descriptor spaces used in cheminformatics are high-dimensional, these approaches cannot be used. However, for most high-dimensional descriptor spaces, the compounds will only have a small number of the entire set of descriptors. In such cases, the similarity between a compound and all the database compounds can be computed using efficient methods based on inverted indices.²¹ Even though this does not change the worst case quadratic complexity, they dramatically reduce the amount of time required to find the k -nearest neighbors of each compound. Our experience has been that the NG and MG graphs of a database with over 224 000 compounds (corresponding to the MLSMR compound library obtained from PubChem²⁷) can be constructed in under three hours on modern workstations. Since these graphs need to be computed only once for the database and then be used to answer different queries, their computational requirements are reasonable.

6. Related Work

Many methods have been proposed for ranked-retrieval and scaffold-hopping. These can be divided into two groups. The first contains methods that rely on better designed descriptor-space representations, whereas the second contains methods that are not specific to any descriptor-space representation but utilize different retrieval strategies to improve the overall performance.

Among the first set of methods, 2D descriptors such as path-based fingerprints,^{8,25} dictionary based keys,^{9,26} and, more recently, extended connectivity fingerprints (ECFP) and¹¹ graph fragments (GF)^{13,14} have all been successfully applied for the retrieval problem. Pharmacophore based descriptors such as ErG⁶ have been shown to outperform simple 2D topology based descriptors for scaffold-hopping.^{6,28} Lastly, descriptors based on 3D structure or conformations of the molecule have also been applied successfully for scaffold-hopping.^{5,28}

The second set of methods include the turbo search schemes (TurboSumFusion and TurboMaxFusion)⁴ and the structural unit analysis based techniques⁷ all of which utilize automatic relevance feedback²¹ ideas. These have been shown to be effective for both scaffold-hopping and ranked-retrieval. The turbo search techniques operate as follows. Given a query q , they start by retrieving the compounds from the database. Let A be the $(k + 1)$ -size set that contains q and the compounds. For each compound $c \in A$, all the compounds in the database are ranked in decreasing order based on their similarity to c , leading to $k + 1$ ranked lists. These lists are used to obtain the final similarity of each compound with respect to the initial query. In particular, in TurboMaxFusion, the similarity between q and a compound c is equal to the similarity corresponding to the maximum similarity of c in the $k + 1$ lists, whereas in TurboSumFusion, the similarity is equal to the sum of all the similarities in these rankings. Thus, these techniques utilize the retrieval strategies similar to those described in sections 5.1.3 and 5.1.2. Similar methods based on consensus scoring, rank averaging, and voting have also been investigated.²⁸

7. Materials

7.1. Data Sets

We used data sets that contain compounds that bind to six different biomolecular targets: COX2 (cyclooxygenase 2), CDK2 (cyclin-dependent kinase 2), FXa (coagulation factor Xa), PDE5 (phosphodiesterase 5), A1A (alpha-1A adrenoceptor), and MAO (monoamineoxidase). Each of these sets represent a different activity class.

The data sets for the first five targets are obtained from refs 29 and 30. The entire set consists of 2142 compounds and there are 50 active compounds for each one of the targets (250 in total). The rest of the compounds are “decoys” (inactive) obtained from the National Cancer Institute diversity set. For each target, we constructed a data set that contains its 50 active compounds and all the decoys. These data sets are termed as COX2, CDK2, PDE5, FXa, and A1A. Note that the decoys are untested molecules in this data set. Therefore, they are not necessarily inactive but rather have low direct similarity to the five activity classes in this data set.

The data set of the sixth target was derived from refs 13, 14, and 31, and after removing compounds with impossible Kekule forms and valence errors, it contains 1458 compounds. The compounds in this data set have been categorized into four different classes, 0, 1, 2, and 3 based on their levels of activity, with 0 indicating no activity. For our experiments, we treat all the compounds that have nonzero activity level (268 compounds) as active.

7.2. Definition of Scaffold-Hopping Compounds

Molecular scaffold is a widely cited concept and is used to evaluate the performance of a method with respect to its scaffold-hopping ability. However the definition of a scaffold-hop is highly subjective with numerous papers using different criteria to define what constitutes a scaffold-hop.^{4,7,28,32}

One of the approaches that is used to define and evaluate scaffold-hopping performance of a method is to “cluster” all the active compounds (global clustering) in a database and assign a structural class to each cluster. The scaffold-hopping performance of a method is then evaluated by identifying how many active compounds from different clusters (structural classes) are obtained.⁶ The assumption in this approach is that compounds with different scaffolds are more likely to be in different clusters and the method that retrieves actives from many different clusters among the top hits is good for scaffold-hopping. This evaluation strategy has the drawback that using global clustering masks the local effects and the final clusters depend on clustering parameters.^{32,33}

Another approach is to use mean pairwise similarity to assess the diversity of a set of actives from a given activity class (target class like kinase C inhibitors, COX2 inhibitors etc). If a method performs well for the class that has high diversity then it is good for scaffold-hopping.⁴ This approach depends on how the diversity is evaluated. Moreover, there is no baseline measure to objectively define absolute low or high diversity class.³² Other methods rely on carefully chosen active class compounds for a given target such that they belong to different chemotypes (maximum diversity).^{7,28} Again, the drawback of such approach that it is highly subjective and depends on the viewpoint of the domain expert.

In this paper, we define scaffold-hops by using an approach that directly relies on the scaffold-hopping problem definition (section 3). In particular, for a given query q , the active compounds are ranked based on their structural similarity to q , and the lowest 50% of them are defined to be the scaffold-hops for q . Thus, this approach identifies a set of scaffold-hopping compounds that are specific to each query and represent the 50% most dissimilar active compounds to the query. We use the 2048-bit path-based fingerprint generated by Chemaxon's screen program²⁵ for measuring the structural similarity between a query and an active compound although any other path-based fingerprint like daylight fingerprints could be used. These fingerprints are well-designed to capture structural similarity between two compounds.⁶ By directly defining scaffold-hops as described above, our approach captures compounds that are structurally different from the query and therefore more likely to be a scaffold-hop. We believe that this is a more objective approach to define scaffold-hops for a given compound. However, we note that this definition is not specifically designed to capture a medicinal chemist's view of a scaffold-hop.

7.3. Experimental Methodology

All the experiments were performed on dual core AMD Opterons with 4 GB of memory. We used the descriptor-spaces GF, ECZ3, and ErG (described in section 4) for evaluating the methods introduced in this paper. Each method is tested against six data sets (section 7.1) using three different descriptor-spaces (section 4) leading to a total of 18 different combinations of data sets and descriptor-spaces. We will refer to them as 18 different problems.

We use the Tanimoto similarity^{3,15,34} for all direct similarity calculations. The Tanimoto similarity function is given by

$$\text{dsim}(c_i, c_j) = \frac{\sum_{k=1}^n c_{ik} c_{jk}}{\sum_{k=1}^n (c_{ik})^2 + \sum_{k=1}^n (c_{jk})^2 - \sum_{k=1}^n c_{ik} c_{jk}} \quad (5)$$

where c_{ik} and c_{jk} are the values for the k th dimension in the n -dimensional descriptor-space representation for the compounds c_i and c_j , respectively. This similarity function was selected because it has been shown to be an effective way of measuring the similarity between chemical compounds^{15,34} for ranked-retrieval and is the most widely used similarity function in cheminformatics.

For each data set, we used each of its active compounds as a query and evaluated the extent to which the various methods lead to effective retrieval of the other active compounds and scaffold-hops.

We varied the parameter values for the methods described in and obtained results for each of the four different sets of values. For the nearest-neighbor methods which depend on the number of neighbors, we used $k = 4, 6, 8$, and 10 for the BestSimNG, BestSumNG, and BestMaxNG and $k = 12, 16, 20$, and 24 for the BestSimMG, BestSumMG, and BestMaxMG schemes. For

NG, the value of k less than 4 leads to graphs with many connected components whereas for MG this value is 12. Hence, we decided not to use values below these thresholds. Note that the threshold for NG is less than that of MG because the criterion for an edge to exist between two nodes of the neighborhood graph is stricter for MG as opposed to NG (section 5.1). We also derive results for the two parameter combination schemes of the graph-based methods (section 5.2) using all of these parameter values. For example, BestSimNG has six sets of results corresponding to the four parameter values ($k = 4, 6, 8$, and 10) and the two combination schemes BestSimNGsum and BestSimNGmax. Similarly we derive six sets of results for other schemes as well.

We compared our schemes against TurboMaxFusion and TurboSumFusion.⁴ For both of these methods, we used $k = 5, 10, 15$, and 20 . These values gave the best results and the results degraded as k was further increased. For a particular value of k , for example $k = 5$, we denote these schemes by TurboMaxFusion5 and TurboSumFusion5. The notation for the two schemes for other values of parameter k are derived in similar fashion.

7.4. Standard Retrieval

For each problem (data set and descriptor-space combination), we obtain a baseline performance by ranking all the compounds with respect to each active compound using the standard retrieval strategy described in section 5.1.1 and the Tanimoto similarity measure (direct similarity measure). We will refer to this scheme as *standard retrieval* and denote it by StdRetTM.

7.5. Performance Assessment Measures

We measure ranked-retrieval and scaffold-hopping performance using *uninterpolated precision*.³⁵ This is calculated as follows. For each active that appears in the top 50 retrieved compounds we compute the precision value. For ranked-retrieval, this is defined as the ratio of the number of actives retrieved over the number of compounds retrieved thus far. For scaffold-hopping it is defined as the number of scaffold-hops retrieved over the number of compounds retrieved thus far. For both ranked-retrieval and scaffold-hopping we sum all their precision values and normalize them by dividing them with 50. This is called the total uninterpolated precision for a query. Similar values are obtained for all the queries for a data set and the total uninterpolated precision is the average of all these values. Note that the total uninterpolated precision captures the number of active compounds (scaffold-hops) for each query as well as the position (rank) information of the actives (scaffold-hops).

To compare the ranked-retrieval performance of two methods, we evaluate their relative performance over all the 18 problems. This is achieved as follows. Let r_i and q_i represent the ranked-retrieval performance achieved by methods r and q on the i th problem respectively. We calculate the log-ratio, $\log_2(r_i/q_i)$, for every problem and take the average of these values. We call this quantity the *average relative performance* or of r with respect to q . On the average, if the is less than zero, r performs worse than q for the task of ranked-retrieval whereas if the is greater than zero, r performs better than q for that task. Note that the reason that we use log-ratios as opposed to simply the ratios is that the distribution of the ratios of two random variables is not symmetric whereas the distribution of their log-ratios is normally distributed. This allows us to compute their average and compare them in an unbiased way. We also compare the scaffold-hopping performance of two methods against each other in a similar fashion. We assess whether the for a given pair of methods is statistically significant using the student's t test,³⁶ which is well-suited to assess statistical significance of a sample of values drawn out of a normal distribution.

8. Results

In this paper, we choose to present the best results obtained for each of the methods described section 5.1 and TurboSumFusion and TurboMaxFusion methods described in section 6 (the complete set of results are available online¹). For each of the three retrieval strategies (i.e., best-sim, best-sum, and best-max) we only present the results obtained by the best-max combination scheme (section 5.2) for the NG and MG graphs. This combination method performed quite well for both scaffold-hopping and ranked-retrieval and achieved better results than those obtained by the schemes using a single value of k or those obtained by the best-sum combination method. For the fusion-based schemes, we present the results obtained for $k = 5$ as it gave the best results for the scaffold-hopping and ranked-retrieval task among the values tested.

8.1. Overall Performance Assessment

Tables 1 and 2 compare the performance of all the methods in a pairwise fashion for scaffold-hopping and ranked-retrieval, respectively. In each of these tables, we present two statistics. The first is the ARP of the row method (r) with respect to the column method (q) as described in section 7.5. The second statistic, shown immediately below the ARP value in parenthesis, is its p -value obtained from the student's t test. Note that for the remainder of this section we will consider the ARP of the two methods to be statistically significant if $p \leq 0.01$.

Comparing the performance of the nearest-neighbor methods, we observe that all of them show good performance for scaffold-hopping as well as ranked-retrieval. Among them, the best performing method is BestSumMGmax. It achieves the best balance between the ranked-retrieval and scaffold-hopping performance. Moreover, the results comparing best-sim retrieval strategies (BestSimNGmax and BestSimMGmax) to the corresponding best-sum and best-max retrieval strategies (BestSumNGmax, BestSumMGmax, BestMaxNGmax, and BestMaxMGmax) show that the latter are statistically better in all cases for the scaffold-hopping task. The best-sum and best-max retrieval methods also outperform the best-sim strategy for ranked-retrieval (although most of these differences are not statistically significant). These results indicate that automatic relevance feedback based strategies (best-sum and best-max) improve scaffold-hopping performance over the best-sim retrieval strategy without degrading the ranked-retrieval performance.

Comparing the results obtained by using the nearest-neighbor graphs over those obtained by using the mutual-nearest-neighbor graphs we see that the relative performance of the various retrieval strategies on these two graphs varies. For example, BestMaxNGmax is worse than BestMaxMGmax for scaffold-hopping but better for ranked-retrieval. However, most of these differences are not statistically significant.

Comparing the scaffold-hopping performance of the graph-based methods with that of StdRetTM (baseline scheme), we see that the former are on average 117% better. For example, BestSimMGmax outperforms StdRetTM by as much as 120%. Note that since the retrieval strategy in both of these methods is the same (best-sim), this particular result indicates that the rather dramatic improvement in performance is a direct consequence of the indirect similarity used by the BestSimMGmax scheme. The performance advantage of the graph-based methods over StdRetTM also carry for the problem of ranked-retrieval for which they always achieve better results (ranging from 3% to 26% better). Moreover, these improvements are statistically significant for the results obtained by the best-sum and best-max retrieval strategies.

Finally, comparing the scaffold-hopping performance of the graph-based methods with that of the two fusion-based schemes, we observe that the graph-based schemes perform consistently better (ranging from 6% to 125% better). In addition, for the best-sum and best-max nearest-

neighbor schemes these performance differences are also statistically significant. As it was the case with StdRetTM, the performance advantage of the graph-based schemes over the fusion-based schemes holds for ranked-retrieval as well. However, the gains are somewhat smaller than those achieved for scaffold-hopping.

8.2. Qualitative Comparisons

To illustrate the differences between StdRetTM, TurboMaxFusion5, and BestSumMGmax, Table 3 shows the active compounds that were present in the top-25 compounds retrieved by each of these methods on the A1A data set and ECZ3 descriptor space using Indoramin as the query compound (compound id 50). For each of the compounds retrieved, Table 3 also shows the chemical class that this compound belongs to. These chemical classes were obtained from the literature as described in ^{ref 30} and are available in the online supplement of ^{ref 30}.

These results show that although all three methods retrieve compounds belonging to classes 3, 6, and 7, the first two methods retrieved far fewer compounds (six and seven actives for StdRetTM and TurboMaxFusion5, respectively) as opposed to BestSumMGmax that retrieved 19 actives. Moreover, BestSumMGmax also retrieves compounds from two additional chemical classes (2 and 11) than the other two methods. The structure of some of these compounds is shown in Figure 3. Note that BestSumMGmax not only retrieves a larger number of active compounds, but it is also able to retrieve active compounds that have much lower direct similarity to the query (e.g., compounds 47, 24, 14, and 8) than the other two schemes. These low similarity compounds tend to have different scaffolds (as illustrated in Figure 3), which explains the improvements in scaffold-hopping that is achieved by BestSumMGmax and the other graph-based schemes. The above trends also holds for most of the other queries and data sets that we analyzed. The retrieved compounds shown in Figure 3 are all arylpiperazines although they have a different scaffold from the Indoramin query molecule.

8.3. Performance of Descriptor-Spaces

Our discussion so far focused on evaluating the average performance of the different methods across the various descriptor-space representations and data sets. In this section, we analyze the performance of the methods on the individual descriptor-spaces and data sets. We limit our evaluation to only the BestSumMGmax method as this methods achieve the best scaffold-hopping and ranked-retrieval performance among the graph-based methods.

The results of these evaluations are shown in Figures 4 and 5, which compare the performance of StdRetTM and TurboMaxFusion5 respectively against BestSumMGmax. In these figures, the left *Y*-axis represents uninterpolated precision values for ranked-retrieval, whereas the right *Y*-axis represents uninterpolated precision values for scaffold-hopping.

The results in Figure 4 show that for scaffold-hopping, BestSumMGmax performs consistently better than StdRetTM for all the descriptor-space and data set combinations. However, the actual gains are data set and descriptor-space dependent. For example, the gains are particularly high for the FXa, A1A, and COX2 data sets using the ErG descriptor-space and for FXa, COX2 and PDE5 data sets using the ECZ3 descriptor-space. Similar trends can be observed with the ranked-retrieval results, with BestSumMGmax outperforming StdRetTM in most of the 18 problems with FXa, COX2 and A1A showing maximum gains and only MAO data set showing a significant performance degradation in terms of ranked-retrieval task in all the three descriptor-spaces.

Finally, the results in Figure 5 show that for scaffold-hopping, BestSumMGmax outperforms TurboMaxFusion5 in most data set and descriptor-space combinations. Moreover, the trends

are similar to than of Figure 4, in that, BestSumMGmax is particularly better for the FXa, A1A, and COX2 data sets.

9. Discussion and Conclusion

In this paper, we focused on the problem of improving the scaffold-hopping performance of the ranked-retrieval systems used for chemical compound databases. We introduced a number of methods that utilize the compound's nearest neighbor graph to indirectly determine their pairwise similarity and devised various retrieval strategies that combine these indirect similarity measurements with retrieval strategies originally developed for automatic relevance feedback. Our results showed that, on the average, these nearest-neighbor graph-based methods consistently and quite often substantially outperform previously developed rank-retrieval and scaffold-hopping approaches. However, we also noticed that these schemes did not lead to major performance gains for those queries whose k -nearest neighbors are mostly inactive compounds. In these cases the number of common neighbors, and consequently the number of two-edge paths, between the query and other active compounds is usually very small. As a result, indirect similarity can not amplify and/or re-enforce direct similarity measures, leading to no measurable improvements.

Martin et al.³⁹ have analyzed the relationship between structural similarity and common biological activity and found that using Unity³⁸ fingerprints, compounds with Tanimoto similarity of ≥ 0.85 have only a 30% chance of being active against the same target. This suggests that a compound that is two edges away from an active compound in the ≥ 0.85 similarity graph, will have only a 9% chance of being active. This value is quite small and it appears to contradict the results obtained in our study in which the use of indirect similarity measures led to substantial performance gains. However, this is not the case. In our approach, the indirect similarity between a pair of compounds is a function of the number of their common neighbors (eq 1), which takes into account multiple two-edge paths. These multiple paths increase the 9% chance derived from a single path and is the reason why our approach is able to improve both the retrieval and the scaffold-hopping performance.

In summary, our results suggest that the nearest-neighbor graph formed by the query and the compounds in the database contains additional information on how the various compounds are related that goes well beyond the information that can be extracted by looking at their direct similarities and the set of the structural descriptors that they share. We believe that this information can be further exploited to obtain additional performance improvements for the ranked-retrieval and scaffold-hopping problems and to also address other problems arising in cheminformatics.

Supplementary Materials

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

This work was supported by NSF ACI-0133464, IIS-0431135, NIH RLM008713A, and by the Digital Technology Center at the University of Minnesota.

References and Notes

1. Eglen, R.M.; Schneider, G.; Bohm, H.J. High Throughput Screening and Virtual Screening: Entry Points to Drug Discovery. In: Bohm, H.J.; Schneider, G., editors. *Virtual Screening for Bioactive Molecules*. 10. Wiley-VCH; Weinheim Germany: 2000. p. 59-79.

2. Bravi, G.; Gancia, E.; Green, D.; Hann, VS.; Mike, M. Modelling structure-activity relationship. In: Bohm, HJ.; Schneider, G., editors. *Virtual Screening for Bioactive Molecules*. 10. Wiley-VCH; Weinheim Germany: 2000. p. 81-113.
3. Barnard, JM.; Downs, GM.; Willett, P. Descriptor-Based Similarity Measures for Screening Chemical Databases. In: Bohm, HJ.; Schneider, G., editors. *Virtual Screening for Bioactive Molecules*. 10. Wiley-VCH; Weinheim Germany: 2000. p. 59-79.
4. Hert J, Willett P, Wilton D. New methods for ligand based virtual screening: Use of data fusion and machine learning to enhance the effectiveness of similarity searching. *J Chem Inf Model* 2006;46:462–470. [PubMed: 16562973]
5. Saeh JC, Lyne PD, Takasaki BK, Cosgrove DA. Lead hopping using svm and 3d pharmacophore fingerprints. *J Chem Inf Model* 2005;45:1122–113. [PubMed: 16045307]
6. Stiefl N, Watson IA, Baumann K, Zaliani A. Erg: 2d pharmacophore descriptor for scaffold hopping. *J Chem Inf Model* 2006;46:208–220. [PubMed: 16426057]
7. Wolohan PN, Akella LB, Dorfman RJ, Nell PG, Mundt SM, Clark RD. Structural units analysis identifies lead series and facilitates scaffold hopping in combinatorial chemistry. *J Chem Inf Model* 2005;46:1188–1193. [PubMed: 16711738]
8. Daylight Toolkit. Daylight Chemical Information System; Aliso Viejo, CA: 2007.
9. MDL Information Systems: San Ramon, CA, 2007.
10. Deshpande M, Kuramochi M, Wale N, Karypis G. Frequent substructure-based approaches for classifying chemical compounds. *IEEE Trans Knowledge Data Eng* 2005;17:1036–1050.
11. Hert J, Willett P, Wilton D, Acklin P, Azzaoui K, Jacoby E, Schuffenhauer A. Comparison of topological descriptors for similarity-based virtual screening using multiple bioactive reference structures. *Org Biomol Chem* 2004;2:3256–3266. [PubMed: 15534703]
12. Rogers D, Brown R, Hahn M. Using extended-connectivity fingerprints with laplacian-modified bayesian analysis in high-throughput screening. *J Biomol Screen* 2005;10:682–686. [PubMed: 16170046]
13. Wale, N.; Karypis, G. Comparison of descriptor spaces for chemical compound retrieval and classification. *IEEE ICDM; Proceedings of the 6th Annual IEEE International Conference on DataMining; Hong Kong. Dec. 2006 p. 678-689.*
14. Wale N, Watson IA, Karypis G. Comparison of descriptor spaces for chemical compound retrieval and classification. *J Knowledge Inf Syst.* August 23rd;published online
15. Willett P. Chemical similarity searching. *J Chem Inf Comput Sci* 1998;38:983–996.
16. Gillet VJ, Willett P, Bradshaw J. Similarity searching using reduced graphs. *J Chem Inf Comput Sci* 2003;43:338–345. [PubMed: 12653495]
17. Harper G, Bravi GS, Pickett SD, Hussain J, Green DV. The reduced graph descriptor in virtual screening and data-driven clustering of high-throughput screening data. *J Chem Inf Comput Sci* 2004;44:45–56.
18. Kearsley SK, Sallamack S, Fluder EM, Andose JD, Mosley RT, Sheridan RP. Chemical similarity using physiochemical property descriptors. *J Chem Inf Comput Sci* 1996;36:118–127.
19. Fouss F, Pirotte A, Renders J, Saerens M. Random walk computation of similarities between nodes of a graph with application to collaborative filtering. *IEEE Trans Knowledge Data Eng* 2007;19:355–369.
20. Teufel B, Schmidt S. Full text retrieval based on syntactic similarities. *Inf Syst* 1988;13:65–70.
21. Baeza-Yates, R.; Ribeiro-Neto, B. *Modern Information Retrieval*. 1st. Addison Wesley; New York: 1999. Modelling.
22. Krauthgamer, R.; Lee, JR. Navigating nets: simple algorithms for proximity search. *ACM-SIAM; Proceedings of the Fifteenth Annual Society for Industrial and Applied Mathematics 2004, Symposium on Discrete Algorithms; New Orleans, Louisiana. Jan 11–14. 2004 p. 798-807.*
23. Hildrum, K.; Kubiawicz, J.; Ma, S.; Rao, S. A note on the nearest neighbor in growth-restricted metrics. *ACM-SIAM; Proceedings of the Fifteenth Annual Society for Industrial and Applied Mathematics 2004, Symposium on Discrete Algorithms; New Orleans, Louisiana. Jan 11–14. 2004 p. 560-561.*

24. Indyk, P. Nearest neighbors in high dimensional spaces. In: Goodman, JE.; O'Rourke, J., editors. Handbook of Discrete and Computational Geometry. 2nd. CRC Press; Boca Raton FL: 2004. p. 877-892.
25. Screen, version 3.1.1. ChemAxon; Budapest, Hungary: 2005.
26. Digital Chemistry: Leeds, U.K., 2007.
27. The PubChem Project. [Nov 2007]. <http://pubchem.ncbi.nlm.nih.gov>
28. Zhang Q, Muegge I. Scaffold hopping through virtual screening using 2d and 3d similarity descriptors: Ranking, voting and consensus scoring. *J Med Chem* 2006;49:1536–1548. [PubMed: 16509572]
29. cheminformatics.org. [July 2005]. <http://www.cheminformatics.org/>
30. Jorissen RN, Gibson MK. Virtual screening of molecular databases using support vector machines. *J Chem Inf Model* 2005;45:549–561. [PubMed: 15921445]
31. Brown R, Martin Y. Use of structure-activity data to compare structure-based clustering methods and descriptors for use in compound selection. *J Chem Inf Comput Sci* 1996;36:576–584.
32. Brown N, Jacoby E. On scaffolds and hopping in medicinal chemistry. *Mini Rev Med Chem* 2006;6:1217–1229. [PubMed: 17100633]
33. Schnecke V, Bostrom J. Computational chemistry-driven decision making in lead generation. *Drug Discov Today* 2006;11:43–50. [PubMed: 16478690]
34. Whittle M, Gillet VJ, Willett P. Enhancing the effectiveness of virtual screening by fusing nearest neighbor list: A comparison of similarity coefficients. *J Chem Inf Comput Sci* 2004;44:1840–1848. [PubMed: 15446844]
35. Hearst, M.; Pedersen, J. Reexamining the cluster hypothesis: Scatter/gather on retrieval results. *ACM SIGIR; Proceedings of the 19th Annual International Conference; Zurich. Jun. 1996 p. 7684*
36. Bland, JM. An introduction to medical statistics. 3rd. Oxford University Press; New York: 2000. Comparing the means of small samples.
37. Samet, H. Foundations of Multidimensional and Metric Data Structures. 1st. Elsevier; San Francisco: 2006. Multidimensional Point Data.
38. Unity Chemical Information Software. Tripos Associates; St Louis, MO: 2002.
39. Martin YC, Kofron JL, Traphagen LM. Do Structurally Similar Molecules have Similar Biological Activity. *J Med Chem* 2002;45:4350–4358. [PubMed: 12213076]

Figure 1.

(a) Part of a directed graph in which there is an edge from c_i to c_j if c_j occurs in the k -nearest-neighbors of c_i . (b) k -Nearest-neighbor graph (NG) created when c_j occurs in the k -nearest-neighbors of c_i or vice-versa. (c) k -Mutual-nearest-neighbor graph (MG) created when c_j occurs in the k -nearest-neighbors of c_i and vice versa.

Figure 2.

Order of the five most similar compounds to query q according to the three retrieval strategies Best-Sim, Best-Sum, and Best-Max. Each of the three strategies retrieves c_6 as the first compound but then diverge due to the different methods employed by them.

Figure 3. Subset of the α_{1A} inhibitors (A1A data set) retrieved using BestSumMGmax and missed by StdRetTM and TurboMaxFusion5 with Indoramin used as query.

Figure 4.

StdRetTM versus BestSumMGmax. Ranked-retrieval precision is defined as sum of the ratios of the number of actives retrieved over the number of compounds retrieved thus far up to rank 50. This number is then normalized by 50. Scaffold-hopping precision is defined as sum of the ratios of the number of actives retrieved over the number of compounds retrieved thus far up to rank 50. This number is then normalized by 50.

Figure 5.

TurboMaxFusion5 versus BestSumMGmax. Ranked-retrieval precision is defined as sum of the ratios of the number of actives retrieved over the number of compounds retrieved thus far up to rank 50. This number is then normalized by 50. Scaffold-hopping precision is defined as sum of the ratios of the number of actives retrieved over the number of compounds retrieved thus far up to rank 50. This number is then normalized by 50.

Table 1

Performance for Scaffold-Hopping^a

	StdRefTM	TurboSumFusion5	TurboMaxFusion5	BestSimNGmax	BestSumNGmax	BestMaxNGmax	BestSimMGmax	BestSumMGmax	BestMaxMGmax
StdRefTM		-0.52 (0.042)	-0.87 (0.008)	-0.94 (0.001)	-1.62 (0.000)	-1.69 (0.000)	-1.15 (0.001)	-1.82 (0.000)	-1.78 (0.000)
TurboSumFusion5	0.52 (0.042)		-0.34 (0.082)	-0.42 (0.017)	-1.10 (0.000)	-1.17 (0.000)	-0.63 (0.004)	-1.30 (0.000)	-1.25 (0.000)
TurboMaxFusion5	0.87 (0.008)	0.34 (0.082)		-0.08 (0.577)	-0.75 (0.001)	-0.83 (0.000)	-0.28 (0.017)	-0.96 (0.000)	-0.91 (0.000)
BestSimNGmax	0.94 (0.001)	0.42 (0.017)	0.08 (0.577)		-0.67 (0.001)	-0.75 (0.001)	-0.20 (0.069)	-0.88 (0.001)	-0.83 (0.001)
BestSumNGmax	1.62 (0.000)	1.10 (0.000)	0.75 (0.001)	0.67 (0.001)		-0.07 (0.561)	0.47 (0.014)	-0.20 (0.181)	-0.16 (0.404)
BestMaxNGmax	1.69 (0.000)	1.17 (0.000)	0.83 (0.000)	0.75 (0.001)	0.07 (0.561)		0.54 (0.006)	-0.13 (0.327)	-0.09 (0.615)
BestSimMGmax	1.15 (0.001)	0.63 (0.004)	0.28 (0.017)	0.20 (0.069)	-0.47 (0.014)	-0.54 (0.006)		-0.67 (0.001)	-0.63 (0.001)
BestSumMGmax	1.82 (0.000)	1.30 (0.000)	0.96 (0.000)	0.88 (0.001)	0.20 (0.181)	0.13 (0.327)	0.67 (0.001)		0.04 (0.692)
BestMaxMGmax	1.78 (0.000)	1.25 (0.000)	0.91 (0.000)	0.83 (0.001)	0.16 (0.404)	0.09 (0.615)	0.63 (0.001)	-0.04 (0.692)	

^aThe top entry in each cell corresponds to the average of the log2 ratios of the unintepolated precision of the row method to the column method for the 18 problems. The number below this entry, in parenthesis, corresponds to the *p*-value obtained from the student's *t* test for that entry.

Table 2

Performance for Ranked-Retrieval^a

	StdRetTM	TurboSumFusion5	TurboMaxFusion5	BestSimNGmax	BestSumNGmax	BestMaxNGmax	BestSimMGmax	BestSumMGmax	BestMaxMGmax
StdRetTM		-0.02 (0.576)	0.07 (0.083)	-0.06 (0.362)	-0.23 (0.038)	-0.18 (0.063)	-0.10 (0.150)	-0.27 (0.040)	-0.11 (0.318)
TurboSumFusion5	0.02 (0.576)		0.09 (0.003)	-0.04 (0.471)	-0.21 (0.029)	-0.16 (0.043)	-0.08 (0.195)	-0.25 (0.036)	-0.09 (0.346)
TurboMaxFusion5	-0.07 (0.083)	-0.09 (0.003)		-0.13 (0.085)	-0.30 (0.008)	-0.26 (0.011)	-0.18 (0.030)	-0.34 (0.012)	-0.19 (0.119)
BestSimNGmax	0.06 (0.362)	0.04 (0.471)	0.13 (0.085)		-0.17 (0.005)	-0.12 (0.030)	-0.04 (0.294)	-0.20 (0.015)	-0.05 (0.431)
BestSumNGmax	0.23 (0.038)	0.21 (0.029)	0.30 (0.008)	0.17 (0.005)		0.05 (0.257)	0.13 (0.059)	-0.04 (0.569)	0.12 (0.065)
BestMaxNGmax	0.18 (0.063)	0.16 (0.043)	0.26 (0.011)	0.12 (0.030)	-0.05 (0.257)		0.08 (0.233)	-0.08 (0.305)	0.07 (0.289)
BestSimMGmax	0.10 (0.150)	0.08 (0.195)	0.18 (0.030)	0.04 (0.294)	-0.13 (0.059)	-0.08 (0.233)		-0.16 (0.020)	-0.01 (0.867)
BestSumMGmax	0.27 (0.040)	0.25 (0.036)	0.34 (0.012)	0.20 (0.015)	0.04 (0.569)	0.08 (0.305)	0.16 (0.020)		0.15 (0.012)
BestMaxMGmax	0.11 (0.318)	0.09 (0.346)	0.19 (0.119)	0.05 (0.431)	-0.12 (0.065)	-0.07 (0.289)	0.01 (0.867)	-0.15 (0.012)	

^aThe top entry in each cell corresponds to the average of the log2 ratios of the uninterpolated precision of the row method to the column method for the 18 problems. The number below this entry, in parenthesis, corresponds to the *p*-value obtained from the student's *t* test for that entry.

Table 3
Number of Actives Retrieved in Top 25 Ranks Using Different Methods for Indoramin (Class 13) As Query^a

no. of pos	StdRefTM				TurboMaxFusion5				BestSumMGmax			
	rank	id	dsim	class	rank	id	dsim	class	rank	id	dsim	class
1	3	30	0.8494	7	4	30	0.8494	7	2	16	0.8204	3
2	7	35	0.8320	9	7	16	0.8204	3	3	15	0.8100	3
3	12	29	0.8264	7	8	29	0.8264	7	4	49	0.8023	12
4	14	16	0.8204	3	9	15	0.8100	3	5	41	0.7768	11
5	16	25	0.8185	6	11	25	0.8185	6	6	13	0.7798	2
6	23	15	0.8100	3	18	49	0.8023	12	7	48	0.7867	12
7					19	35	0.8320	9	8	43	0.7759	12
8									9	46	0.7852	12
9									10	29	0.8264	7
10									11	44	0.7597	12
11									12	8	0.7129	2
12									13	14	0.7050	2
13									15	24	0.7024	6
14									17	25	0.8185	6
15									18	45	0.7838	12
16									19	26	0.7295	7
17									20	47	0.6953	12
18									21	30	0.8494	7
19									23	42	0.7141	12

^aThe table shows the rank, compound identifier (id), direct Tanimoto similarity (dsim), and the chemical class of compounds retrieved for this query. For BestSumMGmax, the table also shows the indirect similarity values (isim) between Indoramin and the actives retrieved.