SCISSORS: A Linear-Algebraical Technique to Rapidly Approximate Chemical Similarities

Imran S. Haque[†] and Vijay S. Pande^{*,†,‡}

Department of Computer Science and Department of Chemistry, Stanford University, Stanford, CA

Received January 7, 2010

Algorithms for several emerging large-scale problems in cheminformatics have as their rate-limiting step the evaluation of relatively slow chemical similarity measures, such as structural similarity or three-dimensional (3-D) shape comparison. In this article we present SCISSORS, a linear-algebraical technique (related to multidimensional scaling and kernel principal components analysis) to rapidly estimate chemical similarities for several popular measures. We demonstrate that SCISSORS faithfully reflects its source similarity measures for both Tanimoto calculation and rank ordering. After an efficient precalculation step on a database, SCISSORS affords several orders of magnitude of speedup in database screening. SCISSORS furthermore provides an asymptotic speedup for large similarity matrix construction problems, reducing the number of conventional slow similarity evaluations required from quadratic to linear scaling.

INTRODUCTION

A fundamental problem in cheminformatics is the calculation of similarity between two given molecules. As a consequence, a large variety of similarity techniques exists.¹ These measures have as their underlying information content a variety of molecular properties, including various encodings of molecular substructure,^{2–4} volume,^{5,6} and surface similarity⁷ and of electrostatic similarity.⁸ While many of the more complicated techniques are able to uncover relevant chemical similarities not found by simpler methods,^{9,10} they are often computationally expensive to evaluate.

Many important algorithms in cheminformatics contain as a critical subroutine these pairwise similarity comparisons. For example, a database search against a single query (without filters) amounts to the evaluation of a similarity measure once for each database molecule. More complicated algorithms, such as those for clustering¹¹ or network construction,¹² may require the evaluation of a number of similarities quadratic in the size of the database, rather than linear. Evaluation of similarities can be a bottleneck, limiting performance as well as the size of problems that can be considered. While fingerprint-style methods have been developed to approximate these similarity measures,¹³ they lack rigorous justifications of their accuracy.

A significant continuing trend in cheminformatics is the increasing size of virtual chemical databases. Public libraries listing known chemical matter, such as PubChem (31 million molecules) and ZINC¹⁴ (34 million molecules), are routinely used in database searches. Continuing advances in both computational power and storage space enable the use of even larger exhaustive and combinatorial databases. GDB-13¹⁵ is an exhaustive database enumerating all 970 million possible compounds composed of 13 or fewer heavy atoms (C, N, O, S, and Cl), according to simple stability and synthesizability criteria. Virtual combinatorial libraries can similarly reach or exceed the 10⁹ molecule mark, even with

as few as three or four points of substitution. In the limit, it is believed that as many as 10⁶⁰ molecules are potentially synthesizable.¹⁶ The combination of rapidly growing chemical libraries with computationally difficult similarity metrics suggests a need for dramatically faster methods of calculating chemical similarity.

In this article, we present SCISSORS [“SCISSORS Calculates Interpolated Shape Signatures over Rapid Overlay of Chemical Structures (ROCS) Space”], a mathematical technique that allows extremely rapid approximate evaluation of compound similarity for a variety of different similarity measures. In one application of interest, the calculation of an all-pairs similarity matrix over a large number of compounds, SCISSORS offers an asymptotic speedup in terms of the number of similarity calculations required from quadratic to linear scaling. We begin with a mathematical derivation of the SCISSORS algorithm and demonstrate its accuracy at Tanimoto estimation and generalizability to various similarity measures. In the discussion, we present links to a method from machine learning known as kernel principal components analysis (PCA) and show how these links justify optimality properties for SCISSORS and offer interesting insights into the nature of chemical shape space. Finally, we present a method by which SCISSORS allows asymptotic speedup (in terms of number of slow similarity evaluations) on the quadratic-scaling problem of similarity matrix construction.

METHODS

One commonality among otherwise diverse techniques to calculate chemical similarity is that, regardless of the underlying method, similarities are often returned as a Tanimoto score. While Tanimotos are often treated simply as real-valued similarities in the range [0,1], they have additional mathematical structure which can be profitably used; this is seen in the definition of the Tanimoto T_{AB} between vectors **A** and **B**:

$$T_{AB} \equiv \frac{\langle \mathbf{A}, \mathbf{B} \rangle}{\langle \mathbf{A}, \mathbf{A} \rangle + \langle \mathbf{B}, \mathbf{B} \rangle - \langle \mathbf{A}, \mathbf{B} \rangle} \quad (1)$$

* Corresponding author. E-mail: pande@stanford.edu.

[†] Department of Computer Science.

[‡] Department of Chemistry.

Equation 1 can be rearranged to obtain an expression for the inner product between two vectors in terms of their Tanimoto:

$$\langle \mathbf{A}, \mathbf{B} \rangle = \frac{T_{AB}}{1 + T_{AB}} (\langle \mathbf{A}, \mathbf{A} \rangle + \langle \mathbf{B}, \mathbf{B} \rangle) \quad (2)$$

Using this equation, the derivation of the SCISSORS method proceeds with two assumptions. We first assume that molecules can be represented as vectors in some (unspecified) inner product space. Given this assumption and a similarity measure on two molecules f and g , one can identify these molecules with vectors $\mathbf{A}$ and $\mathbf{B}$ in eq 2. It then makes sense to ask for the inner product between the two molecules. Since the Tanimoto between the two is known, we must somehow calculate the “magnitude” of each molecule (the $\langle \mathbf{A}, \mathbf{A} \rangle$ terms).

In the absence of further information, a parsimonious decision is to arbitrarily set all magnitudes equal to 1 (that is, $\langle \mathbf{A}, \mathbf{A} \rangle = \langle \mathbf{B}, \mathbf{B} \rangle = 1$); this is equivalent to assuming that all molecules should lie on the surface of a hypersphere of arbitrary dimension N . Alternatively, for certain measures we have more information and can assign meaningful values to the magnitude (see the Discussion Section for an example). For the purpose of the derivation, we take the magnitudes as given (either from parsimony or knowledge of the measure). This implies that it is possible to calculate the inner product between two molecules given their pairwise Tanimoto score.

We further assume that the inner product space in question is $\mathcal{R}^N$, the space of vectors with N real-number coordinates. Under this assumption, it becomes meaningful to try to map molecules to vectors in $\mathcal{R}^N$ in such a way as to preserve their inner products. Given a “basis” set of molecules $\mathbf{B} = \mathbf{B}_1, \dots, \mathbf{B}_k, k \geq N$, we require that similarities be computed between all pairs in $\mathbf{B}$. These similarities are then transformed to inner products (using a method appropriate for the measure in question) and placed into a matrix $\mathbf{G}$ such that G_{ij} contains the inner product between molecules $\mathbf{B}_i$ and $\mathbf{B}_j$ (a matrix of this form is known in the statistical literature as a “Gram matrix”). We wish to find a factorization of $\mathbf{G}$ of the form $\mathbf{G} = \mathbf{M}\mathbf{M}^T$, where $\mathbf{M}$ is an $N \times N$ matrix containing the molecule vectors along its rows.

Since the elements of $\mathbf{G}$ are constructed from molecular similarities, $\mathbf{G}$ is real and symmetric and, therefore, has a spectral decomposition into eigenvectors $\mathbf{V}$ and eigenvalues (in diagonal matrix $\mathbf{D}$). From this, the desired factorization follows by simple algebraic manipulation:

$$\begin{aligned} \mathbf{G} &= \mathbf{M}\mathbf{M}^T \\ &= \mathbf{V}\mathbf{D}\mathbf{V}^T \\ &= \mathbf{V}\mathbf{D}^{1/2}\mathbf{D}^{1/2}\mathbf{V}^T \\ &= (\mathbf{V}\mathbf{D}^{1/2})(\mathbf{V}\mathbf{D}^{1/2})^T \\ \mathbf{M} &= \mathbf{V}\mathbf{D}^{1/2} \end{aligned} \quad (3)$$

The method described in eq 3 is also used to derive vector coordinates from interpoint distances in classical multidimensional scaling (MDS).¹⁷ It has several properties which make it an appealing choice relative to other matrix factorizations (such as the Cholesky decomposition, which also finds matrix $\mathbf{M}$ such that $\mathbf{G} = \mathbf{M}\mathbf{M}^T$ ¹⁸). First, the MDS reconstruction of molecule vectors in N dimensions is

guaranteed to most closely approximate the given inner products, in a least-squares sense. Furthermore, the reconstruction dimension can easily be chosen by sorting the eigenvalues in order of decreasing value and setting to zero all eigenvalues after the first K (for a K -dimensional reconstruction). Our method is also very closely related to a technique known as kernel PCA;^{19,20} this relationship has interesting consequences explored in the Discussion Section.

One mathematical detail skipped in the derivation of eq 3 relates to the square root performed on the eigenvalues. If the measure described by the inner products in $\mathbf{G}$ is a Euclidean metric, then $\mathbf{G}$ will be positive semidefinite (i.e., it will have only zero- or positive-valued eigenvalues),¹⁷ and the square root will exist. However, there is no guarantee that a generic molecular similarity measure will be a Euclidean metric (and several popular measures, including ROCS, are not). Because of the optimality properties of MDS, however, the impact that these negative eigenvalues have on the approximation depends on their magnitude relative to the positive eigenvalues. If no negative eigenvalue has magnitude within the top K , then it would not be used in the K -dimensional reconstruction and does not affect the solution. We demonstrate in our results that for many popular measures, negative eigenvalues are of sufficiently small magnitude that they do not affect the reconstruction and that these non-Euclidean measures can, therefore, be well-approximated by Euclidean inner products. Furthermore, in the next section, we describe a correction factor that can be added to the basic MDS formula to allow interpolation of coordinates associated with negative eigenvalues.

Correction Factors. One drawback of our inference procedure is that the optimization procedure (eq 3) does not respect the vector magnitude assumptions made at the beginning in the first step. In particular, given any molecule $\mathbf{A}$ in the basis set with square magnitude $\langle \mathbf{A}, \mathbf{A} \rangle = x$ and its low-dimensional vector approximation $\hat{\mathbf{A}}$, the approximated square magnitude $\langle \hat{\mathbf{A}}, \hat{\mathbf{A}} \rangle$ will always be smaller than x , because we have ignored the magnitude associated with higher (less significant) dimensions.

As will be shown in the Results Section (Figure 4), this loss of magnitude can cause problems with the absolute accuracy of resulting Tanimoto values (though rank orderings are often preserved). One possible strategy for solving this problem would be to replace eq 3 with a constrained optimization method; however, the simplicity and the ubiquity of implementation of a method based on the spectral decomposition suggest that a procedure to correct the results of the eigenvector decomposition would be useful. Consequently, we choose not to examine constrained optimization methods, although this may be an interesting direction for future work.

A simple correction technique is to simply scale each result vector to force its magnitude to 1. Our testing showed this to be ineffective (data not shown). However, one correction technique, which we call the slack coordinate method, did improve accuracy and furthermore allows us to represent non-Euclidean geometries.

Slack Coordinate Method. Without loss of generality, we consider only the case in which we wish to correct all vector magnitudes to unity. Given real vectors $\mathbf{A}$ and $\mathbf{B}$ of arbitrary dimension n (i.e., $\mathbf{A}, \mathbf{B} \in \mathcal{R}^n$) with magnitude less than 1, we construct vectors $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{B}}$ of unity magnitude by

augmenting them with one additional “slack” coordinate, as follows:

$$\begin{aligned}
 \tilde{\mathbf{A}} &= [a_0 a_1, \dots, a_{n-1} a_n]^T \\
 \tilde{\mathbf{B}} &= [b_0 b_1, \dots, b_{n-1} b_n]^T \\
 \|\tilde{\mathbf{A}}\| &= \tilde{\mathbf{A}}^T \tilde{\mathbf{A}} = 1 \\
 \sum_{i=0}^n a_i^2 &= 1 \\
 a_n^2 &= 1 - \sum_{i=0}^{n-1} a_i^2 \\
 a_n &= \pm \sqrt{1 - \sum_{i=0}^{n-1} a_i^2} \\
 b_n &= \pm \sqrt{1 - \sum_{i=0}^{n-1} b_i^2}
 \end{aligned} \quad (4)$$

An interesting aspect of this correction is that the slack coordinate is incompletely specified, as the sign is indeterminate. It is possible to derive an analytic correction for the resulting Tanimoto. For convenience, we define variables k_1 and k_2 , corresponding to the numerator and denominator, respectively, of the uncorrected Tanimoto equation:

$$T_{AB} = \frac{\sum_{i=0}^{n-1} a_i b_i}{\sum_{i=0}^{n-1} a_i^2 + \sum_{i=0}^{n-1} b_i^2 - \sum_{i=0}^{n-1} a_i b_i} = \frac{k_1}{k_2} \quad (5)$$

We can then define the corrected Tanimoto in terms of these variables, and simple algebraic transformations allow us to express the correction as a function of the uncorrected Tanimoto and the (uncorrected) magnitudes of $\mathbf{A}$ and $\mathbf{B}$:

$$\begin{aligned}
 \|\mathbf{A}\|^2 &= \sum_{i=0}^{n-1} a_i^2 \\
 \|\mathbf{B}\|^2 &= \sum_{i=0}^{n-1} b_i^2 \\
 T_{\tilde{\mathbf{A}}\tilde{\mathbf{B}}} &= \frac{\sum_{i=0}^n a_i b_i}{\sum_{i=0}^n a_i^2 \sum_{i=0}^n b_i^2 - \sum_{i=0}^n a_i b_i} \\
 T_{\tilde{\mathbf{A}}\tilde{\mathbf{B}}} &= \frac{k_1 + a_n b_n}{k_2 + a_n^2 + b_n^2 - a_n b_n} \\
 T_{\tilde{\mathbf{A}}\tilde{\mathbf{B}}} &= \frac{k_1 \pm \sqrt{1 - \|\mathbf{A}\|^2} \sqrt{1 - \|\mathbf{B}\|^2}}{(2 - k_1) \mp \sqrt{1 - \|\mathbf{A}\|^2} \sqrt{1 - \|\mathbf{B}\|^2}} \\
 k_1 &= \frac{T_{AB}(\|\mathbf{A}\|^2 + \|\mathbf{B}\|^2)}{(1 + T_{AB})} \\
 k_3 &= \pm \sqrt{1 - \|\mathbf{A}\|^2} \sqrt{1 - \|\mathbf{B}\|^2} \\
 k_4 &= k_1 + k_3 \\
 T_{\tilde{\mathbf{A}}\tilde{\mathbf{B}}} &= \frac{k_4}{2 - k_4}
 \end{aligned} \quad (6)$$

Although the sign of the slack coordinate ($\pm\sqrt{1 - \|\mathbf{A}\|^2}$ or $\pm\sqrt{1 - \|\mathbf{B}\|^2}$ for molecules $\mathbf{A}$ and $\mathbf{B}$, respectively) is unspecified, the analytic correction term k_3 is not directly dependent on the particular signs of the slack coordinates. Rather, if $\mathbf{A}$ and $\mathbf{B}$ are real vectors, k_3 has sign determined by whether $\mathbf{A}$ and $\mathbf{B}$ have the same or opposite signs in their slack coordinate. The case in which k_3 is always negative (that is, all pairs of vectors $\mathbf{A}$ and $\mathbf{B}$ have opposite signs in their slack coordinates) turns out to be especially useful. However, this case is not realizable if we restrict our vectors $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{B}}$ to lie in $\mathcal{R}^{n+1}$. As a demonstration, consider three vectors $\tilde{\mathbf{A}}$, $\tilde{\mathbf{B}}$, and $\tilde{\mathbf{C}}$ with slack coordinates a_n , b_n , and c_n . Without loss of generality, let a_n be negative and b_n be positive. For c_n to always have an opposite sign to any molecule it was compared to, it would have to be positive when molecule $\mathbf{C}$ is compared with molecule $\mathbf{A}$ and negative when compared to $\mathbf{B}$; this is a contradiction.

Although this style of correction term cannot be implemented with real vectors, by generalizing the vector space, this result is easily understood. The case in which every pair of vectors is associated with a negative correction term k_3 corresponds to vectors lying in $\mathcal{R}^n \times I$, where the slack coordinate is always positive in sign and purely imaginary (e.g., in eq 4, let $a_n = i\sqrt{1 - \sum_{i=0}^{n-1} a_i^2}$, with $i = \sqrt{-1}$). Strictly, this interpretation requires that eq 4 be modified so that the Hermitian transpose is used but does not affect the rest of the derivation. The geometry implied by this vector space is known as pseudo-Euclidean, in contrast with the Euclidean geometry of $\mathcal{R}^{n+1}$, and has the capability to represent curved spaces. In particular, imaginary coordinates of this form can be used to represent the curved dimensions induced by negative eigenvalues in the Gram matrix.²¹ The slack coordinate correction, by generalizing from real to complex vectors, allows approximation of non-Euclidean distance measures.

Fast Approximation of New Vectors. The above procedure for deriving molecule vectors requires computation of $O(k^2)$ molecular similarities. It, therefore, does not scale to computing vectors over large libraries of millions of compounds. However, given a fixed “basis” set of molecules, it is possible to approximate vectors for new molecules in a fixed amount of time for each new molecule by least-squares. Specifically, given a “library” molecule (that is, one not originally in the basis set) L_i and the matrix $\mathbf{M}$ of basis vectors, a vector for L_i can be constructed by the following procedure:

1. Compute similarities between L_i and all basis molecules B_j .
2. Transform similarities to inner products G_{ij} .
3. Store inner products in a vector $\mathbf{T} = [G_{i1}, G_{i2}, \dots, G_{ik}]^T$.
4. Solve equation $\mathbf{M}\tilde{\mathbf{x}} = \mathbf{T}$ by linear least-squares.

The resulting vector $\tilde{\mathbf{x}}$ contains the vector that best approximates the given inner products, in a least-squares sense, and therefore is the molecule vector for molecule L_i . This procedure allows the approximation of vectors for a large library of molecules in time linear in the size of the library. It is, furthermore, embarrassingly parallel once the basis set has been computed, making it a computationally attractive method for operation on large libraries.

RESULTS

Full details of the testing methods used in this work, including exact molecule counts, software versions used, and software settings, are included in the Appendix.

Similarity Measures. To demonstrate the generality of the SCISSORS technique, we have tested it against similarity measures of various types, as implemented in the OpenEye Scientific Software (Santa Fe, NM) oechem, oeshape, and oect toolkits. Representing substructure-fingerprint-based methods is LINGO, a similarity measure which counts the number of substrings in common between the SMILES representations of a pair of molecules.² ROCS,¹⁰ a three-dimensional (3D) shape comparison method which finds the maximal volume overlap between Gaussian representations of molecule shape, represents 3D shape methods. A variation on shape-based ROCS, known as the ROCS Color Tanimoto, uses a similar volume overlap optimization on virtual “color” atoms representing chemical functionalities. We use this color ROCS to demonstrate the generalizability of SCISSORS to 3D methods that include distinct atom types, not all of which interact with one another. Finally, ZAP,⁹ a Poisson–Boltzmann solver which calculates and compares electrostatic fields around molecules, represents electrostatic similarity methods.

Molecules Tested. We draw molecules for testing from a diverse set of collections of drug and drug-like molecules. Most of our tests are performed on a 57 253 molecule database, hereafter named “Maybridge+BBP”, constructed from the union of the Maybridge Screening Collection and a blood-brain barrier penetration data set.²² For Maybridge+BBP, 3D conformers were generated using OpenEye’s OMEGA software. When calculating Tanimotos using ROCS, multiconformer (200 conformer maximum) fit molecules were aligned against single-conformer references, and the maximum Tanimoto over any pair of conformations for the molecule pair was used. Detailed configuration parameters passed to ROCS and OMEGA are included in the Appendix.

In the test for generalizability of SCISSORS, we used the 1665 molecules tested in PubChem Assay #677, an assay for novel allosteric modulators of the M1 muscarinic receptor. To test the sensitivity of the SCISSORS technique against the source of molecules for the basis set, we draw basis molecules from the Maybridge+BBP set and also from the Asinex Gold Collection and the Maybridge Fragment Library, commercial libraries used for compound and fragment screening. Finally, the virtual screening test draws its molecules from the Directory of Useful Decoys, version 2,²³ a computational drug-discovery data set specifically designed to confound methods based on common physical parameters, such as molecular weight and cLogP. For all 3D similarity measures other than the 57 253 molecule Maybridge+BBP set, molecules were assigned single conformers by OpenEye’s OMEGA software. Charges for all molecules were assigned using the AM1-BCC charge model implemented in OpenEye’s oequacpac toolkit.

Differences in setup (i.e., multiconformer vs single conformer) between data sets are due to the fact that we made use of existing databases already in use for other drug discovery projects. These differences are not significant to the content of the paper. Insofar as the goal of SCISSORS is to reproduce the results of existing similarity measures,

we demonstrate comparable performance to the source similarity measure regardless of how the molecules were prepared for the source measures.

Basis Size and Dimensionality. To evaluate SCISSORS’ ability to approximate molecular similarities on data sets of varying size, we used SCISSORS to approximate the ROCS shape Tanimoto for test sets of 19 000 and 50 000 molecules drawn at random from Maybridge+BBP. This was performed multiple times over a large number of basis sizes and dimensionalities (in all cases, the 19 000 or 50 000 library molecules were disjoint from the basis molecules). Tanimotos were computed using multiconformer query molecules matched against single-conformer references; the Tanimoto of the best-matching query conformer against the reference was used as the Tanimoto for the molecule pair. Detailed settings for ROCS are provided in the Appendix. As a measure of performance, we calculated the root-mean-square (RMS) deviation between the SCISSORS and ROCS Tanimoto over all pairings of these 19 000 molecules (Figure 1; $19\,000 \times 19\,000 = 361$ million Tanimotos) or 50 000 molecules (Figure 2; $50\,000 \times 50\,000 = 2500$ million Tanimotos).

Two trends are evident from the data. First, the error in the SCISSORS approximation of the shape Tanimoto falls rapidly at first with the addition of more dimensions and then levels off around 75–100 dimensions. Most of the error falloff occurs in the first few dimensions, with gains beyond 30 dimensions coming much more slowly. A similar trend also appears for basis size; a large error is present for small basis sizes (especially in high-dimensional approximations), as a consequence of overfitting the vector model, but this falls off rapidly as the basis size increases. This result is stable even as we scale to a data set more than double the size. The test set comprising 50 000 molecules required neither more eigenvectors nor a larger basis set to achieve root-mean-square error (RMSE) comparable to that found on the 19 000 molecule test set. Ultimately, SCISSORS is able to approximate the ROCS shape Tanimoto to within approximately 0.08 RMSE around 75–100 dimensions and a basis size of 500–1000, for both the 19 000 and 50 000 molecule test sets.

Basis Selection Strategies. To test the sensitivity of the SCISSORS technique to how the basis set is chosen, we compared randomly chosen basis sets with basis sets chosen such that basis molecules corresponded to “representative” shapes in the data set. To choose representative molecules, we clustered the Maybridge set using the dbclust algorithm¹¹ at several Tanimoto cutoff values; the cluster centers chosen were used as the representative shapes. For each measured basis size N , we ran SCISSORS in 75 dimensions using a basis consisting of N randomly chosen molecules or the N cluster centers with largest number of neighbors (ties broken at random). The resulting library RMSEs are plotted in Figure 3.

For any basis set larger than 200 molecules, it was possible to find a SCISSORS approximation with lower RMS error than that resulting from the random basis; however, in all cases such error reduction was extremely small, on the order of 0.005 Tanimoto units. Interestingly, the optimal clustering Tanimoto value differed by basis size. The optimal Tanimoto (that is, the cluster set with lowest RMSE) increased with increasing basis size, suggesting that, as more basis molecules

Figure 1. RMS Tanimoto error versus basis size and dimensionality of approximation for standard SCISSORS technique on ROCS Shape Tanimoto. The 19 000 library molecules (361 million total Tanimotos) used to calculate root-mean-square error (RMSE). Each bar averaged over at least 10 random bases of given size.

are added, it is somewhat advantageous to select tighter clusters as basis elements.

Corrected Tanimotos. In Figure 4 we examine the effect of the slack coordinate correction term introduced earlier. Because the purpose of the correction term is to account for eigenspectrum density truncated in a low-dimensional approximation, we examine two cases, 10- and 75-dimensional approximations. Figure 4 is a density plot of SCISSORS standard and corrected Tanimotos versus ROCS Tanimotos using the 19 000 molecule test subset of Maybridge+BBP also used in Figure 1. In Figure 4, density contours that are symmetric about $Y = X$ indicate that the plotted Tanimoto values have an unbiased error with respect to the original Tanimoto value, which aids in the preservation of rank ordering.

As expected from the earlier results, in 75 dimensions, the standard SCISSORS Tanimoto values are tightly clustered along the $Y = X$ line, and the contours of the density plot continue to be aligned along this $Y = X$ axis even out to low densities, indicating an overall preservation of rank ordering. In contrast, in 10 dimensions, the SCISSORS Tanimotos, as expected, almost always overestimate the true Tanimoto. Furthermore, the contours are no longer ellipses symmetric about $Y = X$ (or a parallel line), which indicates poorer rank-ordering performance.

As expected, the slack coordinate correction leaves the high-dimensional case largely unaffected, with the major exception of a large group of errors at ROCS Tanimoto = 1.0. These are caused by self-comparisons; the slack coordinate method does not respect the constraint that the Tanimoto of a vector against itself ought to be 1. However, in the low-dimensional case, the slack correction makes a significant difference. Leaving aside the known artifact at ROCS Tanimoto = 1.0, the overall shape of the density contours is shifted below the diagonal but is symmetric. This indicates that, with the exception of a constant shift, the slack-corrected values are much more accurate than the uncorrected low-dimensional values.

Virtual Screening. To illustrate the utility of SCISSORS-derived Tanimoto values in virtual screening, we used SCISSORS on the ROCS-shape Tanimoto to do a virtual screen on the Directory of Useful Decoys (DUD) test set. DUD is a collection of “systems” consisting of target proteins and associated small molecules. For each system, a certain number of molecules known to bind the target are designated as “ligands”; the rest of the molecules (36 times the number of ligands) are “decoys” with similar physical properties believed to not bind the target. We used a previously published protocol²⁴ to evaluate performance and statistical significance; for clarity, we summarize the protocol below.

Figure 2. RMS Tanimoto error versus basis size and dimensionality of approximation for standard SCISSORS technique on ROCS Shape Tanimoto. The 50 000 library molecules (2500 million total Tanimotos) used to calculate RMSE. Each bar averaged over at least 10 random bases of given size.

For each protein system in DUD, we used ROCS and SCISSORS (with three different basis sets) to calculate the Tanimoto of every ligand molecule against every other molecule in the system. For each ligand molecule, we ranked all other molecules in order of decreasing Tanimoto similarity, modeling a screening experiment in which given one active compound, one wishes to find other actives in a pool of inactive molecules. All molecules used were single conformer; settings for ROCS are given in the Appendix. All SCISSORS Tanimoto values were computed in 10 dimensions as uncorrected SCISSORS Tanimotos. Basis sets were drawn from the Maybridge Screening Collection, the Asinex Screening Collection, and the Maybridge Fragment Library to illustrate the effects of different basis choices on performance.

The ROC area under curve (AUC) for this ranking was calculated according to the method in Clark.²⁵ Ligand molecules were considered true positives; false positives were any decoys ranked higher than a true ligand. A bootstrapping procedure²⁴ was used to estimate the 68 and 95% confidence intervals on the calculated AUC values. These AUCs and confidence intervals are illustrated in Figure 5.

For almost all systems, the performance of the different SCISSORS basis sets are statistically indistinguishable from

each other and from that of the original ROCS Tanimoto. In particular, the Maybridge and Asinex basis sets, which are drawn from different libraries but which likely have overall similar characteristics (as both are screening libraries) perform almost identically. The Maybridge Fragment basis is an outlier on several systems; it is likely that this is because these molecules are significantly smaller (in volume/mass) than the drug-like molecules in DUD, indicating that at least rudimentary physical property matching is a prudent step in the use of SCISSORS. It is interesting, however, that there were so few outliers using the fragment basis set, as similarity measures often fail when comparing fragments to elaborated molecules; we have not explored the reasons for this performance in detail.

Similarity Measure Generalizability. To assess the generalizability of SCISSORS to chemical similarity measures other than the ROCS shape Tanimoto, we evaluated the performance of SCISSORS on four similarity measures: ROCS shape and color, LINGO, and ZAP's ET, or electrostatic, Tanimoto. All Tanimotos were originally evaluated on a collection of 1665 small molecules from a screen for allosteric modulators of the M1 muscarinic receptor; for SCISSORS, both basis and library molecules were drawn from this same pool. All tests were conducted with 1000

Figure 3. RMS Tanimoto errors for SCISSORS approximations to ROCS shape Tanimoto constructed from randomly selected basis sets and basis sets chosen by molecule clustering. All Tanimotos evaluated over 361 million molecule pairs.

Figure 4. Density plots of SCISSORS and slack-corrected SCISSORS Tanimoto approximations of ROCS Shape Tanimoto for 361 million molecule pairs from Maybridge. Contours are labeled as $\log_{10}$ of contour height. Light gray $y = x$ line drawn for reference.

Figure 5. ROC AUC values on each system of the DUD test set. Reported are the mean AUC averaged over each ligand in the system (line), the 68% confidence interval on the AUC (box), and the 95% confidence interval on the AUC (whiskers). Within each system, results are reported for ROCS (G, gray) and SCISSORS using basis sets from the Maybridge Screening Collection (M, red), the Asinex Screening Collection (A, orange), and the Maybridge Fragment Library (F, blue). All SCISSORS approximations were done in 10 dimensions; Maybridge and Asinex sets used a 500 molecule basis set; and the Maybridge Fragment set was 473 molecules, the size of the entire library.

library molecules, with RMSE evaluated over an all-pairs ($1000 \times 1000 = 1$ million Tanimotos) similarity comparison. For the electrostatic Tanimoto, molecules were first overlaid using ROCS before electrostatic field calculations.

Figure 6 plots the SCISSORS Tanimoto RMSE for each similarity measure in the large basis (600 molecules) limit as a function of dimensionality. Three trends are particularly notable. First, on this data set, all the similarity measures show nonmonotonic error behavior as the number of dimensions increases, most likely reflecting eventual overfitting (an insufficient number of basis molecules for further dimensions). Second, SCISSORS generalizes well to approximating both the ROCS color and the LINGO Tanimotos, with RMSEs under 0.1 in both cases. Finally, ZAP's electrostatic Tanimoto appears to be inapproximable by the SCISSORS method.

Figure 7 explores why only certain measures are approximable by SCISSORS. Each subplot plots the value of the top 300 eigenvalues from a basis set of size 1000 computed for each of the four similarity measures, on a log scale. For all three well-approximated measures, the magnitude of the 100th eigenvalue has fallen off more than 2 log units relative to the first (largest) one, indicating that high dimensions rapidly become less important to the approximation. In contrast, the ET Tanimoto's eigenspectrum falls off much more slowly; even after 300 dimensions, it has not fallen 2 log units. It is this characteristic of the eigenspectrum (insufficiently rapid falloff) that makes a low-dimensional linear approximation to the ET Tanimoto impossible.

Tanimoto Computation Speed. The precalculation step of SCISSORS is extremely efficient compared to the evaluation of the source similarity measure. We benchmarked the

DISCUSSION

Figure 6. RMS Tanimoto error from SCISSORS approximations to several molecular similarity measures, as a function of dimensionality of approximation. All tests conducted with a basis size of 600 molecules and a test set size of 1000 molecules.

precalculation step for calculating molecule vectors on a 20 480 molecule subset of Maybridge+BBP, using a 500 molecule basis set drawn from the same database. Timings for each phase of SCISSORS precalculation on one core of a 3 GHz Intel Xeon-based (Core 2 architecture) Mac Pro are presented in Table 1 (further machine details are presented in the Appendix). The precalculation time is dominated by ROCS calculations; the spectral decomposition and least-squares inference for SCISSORS take less than 5 s, compared to nearly 11 000 s for the basis versus basis and library versus basis ROCS computations.

A significant benefit of SCISSORS is that it calculates accurate estimates of molecular similarity extremely rapidly once library vectors have been inferred. This is particularly true in the case of the ROCS shape and color similarities, which are industrially important but expensive to calculate. We measured the OpenEye ROCS implementation's performance at approximately 1000 comparisons/s per core on a 3 GHz Mac Pro, while performing the basis–basis and library–basis ROCS comparisons for the above benchmark. In contrast, our C code to evaluate SCISSORS Tanimotos (compiled with gcc 4.0.1 and GotoBLAS2)²⁶ is capable of computing approximately 50 million similarities/s (in 64 dimensions) on one core of the same CPU. With precomputed library vectors, then, SCISSORS is around 50 000 times faster than a direct calculation of Gaussian overlap.

Indeed, SCISSORS is fast enough that for very large similarity matrix problems, it can be faster to recalculate similarities on the fly from SCISSORS vectors than to precalculate all similarities and read them back from disk. This leads to a significant savings in storage, as the storage required for the molecule vectors scales as $O(N)$, while storage for a similarity matrix scales as $O(N^2)$. The combination of high computation speed and reduced storage makes similarity calculations (for large scale virtual screening and data analysis) computationally feasible for data sets with N on the order of 10^5 to 10^7 . While SCISSORS does not solve other issues involved with large data sets (e.g., file handling and molecule preparation), it is able to reduce both computational and total data storage burdens for large-scale problems. In terms of CPU time alone, it can make difficult problems easy and impossible problems difficult.

Connections to Previous Work. *Previous Methods to Characterize Chemical Space.* Several methods have been applied in past work to characterize chemical space by learning relevant dimensions along which chemicals vary to find vector embeddings of molecules. In this section we discuss past work which has used multidimensional scaling and other similar techniques to find vector embeddings of molecules based on particular similarity representations.

Multidimensional scaling has previously been applied to chemical data sets in order to reduce the dimensionality of binary molecular fingerprints.^{27,28} This approach has been used for combinatorial library design; by computing coordinates for each molecule and treating each dimension as a property, it is possible to maximize the distance among molecules to select a maximally diverse library.²⁸ However, using MDS directly in this manner is not scalable to large compound sets, as it requires a full N^2 similarity matrix.

There have also been prior efforts to develop scalable versions of MDS-based methods. One method, due to Nicholls,²⁹ focuses on similarity *metrics*; that is, on similarity measures that compute a distance between pairs of molecules, such that these distances obey the triangle inequality. The Nicholls method, after computing coordinates for a basis set of molecules using MDS, computes coordinates for a new nonbasis molecule by triangulation based on distances to each of the basis molecules.

The principal difference between SCISSORS and these older MDS-based methods is our use of least-squares approximation to calculate coordinates for nonbasis molecules. Least squares allows SCISSORS to scale well to large data sets, unlike the original methods using MDS for dimensionality reduction. Furthermore, because we approximate inner products rather than distances and use least-squares rather than triangulation to calculate coordinates, we do not assume that the approximated similarity measures are true metrics. In particular, the use of least-squares allows SCISSORS to calculate approximate coordinates even in the nonmetric source similarity case (in which the triangulation equations would have no consistent solution). The SCISSORS similarity is a metric (or pseudometric, in the case of the slack-corrected Tanimoto) approximation to the source measure.

Raghavendra and Maggiora³⁰ developed a symmetric-orthogonalization method (similar in some respects to SCISSORS) to calculate vector embeddings of molecules in an arbitrary similarity space. Like SCISSORS, the Raghavendra and Maggiora (RM) method calculates a pairwise similarity matrix on a small set of basis molecules, using this to find vectors for the basis, and then projects library molecule vectors onto this space by transforming library–basis similarity values. It is presented only as a method to describe chemical space and not as a method to approximate similarities. Several mathematical limitations impact the applicability of this method for both purposes.

One critical difference is in the treatment of the basis similarity matrix, the entries of which are Tanimotos. The RM method treats these entries as though they were inner products; however, as eq 1 shows, the Tanimoto is a function of an inner product but not an inner product itself. This inaccuracy will distort vector spaces inferred by the method.

Figure 7. Eigenvalue spectra for Gram matrices resulting from four similarity measures on 1000 molecules from the M1 muscarinic data set.

Table 1. SCISSORS Precalculation Step Timings for 20 480 Library Molecules^a

step	time (sec)
basis matrix ROCS calculation	262
basis vector calculation (eigendecomposition)	0.33
library vs basis ROCS calculation	10 732
library vector calculation (least-squares)	4

^a Using a 500 molecule basis set and one core of a 3 GHz Intel Xeon-based Mac Pro.

Furthermore, this method is unable to handle molecular similarities which do not result in positive definite Gram matrices (i.e., non-Euclidean measures). SCISSORS handles these in two ways. First, the use of the slack correction allows approximation of dimensions associated with the negative eigenvalues in a nonpositive definite similarity measure. Second, a typical use of SCISSORS will retain fewer eigenvalues/eigenvectors (SCISSORS dimensions) than can be computed from the matrix. This allows a positive definite approximation of nonpositive definite measures. It also improves the statistical reliability of our vector approximations. The RM method projects library molecules onto the full number of dimensions computed from the basis matrix (i.e., number of dimensions = number of basis molecules). However, as shown in Figures 1 and 2, fitting to a number

of dimensions similar to the number of molecules leads to large approximation errors from overfitting. Thus, we argue that the SCISSORS method, more faithful to the mathematics behind similarity coefficients, is more flexible by allowing approximation of non-Euclidean measures, and is more accurate because it adds the ability to avoid overfitting.

Connections to Kernel PCA. Kernel principal components analysis (kernel PCA) is a nonlinear variant of PCA that, given data in some input vector space, finds the principal components of the data in a different vector space known as the feature space.^{19,20} The dimensionality of this feature space is often very high (possibly infinite). The computation is achieved by the use of kernel functions, which map a pair of points in the input space to a real number that is interpretable as the inner product of the two points in the feature space, without requiring an explicit expansion of the points into the feature space.

In the Methods, we derived the SCISSORS algorithm from classical multidimensional scaling. The method used in SCISSORS to derive basis vectors is also closely related to kernel PCA (which is itself closely linked to multidimensional scaling).³¹ In particular, by transforming chemical similarity Tanimoto values to inner products, we treat transformed chemical similarity functions as kernels that map molecules from some hypothetical “molecule space”

into a feature space induced by the particular similarity function. Because the mathematical methods underlying SCISSORS are essentially identical to those of kernel PCA, several optimality properties of kernel PCA²⁰ carry over to SCISSORS. As was previously argued using the derivation from multidimensional scaling, a SCISSORS/kernel PCA expansion of a molecule in k dimensions of a chemical feature space has a minimum mean-squared approximation error, compared to any other choice of k directions (e.g., as might be derived from a truncation of a Cholesky decomposition).

Furthermore, the k directions calculated from kernel PCA have minimal representation entropy. As a consequence, no other orthogonal basis transformation can reproduce the given inner product matrix with the same accuracy in fewer bits. This property provides a rigorous bound on the accuracy of vectors calculated by SCISSORS; no other method evaluating chemical similarity inner products by dot products between real vectors can achieve a better coverage of the variance in the training (basis) matrix with fewer bits than used by SCISSORS.

There are several key differences between SCISSORS and kernel PCA related to the treatment of the Gram matrix. First, following multidimensional scaling, we do not rescale the eigenvectors obtained in the diagonalization of the Gram matrix. Additionally, most commonly used chemical similarity measures, when transformed from Tanimoto to inner products, do not induce the type of kernel usually used with a kernel PCA approach (known as a Mercer kernel), as their Gram matrices usually have negative eigenvalues. More importantly, prior to diagonalizing the Gram matrix, we do not center the data to have zero mean. As a consequence, the first eigenvector reflects the mean value in the feature space.³² As will be shown in the next section, this is key to partial interpretability of coordinates derived from SCISSORS.

Interpretation of Chemical Similarities. Besides its use in constructing vector approximations, the SCISSORS inference procedure also gives insight into the nature of the chemical space induced by a given similarity measure. In particular, for well-approximated measures for which a molecule–molecule inner product can be defined, we can interpret the meaning of the first vector coordinate and gain an understanding of the measure. Our example here is the Gaussian volume overlap measure used by ROCS.

ROCS defines its “shape Tanimoto” as eq 7, in which O_{xy} represents the value of the overlap volume between molecules x and y , which have been rotated and translated to maximize their volume overlap.

$$T_{AB} = \frac{O_{AB}}{O_{AA} + O_{BB} - O_{AB}} \quad (7)$$

This equation has clear similarity to the conventional vector Tanimoto (eq 1); by analogy, for ROCS we can treat the maximum overlap volume between a pair of molecules as the inner product between these two molecules. Recall that for a similarity measure to be well-approximated by SCISSORS, we require that the eigenvalue spectrum of the similarity measure falls off sufficiently rapidly, such that each successive dimension is significantly less important than the one preceding it.

Figure 8. Interpretation of SCISSORS coordinates. All data from 128 371 Tanimotos computed on DUD test set; 500 random molecules from Maybridge used as basis.

Consider such a measure on molecules A and B and their vector approximations $\hat{A} = [a_0, \dots, a_n]$ and $\hat{B} = [b_0, \dots, b_n]$. Because SCISSORS is able to approximate the measure well, it is true that $\langle A, B \rangle \approx \langle \hat{A}, \hat{B} \rangle$. The assumption that the measure is approximated well implies (by the discussion in section on generalizability) that the eigenvalue spectrum falls off rapidly. This in turn means that the first coordinate will be the dominant component of the inner product (since further coordinates are decreasingly important), or formally, that $\langle A, B \rangle \approx \langle \hat{A}, \hat{B} \rangle \approx a_0 b_0$. In the special case in which $A = B$ (the self-similarity of molecule A), $\langle A, A \rangle \approx a_0^2$.

For the ROCS similarity measure, this implies that the first coordinate of the vector for a given molecule, and therefore the first principal coordinate of the shape space induced by the measure, is approximately equal to the square root of the molecular volume. This is illustrated in Figure 8a, which plots the first coordinate of many molecule vectors inferred by SCISSORS on the ROCS shape overlap quantity against their molecular volumes. The lines show that the data are well-fit by a square-root function multiplied by a constant that represents the relative importance of molecular volume to the whole similarity measure. This demonstrates that shape-based virtual screening is largely controlled by differences in molecular volume, which correlates very closely with molecular weight (Figure 8b).

As explained by analogy to kernel PCA, this first coordinate (which for ROCS is approximately equal to the

square root of the molecular volume) reflects the mean location of the molecules in the feature space learned from the similarity measure. We do not know whether a similar chemically interpretable understanding of the second and later eigenvectors (projections onto which are the molecular coordinates) exists. It has been demonstrated³³ that the eigenvectors calculated by kernel PCA converge, in the large-basis limit, to eigenfunctions of a linear operator on the underlying kernel. These in a sense reflect a functional basis for the kernel (or similarity measure). However, in general, coordinates in a kernel feature space as calculated by kernel PCA are not required to have pre-images in the input space;²⁰ making matters worse, in the SCISSORS case (similarity kernels operating on molecules), it is not even clear what the input space is! Therefore, we do not have a mathematical derivation for the chemical meaning behind particular coordinates. While it is possible that empirical examination of large molecule sets and their SCISSORS coordinates may uncover such interpretations, interpretability is not crucial to our target application (fast similarity estimation), and we have not carried out such an analysis.

Insights into Chemical Space. An interesting result of the SCISSORS analysis is that relatively few basis molecules (on the order of 500–1000) were sufficient to achieve accurate estimation of shape Tanimotos for compounds in Maybridge and DUD and that the source of these molecules (Maybridge, Asinex, or Maybridge Fragment) made little difference for virtual screening. This finding is paralleled by previous results. Haigh et al. demonstrated that a reference shape set of 3119 molecules was sufficient in their “shape fingerprints” approach to reproduce shape Tanimotos and that these reference shapes were transferable across databases.¹³ Fontaine et al. found that a database of 2458 reference shapes was sufficient to cover the shape space of a one million compound single-conformer data set in their “alignment recycling” approach.³⁴ Thus, our results corroborate past data showing that, although chemical space is very large, the shape space of relevant drug-like molecules may be fairly small, with only a few hundred to a few thousand basis shapes providing adequate coverage of the shape space.

Applications: Fast Library Screening and Clustering with SCISSORS. The combination of learning vectors on a small basis set with fitting of new vectors in linear time using least-squares makes possible highly accelerated library screening. The target scenario for such an application involves screening based on a similarity measure that is expensive to calculate relative to the calculation of a low-dimensional (order of 10 to 100) dot product. The SCISSORS-accelerated screening protocol follows these steps:

1. Select 500–1000 molecules at random to serve as a basis set for the library.
2. Use SCISSORS to infer vectors for basis set by eigenvalue decomposition and for library by least-squares.
3. For each similarity search against the library, run slow similarity method for query molecule versus each basis molecule and learn vector.
4. Screen query molecule against full library by SCISSORS vector Tanimoto.

The advantage to the SCISSORS protocol is that, after vectors have been inferred for a given library of molecules, each new screen only requires that the source similarity measure be run against the basis set, not the whole library.

The majority of the comparisons can be done using the SCISSORS Tanimotos, which are thousands of times faster to evaluate.

In particular, clustering and other approaches that require the computation of a full $O(N^2)$ similarity matrix can be dramatically accelerated by SCISSORS. Assuming a library of size N and a basis set of size k ($k \ll N$), using SCISSORS to compute the Tanimoto similarities for the matrix requires only $O(kN)$ evaluations of the slow similarity measure, rather than $O(N^2)$. Since k can be chosen as a fixed constant, SCISSORS is asymptotically faster in slow similarity evaluations than a straightforward similarity construction by a factor of N/k , potentially many thousands of times. This, combined with the asymptotic reduction in storage space required to quickly retrieve a similarity matrix (discussed in the Results), demonstrates that SCISSORS enables huge reductions in computational cost for large-scale chemical matrix problems.

CONCLUSION

The heavy computational burden imposed by popular chemical similarity measures, combined with the growing size of chemical databases, necessitates faster techniques for similarity comparison to enable large-scale cheminformatics analyses. We have described SCISSORS, a generic method for rapid estimation of chemical similarities that is applicable to several popular similarity measures, including ones based on substructure comparison (LINGO)² and three-dimensional shape and chemical complementarity.^{5,10} SCISSORS is able to achieve high accuracy with an efficient precalculation procedure and a relatively small molecular basis set. Given precalculated molecular vectors, our method achieves several thousand-fold speedup relative to running the source similarity metric for ROCS. We have further described a method that allows construction of a molecular similarity matrix (an important intermediate step in several algorithms) with only a linear number of evaluations of the source similarity measure, rather than the quadratic number normally required.

There are several avenues for further exploration of SCISSORS and related methods. In particular, it would be interesting to examine the performance of previously described molecular-shape basis sets^{13,34} as basis sets for SCISSORS. There is also scope for exploration of methods that can better capture the curvature of the feature spaces induced by certain similarity measures.

Our algorithm enables cheminformatics calculations to be performed with asymptotically lower costs, both in terms of computation (number of slow similarity evaluations) and storage (keeping molecule vectors on disk rather than full similarity matrices). As a consequence, we anticipate that SCISSORS will allow a new class of extreme-scale cheminformatics applications combining very large data sets with similarity measures formerly considered inaccessible for large-compound collections.

ACKNOWLEDGMENT

The authors acknowledge funding from NSF award CNS-0619926 for computer resources and NIH PN1 EY016525-02.

APPENDIX: DETAILED METHODS

Molecular Databases and Preprocessing. The following databases were used:

1. Maybridge+BBP: 56 841 of 56 842 molecules from Maybridge Screening Collection (cromoglicic acid caused Omega to crash and was excluded) + 417 molecules from a blood-brain barrier penetration data set.²² Molecules from Maybridge were provided as 2D SDF files; BBP molecules were provided as SMILES.
2. Directory of Useful Decoys, release 2: 128 371 molecules provided as 3D MOL2 files with embedded charges.
3. Asinex Gold Collection: 233 795 molecules provided in 2D SDF format.
4. Maybridge Fragment Library: 473 molecules provided in 3D MOL2 format with embedded charges.
5. PubChem Assay #677 (discovery of novel allosteric modulators of the M1 muscarinic receptor: antagonist confirmation screen): 1665 molecules provided in 2D SDF format.

For standardization, all molecules (even if provided with 3D geometry and charge information) were converted to 2D uncharged OEBinary format using the OpenEye OEChem toolkit and then processed through the same pipeline. Molecules consisting of multiple disconnected chains (e.g., acids with counterions) were reduced to a single chemical chain using the function “OETheFunctionFormerlyKnownAsStripSalts” to keep only the largest single chain. Molecules containing atoms other H, C, N, O, F, Si, P, S, Cl, Br, and I were excluded, as these are the only atoms allowed in the MMFF94s force field used by OMEGA. Protonation, tautomer, and stereochemical state (other than invertible nitrogens) were maintained as given in the input data.

From the 2D representation, the OpenEye Omega toolkit was used to build one 3D conformer, using all default settings except for MaxConfGen = 50 and MaxConfs = 1. This representative conformer was used to assign charges, using the AM1-BCC charge model in the OpenEye OEProton toolkit. Final 3D conformers were then built using Omega with the following parameters:

1. Bondi van der Waals radii
2. fromCT = False
3. EnumNitrogen = True
4. EnumRing = True
5. BuildForceField = SearchForceField = mmff94s_Trunc
6. EnergyRange = “5.0, 10.0, 15.0, 20.0, 25.0”
7. MaxConfGen = 3000
8. MaxConfRange = 50, 100, 200 (multiconformer molecules only)
9. MaxConfs = 1 (single-conformer molecules only)
10. RangeIncrement = 3
11. RMSRange = “0.5,0.75,1.0,1.5”

For the Maybridge+BBP data set, both multi- and single-conformer molecules were built; for all others, only single-conformer representations were used. These data sets were already in use for other ongoing drug discovery projects and thus were not rebuilt for this study.

All molecule files used as input to ROCS, LINGO, ZAP, and the SCISSORS training routines were stored as 3D OEBinary files.

Software Versions. The Maybridge+BBP data set was prepared using OEChem 1.5.1 for file parsing, OEOmega 2.2.1 for conformer generation, and OEProton 1.2.1 for charge assignment. All other data sets were prepared using

OEChem 1.6.1, OEOmega 2.3.0, and OEQuacPac (the renamed OEProton) 1.3.1.

ROCS shape and color similarities were calculated by the OpenEye OEShape toolkit, version 1.7.0. LINGO similarities were calculated using OEChem 1.6.1. ZAP ET similarities were calculated using OEZap 2.1.0.

Benchmarking of SCISSORS similarities was done using an implementation written in Python and C, using Python 2.5.2, Numpy 1.1.0, and GotoBLAS2 1.13. Most calculations were performed using an older implementation of SCISSORS written in pure Python and Numpy; performance benchmarking was done using a newer implementation of SCISSORS Tanimotos written in C using GotoBLAS2. C code was compiled using gcc 4.0.1.

Software Settings for Similarity Calculations. ROCS shape and color similarities for all accuracy tests, except the virtual screening test, were calculated using the Analytic overlap method, the Implicit Mills–Dean color forcefield, with color optimization enabled, ignoring hydrogens and using exact atomic radii (rather than coercing all radii to the carbon radius). For each pair of molecules, the best overlay was chosen from the set of all conformer pairs (for multiconformer matches) and the ROCS starting positions by selecting the overlay with the highest Combo score. The shape and color Tanimoto values from this overlay were used as the representative Tanimoto values for that particular pair.

For virtual screening, ROCS was run on the DUD molecules with the Grid overlap method, ignoring hydrogens, and with no color force field (as only the shape Tanimoto was being approximated).

For the performance benchmark (in which we measured the time to construct a SCISSORS representation of a 20 480 molecule library), ROCS/OEShape was run in Analytic overlap mode with no color force field (shape only) and with default settings otherwise.

LINGO similarities were calculated using the OELingoSim module in OEChem. Aromaticity (using OpenEye aromaticity model) and chirality were perceived on molecules from their 3D structures; LINGO similarities were then calculated based on an isomeric canonical SMILES representation.

ZAP electrostatic Tanimotos were calculated using the OEET module in OEZap. Molecules were first overlaid into their optimal overlaid pose using ROCS, using the same settings as for shape/color similarity. OEET’s default settings for Tanimoto similarity were used to calculate the electrostatic similarity, based on the AM1-BCC charges calculated during molecule setup.

Performance Benchmarking Setup. Performance benchmarking for ROCS and SCISSORS was run on a Mac Pro with 2 × 3 GHz Quad-Core Intel Xeon CPUs (Core 2 architecture) and 16GB of DDR2 667 MHz fully buffered memory, running OS X 10.5.8. Both the ROCS and SCISSORS benchmarking scripts were written in Python and executed on a 32-bit build of CPython 2.5.2. All benchmarks were conducted as single-CPU tests.

ROCS was called from within a Python script using the OpenEye OEShape toolkit, and greater than 97% of the execution time was spent inside the OEBestOverlay.Overlay function (which calculates the optimal overlap).

SCISSORS operations, other than Tanimoto calculations, were performed using standard operations in Numpy 1.1.0, as downloaded in binary form from numpy.scipy.org. The

C module for SCISSORS Tanimoto calculation was built against GotoBLAS2 1.13, using gcc 4.0.1 with the following performance-relevant compiler options: g++ -arch i386 -fno-strict-aliasing -Wno-long-double -no-cpp-precomp -fno-common -DNDEBUG -msse -msse2 -msse3 -march=core2 -mfpmath=sse -O3. SCISSORS Tanimoto calculation times were computed over $4096 \times 262\,144$ Tanimotos, calculated in blocks of $4096 \times 32\,768$, in 64 dimensions.

REFERENCES AND NOTES

- (1) Nikolova, N.; Jaworska, J. Approaches to Measure Chemical Similarity - a Review. *QSAR Comb. Sci.* **2003**, *22*, 1006–1026.
- (2) Vidal, D.; Thormann, M.; Pons, M. LINGO, an Efficient Holographic Text Based Method To Calculate Biophysical Properties and Inter-molecular Similarities. *J. Chem. Inf. Model.* **2005**, *45*, 386–393.
- (3) Carhart, R. E.; Smith, D. H.; Venkataraghavan, R. Atom pairs as molecular features in structure-activity studies: definition and applications. *J. Chem. Inf. Comput. Sci.* **1985**, *25*, 64–73.
- (4) Durant, J. L.; Leland, B. A.; Henry, D. R.; Nourse, J. G. Reoptimization of MDL Keys for Use in Drug Discovery. *J. Chem. Inf. Comput. Sci.* **2002**, *42*, 1273–1280.
- (5) Grant, J. A.; Gallardo, M. A.; Pickup, B. T. A fast method of molecular shape comparison: A simple application of a Gaussian description of molecular shape. *J. Comput. Chem.* **1996**, *17*, 1653–1666.
- (6) Zauhar, R. J.; Moyna, G.; Tian, L.; Li, Z.; Welsh, W. J. Shape signatures: a new approach to computer-aided ligand- and receptor-based drug design. *J. Med. Chem.* **2003**, *46*, 5674–5690.
- (7) Jain, A. N. Ligand-Based Structural Hypotheses for Virtual Screening. *J. Med. Chem.* **2004**, *47*, 947–961.
- (8) Nicholls, A.; Grant, J. A. Molecular shape and electrostatics in the encoding of relevant chemical information. *J. Comput.-Aided Mol. Des.* **2005**, *19*, 661–686.
- (9) Muchmore, S. W.; Souers, A. J.; Akritopoulou-Zanze, I. The use of three-dimensional shape and electrostatic similarity searching in the identification of a melanin-concentrating hormone receptor 1 antagonist. *Chem. Biol. Drug Des.* **2006**, *67*, 174–176.
- (10) Rush, T. S.; Grant, J. A.; Mosyak, L.; Nicholls, A. A shape-based 3-D scaffold hopping method and its application to a bacterial protein-protein interaction. *J. Med. Chem.* **2005**, *48*, 1489–1495.
- (11) Butina, D. Unsupervised Data Base Clustering Based on Daylight's Fingerprint and Tanimoto Similarity: A Fast and Automated Way To Cluster Small and Large Data Sets. *J. Chem. Inf. Comput. Sci.* **1999**, *39*, 747–750.
- (12) Tanaka, N.; Ohno, K.; Niimi, T.; Moritomo, A.; Mori, K.; Orita, M. Small-World Phenomena in Chemical Library Networks: Application to Fragment-Based Drug Discovery. *J. Chem. Inf. Model.* **2009**, *12*, 2677–2686.
- (13) Haigh, J. A.; Pickup, B. T.; Grant, J. A.; Nicholls, A. Small Molecule Shape-Fingerprints. *J. Chem. Inf. Model.* **2005**, *45*, 673–684.
- (14) Irwin, J. J.; Shoichet, B. K. ZINC - A Free Database of Commercially Available Compounds for Virtual Screening. *J. Chem. Inf. Model.* **2005**, *45*, 177–182.
- (15) Blum, L. C.; Reymond, J.-L. 970 Million Druglike Small Molecules for Virtual Screening in the Chemical Universe Database GDB-13. *J. Am. Chem. Soc.* **2009**, *131*, 8732–8733.
- (16) Bohacek, R. S.; McMartin, C.; Guida, W. C. The art and practice of structure-based drug design: A molecular modeling perspective. *Med. Res. Rev.* **1996**, *16*, 3–50.
- (17) Cox, T. F.; Cox, M. A. A. *Multidimensional Scaling*; Chapman & Hall: London, U.K., 1994.
- (18) Golub, G. H.; Van Loan, C. F. *Matrix Computations*, 3rd ed.; The Johns Hopkins University Press: Baltimore, MD, 1996.
- (19) Schölkopf, B.; Smola, A.; Müller, K.-R. Nonlinear Component Analysis as a Kernel Eigenvalue Problem. *Neural Comput.* **1998**, *10*, 1299–1319.
- (20) Schölkopf, B.; Smola, A.; Müller, K.-R. Technical Report 44. *Nonlinear Component Analysis as a Kernel Eigenvalue Problem*; Max-Planck-Institut für biologische Kybernetik: Tübingen, Germany, 1996.
- (21) Pekalska, E. M. Ph D. Thesis, Wroclaw University: Wroclaw, Poland, 2005.
- (22) Li, H.; Yap, C. W.; Ung, C. Y.; Xue, Y.; Cao, Z. W.; Chen, Y. Z. Effect of Selection of Molecular Descriptors on the Prediction of Blood-Brain Barrier Penetrating and Nonpenetrating Agents by Statistical Learning Methods. *J. Chem. Inf. Model.* **2005**, *45*, 1376–1384.
- (23) Huang, N.; Shoichet, B. K.; Irwin, J. J. Benchmarking Sets for Molecular Docking. *J. Med. Chem.* **2006**, *49*, 6789–6801.
- (24) Haque, I. S.; Pande, V. S. PAPER - Accelerating Parallel Evaluations of ROCs. *J. Comput. Chem.* **2010**, *31*, 117–132.
- (25) Clark, R. D.; Webster-Clark, D. J. Managing bias in ROC curves. *J. Comput.-Aided Mol. Des.* **2008**, *22*, 141–146.
- (26) Goto, K.; Van De Geijn, R. High-performance implementation of the level-3 BLAS. *ACM Trans. Math. Software* **2008**, *35*, 1–14.
- (27) *Practical Application of Computer-Aided Drug Design*, 1st ed.; Charifson, P. S., Ed.; Marcel Dekker: New York, 1997.
- (28) Martin, E. J.; Blaney, J. M.; Siani, M. A.; Spellmeyer, D. C.; Wong, A. K.; Moos, W. H. Measuring Diversity: Experimental Design of Combinatorial Libraries for Drug Discovery. *J. Med. Chem.* **1995**, *38*, 1431–1436.
- (29) Nicholls, A. Method for determining a shape space for a set of molecules using minimal metric distances. U.S. Patent 7,110,888, 2006.
- (30) Raghavendra, A. S.; Maggiora, G. M. Molecular Basis Sets - A General Similarity-Based Approach for Representing Chemical Spaces. *J. Chem. Inf. Model.* **2007**, *47*, 1328–1340.
- (31) Williams, C. K. I. On a Connection between Kernel PCA and Metric Multidimensional Scaling. *Mach. Learn.* **2002**, *46*, 11–19.
- (32) Jenkins, O. C. Technical Report CRES-03–010. *Localization from Pairwise Distance Relationships using Kernel PCA*; USC Center for Robotics and Embedded Systems: Los Angeles, USA, 2003.
- (33) Bengio, Y.; Delalleau, O.; Roux, N. L.; Païement, J.-F.; Vincent, P.; Ouimet, M. Learning Eigenfunctions Links Spectral Embedding and Kernel PCA. *Neural Comput.* **2004**, *16*, 2197–2219.
- (34) Fontaine, F.; Bolton, E.; Borodina, Y.; Bryant, S. H. Fast 3D shape screening of large chemical databases through alignment-recycling. *Chem. Cent. J.* **2007**, *1*, 12.

CI1000136