Published in final edited form as:

J Chem Inf Model. 2011 April 25; 51(4): 760–776. doi:10.1021/ci100223t.

Data-Driven High-Throughput Prediction of the 3D Structure of Small Molecules: Review and Progress

Alessio Andronico[†], Arlo Randall[†], Ryan W. Benz[‡], and Pierre Baldi^{*}

School of Information and Computer Sciences, Institute for Genomics and Bioinformatics, University of California, Irvine, Irvine, CA 92697-3435, USA

Abstract

Accurate prediction of the 3D structure of small molecules is essential in order to understand their physical, chemical, and biological properties including how they interact with other molecules. Here we survey the field of high-throughput methods for 3D structure prediction and set up new target specifications for the next generation of methods. We then introduce COSMOS, a novel data-driven prediction method that utilizes libraries of fragment and torsion angle parameters. We illustrate COSMOS using parameters extracted from the Cambridge Structural Database (CSD) by analyzing their distribution and then evaluating the system's performance in terms of speed, coverage, and accuracy. Results show that COSMOS represents a significant improvement when compared to the state-of-the-art, particularly in terms of coverage of complex molecular structures, including metal-organics. COSMOS can predict structures for 96.4% of the molecules in the CSD [99.6% organic, 94.6% metal-organic] whereas the widely used commercial method CORINA predicts structures for 68.5% [98.5% organic, 51.6% metal-organic]. On the common subset of molecules predicted by both methods COSMOS makes predictions with an average speed per molecule of 0.15s [0.10s organic, 0.21s metal-organic], and an average RMSD of 1.57Å [1.26Å organic, 1.90Å metal-organic], and CORINA makes predictions with an average speed per molecule of 0.13s [0.18s organic, 0.08s metal-organic], and an average RMSD of 1.60Å [1.13Å organic, 2.11Å metal-organic]. COSMOS is available through the ChemDB chemoinformatics web portal at: <http://cdb.ics.uci.edu/>.

Introduction

Accurate prediction of the 3D structure of small molecules is essential in order to understand their physical, chemical, and biological properties including how they react or interact with other molecules. The fundamental importance of 3D structure prediction on a large-scale has increased dramatically with the recent development and free availability of large, public, databases of small molecules^{1–4} containing millions of compounds coupled with the growing interest in the design of new molecules^{5,6} and the exploration of the astronomically-sized space of virtual molecules that are not readily available but could be synthesized from existing building blocks.⁷ In contrast, the current version of the Cambridge Structural Database (CSD),^{6,8,9} the main repository of experimentally-determined small molecule structures which is available only through commercial licenses, contains less than half a million structures. Thus the overwhelming majority of small molecules is not represented in the CSD and there is an important need for methods capable of accurately predicting molecular structures on a high-throughput scale.

^{*}To whom correspondence should be addressed, Department of Biological Chemistry, UCI, pfbaldi@ics.uci.edu.

[†]These authors contributed equally.

[‡]Currently with Applied Proteomics, Glendale, CA.

The study of computational 3D structure generation has a long history, going back more than 30 years.¹⁰ Current methods for predicting 3D structures suffer from problems of speed, coverage, or accuracy. On one hand, quantum mechanics or molecular mechanics methods^{11–14} are often capable of yielding accurate structures. However these methods often depend on a good starting structure and typically require on the order of minutes to hours to generate a low energy structure on current computers. So while these may be the methods of choice for a fine-grained analysis of a small set of molecules of interest, they are not suitable for sifting through very large sets of molecules. On the other hand, approaches based on simple rules or templates are much faster, can be used on a large-scale, and have led in some cases to widely used commercial products.^{10,15} However, these methods suffer from issues of coverage or accuracy, and fail to produce models for some complex organic molecules. In addition, current programs employing these methods typically fail on the majority of metal-organic molecules in the CSD. This shortcoming is critical because metal-organic molecules are used in a variety of important medical and industrial applications. Furthermore, because the best methods in this class are embedded in commercial products, they are not open or freely available to the academic research community. Thus there is a clear need for a new system for predicting the 3D structures of small molecules that is fast, accurate, with good coverage, and open and free for the academic research community.

In the somewhat related field of protein structure prediction, successful methods have been developed that rely heavily on the freely available Protein Data Bank (PDB)¹⁶ in one way or another, for instance by building models using homology to known structures, by including statistical (or knowledge-based) potential terms in the energy function to be minimized, or by extracting libraries of fragments (and their coordinates) from the PDB and using them as building blocks that are then pieced together by some deterministic or stochastic algorithm.^{17–19} This suggests the possibility of using a similar approach for small molecule structure prediction, by relying on a library of fragment parameters that can be used to reconstruct more complex 3D structures. Thus the goal here is to present this idea in detail and build the first version of a system capable of predicting 3D structures rapidly, accurately, with good chemical coverage, that can scale to large databases or combinatorial libraries, and that is as free and open for academic research as possible. Here we illustrate the general approach using a fragment library extracted from the CSD, but it should be clear that the approach can be applied using any fragment library.

Here we present the methods and benchmark the performance of a novel data-driven system called COSMOS (COordinate of Small MOleculE S) for the high-throughput prediction of small molecule 3D structures. In essence, the system decomposes any molecule into fragments that are matched to libraries of fragments extracted for illustration purposes from the CSD, and assembles the resulting matching fragments into a 3D model. The presentation follows a broader review perspective that also covers key issues related to small molecule structure prediction. The rest of the article is organized as follows. First, a short summary section on the growth and composition of the CSD is presented. This is followed by a brief review section on the current state-of-the-art in the field and a set of target specifications for the next generation of prediction methods are proposed. Then a section on fundamental methodological issues for the field discusses structural ambiguities and assessment protocols. This is followed by a section that presents our main results obtained by using a fragment library extracted from the CSD, with a focus on the statistical distributions of various fragment types. Then the fragment matching methods and the procedure for building molecules from fragments are described in their own sections. In the final section the benchmark performance of the new prediction method is assessed using various metrics for coverage, speed, and accuracy.

Growth and Composition of the Cambridge Structural Database

The Growth of the CSD

Because here we illustrate our approach using the CSD, we first briefly review several statistics regarding the CSD, starting with its growth. Taking snapshots of the number of entries with error-free 3D coordinates in the CSD in recent decades demonstrates the significant growth: 25,016 structures in 1980, 84,068 structures in 1990, 220,013 structures in 2000, and 425,122 in the 2009 release of the CSD (CSD09) (Figure 1). In fact, the CSD has been growing exponentially, roughly doubling every 7–8 years.

The Composition of the CSD: Importance of Organometallic Molecules

The growth within the CSD has been increasingly dominated by organometallic (aka metal-organic) molecules, with their representation growing from 44.9% in 1980 to 53.8% in 1990, 60.1% in 2000, and 61.7% in 2009. The increasing numbers of organometallic molecules crystallized in recent years results probably from multiple factors including the difficulty of characterizing their structures by non-X-ray methods and their increasing relevance in several areas of research. Organometallics already play essential roles in medicine, chemical research, and industry. As a broad class, organometallic small molecules are much more complex and structurally diverse than organic small molecules, thus, they can exhibit functional properties that cannot be realized by any organic molecule. Their complexity and diversity also present a significant challenge for structure prediction methods and prediction failure rates are much higher on organometallics. Thus, there is a significant need for new prediction methods that can accurately predict the structures of organometallic molecules.

While organic small molecules consist only of the following atoms: H, B, C, N, O, F, P, S, Cl, Br, and I, organometallics contain combinations of metal atoms and atoms from the organic subset. Organometallics have a much larger set of atomic building blocks and can exhibit much more complex connectivity between atoms than organic molecules. Two simple measures of molecular complexity are the number of atoms and the maximum coordination number (*CN*) observed among any of its atoms. The average number of heavy atoms in the organic molecules in the CSD09 is 20.0 and the average maximum *CN* is 3.5, while the average number of heavy atoms in organometallics is 34.5 and the average maximum *CN* is 6.1. Less than 1% of organic molecules have a maximum *CN* of more than four, in contrast more than 70% of organometallic molecules have a maximum *CN* of more than four.

From a significance standpoint, organometallics have important applications in chemical research and industry, for instance as catalysts in the large scale production of chemicals used in pharmaceuticals, food products, and fuel additives. Organometallic small molecules are also playing an increasing role in medicine. Most drugs are organic small molecules, however, 45 of the 1,382 drugs (3.2%) in the Drug-Bank database^{20,21} listed as approved are organometallic, and 145 of the 3,203 drugs (4.5%) listed as experimental are organometallic. Only three of the approved organometallic drugs (0.22%) have a maximum *CN* of more than four, and two of them are only slightly different versions of vitamin B-12 (Cyanocobalamin and Hydroxocobalamin). There are 24 experimental organometallic drugs (0.75%) with a maximum *CN* of more than four. A variety of organometallic molecules have been shown to have anti-cancer activity.^{22,23} Organometallic molecules also act as delivery systems for vital minerals such as iron, magnesium, and zinc. In another medical application, a variety of organometallic molecules have been shown to have anti-cancer activity.^{22,23}

Current State-of-the-Art in Small Molecule 3D Structure Prediction and Specifications

State-of-the-Art

The 3D structure of small molecules can be predicted fairly accurately using quantum mechanics (QM) or molecular mechanics (MM) approaches.^{11–14,24} However, these methods can still produce imperfect models and suffer from two drawbacks: they often require a good starting configuration and they require a non-trivial computational time. MM may require from seconds to minutes for a small organic molecule and scale as the square of the number of atoms (n^2) involved. For metal-organic molecules, in general there are no proper MM force fields and thus one must typically resort to QM methods which may take from minutes to hours and scale as the cube (n^3) or higher. In either case, these durations are not compatible with the prediction of the structures of millions or billions of molecules. Furthermore, both QM and MM methods benefit from a good starting point. As a result, both QM and MM methods ought to be viewed as refined methods that can be used, in conjunction with the high-throughput methods described here, to produce refined structures for a small number of suitably selected molecules of interest. In particular, QM and MM methods can be used to build libraries of basic fragments with even greater coverage than what can be extracted from the CSD.

For the rapid, high-throughput, prediction of molecular 3D structures, several methods have been proposed over the years ranging from methods based on distance geometry to rule based methods.^{10,25–30} Some of these methods have been incorporated into widely used commercial product for general coordinate generation including Concord (<http://www.tripos.com>^{27,31}), CORINA (<http://www.mol-net.de/>^{10,15}), LigPrep (<http://www.schrodinger.com/products/14/10/>), and Omega (<http://www.eyesopen.com>). Probably the most widely used of these programs is CORINA, which has been licensed in the past to predict 3D structures for some of the molecules in the main large public databases of small molecules such as PubChem (<http://pubchem.ncbi.nlm.nih.gov/>^{4,32}) and ChemDB.^{2,3} Thus we have obtained a license for CORINA for evaluation and comparison purposes. While commercial methods have demonstrated some success in the field of small molecule structure prediction, there has been a general lack of comparable programs that are freely available to the academic community. In particular, public chemical databases such as PubChem, ZINC, and ChemDB could benefit from a high quality free method. A few free programs have emerged, such as Open Babel (http://openbabel.org/wiki/Main_Page³³) and FROG (<http://bioserv.rpbs.univparis-diderot.fr/cgi-bin/Frog2>³⁰). However, the performance gap between these free methods and the commercial methods is significant.

Evaluation

Here we evaluate the coverage, speed, and accuracy of some relevant existing methods on a random subset of 10,000 molecules extracted from the 2009 version of the CSD (CSD09). All timing benchmarks in this paper are carried on a single workstation with AMD Opteron Processor 250 (2.4 Gz processor), 4 GB of RAM, and 1 MB of cache. We tested the free methods Open Babel and FROG, the commercial method CORINA (version 3.46), and a distance geometry (DG) approach similar to^{29,34} developed in-house. The DG method is evaluated alone, without a subsequent energy minimization step, to assess how this general approach performs while retaining its fast conversion speed—adding an energy minimization step significantly increases the DG conversion times. The results of the benchmark tests are summarized in Table 1 and Table 2. Table 1 displays the percentage of molecules for which a prediction was returned (coverage) and the average prediction time for the entire set of 10,000 molecules as well as the organic (3,444 molecules) and metal-organic subsets (6,556 molecules). In order to make a fair comparison of the methods, we identified the common

subset of molecules that all methods were able to produce models for. The common subset includes 2,443 organic molecules and only 47 metal-organic molecules. Table 2 displays the average RMSD, the percentage of models containing atom-atom clashes¹, and the average prediction time of the methods on the common subset.

Coverage—The DG method has perfect coverage of the organic and metal-organic molecules while CORINA covers 98.5% of organics and approximately half of metal-organics (49.6%). Open Babel covers 85.7% of the organics and only 2.4% of the metal-organics. Similarly, FROG covers 86.1% of the organics and only 1.7% of the metal-organics.

Speed—CORINA is by the fastest method with an average prediction time of 0.01s on the common subset of molecules and 0.04s calculated on all the molecules that it was able to make a prediction on. For some additional molecules CORINA attempted to make a prediction, but then failed. When the run-times on these failed molecules are included the average time to produce a model is 0.20s. The DG method had an average prediction time of 0.21s on the common subset molecules and 14.60s when calculated on the entire set of 10,000 molecules. Open Babel had an average time of 23.9s on the common subset and 28.4s over all the molecules it was able to predict. The most recent version of FROG is only available as a web-server (the authors plan to release a downloadable version in the future), thus we could not accurately calculate an average run-time. The average prediction time calculated from the server response times was approximately 7s.

Accuracy—CORINA is the most accurate on both the organic and metal-organic subsets with average RMSD of 0.95Å and 1.18Å respectively. There is a noticeable performance gap between CORINA and the next best method, Open Babel, which has an average of 1.20Å on the organics and 1.34Å on the metal-organics. FROG has slightly higher averages than Open Babel on both subsets: 1.31 Å on the organics and 1.42Å on the metal-organics. DG is significantly worse than these with averages of 1.79Å on organics and 1.88Å on the metal-organics. Regarding atom-atom clashes, less than 1% of the models produced by Open Babel and FROG contain clashes. 2.8% of the CORINA models contain clashes and 10.0% of the DG models contain clashes.

In summary, the commercial method CORINA is the fastest and most accurate method among the ones we tested. In addition, CORINA covers nearly all organic molecules and about half of the metal-organics. This coverage is significantly better than both FROG and Open Babel. The obvious strength of the DG method is that it can make predictions on virtually any molecule; however, the weak accuracy and speed results indicate that it would be a poor choice to use as a stand-alone method. The DG accuracy could be improved by adding an energy minimization step but at the cost of further increasing the run time. Based on the benchmark tests CORINA is clearly one the best existing method for high-throughput small molecule structure prediction, thus, only CORINA is used for all of the comparisons with COSMOS presented in the Results section. CORINA's most glaring weakness is that it performs very poorly on metal-organic molecules when compared to organic molecules, in terms of speed and accuracy, but even more significantly in coverage. One of the key challenges associated with metal-organic molecules is how to handle the molecules with high degrees of connectivity between atoms. For instance, 29% of the metal-organic molecules in the CSD09 have a maximum *CN* of 7 or higher, and CORINA does not attempt prediction on any molecule with maximum *CN* greater than 6.

¹For details about the definition of atom clash see the section: Molecular Fragments, Statistics, and Representations.

Challenges Associated with Metal-Organics

One major difficulty associated with the higher *CN* atoms is that the 1D and 2D molecular representations containing these atoms can correspond to an under-constrained system. The stereo descriptors in these representations can be used to uniquely define the placement of neighboring atoms around an atom with *CN*=6; however, this cannot be applied to atoms with higher *CN*. In some cases the connectivity among the entire set of atoms defines a unique solution in 3D space. However, in other cases the system is under-constrained and many unique 3D structures will satisfy the constraints provided by the input. In this situation, it is still preferable to produce one or more plausible models than to just fail, as long as the issue is communicated to the user of the predictor. Another challenge associated with predicting metal-organic molecules that is specific to data-driven methods, is that, in general, a fragment from a query metal-organic molecule is much less likely to match against a library of fragments than an organic fragment. This results from the fact that the population of metal-organic fragments is much more diverse than the population of organic fragments. The diversity of various types of organic and metal-organic fragments are discussed and analyzed in detail in the section on Molecular Fragments, Statistics, and Representations.

Specifications

Given the results discussed above, the current state-of-the-art in small molecule structure prediction falls short of what will be required by the next generation of high-throughput computational and laboratory experiments that rely on predicted structural models. Here we outline new target specifications for the next generation of high-throughput small molecule structure prediction methods. Based on the current state-of-the-art, for a new prediction method to have an impact in the field, it must excel in one or more (preferably all) of the following areas:

Coverage—The system should cover all organic and metal-organic small molecules.

Speed—The system should be very fast. It should be suitable for high-throughput applications and capable of processing on the order of one million small molecules per hour on a small cluster (10–100 nodes). Thus the average prediction time per molecule should be on the order of 0.1s on a single node.

Accuracy—The system should be very accurate. The average Root Mean Square Deviation (RMSD), a standard measure of performance in 3D structure prediction, should not exceed 1.2Å when computed on a large set of small molecules, consisting of organic and metal-organic molecules. This target accuracy takes into account structural ambiguities discussed in the section on Fundamental Methodological Issues. In addition, the predicted models should have correct coordination geometries and be as free of atom-atom clashes as experimentally-determined models.

Openness—The system should be freely available for academic research purposes.

Fundamental Methodological Issues

General Approach

The basic approach developed in this work consists of three steps: (1) build fragment libraries; (2) for any query molecule, decompose the molecule into elementary fragments and match those against the libraries; and (3) use the matching fragments to produce a 3D structure prediction. In turn, there are several possibilities for implementing each step such as: (1) the nature and size of the fragments (e.g. rigid fragments) and their origin (e.g. CSD

vs other sources); (2) the matching strategy (e.g. exact, partial, fuzzy); and (3) the algorithms for assembling fragments together (e.g. deterministic versus stochastic). For each step, we describe the implementation of the current version of COSMOS. Before addressing these technical issues we briefly discuss molecular representations, structural ambiguities, and assessment methodology.

Inputs and Outputs: SMILES and SDF Representations

Molecules can be represented in many computer formats, such as Simplified Molecular Input Line Entry Specification (SMILES) strings, connectivity matrices (e.g. sdf files), and fingerprints.^{35,36} Programs for converting from one format to another exist, thus in principle one can work with any format. For simplicity and consistency with wide practice, we will use SMILES. SMILES is a chemical language that can be used to specify the connectivity and stereochemistry of a molecule in a straightforward and compact line of text.^{37–39} The language was extended and improved by DAY-LIGHT Chemical Information Systems, Inc. (<http://www.daylight.com>). In SMILES, branches are represented by parentheses and ring closures are indicated by matching integer tags. Tetrahedral and bond stereochemistry can also be defined with SMILES. The SMILES format is supported by most chemoinformatics software tools, and is used extensively to store the 2D structure of molecules in various public databases. For instance, PubChem,^{4,32,40} ChemDB,^{2,3} ZINC,¹ and DrugBank²¹ all utilize SMILES. Internally, COSMOS takes SMILES input and produces 3D structure output using the standard sdf (Structure-Data File) format. However, since these formats can be converted to the other standard representations, users of the system can provide any standard input format and receive any standard output format.

Structural Ambiguities

There are several sources of ambiguity between 1D or 2D molecular representations and 3D structure that must be considered and accounted for during the development of a prediction system. The SMILES string or the bond matrix/graph of a molecule does not always specify a unique 3D structure. Ambiguities can arise even within a single rigid fragment due to the presence of mirror symmetries (enantiomers) or tetrahedral atoms with unspecified stereochemistry (stereoisomers). These ambiguities can combine exponentially when a molecule with multiple rigid fragments is considered. Some degree of stereochemistry can be specified in the SMILES. However this does not always occur, and even when it occurs, the specification may be insufficient to describe a unique structure. These ambiguity problems are common in current large databases of molecules and their associated SMILES representations.

Some of these issues are illustrated in Figure 2 using simple rigid fragments. Figure 2a and Figure 2b display a fragment with two different tetrahedral arrangements. The tetrahedral variation can be explicitly indicated in the corresponding SMILES, in which case the correct unique structure ought to be produced. But if the tetrahedral annotation is omitted from the SMILES, then the SMILES does not specify a unique structure. If not specified, then it can be modeled in two possible ways. In some cases one of the options is clearly superior energetically and thus the ambiguity is not a problem in model building; however, if both options are energetically equal (as in Figure 2a and Figure 2b) then either option is equally likely. A simple solution is to notify the user of the source of ambiguity and to return both models. However, these issues can multiply with the number of tetrahedral atoms within a rigid fragment and with the number of fragments in a molecule. Approximately 28% of the molecules in CSD09 have one or more undefined stereo centers in the SMILES representation when the source SDF file properties alone are used to generate the SMILES

There are other well known sources of ambiguity that cannot be incorporated in the current SMILES notation and result in a SMILES string being associated with multiple 3D structures with identical or similar energies. A clear example of this phenomenon is provided by mirror images, as shown in Figure 2c and Figure 2d. In the rigid fragments from CSD09, we find that 95% of the SMILES of the original and the mirror image are identical. For the vast majority of these the original and the mirror have the same 3D structure due to additional symmetries; however, the RMSD between the original and the mirror is greater than 0.30Å for approximately 10% of these. In these cases, the mirror structure is the enantiomer of the original. Approximately 25% of the molecules in CSD09 have one or more rigid fragments with possible enantiomers. Figure 2c and Figure 2d provide an example of the enantiomeric structures for a single SMILES, and the RMSD between the structures is 2.52Å. The solution to this source of ambiguity is to check the SMILES of the mirror image of the rigid fragment being considered. If it matches the original, then the RMSD between the model and its mirror image can be checked. If the RMSD is significant then both models should be returned and the user should be notified of the reason. Again, however, this can become impractical due to the combinatorial explosion across multiple rigid fragments. Approximately 50% of the SMILES from the molecules in CSD09 have at least one rigid fragment flagged for ambiguity due to undefined stereo centers or enantiomers.

At the level of whole molecules rotatable bonds introduce another source of ambiguity that must be dealt with appropriately. The experimentally determined coordinates of a small molecule in the CSD represents one low-energy conformation, but when the molecule has rotatable bonds it may have many conformations that are approximately equal energetically.

These sources of ambiguity can be handled in a variety of ways both in how a prediction method is designed and during assessment. In cases of ambiguity, returning no prediction is a poor solution. In many applications (e.g. docking⁴¹) having one low energy structure is better than having none. Likewise, QM or MM methods cannot be used to perform further refinements, when no structure is produced. Furthermore, not producing a structure prevents comparison to other predictors that output a structure. The solution we adopt is to: (1) always produce at least one structure, sampled from the set of lowest energy structures described by the same SMILES string; (2) alert the user to the presence of ambiguity; (3) produce also a second structure when the ambiguity is only between two alternatives; and (4) when the source of ambiguity is undefined stereocenters, return all stereoisomers (e.g. up to 64 models). Likewise, these sources of ambiguity can be handled differently during assessment. It is possible, for instance, to perform a restricted form of assessment on the rigid fragments only, or on the ring systems only. This at least removes a major source of ambiguity (rotatable bonds). Another approach is to remove molecules with ambiguity from the assessment entirely or, if the predictors produce multiple predictions, to take the predictions with the lowest RMSD to the target. All these assessment approaches have their advantages and drawbacks; therefore we have utilized multiple assessment measures in this work. In the end, however, it is common practice to report primarily the RMSD between the first model returned by the predictor and the experimentally solved structure, thus, this measure is always included in this work.

Assessment Methodology

Accounting for Redundancy—Speed, accuracy, and coverage of prediction methods can be assessed using the molecules in the CSD. However great care must be exercised to avoid the circularity of using the entire CSD data to predict CSD structures. This is a standard statistical problem that often comes up in machine learning and bioinformatics applications. To address this issue, one must run validation experiments where a subset of

the data is used to make blind predictions on a different set of the data. In particular one can use the molecules contained in one older version of the CSD to predict the new structures that have been deposited in a more recent version. One can further reduce the chance of any circularity by removing any residual redundancy between the training and validation sets, as is routinely done in protein structure prediction. This redundancy-reduction is achieved by retaining in the validation set only molecules that have low similarity to the molecules in the training set. In protein structure prediction, it is typical to retain for instance only proteins that have less than 25% sequence similarity to the sequences in the training set.

As a conservative approach we used data from the 2004 version of the CSD (CSD04) to build the fragment libraries for predicting the structures of new molecules found in the 2009 version of the CSD (CSD09). This is quite stringent since the CSD almost doubled in size from 2004 to 2009. Furthermore, in the Results section, we present the speed and accuracy results obtained by using a set of 10,000 molecules and also a reduced set of molecules that have low similarity with the molecules in our training set.

Similarity is assessed here by comparing chemical fingerprints^{42,43} with the most widely used measure of molecular similarity, the Tanimoto coefficient,^{36,44} which scales from 0 to 1. For our reduced set, we retained only the molecules with Tanimoto coefficient below 0.6.

Accuracy Measures—The average root mean square deviation (RMSD) is the primary measure used to assess the accuracy of predicted models in this work. When a molecule contains topologically equivalent atoms the RMSD is calculated for all possible mappings between the model and the native structure and the minimum resulting value is used for assessment. While the RMSD is a good measure for assessing the overall accuracy of predicted model coordinates, it is not appropriate for every type of evaluation (see Structural Ambiguities section). Specifically, an interesting accuracy measure for metal-organic molecules is the percentage of correctly predicted metal coordination geometries. Thus, we have developed a simple angle-based measure for this purpose.

For each *CN*, 3 through 8, there are two or more metal coordination geometries that occur frequently in the CSD. The specific geometries analyzed in this work are: (*CN*=3) T-Shaped, Trigonal Pyramidal, and Trigonal Planar; (*CN*=4) Square Planar and Tetrahedral; (*CN*=5) Square Pyramidal and Trigonal Bipyramidal; (*CN*=6) Octahedral and Trigonal Prismatic; (*CN*=7) Capped Octahedron, Capped Trigonal Prismatic, and Pentagonal Bipyramidal; and (*CN*=8) Cubic, Dodecahedral, and Square Antiprismatic.⁴⁵ To develop an approach for classifying the coordination geometries we started by generating an idealized mono-centric (central atom and its covalently bonded neighbors only) template structure for each of these geometries. Next, we developed an angle-based similarity metric where each mono-centric fragment is represented as a set of angles (*Mθ*). A given angle in the set, *Mθ_{i,j}*, is the angle that neighbor atoms *i* and *j* form with the central atom, where the central atom is the vertex. The set is composed of the angles calculated using all *CN(CN-1)/2* possible pairs of neighbor atoms. The difference between the model and template angle sets is summarized by the mean difference between their corresponding angles:

$$\overline{\Delta\theta}(M\theta, T\theta) = \left[\sum_j^{CN} \sum_{j<i}^{CN} |M\theta_{i,j} - T\theta_{i,j}| \right] / [CN(CN-1)/2].$$

For the metric to assess the geometry only, and ignore the relative placement of atoms, all of the neighbor atoms are considered equivalent and $\overline{\Delta\theta}$ is calculated for all possible atom mappings between the model and template fragment and only the minimum value is kept. Finally, in order to discover an appropriate threshold for classification we applied the mean difference metric ($\overline{\Delta\theta}$) to each pair of template geometries with the same *CN*. The results of are summarized in Table 3. Based on these results we selected $\overline{\Delta\theta} < 10.0^\circ$ as an upper threshold for classifying a

model as a specific geometry because it provides clear separation between most of the geometries, but also allows for some variability.

To classify the geometry of a fragment we calculate $\overline{\Delta\theta}$ versus all templates with the same CN and identify the best match. If for the best match $\overline{\Delta\theta} < 10.0^\circ$ then the fragment is classified as that geometry, if not then it is considered unclassified. If a CSD fragment can be classified, then the corresponding model fragment is considered correct if it has the same geometry class. If a CSD fragment is unclassified, then the ideal templates are ignored and the corresponding model fragment is compared directly to the CSD fragment. The model fragment is considered correct if $\overline{\Delta\theta}(M\theta, T\theta) < 10.0^\circ$ where $T\theta$ is the angle set of the CSD fragment.

Molecular Fragments, Statistics, and Representations

Fragment Libraries, Statistics, and Representations

Molecular modeling methods utilize many kinds of fragments at varying levels of granularity, from entire molecules at the coarser level, to single pairs of bonded atoms (e.g. C-C) at the finer level. In this work, primarily two types of intermediate fragment representations were utilized: rigid fragments and cyclic fragments. These representations are described in detail below. In the determination of fragments all bonds present in the source files were utilized, including metal pi bonds. In addition, polymeric structures were handled by considering only one copy of the monomer.

Rigid Fragments and Their Torsion Angles—Rigid fragments are molecular substructures obtained by fragmenting a molecule at each of its rotatable bonds, such that the resulting fragments contain no rotatable bonds themselves (Figure 3). For each rotatable bond, the torsion angle and the SMILES for the atoms used to calculate the angle and their first bonded neighbors (“torsion key”) are saved along with the rigid fragments themselves. For a molecule with n rotatable bonds, $n + 1$ rigid fragments and n torsion angles are saved. The molecules in CSD09 fragmented into nearly 2.3 million total rigid fragments, but only 294,856 are unique according to their isomeric SMILES. The ten most frequently occurring organic and metal-organic rigid fragments are presented in Table 4. How the frequencies are distributed among the unique rigid fragments has important implications for the data-driven prediction approach. In fact, the frequency of occurrence of small molecule fragments have been shown to distribute via power-laws.⁴⁶ Figure 4 presents a plot of the CSD fragment frequencies on a log-log scale, which clearly illustrates the power-law phenomenon. The presence of power-laws in the fragment distributions indicates that a relatively small fraction of the rigid fragments occur the vast majority of the time. Although the power-law distribution applies to both organic and metal-organic fragments, there is a significant difference in how the organic fragments distribute compared to the metal-organic (Figure 4b). The top 0.1% most frequently occurring organic rigid fragments occur so often that they account for 81.3% of all occurrences. In contrast, the top 10% most common metal-organic rigid fragments account for only 69.3% of all occurrences. The basic implication of this result for small molecule structure prediction is that the organic rigid fragments of new molecules are much more likely to have a match in the data libraries than the metal-organic rigid fragments.

In this work a single prototype structure is used to represent the class of rigid fragment structures associated with the same SMILES. The prototype is selected by calculating the RMSD between all pairs of structures in the group and keeping the one with the minimum sum of RMSD values versus all other pairs. For the CSD09, this process yields 294,856 prototype rigid fragments structures.

During the fragmentation of the experimental 3D structures, the torsion angles between all connected pairs of rigid fragments are calculated and saved along with SMILES representing the atoms in the rotor bonds and their first bonded neighbors (Figure 3d and Figure 5). Figure 4e shows a plot of the frequency distribution of torsion angles of CSD09 on a log-log scale. The torsion angles clearly distribute via power-law, similar to what was observed for the rigid segments. In CSD09, there are approximately 1.9 million torsion angles with 218,391 unique “torsion keys”. To use this torsion information in COSMOS, we first bin all the torsions for a given torsion key in 10° bins. As an example, Figure 5 displays a schematic diagram of the torsion key [cH]c([s])c([cH])[s] and the resulting histogram of torsion angles. We then proceed to save in COSMOS’s torsion library the single value corresponding to the center of the bin with the highest count. For the torsion key [cH]c([s])c([cH])[s], this single value is 180°. One advantage of this simple univariate approach is that it can be applied consistently regardless of the amount of data available for a given torsion key; however, future versions of COSMOS may utilize multivariate models when appropriate. It has been demonstrated that with sufficient data multivariate distribution models of adjacent torsions can be used to improve the accuracy of model conformations.⁴⁷

Cyclic Fragments—The cyclic fragments are extracted from the rigid fragments by removing all non-ring atoms (Figure 6). This yields complete ring systems that can be further broken into individual rings. The complete ring systems and the individual rings are used to build the library.

In CSD09 there are 966,826 complete ring systems, of which 671,097 (69.4%) consist of a single ring. These yield 163,312 unique complete ring systems, including 9,276 unique single rings. The statistical distribution of the ring systems, as shown in Figure 4c, also follows a power-law, with clear differences between organics and organometallics (Figure 4d). For instance, the top 0.1% most common organic complete ring systems account for 83.5% of all occurrences, and benzene alone accounts for 62.5% of all occurrences. In contrast, the top 10% most common metal-organic complete ring systems account for only 45.0% of all occurrences. The prototype cyclic fragments used to populate the library are selected with the same procedure described above for the rigid fragments.

Atom Clashes—Here we analyze the degree of overlap between non-bonded atoms in the molecules in CSD09. To assign a single overlap score ($\mathcal{O}$) to each structure we use the following definition: $\mathcal{O} = \min_{(i,j)} d(a_i, a_j) - (v_i + v_j)$, where v_i and v_j are the van der Waals radii⁴⁸ of atom a_i and a_j respectively, $d(a_i, a_j)$ is the distance between the two atoms, and the minimum is computed over all pairs of heavy atoms in the molecule separated by more than two bonds. In other words, $\mathcal{O}$ quantifies the single most offensive steric clash in a given structure. This measure is also referred to as “overlap factor” and “close contact ratio” and has been used previously in the assessment of crystal structures and predicted models.^{15,27,31,49}

In Figure 7, we report the cumulative percentage of molecules in the CSD09 versus the overlap factor ($\mathcal{O}$). From Figure 7, it is clear that the vast majority of the molecules in CSD09 have an overlap factor greater than 0.7. In fact, only 0.7% of the CSD molecules have an overlap factor below 0.7. As a result of this analysis, $\mathcal{O} < 0.7$ was selected as the criteria to categorize a predicted model as containing clashes for assessment purposes.

Matching Fragments: Exact Matching, Partial Matching, Fuzzy Matching

Given a query molecule, it is decomposed into fragments, and the fragments are matched against the fragment libraries. The matching procedure can be exact when an exact match is found, partial when only a substructure of the fragment is matched, or fuzzy⁵⁰ which allows

also for variability in the atom types. One tool for implementing fuzzy matching is provided by the SMARTS language (Daylight) which allows one essentially to create fuzzy SMILES strings. An example of fuzzy matching using SMARTS is given in Figure 8. In many cases cyclic fragments with different atom types at one or two positions, but with the same connectivity, have nearly identical 3D structures. For instance, the basic structure of the porphyrin ring system occurs in a variety of molecules, with different atoms at the center of the structure. Figure 8 shows an example of a SMARTS pattern generated from an unmatched complete ring system (porphyrin portion of Chlorophyll A) that could be used to match a complete ring system analog in the dictionary (porphyrin portion of Heme B). The SMARTS language supports flexibility in the definitions of patterns beyond the simple wildcards shown in Figure 8. It contains a variety of built-in pattern matching, for instance [a] matches any aromatic atom, [#5] matches any nitrogen (n or N), and [#6] matches any carbon atom (C or c). In addition, SMARTS allows for the creation of user-defined patterns. COSMOS uses the following definitions for categorical fuzzy matching: [AM] Alkali Metals, [AE] Alkaline Earth Metals, [ML] Metalloids, [HL] Halogens, [NM] Non-Metals, [NG] Noble Gases, [MT] Metals, [TM] Transition Metals, and [AR] Aromatics.

3D Structure Prediction: Building Models from Fragments

To a first approximation, there are two general strategies for constructing 3D models from fragments: (1) a general top-down hierarchical approach, where first large fragments are matched to the libraries as much as possible, followed by smaller fragments; or (2) a general bottom-up approach where only small fragments are matched to the libraries and then the matches are used to reconstruct the molecules. Here the hierarchical top-down approach used by the current version of COSMOS is described together with the corresponding pseudocode.

Algorithm 1.1 describes the method at the highest level: COSMOS takes a query molecule and extracts its rigid fragments and torsion keys at the rotatable bonds, builds a model for each rigid fragment, reconnects the rigid fragments, and sets the torsion angles of the final model. Building models for the rigid fragments is the most fundamental step, and this process is presented in detail in Algorithm 1.2. COSMOS first checks the fragment library for an exact match. If this fails, then the next step depends on whether or not the rigid fragment contains a ring system.

The rigid fragments which contain ring systems are handled by removing all non-ring atoms and using the remaining ring system to search the fragment library using progressively looser matching criteria until a match is found. If the entire ring system cannot be matched in any manner, then the ring system is broken down into individual rings and they are matched against the fragment library (Algorithm 1.2 is applied to each ring). When an individual ring cannot be matched in the library, GEOMETRICRINGBUILD() is called, which constructs the ring in a bottom-up fashion using angles determined by the number of atoms in the ring, their individual properties, and ideal bond lengths.^{8,51} Once all of the individual rings have been modeled, they are reconnected to produce coordinates for the entire ring system. Figure 9 provides two examples of how rings can be stitched together. The first example (Figure 9a) corresponds to a simple tetrahedral join. The second example (Figure 9b) corresponds to a simple planar join. The non-ring atoms which are covalently bonded to atoms in the ring system are placed using PLACETERMINALATOMS(). Each terminal atom is placed using the properties of its source cyclic fragment atom, including stereochemistry, to determine the angle for placement, and ideal bond lengths.^{8,51}

The rigid fragments which contain no rings (branched) are handled by breaking them down into mono-centric fragments. Each mono-centric fragment is matched against the fragment

library and those which are unmatched are built by GEOMETRYMONOCENTRICBUILD(). This function selects a specific mono-centric geometry⁴⁵ based on the properties of the central atom, and if there is no way to discriminate a single geometry then the most prevalent one is chosen. The mono-centric fragment is constructed using the selected ideal mono-centric geometry and ideal bond lengths. Once all of the individual mono-centric fragments have been modeled, they are reconnected by aligning the overlapping bond vectors and the dihedral angles are adjusted to minimize steric hindrance.

When all of the rigid fragments of a molecule have been constructed, they are joined together at their rotor bonds by aligning the overlapping bond vectors. Then, each torsion key is matched against the torsion library, and if a match is found the resulting torsion value is set. If no match is found, then a torsion optimizing procedure built into OEChem is used to set the torsion angle.⁵² Considering the 6,795 molecule test set from the Results section, no match was found for 10.5% of the torsion keys, and 23.5% of the molecules had one or more unmatched torsion keys.

Algorithm 1.1

COSMOS(*Molecule*; *\$FragmentLibrary*; *\$TorsionLibrary*)

```

COSMOS(Molecule; $FragmentLibrary; $TorsionLibrary)
1. mol ← Molecule
2. mol ← mol - mol.GetAtoms()
3. mol ← mol - mol.GetBonds()
4. mol ← mol - mol.GetAngles()
5. mol ← mol - mol.GetTorsions()
6. mol ← mol - mol.GetChiralCenters()
7. mol ← mol - mol.GetMolChiralCenters()
8. mol ← mol - mol.GetMolChiralCenters()
9. mol ← mol - mol.GetMolChiralCenters()
10. mol ← mol - mol.GetMolChiralCenters()
11. mol ← mol - mol.GetMolChiralCenters()
12. mol ← mol - mol.GetMolChiralCenters()
13. mol ← mol - mol.GetMolChiralCenters()
14. mol ← mol - mol.GetMolChiralCenters()
15. mol ← mol - mol.GetMolChiralCenters()
16. mol ← mol - mol.GetMolChiralCenters()
17. mol ← mol - mol.GetMolChiralCenters()
18. mol ← mol - mol.GetMolChiralCenters()
19. mol ← mol - mol.GetMolChiralCenters()
20. mol ← mol - mol.GetMolChiralCenters()
21. mol ← mol - mol.GetMolChiralCenters()
22. mol ← mol - mol.GetMolChiralCenters()
23. mol ← mol - mol.GetMolChiralCenters()
24. mol ← mol - mol.GetMolChiralCenters()
25. mol ← mol - mol.GetMolChiralCenters()
26. mol ← mol - mol.GetMolChiralCenters()
27. mol ← mol - mol.GetMolChiralCenters()
28. mol ← mol - mol.GetMolChiralCenters()
29. mol ← mol - mol.GetMolChiralCenters()
30. mol ← mol - mol.GetMolChiralCenters()
31. mol ← mol - mol.GetMolChiralCenters()
32. mol ← mol - mol.GetMolChiralCenters()
33. mol ← mol - mol.GetMolChiralCenters()
34. mol ← mol - mol.GetMolChiralCenters()
35. mol ← mol - mol.GetMolChiralCenters()
36. mol ← mol - mol.GetMolChiralCenters()
37. mol ← mol - mol.GetMolChiralCenters()
38. mol ← mol - mol.GetMolChiralCenters()
39. mol ← mol - mol.GetMolChiralCenters()
40. mol ← mol - mol.GetMolChiralCenters()
41. mol ← mol - mol.GetMolChiralCenters()
42. mol ← mol - mol.GetMolChiralCenters()
43. mol ← mol - mol.GetMolChiralCenters()
44. mol ← mol - mol.GetMolChiralCenters()
45. mol ← mol - mol.GetMolChiralCenters()
46. mol ← mol - mol.GetMolChiralCenters()
47. mol ← mol - mol.GetMolChiralCenters()
48. mol ← mol - mol.GetMolChiralCenters()
49. mol ← mol - mol.GetMolChiralCenters()
50. mol ← mol - mol.GetMolChiralCenters()
51. mol ← mol - mol.GetMolChiralCenters()
52. mol ← mol - mol.GetMolChiralCenters()
53. mol ← mol - mol.GetMolChiralCenters()
54. mol ← mol - mol.GetMolChiralCenters()
55. mol ← mol - mol.GetMolChiralCenters()
56. mol ← mol - mol.GetMolChiralCenters()
57. mol ← mol - mol.GetMolChiralCenters()
58. mol ← mol - mol.GetMolChiralCenters()
59. mol ← mol - mol.GetMolChiralCenters()
60. mol ← mol - mol.GetMolChiralCenters()
61. mol ← mol - mol.GetMolChiralCenters()
62. mol ← mol - mol.GetMolChiralCenters()
63. mol ← mol - mol.GetMolChiralCenters()
64. mol ← mol - mol.GetMolChiralCenters()
65. mol ← mol - mol.GetMolChiralCenters()
66. mol ← mol - mol.GetMolChiralCenters()
67. mol ← mol - mol.GetMolChiralCenters()
68. mol ← mol - mol.GetMolChiralCenters()
69. mol ← mol - mol.GetMolChiralCenters()
70. mol ← mol - mol.GetMolChiralCenters()
71. mol ← mol - mol.GetMolChiralCenters()
72. mol ← mol - mol.GetMolChiralCenters()
73. mol ← mol - mol.GetMolChiralCenters()
74. mol ← mol - mol.GetMolChiralCenters()
75. mol ← mol - mol.GetMolChiralCenters()
76. mol ← mol - mol.GetMolChiralCenters()
77. mol ← mol - mol.GetMolChiralCenters()
78. mol ← mol - mol.GetMolChiralCenters()
79. mol ← mol - mol.GetMolChiralCenters()
80. mol ← mol - mol.GetMolChiralCenters()
81. mol ← mol - mol.GetMolChiralCenters()
82. mol ← mol - mol.GetMolChiralCenters()
83. mol ← mol - mol.GetMolChiralCenters()
84. mol ← mol - mol.GetMolChiralCenters()
85. mol ← mol - mol.GetMolChiralCenters()
86. mol ← mol - mol.GetMolChiralCenters()
87. mol ← mol - mol.GetMolChiralCenters()
88. mol ← mol - mol.GetMolChiralCenters()
89. mol ← mol - mol.GetMolChiralCenters()
90. mol ← mol - mol.GetMolChiralCenters()
91. mol ← mol - mol.GetMolChiralCenters()
92. mol ← mol - mol.GetMolChiralCenters()
93. mol ← mol - mol.GetMolChiralCenters()
94. mol ← mol - mol.GetMolChiralCenters()
95. mol ← mol - mol.GetMolChiralCenters()
96. mol ← mol - mol.GetMolChiralCenters()
97. mol ← mol - mol.GetMolChiralCenters()
98. mol ← mol - mol.GetMolChiralCenters()
99. mol ← mol - mol.GetMolChiralCenters()
100. mol ← mol - mol.GetMolChiralCenters()

```

Algorithm 1.2

BUILDRIGIDFRAGMENTS(*Fragment*; *\$FragmentLibrary*)

```

BUILDRIGIDFRAGMENTS(Fragment; $FragmentLibrary)
1. mol ← Fragment
2. mol ← mol - mol.GetAtoms()
3. mol ← mol - mol.GetBonds()
4. mol ← mol - mol.GetAngles()
5. mol ← mol - mol.GetTorsions()
6. mol ← mol - mol.GetChiralCenters()
7. mol ← mol - mol.GetMolChiralCenters()
8. mol ← mol - mol.GetMolChiralCenters()
9. mol ← mol - mol.GetMolChiralCenters()
10. mol ← mol - mol.GetMolChiralCenters()
11. mol ← mol - mol.GetMolChiralCenters()
12. mol ← mol - mol.GetMolChiralCenters()
13. mol ← mol - mol.GetMolChiralCenters()
14. mol ← mol - mol.GetMolChiralCenters()
15. mol ← mol - mol.GetMolChiralCenters()
16. mol ← mol - mol.GetMolChiralCenters()
17. mol ← mol - mol.GetMolChiralCenters()
18. mol ← mol - mol.GetMolChiralCenters()
19. mol ← mol - mol.GetMolChiralCenters()
20. mol ← mol - mol.GetMolChiralCenters()
21. mol ← mol - mol.GetMolChiralCenters()
22. mol ← mol - mol.GetMolChiralCenters()
23. mol ← mol - mol.GetMolChiralCenters()
24. mol ← mol - mol.GetMolChiralCenters()
25. mol ← mol - mol.GetMolChiralCenters()
26. mol ← mol - mol.GetMolChiralCenters()
27. mol ← mol - mol.GetMolChiralCenters()
28. mol ← mol - mol.GetMolChiralCenters()
29. mol ← mol - mol.GetMolChiralCenters()
30. mol ← mol - mol.GetMolChiralCenters()
31. mol ← mol - mol.GetMolChiralCenters()
32. mol ← mol - mol.GetMolChiralCenters()
33. mol ← mol - mol.GetMolChiralCenters()
34. mol ← mol - mol.GetMolChiralCenters()
35. mol ← mol - mol.GetMolChiralCenters()
36. mol ← mol - mol.GetMolChiralCenters()
37. mol ← mol - mol.GetMolChiralCenters()
38. mol ← mol - mol.GetMolChiralCenters()
39. mol ← mol - mol.GetMolChiralCenters()
40. mol ← mol - mol.GetMolChiralCenters()
41. mol ← mol - mol.GetMolChiralCenters()
42. mol ← mol - mol.GetMolChiralCenters()
43. mol ← mol - mol.GetMolChiralCenters()
44. mol ← mol - mol.GetMolChiralCenters()
45. mol ← mol - mol.GetMolChiralCenters()
46. mol ← mol - mol.GetMolChiralCenters()
47. mol ← mol - mol.GetMolChiralCenters()
48. mol ← mol - mol.GetMolChiralCenters()
49. mol ← mol - mol.GetMolChiralCenters()
50. mol ← mol - mol.GetMolChiralCenters()
51. mol ← mol - mol.GetMolChiralCenters()
52. mol ← mol - mol.GetMolChiralCenters()
53. mol ← mol - mol.GetMolChiralCenters()
54. mol ← mol - mol.GetMolChiralCenters()
55. mol ← mol - mol.GetMolChiralCenters()
56. mol ← mol - mol.GetMolChiralCenters()
57. mol ← mol - mol.GetMolChiralCenters()
58. mol ← mol - mol.GetMolChiralCenters()
59. mol ← mol - mol.GetMolChiralCenters()
60. mol ← mol - mol.GetMolChiralCenters()
61. mol ← mol - mol.GetMolChiralCenters()
62. mol ← mol - mol.GetMolChiralCenters()
63. mol ← mol - mol.GetMolChiralCenters()
64. mol ← mol - mol.GetMolChiralCenters()
65. mol ← mol - mol.GetMolChiralCenters()
66. mol ← mol - mol.GetMolChiralCenters()
67. mol ← mol - mol.GetMolChiralCenters()
68. mol ← mol - mol.GetMolChiralCenters()
69. mol ← mol - mol.GetMolChiralCenters()
70. mol ← mol - mol.GetMolChiralCenters()
71. mol ← mol - mol.GetMolChiralCenters()
72. mol ← mol - mol.GetMolChiralCenters()
73. mol ← mol - mol.GetMolChiralCenters()
74. mol ← mol - mol.GetMolChiralCenters()
75. mol ← mol - mol.GetMolChiralCenters()
76. mol ← mol - mol.GetMolChiralCenters()
77. mol ← mol - mol.GetMolChiralCenters()
78. mol ← mol - mol.GetMolChiralCenters()
79. mol ← mol - mol.GetMolChiralCenters()
80. mol ← mol - mol.GetMolChiralCenters()
81. mol ← mol - mol.GetMolChiralCenters()
82. mol ← mol - mol.GetMolChiralCenters()
83. mol ← mol - mol.GetMolChiralCenters()
84. mol ← mol - mol.GetMolChiralCenters()
85. mol ← mol - mol.GetMolChiralCenters()
86. mol ← mol - mol.GetMolChiralCenters()
87. mol ← mol - mol.GetMolChiralCenters()
88. mol ← mol - mol.GetMolChiralCenters()
89. mol ← mol - mol.GetMolChiralCenters()
90. mol ← mol - mol.GetMolChiralCenters()
91. mol ← mol - mol.GetMolChiralCenters()
92. mol ← mol - mol.GetMolChiralCenters()
93. mol ← mol - mol.GetMolChiralCenters()
94. mol ← mol - mol.GetMolChiralCenters()
95. mol ← mol - mol.GetMolChiralCenters()
96. mol ← mol - mol.GetMolChiralCenters()
97. mol ← mol - mol.GetMolChiralCenters()
98. mol ← mol - mol.GetMolChiralCenters()
99. mol ← mol - mol.GetMolChiralCenters()
100. mol ← mol - mol.GetMolChiralCenters()

```

Results

We have described a general framework for developing data-driven predictors for the 3D structures of small molecules. Here we describe results obtained with the current implementation where: (1) we use only a top-down approach; (2) fragment libraries are composed of rigid fragments and rings, extracted from the old CSD04; (3) only one prototype per fragment and a single point estimate for each torsion key value are used.

To evaluate COSMOS, we first selected molecules that were unique to CSD09 according to their SMILES. From this subset, we took a random sample of 10,000 molecules to use as a test set (CSD09Test). The set CSD09Test contains 3,597 organic and 6,403 organometallic molecules. Note that CSD09Test is a subset of all the molecules added to CSD after 2004, and that it does not contain any of the molecules that were used to build the fragment libraries.

Coverage, the percentage of molecules predicted among a large set, is the most straightforward assessment of a predictor's ability to generalize. The coverage results of COSMOS and CORINA are summarized in Table 5. COSMOS was able to return a prediction for 96.4% of the molecules with nearly perfect coverage (99.6%) of the organic subset. CORINA performed comparably on organic molecules, but predicted only 51.6% of the organometallic molecules. Of the molecules which CORINA was not able to produce models for, no attempt was made to predict for approximately half of them due to the presence of atoms with $CN > 6$, and the other half failed due to various errors that occurred during the attempted prediction. These results indicate that COSMOS represents a significant improvement in coverage of complex molecules.

To objectively assess COSMOS and CORINA we calculated the accuracy and speed using the subset of molecules in CSD09Test that both methods produced predictions for. This subset consists of 6,795 molecules and contains 52.1% organic molecules. The predicted structures of each method were compared in terms of (1) RMSD values between the generated structures and their corresponding experimental CSD counterparts; (2) percentage of predicted models that contain clashes; and (3) prediction time. The results are summarized in Table 6. When the entire set is considered the average accuracy and average speed of the two methods are quite similar. On the organic subset CORINA is slightly more accurate and it returns fewer models that contain clashes, but it is slower than COSMOS. On the organometallic subset, COSMOS is more accurate and produces fewer models containing atom clashes, but it is slower than CORINA.

In addition to the average accuracy and speed, Table 6 also provides the percentage of molecules with RMSD (1) lower than 0.3 Å; (2) lower than 1.0 Å. The first measure was used in¹⁵ as a valid upper limit to assess the accuracy of a prediction: if two structures have $RMSD < 0.3$ Å, they can be considered equivalent conformations. In our evaluation both COSMOS and CORINA fail to produce a high percentage of models below this stringent threshold. The threshold of 1.0 Å provides a more forgiving measure for considering a prediction successful. Using both thresholds, CORINA is slightly better on the organic subset and COSMOS is significantly better on the metal-organic subset. It is important to note that, in practice, RMSD is length dependent, thus these arbitrary cut-off values are only informative when considering large sets of predictions. For instance, for a target molecule of eight atoms an RMSD value of 1.0 Å would represent a very poor prediction, whereas a result of 1.0 Å for a molecule of 150 atoms would be quite good. To provide a sense of how the RMSD values distribute overall, Figure 10 compares the histograms of the RMSD values calculated from the predicted models of CORINA and COSMOS.

In Table 7 we also show the accuracy results for a reduced set of 4,710 molecules. This test set was obtained by removing from the set of Table 6 all the molecules with high (larger than 0.6) Tanimoto similarity with our "training set" (CSD04). As expected, COSMOS performs slightly worse on the redundancy reduced set; however, the results on this more difficult set still compare favorably with CORINA.

The experimentally determined structure of a small molecule in the CSD represents one low-energy conformation, but when the molecule has rotatable bonds it may have many conformations that are approximately equal energetically. This introduces some noise in the RMSD assessment protocol. Thus to further assess the two methods and remove a major source of ambiguity, we also compared their performance on rigid fragments in isolation. The accuracy and speed results of COSMOS and CORINA calculated on the first occurrence of each unique rigid fragment of CSD09Test are summarized in Table 8. Overall COSMOS performs better on the rigid fragment set than CORINA in terms of accuracy (0.530 Å vs 0.664 Å); however, CORINA is slightly more accurate on the organic rigid fragments, with a

clash percentage (1.3%) comparable to COSMOS (1.2%). COSMOS is more accurate on the organometallic fragments (0.660Å vs 0.896Å) and is faster (0.073s vs 0.097s) overall than CORINA.

Even when rigid segments are evaluated in isolation the RMSD is not the most appropriate measure for assessing some aspects of the models. To assess the correctness of the metal coordination geometries of the predicted models we use the mean angle difference ($\overline{\Delta\theta}$), described in detail in the Assessment section. Here the metric is applied to the mono-centric fragments of CSD09Test and we assess how often COSMOS and CORINA models have the correct metal coordination geometry. Note that mono-centric fragments which are extracted from metal-catenes structures and have incomplete metal coordination spheres are excluded from this assessment.

We initially chose $\overline{\Delta\theta} < 10.0^\circ$ as a threshold to classify a fragment geometry as correct based on the results from comparing the ideal templates using $\overline{\Delta\theta}$ (Table 3); however, when $\overline{\Delta\theta} < 10.0^\circ$ is used for classification 27.5% of the CSD fragments in the test set remain unclassified. When a more lenient threshold of $\overline{\Delta\theta} < 15.0^\circ$ is used only 11.2% of the CSD fragments remain unclassified. Thus, the evaluation was performed using both thresholds and the results, broken down by coordination number (CN) and specific geometry, are summarized in Table 9. Overall, the COSMOS model fragments have the correct metal coordination geometries more often than those of CORINA using $\overline{\Delta\theta} < 10.0^\circ$ (62.6% vs 54.0%) and $\overline{\Delta\theta} < 15.0^\circ$ (71.3% vs 64.7%). For both thresholds COSMOS is more accurate for CN=3, CN=4, and CN=5, while CORINA is more accurate for CN=6. When considering specific geometries, one result which stands out is that for CN=4 CORINA predicts none of the 648 square planar fragments correctly using $\overline{\Delta\theta} < 10.0^\circ$, while COSMOS predicts 42.7% correctly. One example of a rigid fragment where the COSMOS model correctly utilizes square planar geometry and the CORINA model incorrectly utilizes tetrahedral geometry is shown in Figure 11. Similarly, for CN=5 CORINA only predicts 3.7% of the square pyramidal fragments correctly using $\overline{\Delta\theta} < 10.0^\circ$ while COSMOS predicts 46.3% correctly. These results indicate that CORINA is biased to use certain metal coordination geometries and avoid others. While this type of bias may have a negative effect on accuracy in some cases, CORINA's bias towards using octahedral geometry does not hurt overall accuracy for CN=6 because it is the correct geometry for the overwhelming majority of the structures in the test set.

Overall COSMOS compares favorably with CORINA in terms of the correctness of the metal coordination geometries; however, these results demonstrate that both methods have room for improvement and they give a sense of the difficulty of predicting the correct metal geometry when there are multiple plausible geometries. Since CORINA does not handle CN>6, only the COSMOS results are presented for CN=7 and CN=8. For CN=7, 36.2% of the COSMOS model fragments have the correct geometry using $\overline{\Delta\theta} < 10.0^\circ$ and 43.3% are correct with $\overline{\Delta\theta} < 15.0^\circ$. For CN=8, 41.8% of the COSMOS model fragments are correct using $\overline{\Delta\theta} < 10.0^\circ$ and 56.6% are correct with $\overline{\Delta\theta} < 15.0^\circ$.

Conclusion

We have reviewed the field and described a general framework for the rapid generation of accurate 3D coordinates for small molecules that relies primarily on precomputed libraries of fragments, as opposed to rules or physics approximations. The power-law distributions of the rigid segments and torsion angle data found in current databases of small molecules indicates that a very good coverage of chemical space can be achieved with relatively small, pre-computed, libraries and that adding new fragment information will help to increase the

coverage of rarely occurring fragments and torsion angles. The framework was deployed and tested here using fragment parameters obtained only from the CSD04 to predict structures of novel molecules in the CSD09. In the benchmark tests using these fragments, the performance of COSMOS and CORINA were comparable, with a clear advantage for COSMOS on metal-organics. The current version of COSMOS is available for academic research through the ChemDB chemoinformatics web portal at: <http://cdb.ics.uci.edu/> as a web server.

While the current version of COSMOS provides an improvement over the current state-of-the-art, it does not yet fully achieve all the specifications that we believe to be desirable for the next generation of prediction systems. We believe that further improvements of COSMOS are possible in many ways, for instance by modifying the libraries and by developing more complete representations of both the fragment and the torsion angle distributions.

We also hope that future versions of COSMOS, or other similar systems, will be able to achieve the greater degree of data and software openness that is indispensable for real scientific progress in the field. One obstacle in this area may be the closed nature of the CSD which, unlike the PDB, cannot be used without severe restrictions, even for academic research purposes. This is yet another example of the unfortunate state of affairs in chemoinformatics where an overly zealous culture of closeness and secrecy, sometimes related to short term profits, have greatly hampered scientific progress. A culture shift towards open access to critical data, algorithms, and software derived from them—as is pervasively the case in physics or biology—is essential to improve the quality of the science and ultimately lead to technological advances of real economic value that benefit the largest number of taxpayers.

Acknowledgments

Work partly supported by NIH Biomedical Informatics Training grant (LM-07443-01), NSF MRI grant (EIA-0321390), and NSF grant IIS-0513376 to PB, by an NIH/NLM Pathway to Independence Award (K99LM010821) to AR, and by the UCI Institute for Genomics and Bioinformatics. We thank Ramzi Nasr and Jordan Hayes for helping with the calculation of the Tanimoto coefficients and with the setup of the web server. We would like also to acknowledge the OpenBabel project and OpenEye Scientific Software for their free software academic license.

References

1. Irwin JJ, Shoichet BK. ZINC—A Free Database of Commercially Available Compounds for Virtual Screening. *J. Chem. Inf. Comput. Sci.* 2005; 45:177–182.
2. Chen J, Swamidass SJ, Dou Y, Bruand J, Baldi P. ChemDB: a public database of small molecules and related chemoinformatics resources. *Bioinformatics.* 2005; 21:4133–4139. [PubMed: 16174682]
3. Chen J, Linstead E, Swamidass SJ, Wang D, Baldi P. ChemDB Update—Full Text Search and Virtual Chemical Space. *Bioinformatics.* 2007; 23:2348–2351. [PubMed: 17599932]
4. Wang Y, Xiao J, Suzek T, Zhang J, Wang J, Bryant S. PubChem: a public information system for analyzing bioactivities of small molecules. *Nucleic Acids Res.* 2009; 37:W623–W633. [PubMed: 19498078]
5. Lauri G, Bartlett P. CAVEAT: A program to facilitate the design of organic molecules. *J. Comput.-Aided Mol. Des.* 1994; 8:51–66. [PubMed: 8035213]
6. Taylor R. Life-science applications of the Cambridge Structural Database. *Acta Crystallogr., Sect. D: Biol. Crystallogr.* 2002; 58:879–888. [PubMed: 12037325]
7. Dobson CM. Chemical Space and Biology. *Nature.* 2004; 432:824–828. [PubMed: 15602547]
8. Allen F. The Cambridge Structural Database: a quarter of a million crystal structures and rising. *Acta Crystallogr., Sect. B: Struct. Sci.* 2002; 58:380–388.

9. Allen F, Motherwell W. Applications of the Cambridge Structural Database in organic chemistry and crystal chemistry. *Acta Crystallogr., Sect. B: Struct. Sci.* 2002; 58:407–422.
10. Sadowski J, Gasteiger J. From atoms and bonds to three-dimensional atomic coordinates: automatic model builders. *Chem. Rev.* 1993; 93:2567–2581.
11. Frisch, MJ., et al. Gaussian 03, Revision C.02. Wallingford, CT: Gaussian, Inc.; 2004.
12. Brooks B, Bruccoleri RE, Olafson BD, States DJ, Swaminathan S, Karplus M. CHARMM: A program for macromolecular energy, minimization, and dynamics calculations. *J. Comput. Chem.* 1983; 4:187–217.
13. Kalé L, Skeel R, Bhandarkar M, Brunner R, Gursoy A, Krawetz N, Phillips J, Shinozaki A, Varadarajan K, Schulten K. NAMD2: Greater Scalability for Parallel Molecular Dynamics. *J. Comput. Phys.* 1999; 151:283–312.
14. Case D, CHEATHAM T III, Darden T, Gohlke H, Luo R, Merz K Jr, Onufriev A, Simmerling C, Wang B, Woods R. The Amber biomolecular simulation programs. *J. Comput. Chem.* 2005; 26:1668–1688. [PubMed: 16200636]
15. Sadowski J, Gasteiger J, Klebe G. Comparison of Automatic Three-Dimensional Model Builders Using 639 X-ray Structures. *J. Chem. Inf. Model.* 1994; 34:1000–1008.
16. Berman HM, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, Shindyalov IN, Bourne PE. The Protein Data Bank. *Nucleic Acids Res.* 2000; 28:235–242. [PubMed: 10592235]
17. Simons K, Kooperberg C, Huang E, Baker D. Assembly of protein tertiary structures from fragments with similar local sequences using simulated annealing and Bayesian scoring functions. *J. Mol. Biol.* 1997; 268:209–225. [PubMed: 9149153]
18. Chivian D, Kim D, Malmstrom L, Bradley P, Robertson T, Murphy P, Strauss C, Bonneau R, Rohl C, Baker D. Automated prediction of CASP-5 structures using the Robetta server. *Proteins.* 2003; 53:524–533. [PubMed: 14579342]
19. Zhang Y, Kolinski A, Skolnick J. TOUCHSTONE II: a new approach to ab initio protein structure prediction. *Biophys. J.* 2003; 85:1145–1164. [PubMed: 12885659]
20. Wishart D, Knox C, Guo A, Shrivastava S, Hassanali M, Stothard P, Chang Z, Woolsey J. DrugBank: a comprehensive resource for in silico drug discovery and exploration. *Nucleic Acids Res.* 2006; 34:D668–D672. [PubMed: 16381955]
21. Wishart D, Knox C, Guo A, Cheng D, Shrivastava S, Tzur D, Gautam B, Hassanali M. DrugBank: a knowledgebase for drugs, drug actions and drug targets. *Nucleic Acids Res.* 2008; 36(Database Issue):D901–D906. [PubMed: 18048412]
22. Tabassum S, Pettinari C. Chemical and biotechnological developments in organotin cancer chemotherapy. *J. Organomet. Chem.* 2006; 691:1761–1766.
23. Berger I, et al. In Vitro Anticancer Activity and Biologically Relevant Metabolization of Organometallic Ruthenium Complexes with Carbohydrate-Based Ligands. *Chem.–Eur. J.* 2008; 14:9046–9057.
24. Pearlman D, Case D, Caldwell J, Ross W, Cheatham T, DeBolt S, Ferguson D, Seibel G, Kollman P. AMBER, a package of computer programs for applying molecular mechanics, normal mode analysis, molecular dynamics and free energy calculations to simulate the structural and energetic properties of molecules. *Comput. Phys. Commun.* 1995; 91:1–41.
25. Crippen GM. Rapid Calculation of Coordinates from Distance Matrices. *J. Comput. Phys.* 1978; 26:449–452.
26. Crippen, GM.; Havel, TF. Distance Geometry and Molecular Conformations. New York: Wiley; 1988. Chemometrics Series
27. Pearlman R. Rapid generation of high quality approximate 3D molecular structures. *Chem. Des. Auto. News.* 1987; 2:1–6.
28. Gasteiger J, Schuur J, Selzer P, Steinhauer L, Steinhauer V. Finding the 3D structure of a molecule in its IR spectrum. *Fresenius J. Anal. Chem.* 1997; 359:50–55.
29. Xu H, Izrailev S, Agrafiotis DK. Conformational Sampling by Self-Organization. *J. Chem. Inf. Model.* 2003; 43:1186–1191.
30. Leite TB, Gomes D, Miteva MA, Chomilier J, Villoutreix BO, Tuffery P. Frog: a FFree Online druG 3D conformation generator. *Nucleic Acids Res.* 2007; 35:W568–W572. [PubMed: 17485475]

31. Pearlman, RS. 3D QSAR in Drug Design. Kubinyi, H., editor. Leiden: ESCOM; 1993. p. 41-79.
32. Sayers E, et al. Database resources of the National Center for Biotechnology Information. *Nucleic Acids Res.* 2010; 38:D5–D16. [PubMed: 19910364]
33. Guha R, Howard MT, Hutchison GR, Murray-Rust P, Rzepa H, Steinbeck C, Wegner JK, Willighagen EL. The Blue Obelisk–Interoperability in Chemical Informatics. *J. Chem. Inf. Model.* 2006; 46:991–998. [PubMed: 16711717]
34. Agrafiotis DK, Xu H. A self-organizing principle for learning nonlinear manifolds. *Proc. Natl. Acad. Sci. U. S. A.* 2002; 99:15869–15872. [PubMed: 12444256]
35. Gasteiger, J.; Engel, T., editors. *Chemoinformatics: A Textbook*. Wiley; 2003.
36. Leach, AR.; Gillet, VJ. *An Introduction to Chemoinformatics*. Dordrecht, The Netherlands: Springer; 2005.
37. Weininger D, Weininger A, Weininger J. SMILES, A Chemical Language and Information System. 1. Introduction to Methodology and Encoding Rules. *J. Chem. Inf. Comput. Sci.* 1988; 28:31–36.
38. Weininger D, Weininger A, Weininger J. SMILES. 2. Algorithm for Generation of Unique SMILES notation. *J. Chem. Inf. Comput. Sci.* 1989; 29:97–101.
39. Weininger D. SMILES. 3. DEPICT. Graphical depiction of chemical structures. *J. Chem. Inf. Comput. Sci.* 1990; 30:237–243.
40. Wheeler D, et al. Database resources of the National Center for Biotechnology Information. *Nucleic Acids Res.* 2007; 35:D5–D12. [PubMed: 17170002]
41. Lang P, Brozell S, Mukherjee S, Pettersen E, Meng E, Thomas V, Rizzo R, Case D, James T, Kuntz I. DOCK 6: Combining techniques to model RNA–small molecule complexes. *RNA.* 2009; 15:1219–1230. [PubMed: 19369428]
42. Baldi P, Benz RW, Hirschberg D, Swamidass S. Lossless Compression of Chemical Fingerprints Using Integer Entropy Codes Improves Storage and Retrieval. *J. Chem. Inf. Model.* 2007; 47:2098–2109. [PubMed: 17967006]
43. Rogers D, Hahn M. Extended-Connectivity Fingerprints. *J. Chem. Inf. Model.* 2010; 50:742–754. PMID: 20426451. [PubMed: 20426451]
44. Gasteiger J. Chemoinformatics: a new field with a long tradition. *Anal. Bioanal. Chem.* 2006; 384:57–64. [PubMed: 16177914]
45. Crabtree, RH. *The Organometallic Chemistry of the Transition Metals*. New York, NY: Wiley; 1988. p. 32-33.
46. Benz RW, Swamidass SJ, Baldi P. Discovery of Power-Laws in Chemical Space. *J. Chem. Inf. Model.* 2008; 48:1138–1151. [PubMed: 18522387]
47. Feuston B, Miller M, Culberson J, Nachbar R, Kearsley S. Comparison of knowledge-based and distance geometry approaches for generation of molecular conformations. *J. Chem. Inf. Comput. Sci.* 2001; 41:754–763. [PubMed: 11410056]
48. Bondi A. Van der Waals Volumes and Radii. *J. Phys. Chem.* 1964; 68:441–451.
49. Sadowski, J. *Molecular Drug Properties: Measurement and Prediction*. Mannhold, R., editor. Weinheim, Germany: Wiley-VCH Verlag GmbH & Co. KGaA; 2008. p. 157-181.
50. Wipke W, Hahn M. Analogy in computer-assisted design. *Prog. Clin. Biol. Res.* 1989; 291:141–145. [PubMed: 2726836]
51. Cordero B, Gómez V, Platero-Prats A, Revés M, Echeverría J, Cremades E, Barragán F, Alvarez S. Covalent radii revisited. *Dalton Trans.* 2008; 21:2832–2838. [PubMed: 18478144]
52. Halgren T. Merck molecular force field. I. Basis, form, scope, parameterization, and performance of MMFF94. *J. Comput. Chem.* 1996; 17:490–519.

Figure 1.
CSD Molecules Growth over four decades. Green = organics. Red = metal-organics.

Figure 2.

Rigid Fragment Structural Ambiguities. The rigid fragments in **a** and **b** can be distinguished in their SMILES by including stereochemistry information. **a**: Stereo SMILES: Cc1ccc-2c(c1)[C@]3(c4c2cccc4)c5ccccc5-c6c3cc(cc6)C. **b**: Stereo SMILES: Cc1ccc-2c(c1)[C@@]3(c4c2cccc4)c5ccccc5-c6c3cc(cc6)C. If the stereochemistry of the central tetrahedral atom is not specified, then the SMILES calculated from the rigid fragments in both **a** and **b** are identical and equal to: Cc1ccc-2c(c1)C3(c4c2cccc4)c5ccccc5-c6c3cc(cc6)C. However, the RMSD between the two structures is 3.46Å. The fully specified SMILES for the rigid fragments in **c** and **d** is Cc1cc2cc(c1CCc3ccc(c(c3)Br)CC2)Br. The rigid fragment in **c** is the enantiomeric structure of the rigid fragment in **d**. The RMSD between the two structures is 2.52Å.

Figure 3.

Fragmentation at Rotatable Bonds. Small molecule structures can be divided into independent fragments at each rotatable bond. **a:** The rotatable bond is highlighted and the four atoms used to define the torsion angle are labeled with capital letters. **b:** The same molecule rotated 90° degrees on the vertical axis. **c:** Schematic representation with additional atoms bound to the torsion bond atoms labeled with lower-case letters. These atoms, in addition to the torsion atoms, are used to define the torsion key SMILES. **d:** Resulting fragments and torsion key SMILES. The original SMILES string for the molecules is: c1ccc2c(c1)ccc(c2c3c4ccccc4ccc3C(=O)O)C(=O)O.

Figure 4.

Distributions of rigid segments, cyclic fragments, and torsion keys on a log-log scale. The distributions were derived by identifying the unique SMILES for rigid segments, cyclic fragments, and torsion keys, then counting the number of occurrences of each. The probability of each item is plotted versus its frequency rank, such that a rank of 1 corresponds to the most frequently occurring item. The calculated power-law distribution exponents are indicated by $|\alpha|$. **a:** all rigid segments, $|\alpha| = 1.8$. **b:** organic rigid segments, $|\alpha| = 1.7$; organometallic rigid segments, $|\alpha| = 1.9$. **c:** all cyclic fragments, $|\alpha| = 1.9$. **d:** organic cyclic fragments, $|\alpha| = 1.7$; organometallic cyclic fragments, $|\alpha| = 2.3$. **e:** torsion keys, $|\alpha| = 1.9$.

Figure 5.

Torsion angle bins for the torsion key SMILES: [cH]c([s])c([cH])[s]. This torsion key SMILES occurs 485 times in CSD09. **a:** schematic and SMILES representations of torsion key. **b:** counts of the observed torsion angles in 10° bins. The center of the most populated bin, from 175° to -175°, is 180°.

Figure 6.

Cyclic Fragmentation. **a:** Starting rigid fragment (caffeine). **b:** Terminal atoms are removed from the molecule (caffeine) leaving only the complete ring system (purine). **c:** Subrings are extracted from the complete ring system (imidazole and pyrimidine). Using this procedure, three unique cyclic structures (purine, imidazole, and pyrimidine) can be added to the cyclic fragment library.

Figure 7. Cumulative probability distribution of overlap factors in CSD09. The black, green, and red curves represent the entire CSD, the subset of organic molecules, and the subset of metal-organic molecules, respectively.

Figure 8. Simple Fuzzy Matching. Matching the porphyrin ring system of Chlorophyll A with the analogous porphyrin ring system of Heme B using the SMARTS pattern language. The asterisk represents a wildcard in SMARTS, so it will match any atom at that position.

Figure 9. Fragmenting and Reconnecting Subfragments of Cyclic Fragments. The process of breaking down a cyclic fragment into subfragments at simple connections points, matching the subfragments in the library, and reconnecting the fragments, is shown for two types of connection points. **a:** Simple tetrahedral joins occur when a single atom is the only connection point between two subfragments of a rigid fragment. These connection points are assembled according to ideal tetrahedral geometry of the connecting tetrahedral atom (circled in red). **b:** Simple planar joins occur when two atoms join subfragments, and these atoms plus all of their first bonded neighbors are in the same plane. Planar connections are assembled according to ideal planar geometry. The example shows a molecule with two planar joins, the first with connection atoms circled in red and green, and the second with connection atoms circled in orange and blue.

Figure 10.
Histograms of the RMSD values obtained from the 3D structures predicted with COSMOS (green) and with CORINA (red).

Figure 11.

Square planar versus tetrahedral modeling of a metal-organic rigid fragment. **a:** Rigid fragment from the CSD with square planar metal coordination geometry. **b:** CSD rigid fragment rotated 90° degrees on the vertical axis. **c:** Model predicted by COSMOS with correct square planar metal coordination geometry. **d:** COSMOS model rotated 90° degrees on the vertical axis. **e:** Model predicted by CORINA with incorrect tetrahedral metal coordination geometry. **f:** CORINA model rotated 90° degrees on the vertical axis.

Table 1

Comparison of four programs using a random sample of 10,000 molecules from CSD09. Time represents the average prediction time per molecule in seconds, computed over all molecules for which the program outputs a prediction. Times shown in parenthesis represent the average prediction times per molecule, including all molecules for which a prediction is attempted. Accurate prediction time for FROG is not available (NA) because it is only available as a Web server. Coverage (Cov.) is expressed as a percentage of the total number of molecules in the corresponding class.

	All (10,000)		Organic (3,444)		Metal-Organic (6,556)	
	Time (s)	Cov. (%)	Time (s)	Cov. (%)	Time (s)	Cov. (%)
DG	14.6	100.0	8.16	100.0	18.0	100.0
FROG	NA	30.8	NA	86.1	NA	1.7
Open Babel	28.4	31.1	27.5	85.7	45.1	2.4
CORINA	0.04 (0.20)	66.4	0.02 (0.05)	98.5	0.07 (0.31)	49.6

Table 2

Comparison of the programs reported in Table 1 using the common subset of molecules for which all the four methods were able to generate 3D structures. Average root mean square deviation (RMSD) per molecule is in Angstroms. The clash percentage column reports the percentage of returned models containing atom clashes. Time represents the average prediction time per molecule in seconds. Accurate prediction time for FROG is not available (NA) because it is only available as a Web server.

	RMSD (Å)	Clash %	Time (s)
Common Set (2,490 Molecules)			
DG	1.79	10.0	0.21
FROG	1.32	0.6	NA
Open Babel	1.20	0.8	23.9
CORINA	0.95	2.8	0.01
Organic (2,443 Molecules)			
DG	1.79	10.2	0.21
FROG	1.31	0.6	NA
Open Babel	1.20	0.9	23.7
CORINA	0.95	2.8	0.01
Metal-Organic (47 Molecules)			
DG	1.88	2.1	0.28
FROG	1.42	4.3	NA
Open Babel	1.34	0.0	34.2
CORINA	1.18	2.1	0.01

Table 3

Comparison of the metal geometry templates using the mean angle difference ($\overline{\Delta\theta}$) assessment measure.

	Mean Angle Difference ($\overline{\Delta\theta}$)
Coordination Number 3	
T-Shaped vs Trigonal Pyramidal:	35.7°
T-Shaped vs Trigonal Planar:	40.0°
Trigonal Pyramidal vs Trigonal Planar:	13.0°
Coordination Number 4	
Square Planar vs Tetrahedral	36.5°
Coordination Number 5	
Square Pyramidal vs Trigonal Bipyramidal	12.0°
Coordination Number 6	
Octahedral vs Trigonal Prismatic	22.5°
Coordination Number 7	
Capped Octahedron vs Capped Trigonal Prismatic	21.6°
Capped Octahedron vs Pentagonal Bipyramidal	14.6°
Capped Trigonal Prismatic vs Pentagonal Bipyramidal	14.4°
Coordination Number 8	
Cubic vs Dodecahedral	19.0°
Cubic vs Square Antiprismatic	15.6°
Dodecahedral vs Square Antiprismatic	8.0°

Table 4

Ten most frequently occurring organic and metal-organic rigid fragments. The percentage of total calculation considers all organic and metal-organic rigid fragments, and the total number of rigid fragments is 2,273,659.

Top Ten Organic Rigid Fragments			Top Ten Metal-Organic Rigid Fragments		
SMILES	2D	% of Total	SMILES	2D	% of Total
<chem>c1ccc(cc1)P</chem>		7.13%	<chem>C[Si](C)(C)C</chem>		0.83%
<chem>CCC</chem>		6.23%	<chem>C[Si](C)(C)N</chem>		0.49%
<chem>COC</chem>		5.00%	<chem>CC(C)C[SiH3]</chem>		0.37%

Top Ten Organic Rigid Fragments			Top Ten Metal-Organic Rigid Fragments		
SMILES	2D	% of Total	SMILES	2D	% of Total
<chem>Cc1ccccc1</chem>		4.31%	<chem>c1cccc(cc1)[SHH3]</chem>		0.36%
<chem>CC(=O)O</chem>		2.92%	<chem>CP(C)(C)[Pt]</chem>		0.23%
<chem>CCO</chem>		2.73%	<chem>CO[Cu]</chem>		0.22%

J Chem Inf Model. Author manuscript; available in PMC 2012 April 25.

Table 5

Coverage of COSMOS compared to CORINA. The percentages refer to the set of 10,000 molecules randomly extracted from CSD09 used in our tests.

	COSMOS	CORINA
% of Molecules Predicted	96.4	68.5
% of Organic Molecules Predicted	99.6	98.5
% of Organometallic Molecules Predicted	94.6	51.6

Table 6

Comparison of COSMOS and CORINA in terms of speed (mean time per prediction) and accuracy (RMSD value and percentage of molecules with atom clashes). The results are broken down by set of molecules: (1) molecules where both methods generated a 3D structure; (2) subset of organic molecules; and (3) subset of organometallic molecules. The CSD column reports the percentage of original CSD molecules with atom clashes.

	COSMOS	CORINA	CSD
Common Set (6,795 Molecules)			
Time (s)	0.151	0.134	
RMSD (Å)	1.566	1.600	
Clash %	14.1	18.0	0.2
% of Molecules with RMSD < 0.3	12.2	10.3	
% of Molecules with RMSD < 1.0	37.0	35.4	
Organic (3,538 Molecules)			
Time (s)	0.099	0.182	
RMSD (Å)	1.258	1.133	
Clash %	6.6	4.0	0.1
% of Molecules with RMSD < 0.3	17.4	18.5	
% of Molecules with RMSD < 1.0	46.4	54.4	
Metal-Organic (3,257 Molecules)			
Time (s)	0.208	0.082	
RMSD (Å)	1.899	2.108	
Clash %	22.2	33.3	0.3
% of Molecules with RMSD < 0.3	6.6	1.4	
% of Molecules with RMSD < 1.0	26.8	14.7	

Table 7

Comparison of COSMOS and CORINA for the reduced set of molecules with low similarity with molecules in the “training set” of COSMOS (CSD04). The CSD column reports the percentage of original CSD molecules with atom clashes. The results are broken down by set of molecules as in Table 6.

	COSMOS	CORINA	CSD
Common Set (4,710 Molecules)			
RMSD (Å)	1.729	1.728	
Clash %	17.5	20.9	0.3
Organic (2,163 Molecules)			
RMSD (Å)	1.384	1.207	
Clash %	8.4	4.6	0.1
Metal-Organic (2,547 Molecules)			
RMSD (Å)	2.022	2.171	
Clash %	25.2	34.7	0.4

Table 8

Comparison of COSMOS and CORINA in terms of speed (mean time per prediction) and accuracy (RMSD value and percentage of molecules with atom clashes) for rigid segments only. The results are broken down by set of fragments: (1) entire set of rigid segments; (2) subset of organic rigid segments; and (3) subset of organometallic rigid segments. The CSD column reports the percentage of original CSD rigid segments with atom clashes.

	COSMOS	CORINA	CSD
Common Set (9,027 Rigid Segments)			
Time (s)	0.073	0.097	
RMSD (Å)	0.530	0.664	
Clash %	5.2	6.0	0.1
Organic (3,489 Rigid Segments)			
Time (s)	0.070	0.184	
RMSD (Å)	0.323	0.296	
Clash %	1.2	1.3	0.0
Metal-Organic (5,538 Rigid Segments)			
Time (s)	0.075	0.042	
RMSD (Å)	0.660	0.896	
Clash %	7.6	9.0	0.1

Table 9

Comparison of COSMOS and CORINA in terms of correctness of metal coordination geometries using the mono-centric fragments extracted from CSD09Test. The CSD column indicates the number of mono-centric fragments of each geometry, and the COSMOS and CORINA columns indicate the percentage of the corresponding predicted model fragments with the correct geometry. The criteria used to classify the fragment geometries are described in the Assessment section. For $CN=7$ and $CN=8$ no results are available because CORINA does not handle $CN>6$.

	$\overline{\Delta\theta}<10.0^\circ$			$\overline{\Delta\theta}<15.0^\circ$		
	CSD	COSMOS	CORINA	CSD	COSMOS	CORINA
Coordination Number 3						
Trigonal Planar	111	70.3%	79.3%	133	77.4%	82.0%
T-Shaped	41	26.8%	17.1%	61	41.0%	14.8%
Trigonal Pyramidal	73	64.4%	41.1%	99	64.6%	46.5%
Unclassified	190	57.4%	22.1%	122	73.8%	51.6%
Total	415	59.0%	40.2%	415	68.0%	54.7%
Coordination Number 4						
Square Planar	647	42.7%	0.0%	698	48.7%	1.0%
Tetrahedral	872	86.7%	86.8%	1061	88.7%	93.7%
Unclassified	483	45.5%	25.7%	243	61.3%	39.9%
Total	2,002	62.5%	44.0%	2,002	71.4%	54.8%
Coordination Number 5						
Square Planar	294	46.3%	3.7%	369	50.9%	6.5%
Tetrahedral	214	52.6%	80.4%	363	55.1%	86.8%
Unclassified	319	33.9%	25.1%	95	56.8%	38.9%
Total	827	43.2%	31.8%	827	53.4%	45.5%
Coordination Number 6						
Octahedral	1,212	82.4%	94.4%	1,406	85.6%	97.7%
Trigonal Prismatic	16	25.0%	0.0%	71	26.8%	2.8%
Unclassified	308	44.5%	44.5%	59	55.9%	45.8%
Total	1,536	74.2%	83.4%	1,536	81.7%	91.3%
Coordination Number 7						
Capped Octahedron	33	9.1%	NA	213	48.4%	NA
Capped Trigonal Prism.	11	27.3%	NA	32	40.6%	NA

	$\overline{\Delta\theta} < 10.0^\circ$			$\overline{\Delta\theta} < 15.0^\circ$		
	CSD	COSMOS	CORINA	CSD	COSMOS	CORINA
Pentagonal Bipyramidal	135	25.9%	NA	178	33.7%	NA
Unclassified	372	42.7%	NA	128	49.2%	NA
<i>Total</i>	551	36.3%	NA	552	43.4%	NA
Coordination Number 8						
Cubic	7	57.1%	NA	22	45.5%	NA
Dodecahedral	164	30.5%	NA	194	53.1%	NA
Square Antiprismatic	124	32.3%	NA	137	35.8%	NA
Unclassified	388	49.5%	NA	330	67.9%	NA
<i>Total</i>	683	41.9%	NA	683	56.5%	NA