Available online at www.sciencedirect.com
ScienceDirect
Journal of Molecular Graphics and Modelling 25 (2007) 700–710
Journal of Molecular Graphics and Modelling
www.elsevier.com/locate/JMGM

Analyzing the performance of conformational search programs on compound databases

Markus J. Loferer [a], István Kolossváry [b,1], András Aszódi [a,*]

[a] *In Silico Sciences Unit, Informatics and Knowledge Management, Novartis Institutes for Biomedical Research Vienna, Brunnerstrasse 59, A-1235 Vienna, Austria*

[b] *Budapest University of Technology and Economics, Faculty of Chemical Engineering, Department of Chemical Information Technology, Szt. Gellért tér 4, H-1111 Budapest, Hungary*

Received 11 January 2006; received in revised form 17 May 2006; accepted 18 May 2006
Available online 26 May 2006

Abstract

We have studied the sampling performance of conformational search programs using geometric and energetic criteria. Ideally, a conformational search algorithm should identify the largest possible number of low-energy structures (energy criterion) covering the widest possible range of molecular shapes (geometric criterion). Geometric analysis consisted in comparing the distribution of conformations within the generated ensembles by multidimensional scaling and by analysing the eigenvalue structure of the pairwise coordinate covariance matrices. The energetic comparison was carried out by assessing the energy distribution of conformers after minimizing them all using the same semi-empirical quantum mechanics optimization protocol. The present investigation focused on five conformational search programs: DGEOM, QXP, ROTATE, LMOD and OMEGA. We have applied these methodologies to a maximally diverse 604-compound subset of the LeadQuest library. The program LMOD performs best according to the energetic criterion, whereas a wider range of geometrically diverse conformations is sampled by the other programs, at the cost of higher median conformer energies. In terms of speed, OMEGA is fastest. We recommend the use of LMOD or OMEGA for high-quality conformational search applications.

Keywords: Conformational search; Conformer generation; Sampling efficiency; Conformational ensemble; Compound databases

1. Introduction

The fast and reliable calculation of low-energy conformer ensembles are crucially important in a number of computational chemistry applications, such as pharmacophore searches, building protein models and the analysis of ligand–receptor interactions. Exploring the conformational landscape of a molecule is thus very important for pharmaceutical research. Considerable efforts have therefore been invested in designing efficient conformational search algorithms.

In the following we give only a very short high-level overview of the techniques applied in this important field, and advise the interested reader to consult reviews, e.g. [1] for further details.

Most conformational search methods can be classified into one of the major categories of *stochastic* and *deterministic* sampling approaches, or into a hybrid class that combines deterministic and stochastic elements in the algorithm.

Deterministic approaches attempt to enumerate the full set of low-energy conformations by performing a systematic conformer space search. A classical implementation is the ''brute-force'' torsion angle search whereby all rotatable bonds in a molecule are rotated by a prescribed angle until all possible combinations have been tried. Due to the exponentially growing number of torsion angle combinations (known as the ''combinatorial explosion'' phenomenon) this approach is feasible only for very small molecules, and all practical implementations resort to various heuristics to restrain the space of accessible conformations. These heuristics include pruning techniques whereby branches of the search path are excluded (taboo search, branch-and-bound algorithms, etc.), or

* Corresponding author. Tel.: +43 1 86634 452; fax: +43 1 86634 727.
E-mail address: aszodi@inode.at (A. Aszódi).
[1] Present address: D.E. Shaw Research and Development, LLC, 120 West 45th Street, New York, NY 10036, USA.

1093-3263/$ – see front matter
doi:10.1016/j.jmgm.2006.05.008

the use of predefined libraries of small structural fragments generated by optimization [2] or obtained from experimental data (''knowledge-based'' searches).

Stochastic sampling methods explore the conformation space by incorporating a random element into the search process. Perhaps the most important methods of this class are the various Monte Carlo-type (MC) simulations and genetic algorithms. Mode-following and ''minimum jumping'' approaches also implement a random walk on the molecular energy hypersurface. Obviously there is no guarantee that all low-energy conformers will be identified within a finite simulation run.

Finally, hybrid approaches include molecular dynamics (MD) simulations where the random element is represented by the initial distribution of the atomic velocity vectors, and the MC/MD combination, where each MC step, designed to improve sampling quality by forcing ''jumps'' in conformational space, is followed by a short deterministic MD run. Distance-geometry techniques, where a random set of interatomic distances is selected in such a way that a prescribed set of distance restraints are satisfied, also belong to this category because once the full distance matrix is specified the creation of three-dimensional coordinates (''embedding'') proceeds in a deterministic fashion.

Despite all these efforts, the conformational search problem still remains challenging for complex molecules containing a large number of rotatable bonds and flexible rings. In practical modelling work, one often encounters the situation where in the words of Allinger ''The number of conformations becomes so large that a complete analysis becomes very laborious. The results depend not so much on the force field as on the intuition of the person doing the calculation and which starting geometries were used for the energy minimizations'' [3].

Given this state of affairs, it is highly desirable to conduct systematic investigations of sampling properties, but these seem to be rather thinly spread in literature. Studies mainly focus on the retrieval of bioactive conformers [4,5] or on the test of few molecules with defined energy hypersurfaces [6–9]. Although these studies provide useful insights and suggest reasonable approaches to overcome the limitations, to our knowledge no attempt has yet been undertaken to examine systematically the *overall* performance of conformational search algorithms when they are applied to large collections of pharmaceutically relevant molecules. In this report we set out to carry out such an analysis, comparing the sampling performance of several conformational search algorithms.

Our analysis focuses on the geometric, energetic and speed aspects of conformation generation. The geometric assessment was concerned with the extent and uniformity of sampling. We define ''extent'' as the size of conformer space occupied by the generated conformers, whereas ''uniformity'' measures the unbiasedness of the sampling, i.e. whether there is a certain ''preference'' towards subregions of the sampled region of conformational space. The required statistical methods are explained in detail in the next section. First, generalized procrustes analysis (GPA), a method for optimal conformer superimposition is presented. Distances between the superimposed conformers are calculated and multidimensional scaling (MDS) is used to embed the conformers as points in an abstract conformation space where the extent of sampling can be measured. To assess uniformity, the covariances between the superimposed conformer structures are examined. The energetic analysis is based on the distribution of conformer energies. Thus, an assignment of force-field parameters to all molecules was required and statistical analysis on median energies obtained from all conformers of a molecule performed.

Although in physical reality the molecular geometries and energies mutually determine each other, the geometric and energetic criteria describing algorithmically generated conformer ensembles are not completely ''correlated'' so we used them separately to assess the conformation search programs in an unbiased manner. To give a theoretical example: a systematic torsion grid search without bump checks would generate an ensemble that samples the conformational space perfectly but the energy distribution would be totally unacceptable; on the other hand a perfect global minimizer would find the lowest energy possible with zero-extent sampling (assuming one global minimum). All real programs implement better or worse compromises between sampling extent and energy distribution and therefore it is justified to assess both aspects separately.

2. Methods and computational procedure

2.1. Theory

We investigated the performance of $N_P = 5$ conformation search programs applied in batch mode to a database containing $N_M = 604$ molecules. Each algorithm generated a set of conformers for each molecule in the database, whereby the number of conformers $N_{conf}(p, m)$ was in general different for each conformational ensemble of molecule m generated by program p for reasons explained below. The total size of our raw dataset containing N_{total} conformations was thus

$$N_{total} = \sum_{p=1}^{N_P} \sum_{m=1}^{N_M} N_{conf}(p, m)$$

and we carried out analyses across programs and/or compounds. On average 921 conformers were generated per molecule, which resulted in the production of approximately 2.8 million conformers for all molecules and programs. The relationships of the conformer sets are best explained on a schematic figure (Fig. 1).

It is by no means straightforward to come up with a scientifically sound and quantitative measure of conformational search performance. Ideally, a good algorithm is expected to find all physically correct low-energy conformers in a reasonable amount of time. This qualitative requirement can be decomposed into the following factors:

1. *Sampling*: A good algorithm is expected to sample efficiently, i.e. to visit the largest possible hypervolume in

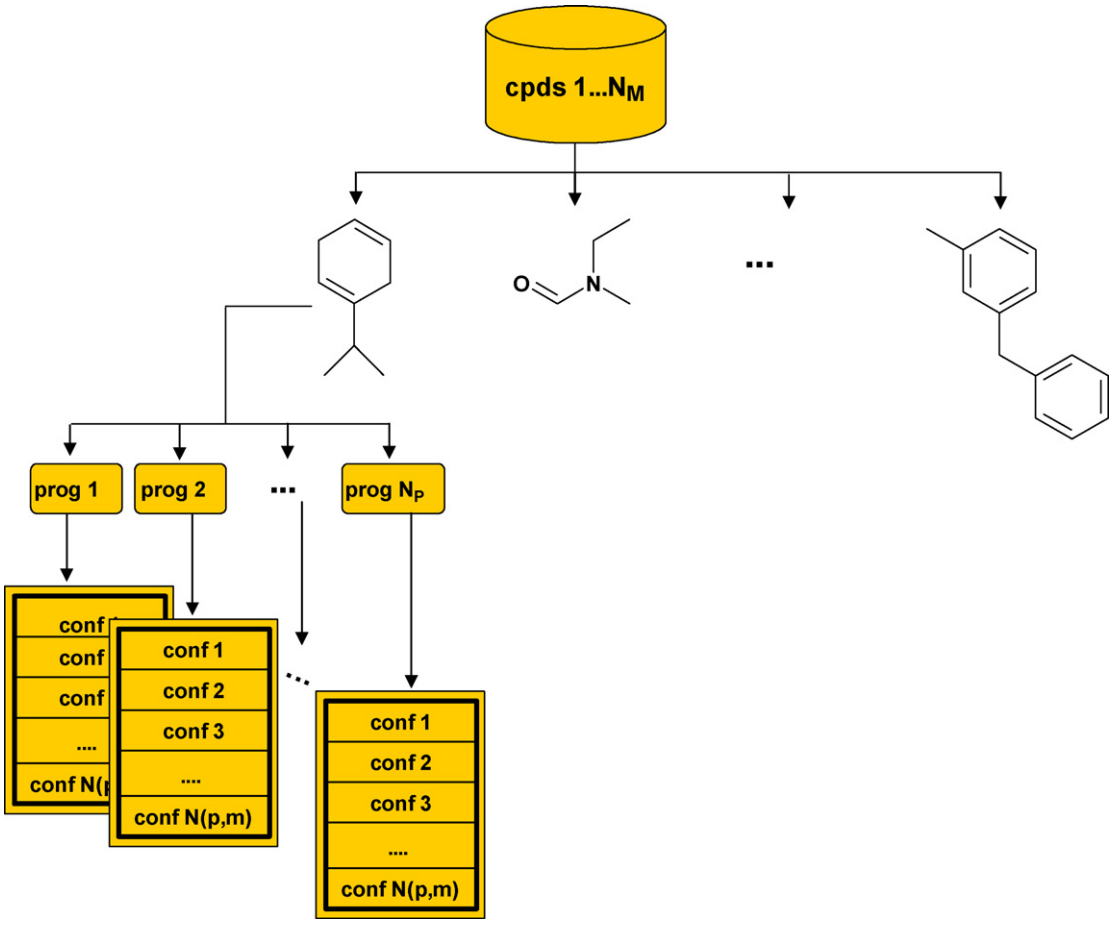

Fig. 1. Overview of the analysis process. For each of the N_M probe molecules a conformer ensemble containing $N(p, m)$ structures is produced by each of the N_P programs.

conformational space in an approximately uniform manner. Assessment of sampling efficiency is based on analysing the dissimilarities between individual conformations in the generated ensemble using various geometric criteria.

2. *Energetics*: The search must be conducted using a realistic model of molecular energetics. For practical reasons, most conformational search methods employ some kind of force-field based approximation. While these enable fast energy calculations, they describe physical reality only with limited accuracy.
3. *Execution speed*: In order to be practically applicable, a good algorithm should be implemented efficiently. Execution speed is partly an inherent property of the algorithm itself, and partly a consequence of the quality of the implementation.

In all assessments we were focusing on evaluating the average performance of conformational search algorithms when they are applied to a large dataset containing molecules of varying size and complexity.

2.1.1. Assessment of sampling efficiency

In order to analyse the extent of conformational space covered by a conformer ensemble of a given molecule, we first superposed all conformers within the ensemble using the generalized procrustes analysis method (GPA). This iterative algorithm translates and rotates the conformers relative to each other so that the total sum of squares of pairwise atomic coordinate differences is minimized. The algorithm usually converges quickly [10], and delivers an optimal, unbiased alignment between all members of the ensemble [11]. Only the heavy (non-hydrogen) atoms were considered in the alignments.

Next we calculated the matrix of all pairwise similarities between the aligned conformers in the ensemble using the Riemann distance as a similarity measure. The Riemann distance is the natural distance measure between two optimally aligned geometrical objects in an abstract shape space. This shape space corresponds to the surface of a hypersphere and the Riemann distance (the geodetic length of the main hypercircle on this surface) in it is given by the following expression:

$$d_{\text{Riemann}} = S_1^2 + S_2^2 - 2S_1S_2 \cos \rho(X_1^0, X_2^0)$$

$X_{1(2)}^0$ denotes superimposed coordinates for two conformers 1 (or 2) and $S_{1(2)}$ is a measure of size of conformer 1 (or 2), while ρ is the angle between the two conformers (for a detailed discussion refer to Ref. [12]). When the two objects are similar, the Riemann distance can be approximated by the RMS distance, however, for more dissimilar objects which are likely to

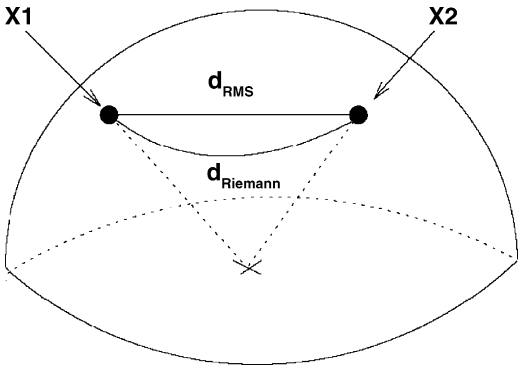

Fig. 2. Two conformers, X_1 and X_2, after optimal superimposition. Each conformer is depicted by a point in conformer shape space. This abstract space has spherical topology (in general the surface of a hypersphere) and distance between the two conformers is correctly measured by the arc length, corresponding to the Riemann distance. The RMS-distance is a linear approximation valid for similar conformers.

be found in well-sampled ensembles the RMS distance is not appropriate any more (Fig. 2).

The optimally aligned conformer ensemble and the corresponding dissimilarity matrix was then subjected to various geometric analyses as described below.

Multidimensional scaling (MDS), also known as *embedding*, is a statistical procedure to represent a set of entities as points in Euclidean space in such a way that the geometric distances between the points correspond to the pairwise dissimilarities between the entities. If the dissimilarities are Euclidean distances derived from an actual n-dimensional point set, then it is possible to recover the original point set by performing metric matrix embedding [13]. In most cases, however, the dissimilarities are not metric Euclidean distances, and therefore an exact embedding is not possible. In such cases an approximation is used that was originally introduced in psychological applications [14]. The basic idea is to force the embedding into a low-dimensional Euclidean space and then minimize the difference between the actual and prescribed distances (d_{ij} and d'_{ij}, respectively). The relative error, which is minimized in this procedure, is called "STRESS" and measures the quality of the embedding:

$$\text{STRESS} = \sqrt{\frac{\sum_{i<j}(d_{ij} - d'_{ij})^2}{\sum_{i<j} d_{ij}^2}}$$

In our case MDS was used to map a conformational ensemble to a set of points in an abstract "conformer space", with each point corresponding to a conformer. We used a 3D conformer space so that we could visualize the extent of sampling (Fig. 3).

To quantify the extent of the sampling we have calculated the volume of the inertial ellipsoid of the point set. This value does not measure whether the points forming this ellipsoid cluster into a number of distinct sets or distribute evenly in accessible conformer space. Thus, the uniformity of sampling had to be tested in a way that is presented below.

Analysis of the covariance structure. The unbiased GPA alignment of conformational ensembles offered another way of assessing sampling. If the sampling is biased, then there are only a few distinct conformer families in the ensemble and the family members are highly correlated to each other, whereas uniform sampling should lead to low correlation between the conformers. This intuitive notion could be quantified as follows. We calculated the covariance of the atomic coordinates between all possible pairs of aligned structures within an N_{conf}-conformer ensemble, and stored the values in a symmetric $N_{\text{conf}} \times N_{\text{conf}}$ covariance matrix. Theoretically, if all realisable conformers are equally well sampled, a certain shape of the eigenvalue spectrum ("scree plot") is predicted, approaching a constant value for all eigenvalues in the limit of large numbers (see Ref. [15] for a detailed derivation). If the sampling is non-uniform then the eigenvalue spectrum of the coordinate covariance matrix contains a few large eigenvalues.

In a series of preliminary computational experiments we always found considerable non-uniformity in the sampling indicated by the eigenvalue spectrum tracing a curve similar to an exponential function as depicted in the right part of Fig. 4. As a simple empirical measure we fitted a

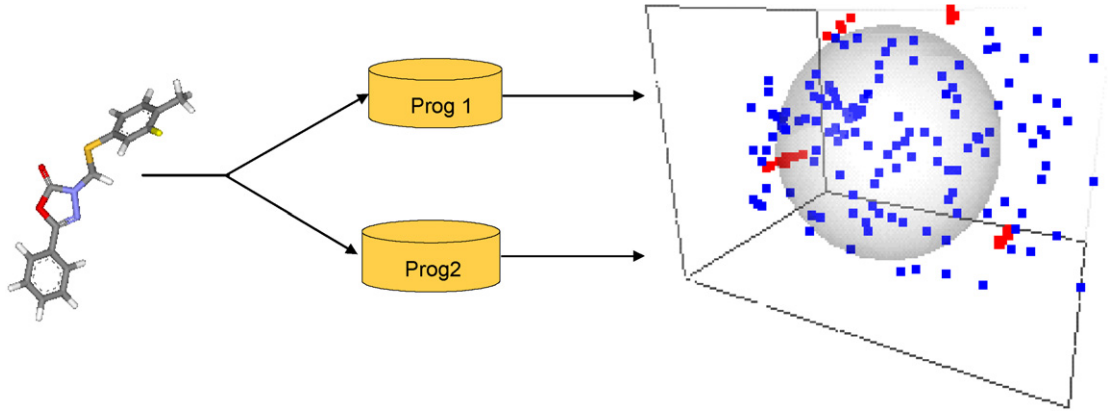

Fig. 3. Schematic overview of the MDS process. A molecule is sampled by two conformer generation programs. The conformers are then embedded as points in 3D space by means of the MDS algorithm. The distribution of the points depends on the distances among the conformers. Program 1 (blue points) samples a larger conformer space more uniformly than program 2 (red points).

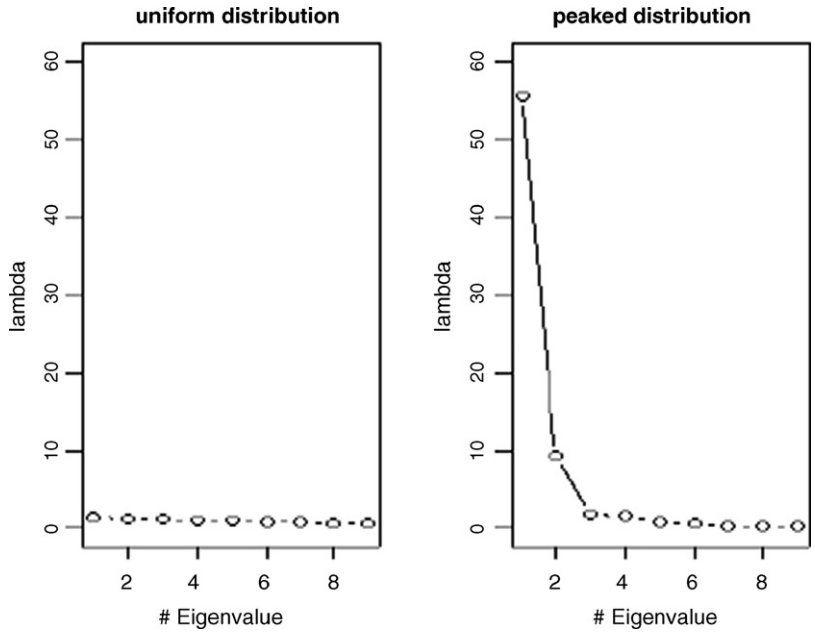

Fig. 4. Eigenvalue spectra of the correlation of a Gaussian random variable. The dimension of the random variables was 10 and Pearson's correlation was calculated for (left) 10,000 and (right) 100 realizations of the random variable. The expected uniform shape is only observed if a high number of tests is performed, whereas (uncorrelated) random numbers produce a eigenvalue spectrum of exponential shape for few numbers tested.

single-exponential model to the eigenvalues and the ''decay constant'' was taken as a descriptor of non-uniformity of the ensemble of sampled conformers. Fig. 4 shows numerical examples for correlations between 10-dimensional vectors sampled from a multinormal distribution, and demonstrates that the distribution of eigenvalues is highly dependent of the number of samples.

2.1.2. Assessment of energetics

We compared the distribution of molecular energies of conformational ensembles generated by the programs tested. The programs were ranked according to how often they delivered the lowest median energy for a molecule in the database. We used the median to describe the central tendency of the energy distributions because this measure is less sensitive to skewness than the mean. A program that always produced the lowest median energy conformer ensemble for all $N_M = 604$ molecules in the dataset would be considered optimal.

2.1.3. Assessment of execution speed

Speed is essential if high-throughput conformational search is desired. Ideally, the time needed for creation of a ''good'' conformer ensemble should be measured. In reality, this is hardly achievable as different conformational search programs generate varying numbers of conformers of varying quality which often require post-processing steps (clustering, etc.). Thus, we have restricted ourselves to the measurement of production times of conformers, no matter whether the conformer ensembles needed further refinement or not. Thus, the average time for one conformer was calculated as the ratio of the total elapsed CPU-time for a run with a single molecule divided by the number of conformers generated during the run.

2.2. Computational details

2.2.1. Conformational search programs

We selected the conformational search programs so that a wide range of methodologies (both deterministic and stochastic) could be analysed. For practical reasons we restricted ourselves to programs that were easily available to us. In the following we provide a brief description of the five programs tested.

DGEOM (© Chiron Corporation, 1995; available from the Quantum Chemistry Program Exchange as program no. QCPE-590 for a modest fee): DGEOM employs distance geometry (metric matrix embedding) to produce conformers satisfying a set of pairwise atom distance restraints which are generated using chemical knowledge [16]. Note that this procedure, although it is also based on MDS, is different from the embedding of conformers we used for assessing sampling efficiency (cf. Section 2.1.1). The algorithm is almost purely geometric, no force-field calculations are involved except for a van der Waals check to filter out steric clashes. We used the default parameter settings except that the maximum number of accepted structures was increased to 2000 conformations, 100

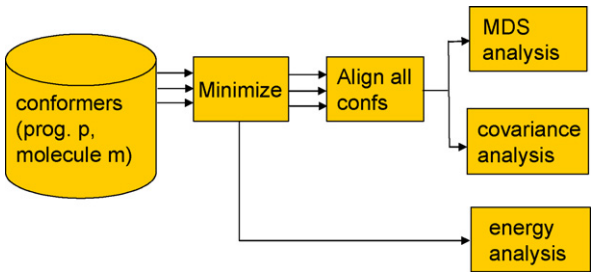

Fig. 5. Workflow for the statistical analysis of a specific molecule and sampling program.

trials per structure were allowed and the acceptance rate (minimum dissimilarity from other structures) was set to 0.3 Å.

QXP (available at Novartis in-house as part of the FLO.01 program suite): QXP employs torsional MC sampling [17]. Rings are sampled by perturbations in Cartesian space. The sampling temperature was set to 900 K and a maximum number of 1000 accepted conformers, two minimizations per search cycle (which uses simplified force fields), and a minimum 30.0° RMS deviation in dihedral angle space for acceptance were set.

ROTATE (Version 1.15, © Molecular Networks GmbH Computerchemie, available to us under a demo license): ROTATE is a knowledge-based system, employing dihedral angles obtained from data of the Cambridge Crystallographic Data Centre. A set of starting structures with different ring conformations was created with CORINA, as these are not sampled by ROTATE. The maximum number of rotatable bonds, which are counted starting from the innermost rotatable bond, was set to 8, and duplicate conformations which by

Table 1
Mean values of sampling extent and sampling diversity for each program

Program	Median MDS 3D volume ($\times 10^3$ Å^3)	Median decay constant
DGEOM	608	−0.52
QXP	468	−0.52
ROTATE	574	−0.54
LMOD	571	−0.62
OMEGA	574	−0.57

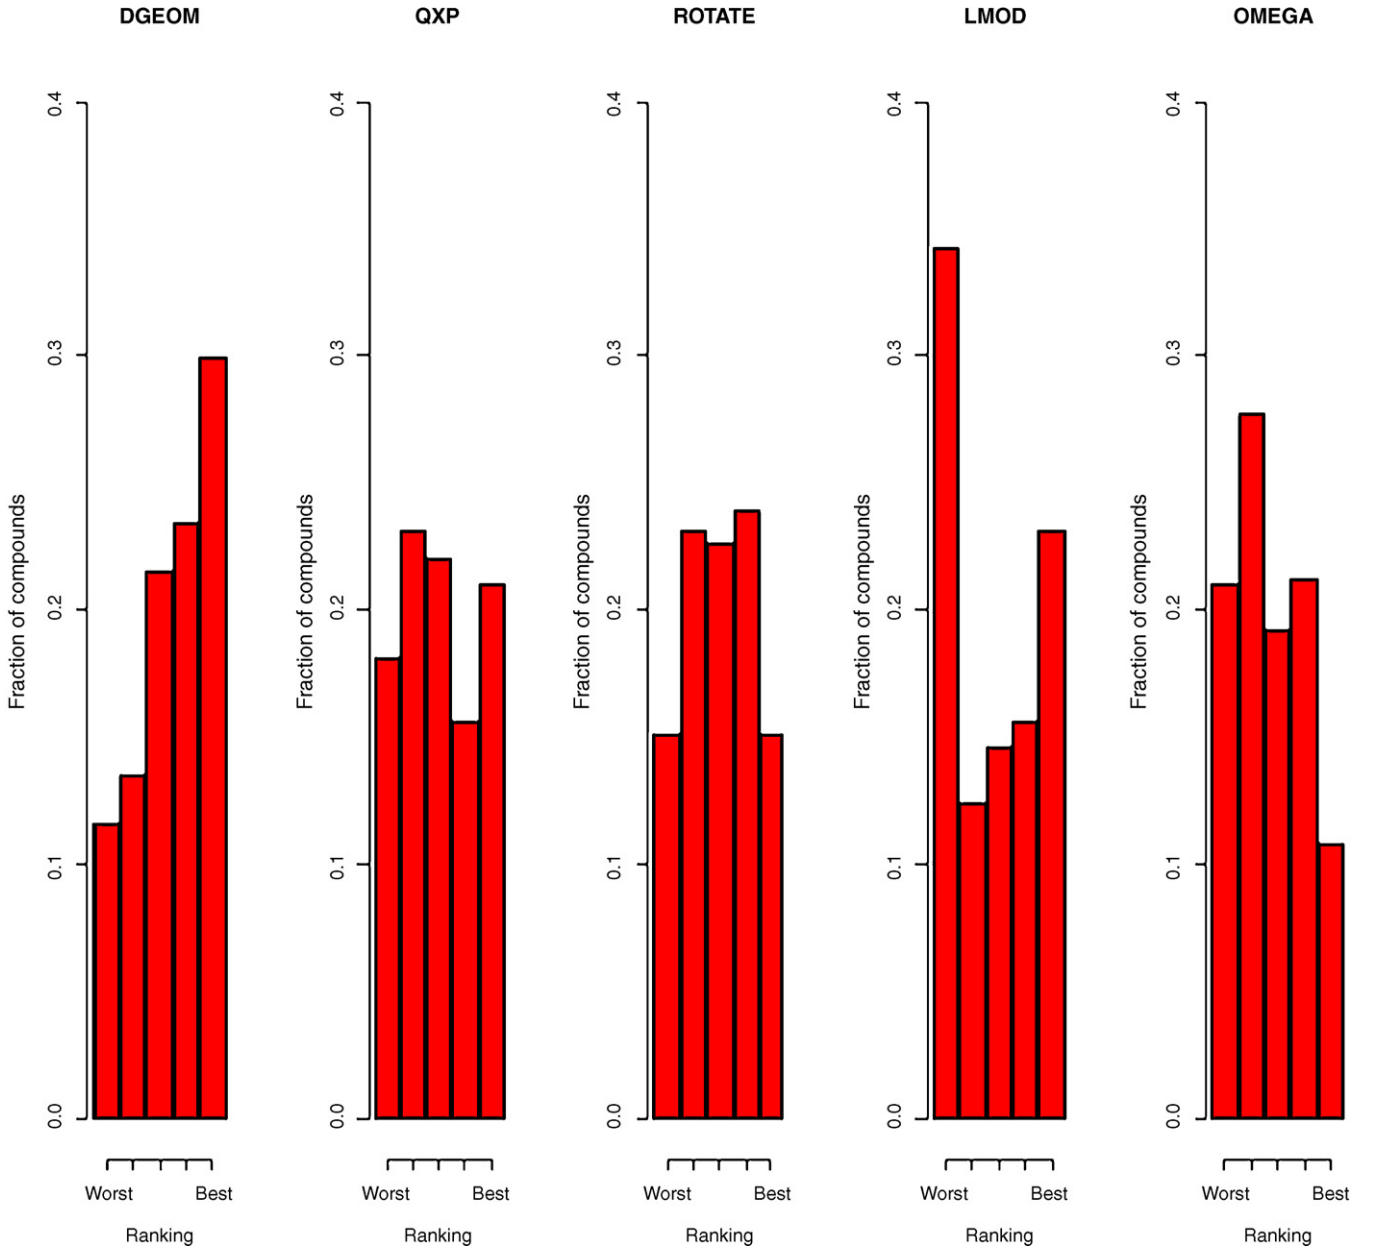

Fig. 6. Ranking of programs by the extent of sampled conformers. For each molecule the embedding volume of the ensemble generated by the five programs were compared and ranked. The histograms show how often a specific program ranked worst, . . ., best. DGEOM is most often ranked the best program producing the largest embedding volumes, followed by LMOD and QXP.

definition differed by less than 30° RMS deviation in torsion angle space were filtered out.

LMOD (© Biokol Research LLC; source code available to the In Silico Sciences unit): This program implements an enhanced version of the low-mode conformational search method that jumps to neighbouring local minima on the potential energy surface by employing an efficient mode-following technique [9]. LMOD operates by using a detailed force field; the current implementation relies on the freely available NAB tool [18,19] which can evaluate potential energy expressions and perform molecular dynamics using the AMBER force-field family. Here we have made use of the GAFF-1.2 force field [20], parameter assignment was carried out by the ANTECHAMBER program [21]. The calculations were performed without explicit solvent, but the dielectric constant was set to the bulk water value ($\varepsilon = 80$) to model electrostatic screening. This approximation was acceptable because of the small sizes of the compounds in the dataset. Two thousand iterations were allowed, and 1500 conformations, ranked by increasing energy, were accepted. The lowest five eigenmodes were searched in both directions and structures with a RMS gradient below 0.1 kcal/mol/Å^2 accepted as a new minimum.

OMEGA (Version 1.8, © OpenEye Software, Inc.; we tested the software using an evaluation license): This program employs a knowledge-based approach but the details of the algorithm unfortunately constitute a trade secret. We used the program in rule-based conformer generation mode and no

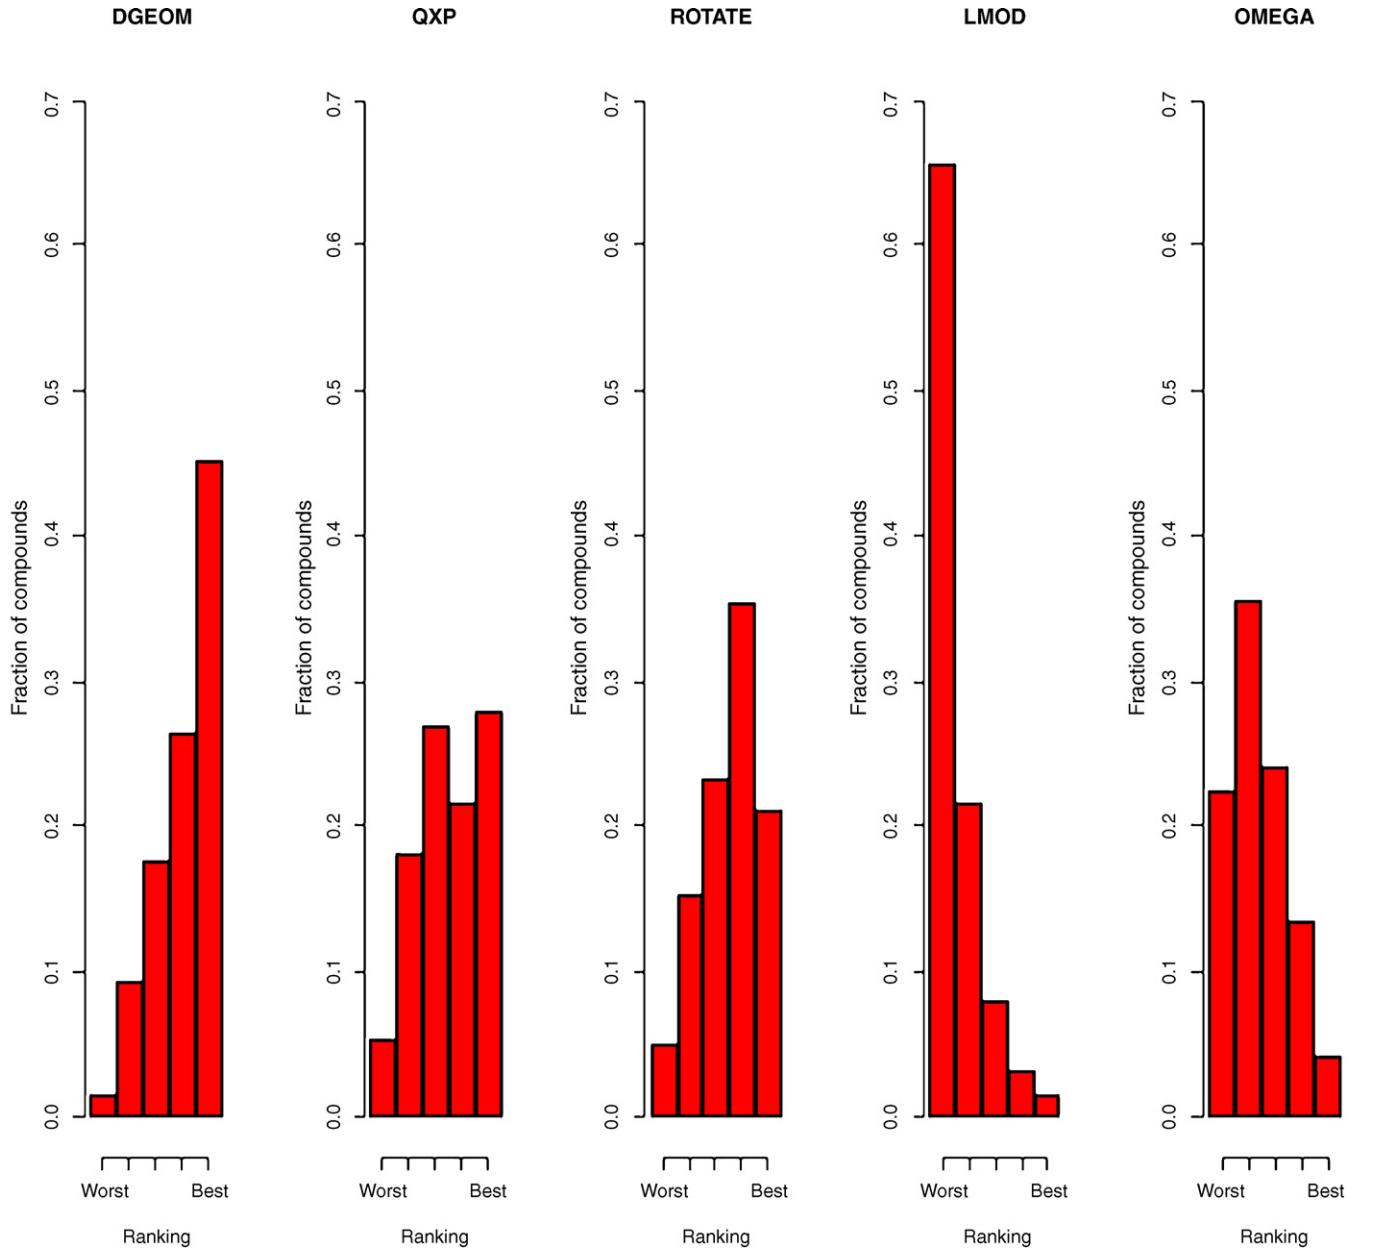

Fig. 7. Comparison of the uniformity of the sampling based on the correlation measure. DGEOM generates most often the most uniform distribution, LMOD the least often.

subsequent energy minimizations were performed. The criterion for duplicate removal was set to a value of 0.6 Å RMSD and a maximum of 2000 stored conformations were allowed.

All programs were run on a Linux cluster at NIBR Vienna under Suse Linux 9.1. Up to 96 1.8 GHz AMD Opteron B CPUs have been used.

2.2.2. Data

The LeadQuest Database (Tripos, Inc.) is an *in silico* library containing 80,000 drug-like molecules. A representative subset of this library was selected by means of the OptiSim algorithm [22]. As the criterion for diversity a Tanimoto coefficient of less then 0.50 based on UNITY fingerprints was chosen. Thus, a maximally dissimilar subset of 604 molecules was obtained (see Supplementary Information).

2.2.3. After-minimization of raw conformers and duplicate removal

Since the programs use different energy minimization strategies and force fields, it is not possible to compare the energy distributions of the conformer sets directly. A possible solution is to after-minimize the conformers using the same energetics and minimization protocol for the result set of each conformation generation program. To this end one would need a kind of ''gold standard'' program. We chose the semiempirical quantum mechanics (QM) package MOPAC (Version 7.0) as an unbiased after-minimization tool because a QM optimizer provides high accuracy and would not prefer any classical force field over another. Note that we do not suggest that QM-based after-minimization should be used in a productive context when exploring the conformational landscapes of many molecules in batch mode; we made use of MOPAC only as an analysis tool.

We used the AM1 Hamiltonian with a gradient RMS convergence criterion (GNORM) of 0.5 kcal/mol/Å^2. The final optimized geometry and energy (''heat of formation'') were used for further processing.

For duplicate removal we clustered the minimized conformers. Coordinate RMS deviation was used as a similarity measure without considering intramolecular local symmetries, and only conformers differing less than 1.0 Å were clustered together. Clustering was performed using the command-line version of the XCLUSTER program (Version 1.7, supplied by Schrödinger, Inc.).

2.2.4. Statistical analysis

In order to compare the five programs in an unbiased manner, the conformer ensembles must contain the same number of conformers for each molecule. Since some of the programs perform conformer clustering to filter out repeatedly occurring almost-identical conformers while others do not, it was not possible to set the ensemble size to a predefined value for all programs prior to the batch runs. As a workaround we took the smallest number of conformers for a given molecule and subsampled the other conformer sets by drawing the same number of conformers randomly. To improve statistical significance, this procedure was repeated until the estimated confidence interval for the total population expectation values fell below 3%.

For all statistical analyses we used the public-domain statistical modelling package R 2.0.1 [23].

3. Results and discussion

3.1. Outline

Fig. 5 shows all the steps described above in a single workflow. As soon as statistical data are generated for all molecules and programs analysis is performed. We employed the following analysis scheme: for each molecule the results of the five programs tested are ranked (rank 1 = worst, rank 5 = best performer). Then the programs are compared among each other in terms of the distribution of these rankings, considering all molecules.

3.2. Geometric analysis

3.2.1. Multidimensional scaling

This test assessed the size of the 3D ellipsoid occupied by the ensembles in an abstract conformational space when the inter-conformer Riemann distance matrix was embedded in $\mathbf{R}^3$. We found that the conformers generated by DGEOM occupied the largest median volume after embedding, followed by OMEGA, ROTATE and LMOD (Table 1). DGEOM was also

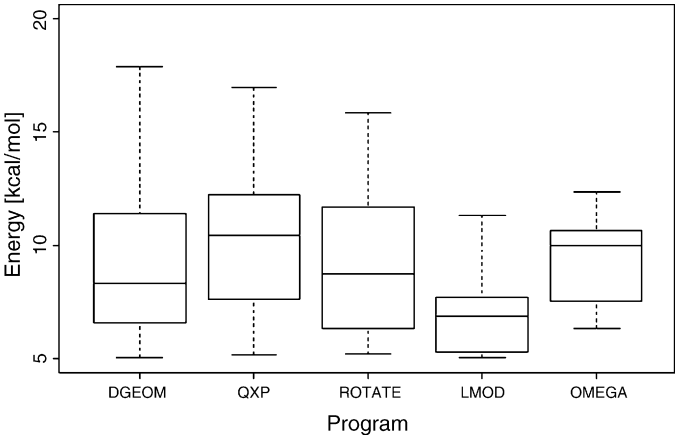

Fig. 8. A typical energy distribution for a set of conformers. The boxed areas contain conformers within the second and third quartiles which are separated by the median line. The whiskers indicate the lower and upper limits of the empirical distribution.

Table 2
Average energy values: the median energies for a specific molecule and program have been averaged on all molecules

Program	Mean energy (kcal/mol)
DGEOM	−38.7
QXP	−39.6
ROTATE	−39.3
LMOD	−40.4
OMEGA	−39.8

most often ranked as the program generating the best (largest) embedding volume for a given molecule (Fig. 6).

3.2.2. Analysis of covariance structure

Fig. 7 displays a ranking of the "decay constant" fitted to the eigenvalue spectra of the coordinate covariance matrices generated by the five programs examined. Again DGEOM performed most often as the best program, followed by QXP and ROTATE. If averaged values instead of a rank-based analysis are considered (Table 1) then QXP and DGEOM give the best results. We note that the differences between the programs are quite small. The average "decay constant" of the worst performing program, LMOD, had a mean value of −0.62, which is 19.2% different from the best, QXP, −0.52.

3.3. Energy analysis

In this analysis, the median energies of the conformer ensembles generated by the programs were compared. We found that the conformer ensembles of "simple" molecules containing just a few rotatable bonds and rigid ring systems had very similar energy spectra, no matter which program generated the conformers (data not shown). On the other hand, large differences of more than 10 kcal/mol in median energies were observed for large, highly flexible molecules. A typical distribution of energy values observed for one molecule is shown in Fig. 8. The energy distributions created by the programs examined are drawn here in form of a boxplot. The energies span a wide range and considerable differences occur.

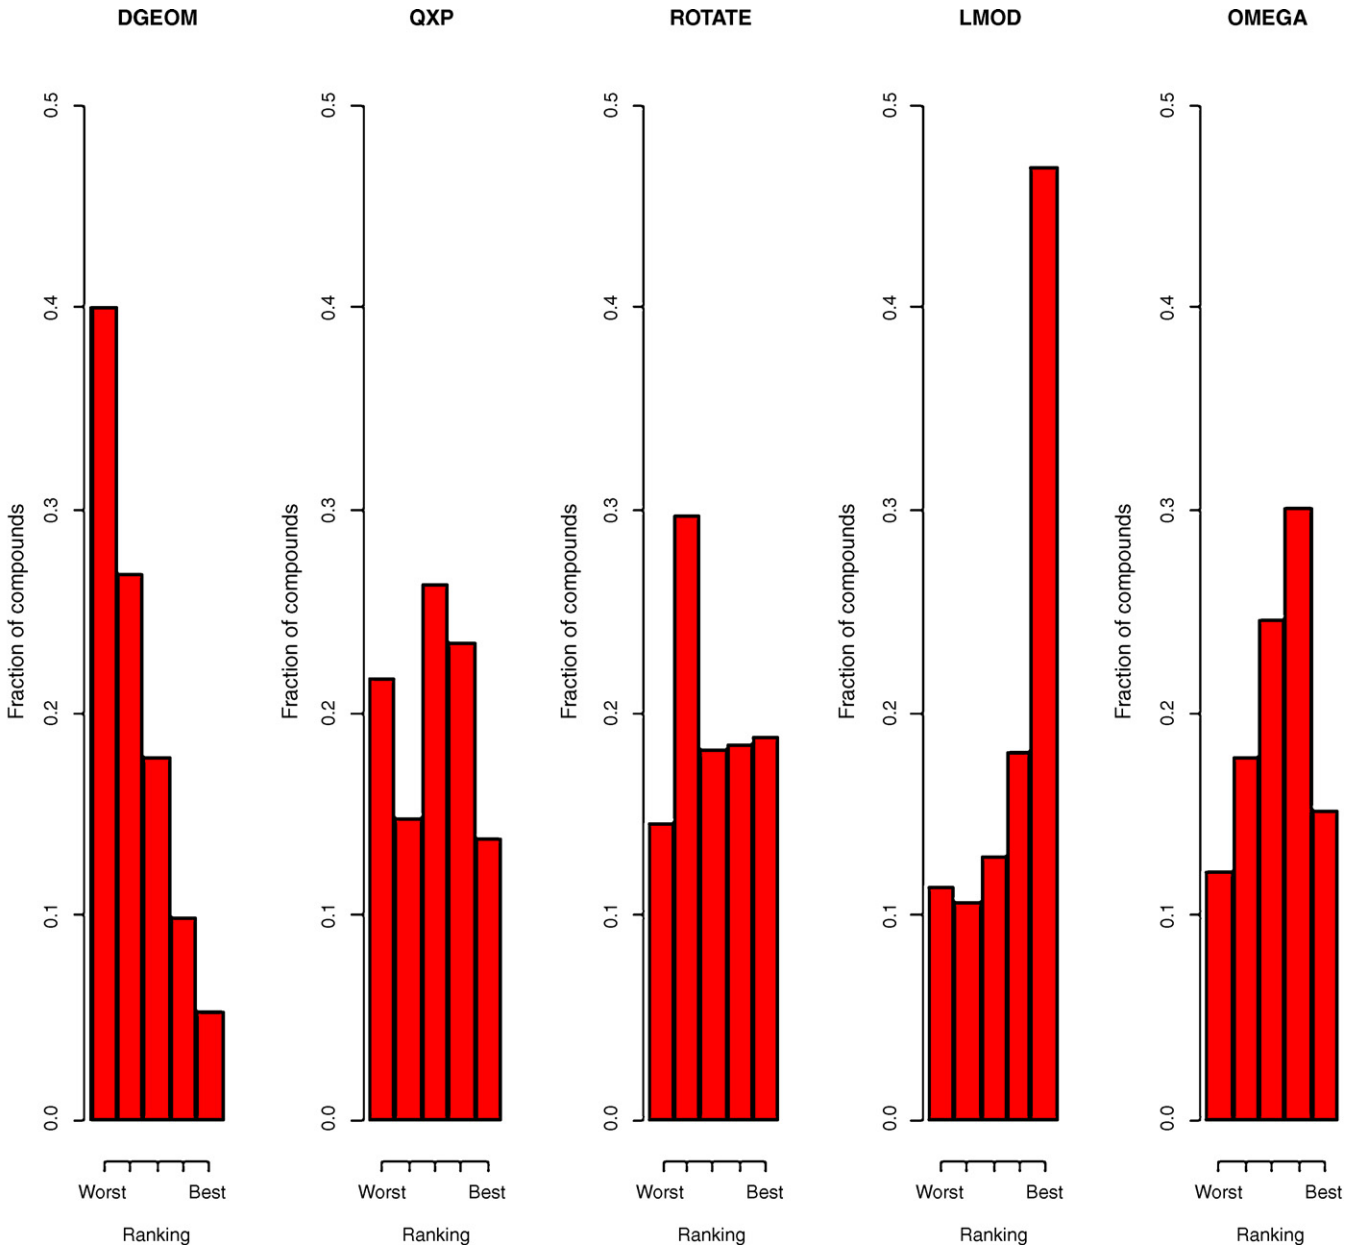

Fig. 9. Ranking of conformer analysis programs by energies. For each molecule the median energies of the five conformer generation programs were compared and ranked. The histograms show how often a specific program ranked worst, . . ., best. The program LMOD creates lowest-energy conformers most often, whereas DGEOM performs worst.

Table 3
Execution times

Program	Average time/conformer/CPU (ms)
DGEOM	76
QXP	213
ROTATE	305
LMOD	996
OMEGA	12

Elapsed time per conformer is measured.

Table 2 shows an overview of the obtained averaged values. Higher energy conformations are mainly observed with programs QXP, ROTATE, DGEOM and OMEGA. The majority of lowest-energy conformations is generated by LMOD. This is evident from Fig. 9, showing that LMOD performs best for most molecules analysed, with the other programs falling significantly behind.

3.4. Execution speed

Measurement of execution speed showed a significant advantage for OMEGA which consumed only 12 ms elapsed CPU-time per conformer on average on the hardware specified above. Also DGEOM generates conformers at a very high speed. The other programs generate an average conformer at a rate of 200–1000 ms. The LMOD program, which uses the most elaborated force field performed slowest (Table 3).

4. Conclusion

The general conclusion from our study is that there is no ''overall winner'', i.e. no single conformational analysis program provides the lowest-energy conformations at the highest possible sampling efficiency. The amount of relevant conformer space of course critically depends on the number of microstates linked to a certain energy (the density of states) and the possibility to attain these conformers (occupation probability). The two characteristics ''energy'' and ''geometry'' are thus intimately related. Interesting results concerning these findings exist for peptides and the protein-folding problem [24,25].

If low-energy conformations are desired, then we recommend LMOD. This program appears to bias the sampling towards the minima of the applied force field, which means that the geometrical extent and diversity of the conformer ensembles is limited compared to the others, as the preferentially sampled lower-energy regions contain fewer minima. Such a biased sampling is obviously a very desirable property if the force field approximates physical reality well. One should note, however, that force field inaccuracies in general and the difficulty of properly modelling solute–solvent interactions in particular can seriously limit the usefulness of conformer ensembles generated in a simple protocol if the structures obtained are not refined further.

On the other hand, DGEOM and, generally, Monte Carlo-based programs such as QXP generate a larger variety of sampled conformations both in terms of sampling volume and uniformity, albeit at the cost of higher median energies. While MC simulations generate (after infinite steps) the appropriate thermodynamic ensemble corresponding to the force field used, the MC-based conformational search programs usually perform only simple bump-checks or employ simplified force fields for efficiency reasons which leads to higher-energy conformers. The performance of knowledge-based programs is highly dependent of the implemented rules and patterns. OMEGA seems to be biased more towards lower energy conformations than ROTATE, which yields better results in terms of geometry properties. OMEGA might constitute a viable option if speed is imperative as it is by far the fastest program and generates conformers at lower energies than the other programs (except LMOD). For all these conformer generation programs our results show that not even after-minimization using a semiempirical quantum mechanical protocol is sufficient to move the sampled ensemble towards the more relevant low-energy regions. Thus, these programs are preferably used whenever low-energy conformations are not essential. This is the case if abstractions of a molecule (scaffolds, pharmacophores) or interactions between chemical entities are studied. In many protein–ligand complexes ligands are not bound in local minimum energy conformations and also exhibit considerable strain-energies of sometimes more than 9 kcal/mol [10]. Then, good performance with respect to geometric properties is more important than energy considerations.

Acknowledgements

The authors wish to thank Dr. Markus Jaritz, Dr. Torsten Schindler and Dr. Dmitri Mikhailov (In Silico Sciences Unit) for their useful comments and suggestions. M.J.L. was supported by a Global Research Management Postdoctoral Fellowship at the Novartis Institutes for Biomedical Research. I.K. acknowledges financial support from a grant of the Hungarian Science Foundation OTKA T042933.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jmgm.2006.05.008.

References

[1] C.H. Schwab, Conformational analysis and searching, in: J. Gasteiger (Ed.), Handbook of Chemoinformatics, Wiley/VCH, Weinheim, 2003, pp. 262–301.
[2] A. Smellie, R. Stanton, R. Henne, S. Teig, Conformational analysis by intersection: CONAN, J. Comp. Chem. 24 (2003) 10–20.
[3] O. Burkert, N.L. Allinger, Molecular mechanics, ACS Monographs, Washington, 1982.
[4] J. Boström, Reproducing the conformations of protein-bound ligands: a critical evaluation of several popular conformational searching tools, J. Comput.-Aided Mol. Des. 15 (2001) 1137–1152.
[5] B.P. Feuston, D.M. Miller, J.C. Culberson, R.B. Nachbar, S.K. Kearsley, Comparison of knowledge-based and distance geometry approaches for generation of molecular conformations, J. Chem. Inform. Comput. Sci. 41 (2001) 754–763.

[6] C.-S. Wang, Efficient algorithm for conformational search of macrocyclic molecules, J. Comp. Chem. 18 (1996) 277–289.
[7] A. Smellie, S.D. Kahn, S.L. Teig, Analysis of conformational coverage. 1. Validation and estimation of coverage, J. Chem. Inform. Comput. Sci. 35 (1995) 285–294.
[8] C. Parish, R. Lombardi, K. Sinclair, E. Smith, A. Goldberg, M. Rappleye, M. Dure, A comparison of the low mode and Monte Carlo conformational search methods, J. Mol. Graph. Model. 21 (2002) 129–150.
[9] I. Kolossváry, W.C. Guida, Low-mode conformational search elucidated: application to C39H80 and flexible docking of 9-Deazaguanine inhibitors into PNP, J. Comp. Chem. 20 (1999) 1671–1684.
[10] J.C. Gower, Generalized procrustes analysis, Psychometrika 40 (1975) 33–50.
[11] P.M. Kroonenberg, W.J. Dunn III, J.J.F. Commandeur, Consensus molecular alignment based on generalized procrustes analysis, J. Chem. Inform. Comput. Sci. 43 (2003) 2025–2032.
[12] I.L. Dryden, K.V. Mardia, Statistical Shape Analysis, John Wiley & Sons, 1998.
[13] G.M. Crippen, T.F. Havel, Distance Geometry and Molecular Conformation, Research Studies Press Ltd., 1988.
[14] R.N. Shepard, The analysis of proximities: multidimensional scaling with an unknown distance function I, Psychometrika 27 (1962) 125–140.
[15] R. Everson, S. Roberts, Inferring the eigenvalues of covariance matrices from limited, noisy data, IEEE Trans. Signal Proc. 48 (2000) 2083–2091.
[16] D.C. Spellmeyer, A.K. Wong, M.J. Bower, J.M. Blaney, Conformational analysis using distance geometry methods, J. Mol. Graph. Mod. 15 (1997) 18–36.
[17] C. McMartin, R.S. Bohacek, QXP: powerful, rapid computer algorithms for structure-based drug design, J. Comput.-Aided Mol. Des. 11 (1997) 333–344.
[18] T.J. Macke, D.A. Case, Modeling Unusual Nucleic Acid Structures, American Chemical Society, Washington, DC, 1998, pp. 379–393.
[19] T.J. Macke, D.A. Case, 2004. http://www.scripps.edu/mb/case/casegr-sh-3.2.html.
[20] J. Wang, R.M. Wolf, J.W. Caldwell, P.A. Kollman, D.A. Case, Development and testing of a general amber force field, J. Comp. Chem. 25 (2004) 1157–1174.
[21] J. Wang, W. Wang, P.A. Kollman, D.A. Case, Automatic atom type and bond type perception in molecular mechanical calculations, J. Mol. Graph. Model (2006) in press.
[22] R.D. Clark, OptiSim: an extended dissimilarity selection method for finding diverse representative subsets, J. Chem. Inform. Comput. Sci. 37 (1997) 1181–1188.
[23] The R Project for Statistical Computing, 2005. http://www.r-project.org.
[24] J.J. Perez, N.B. Centeno, New insights into the multiple minima problem, J. Mol. Struct. (Theochemistry) 370 (1996) 173–179.
[25] A.R. Leach, A.P. Lemon, Exploring the conformational space of protein side chains using dead-end elimination and the A*-algorithm, Proteins Struct. Funct. Gen. 33 (1998) 227–239.