

Global Optimization via a Modified Potential Smoothing Kernel

Alan Grossfield and Jay W. Ponder

Department of Biochemistry and Molecular Biophysics
Washington University School of Medicine
660 S. Euclid Ave.
Box 8231
St. Louis, MO
63110
phone: 314 362 4195
fax: 314 362 7183
email: ponder@dasher.wustl.edu, alan@dasher.wustl.edu

Keywords: global optimization, potential smoothing

Abstract

A new potential smoothing method, the shifted-tophat (or stophat) is presented. This method uses a tophat function as the smoothing kernel, instead of the gaussian used in conventional methods. Stophat smoothing is applied, as part of the Potential Smoothing and Search (PSS) procedure for global optimization, to several biomolecular problems, including polyalanine helices, united-atom and all-atom models of enkephalin, and rigid helices from glycophorin.

I Introduction

Global optimization of molecular structure is one of the classic problems of theoretical chemistry and biophysics. The problem can be stated simply: for a given molecule, what is the conformation with the lowest potential energy? Although simple to state, this problem is quite difficult to solve in a general way. Conformational space grows exponentially with molecular size, making systematic searching difficult for all but the smallest problems. Moreover, the number of local minima usually grows rapidly as well, with the result that straight-forward optimization methods tend to get trapped in local minima and fail to find the global minimum.

A variety of approaches exist to deal with this difficulty. One of the most common is simulated annealing, which uses either Monte Carlo or molecular dynamics to sample conformational space at high temperatures, where even large energy barriers are easily crossed (1). By gradually lowering the temperature, the conformational sampling can be focused on lower energy regions of phase space. In principle, as the temperature approaches 0K, only the global minimum energy structure remains. In practice, simulated annealing is not deterministic, because of the prohibitive computational cost of a sufficiently long trajectory. Rather, even long trajectories have a finite chance of finding several different minima, depending on the topology and roughness of the potential energy surface.

The thermodynamic importance of any specific conformation is determined by the Boltzmann factor $E/k_B T$, where E is the potential energy for the conformation. Thus, simulated annealing can be viewed as a rescaling of the energy surface; high temperatures effectively scale the energies and make the surface flatter and thus easier to search, at the cost of diminishing the favorability of the low energy states. However, the molecular dynamics and

Monte Carlo methods used with simulated annealing generate a Boltzmann distribution at each temperature, rather than proceeding directly toward a minimum. Combined with the exponential dependence of the barrier-crossing time on barrier height and the exponential growth in the number of barriers with system size, this makes simulated annealing inefficient for global optimization.

Potential smoothing methods take an alternative approach to the problem. If global optimization is difficult because the potential surface has many local minima and large barriers, why not simply alter the potential energy surface to eliminate them? Smoothing methods transform the potential surface in such a way that the number of minima is greatly reduced and global optimization is easy. The degree of surface deformation is then gradually decreased, such that the global minimum on the deformed surface can be tracked back to the original undeformed surface.

There are several distinct methods which proceed along these lines, including the diffusion equation method (2; 3; 4; 5), gaussian density annealing (6; 7), the effective diffused potential (8; 9; 10), and gaussian packet annealing (11). In all of these methods, discrete atoms are replaced by gaussian distributions. The interactions between the distributions result in a smoother potential surface with fewer minima.

These methods have a variety of physical interpretations, but there are strong mathematical similarities. For example, in gaussian density annealing, one tracks the classical atomic probability density as a function of temperature; as the temperature gets small, the density converges about a single low energy conformation (6). To make the problem mathematically tractable, the atomic distributions are assumed to be well modeled by single gaussians. Shalloway's packet annealing method further generalizes this approach, allowing the distributions to be represented by the sum of multiple gaussians, which can split and merge over the course of the calculation, to better capture both conformational and vibrational heterogeneity (11).

Vershelde *et al.* interpret the gaussian distributions to be the result of atomic fluctuations (9). By again assuming a gaussian form for the fluctuations, they are able to analytically calculate the entropy associated with them. This entropy is combined with the smoothed potential energy to create an effective potential (10). In the limit that the true atomic fluctuations are well-described by gaussians, the effective potential is identical to the free energy. All of these methods share essentially identical functional forms for the smoothed van der Waal's and electrostatic interactions, as does the orig-

inal diffusion equation method of Piela *et al.*(2). For a further discussion of the theoretical relationships between these methods, see the reviews by Piela and Schelstraete *et al.*(12; 13).

Thus far, we have discussed smoothing methods in terms of the substitution of probability distributions for discrete atomic locations. However, we can equivalently focus instead on the many-dimensional potential energy surface. Calculating the interaction between two atomic distributions is equivalent to averaging the potential energy surface, using the convolution of the two distributions as a weighting function. In the case where the distributions are gaussians, this leads to a weighting function which is also a gaussian. This is not the case with other distributions. For example, representing the atoms with tophat or impulse distributions would be equivalent to averaging the surface using a 2-step impulse function.

In the atomic distribution view of smoothing, the use of gaussians has obvious physical significance. However, if one views smoothing as an empirical transformation applied to the surface, the choice of smoothing kernels becomes a matter of mathematical and computational convenience. If the primary goal of the smoothing is to simplify conformational searching, it may be the case that other functional forms offer advantages over gaussians. For example, the method of bad derivatives uses a tophat function in place of a gaussian, with a number of advantages (14). Most importantly, this form of smoothing is effective for molecular conformational optimization. Also, the tophat width λ exactly determines the length scale of averaging, which aids interpretation of the resulting deformations. The method of bad derivatives, moreover, does not require that the potential function be differentiable, which allows the application of novel scoring functions, such as the Boltzmann weight. However, the method of bad derivatives calculates the energies and forces in the form of a numerical derivative, which becomes computationally expensive for larger systems.

This paper presents the results of *analytically* smoothing the potential energy surface using a tophat smoothing kernel. Section II contains a derivation of gaussian smoothing, tophat smoothing, and shifted-tophat (stophat) smoothing, as applied to molecular mechanics potential functions. Section III describes the potential smoothing and search (PSS) procedure. Section IV describes the application of stophat smoothing to a number of molecular systems (butane, polyalanine, enkephalin, and glycophorin), with emphasis on its efficiency at global optimization. Finally, Section V will compare the physical interpretation and computational performance of stophat smoothing

to gaussian smoothing and simulated annealing.

II Derivations

II-A Gaussian smoothing

Many smoothing methods operate by replacing the potential function with an appropriately weighted average value. Mathematically, the smoothed function is defined via an integral of the potential function, multiplied by a suitable weighting function or kernel, $k(\vec{x})$:

$$\mathcal{F}(\vec{x}_0, \lambda) = \frac{\int F(\vec{x})k(\vec{x}_0, \vec{x}, \lambda) d\vec{x}}{\int k(\vec{x}_0, \vec{x}, \lambda) d\vec{x}} \quad (1)$$

In this notation, \vec{x} is a $(3N - 6)$ -dimensional vector (for an N atom system) representing all degrees of freedom for the system, the integral is over all of phase space, and \mathcal{F} is the average of F , subject to weighting by k .

The most common choice of smoothing kernel is

$$k(\vec{x}_0, \vec{x}, \lambda) = \exp\left(\frac{-(\vec{x} - \vec{x}_0)^2}{\lambda}\right), \quad (2)$$

a Gaussian of width λ centered at \vec{x}_0 . The degree of smoothing is controlled by varying λ ; larger values of λ result in more averaging, producing a smoother \mathcal{F} , while $\lim_{\lambda \rightarrow 0} \mathcal{F} = F$.

Equation 1 is not convenient for molecular calculations involving empirical forcefields where the total potential is calculated as a sum of a large number of terms. However, the integral in Equation 1 can be rewritten as a sum of integrals over the individual terms. In the general case each energy term depends on all of the coordinates \vec{x}_0 , requiring integration over the full $(3N - 6)$ -dimensional space. However, in the most common case of a pairwise potential, each term depends on the three variables specifying the vector between two atoms. The integration over the other dimensions does not affect the average, and can be ignored. If the terms in the potential are pairwise and depend solely on distance, the integral for each term can be reduced to a 1-dimensional integral (15). For clarity, we will summarize the derivation here. We begin by considering the smoothing of a single pairwise potential term $F(\vec{x}_0)$. Since F depends solely on the magnitude of \vec{x}_0 , we can change coordinate systems to place one atom at the origin, and the

other along the z axis, at $(0, 0, x_0)$. In this coordinate system, an arbitrary vector $\vec{x} = (r \cos \theta \sin \phi, r \sin \theta \sin \phi, r \cos \phi)$, where r , θ , and ϕ have their conventional meaning from spherical coordinates. Using this notation, we can rewrite the numerator of the exponential term in Equation 2 as

$$(\vec{x} - \vec{x}_0)^2 = r^2 + x_0^2 - 2rx_0 \cos \phi \quad (3)$$

Substituting this back in to Equation 1 (written in spherical coordinates) and performing the trivial integration over θ yields

$$\mathcal{F}(\vec{x}_0, \lambda) = \frac{2\pi}{(\sqrt{\pi}\lambda)^3} \int_0^\infty dr \int_0^\pi r^2 F(r) \exp\left[-\left(\frac{x_0^2 + r^2}{\lambda}\right)\right] \exp\left[\left(\frac{2x_0r \cos \phi}{\lambda}\right)\right] \sin \phi d\phi \quad (4)$$

Only the final exponential depends on ϕ , so we can calculate it separately as

$$\int_0^\pi \exp\left[\left(\frac{2x_0r \cos \phi}{\lambda}\right)\right] \sin \phi d\phi = \frac{\lambda}{2x_0r} \left\{ \exp\left[\frac{2x_0r}{\lambda}\right] - \exp\left[-\frac{2x_0r}{\lambda}\right] \right\} \quad (5)$$

Substituting this result back into Equation 4, we obtain

$$\mathcal{F}(\vec{x}_0, \lambda) = \frac{1}{x_0\sqrt{\pi}\lambda} \int_0^\infty r F(r) \left\{ \exp\left[-\frac{(r - x_0)^2}{\lambda}\right] - \exp\left[-\frac{(r + x_0)^2}{\lambda}\right] \right\} dr \quad (6)$$

Setting F to be a typical potential energy term (for instance, $1/x$ for the coulomb potential) allows us to recover the standard diffusion equation potential functions (2). This procedure can only be applied to functions F where an analytic solution for the integral in Equation 6 exists. As a result, F must diverge no faster than r^{-1} as $r \rightarrow 0$, or the integrand will diverge at the origin.

II-B Tophat Smoothing

An alternative choice of smoothing kernel is an impulse or tophat function. Specifically,

$$k(\vec{x}_0, \vec{x}, \lambda) = \begin{cases} 1 & \text{for } |\vec{x} - \vec{x}_0| \leq \lambda \\ 0 & \text{otherwise} \end{cases}, \quad (7)$$

which selects out a (hyper)sphere of radius λ surrounding \vec{x}_0 . As above, each pairwise potential term will depend on three variables, allowing us to

neglect the other dimensions in the integral. Substituting this kernel back into Equation 1 and assuming $\lambda \leq |\vec{x}_0|$ gives

$$\mathcal{F}(x_0, \lambda) = \frac{3}{4\pi\lambda^3} \int_{x_0-\lambda}^{x_0+\lambda} dr \int_0^{\cos^{-1}\left(\frac{x_0^2+r^2-\lambda^2}{2x_0r}\right)} d\theta \int_0^{2\pi} F(r)r^2 \sin\theta d\phi \quad (8)$$

Performing the integrals over θ and ϕ yields

$$\mathcal{F}(x_0, \lambda) = \frac{3}{2\lambda^3} \int_{x_0-\lambda}^{x_0+\lambda} r \left(\frac{\lambda^2 - x_0^2 - r^2}{2x_0r} + 1 \right) F(r) dr \quad (9)$$

If $\lambda > |\vec{x}_0|$, the region of integration contains the origin, and the boundaries on the integral are different:

$$\mathcal{F}(x_0, \lambda) = \frac{3}{2\lambda^3} \left[\int_0^{\lambda-x_0} 2r^2 F(r) dr + \int_{\lambda-x_0}^{\lambda+x_0} r \left(\frac{\lambda^2 - x_0^2 - r^2}{2x_0r} + 1 \right) F(r) dr \right] \quad (10)$$

Together, Equations 9 and 10 are analogous to Equation 6. Given a function F , they allow us to derive a tophat-smoothed functional form \mathcal{F} . As with gaussian smoothing, F must diverge no faster than r^{-1} at the origin.

II-C Stophat Smoothing

The shifted tophat (or stophat) smoothing method is derived in much the same way as the tophat method. We proceed as we did before, but rather than selecting a sphere centered at \vec{x}_0 as in Equation 7, we center it at $\vec{x}_0 + \lambda\hat{z}$. Instead of Equation 8, we obtain for all λ

$$\begin{aligned} \mathcal{F}(x_0, \lambda) &= \frac{3}{4\pi\lambda^3} \int_{x_0}^{x_0+2\lambda} dr \int_0^{\cos^{-1}\left(\frac{(x_0+\lambda)^2+r^2-\lambda^2}{2(x_0+\lambda)r}\right)} d\theta r^2 \sin\theta \int_0^{2\pi} d\phi F(r) \\ &= \frac{3}{4\lambda^3} \int_{x_0}^{x_0+2\lambda} -r \left(\frac{x_0^2 + 2x_0\lambda + r^2 - 2r(x_0 + \lambda)}{x_0 + \lambda} \right) F(r) dr \end{aligned} \quad (11)$$

Unlike gaussian and tophat smoothing, stophat remains well-defined even for functions which diverge rapidly as $r \rightarrow 0$. This is because the integral contains no contribution from the potential function for $r < |\vec{x}_0|$.

Stophat smoothing is not strictly an average of the total potential, and cannot be derived from Equation 1. While each individual energy term is averaged according to a tophat kernel, the locus of integration varies for different atom pairs.

II-D Application of Stophat Smoothing to Forcefield Terms

Non-Bonded Terms

Stophat smoothing is directly applied to the non-bonded terms of standard forcefields simply by choosing F appropriately and solving the integral in Equation 11. For example, the stophat smoothed Coulomb function is calculated as

$$\begin{aligned} \mathcal{U}_{\text{coul}}(x_0, \lambda) &= \frac{3}{4\lambda^3} \int_{x_0}^{x_0+2\lambda} -r \left(\frac{x_0^2 + 2x_0\lambda + r^2 - 2r(x_0 + \lambda)}{x_0 + \lambda} \right) \frac{1}{r} dr \\ &= \frac{1}{x_0 + \lambda} \end{aligned} \quad (12)$$

Figure 1 shows the behavior of the Coulomb potential as λ increases.

Unlike gaussian smoothing methods, stophat can be directly applied to the Lennard-Jones 6-12 potential. Solving the integral

$$\mathcal{U}_{\text{LJ}}(x_0, \lambda) = \frac{3}{4\lambda^3} \int_{x_0}^{x_0+2\lambda} -r \left(\frac{x_0^2 + 2x_0\lambda + r^2 - 2r(x_0 + \lambda)}{x_0 + \lambda} \right) \left[\left(\frac{\sigma}{r} \right)^{12} - 2 \left(\frac{\sigma}{r} \right)^6 \right] dr \quad (13)$$

gives

$$\mathcal{U}_6(x_0, \lambda) = \frac{-2\sigma^6}{x_0^3(x_0 + 2\lambda)^3} \quad (14)$$

and

$$\mathcal{U}_{12}(x_0, \lambda) = \frac{\sigma^{12}(15x_0^6 + 90\lambda x_0^5 + 288\lambda^2 x_0^4 + 522\lambda^3 x_0^3 + 648\lambda^4 x_0^2 + 432\lambda^5 x_0 + 128\lambda^6)}{15x_0^9(x_0 + 2\lambda)^9} \quad (15)$$

for the r^{-6} and r^{-12} terms, respectively. Figure 2 demonstrates how the Lennard-Jones function is modified by stophat smoothing. Unlike gaussian smoothing, which tends to increase the effective van der Waal's radius, stophat smoothing moves the repulsive barrier closer to the origin.

Torsions

Torsion terms depend on the cosine of the torsional angle defined by the positions of four atoms. Because the smoothing derivations explicitly assume

that the energy function is pairwise and dependent only on distance, a different approach is needed. We adopt here the convention used by Pappu *et al.*, treating the torsion as a 1-dimensional function of the dihedral angle (16). Moreover, given the symmetry and periodicity of torsional angles, it seems logical to forgo the shifting used in the stophat method, and apply tophat smoothing instead. Accordingly, the smoothing integral is

$$\begin{aligned}
U_{\text{tors}}(\phi_0, \lambda_t) &= \frac{1}{2\lambda_t} \int_{\phi_0 - \lambda_t}^{\phi_0 + \lambda_t} \cos(n\phi) d\phi \\
&= \frac{1}{2n\lambda_t} [\sin(n\phi_0 + n\lambda_t) - \sin(n\phi_0 - n\lambda_t)] \\
&= \cos(n\phi_0) \frac{\sin(n\lambda_t)}{n\lambda_t}
\end{aligned} \tag{16}$$

In this notation, λ_t is the angular smoothing parameter in radians. Clearly, there are a variety of choices one could make to relate λ_t to the linear smoothing parameter λ . Pappu *et al.* chose this scaling factor based on the extent of diffusion along torsional degrees of freedom, a reasonable physical choice given their use of the diffusion equation method (17). A different choice is made here; since stophat smoothing averages over a specific lengthscale, we attempt to match the lengthscale of torsional averaging to that of the non-bonded terms. Specifically, we tracked the distance between the two atoms at the ends of the dihedral (the 1-4 distance), and chose

$$\lambda_t = \frac{\lambda}{\frac{\partial D_{14}}{\partial \phi}} \tag{17}$$

Using the notation shown in Figure 3, we can write the 1-4 distance as

$$D_{14} = \sqrt{d_1^2 + d_2^2 + d_3^2 - 2(d_1d_2 \cos \theta_1 + d_2d_3 \cos \theta_2 - d_1d_3[\cos \theta_1 \cos \theta_2 - \sin \theta_1 \sin \theta_2 \cos \phi])} \tag{18}$$

and the derivative as

$$\frac{\partial D_{14}}{\partial \phi} = \frac{d_1d_3 \sin(\theta_1) \sin(\theta_2) \sin(\phi)}{D_{14}} \tag{19}$$

Using realistic bond lengths and angle values, we computed the probability distribution function for $\frac{\partial D_{14}}{\partial \phi}$ over all values of the dihedral. Based on this distribution, we chose to use $\lambda_t/\lambda = 2.5 \text{ rad}/\text{\AA}$. It should be noted that

the results of optimizations using stophat smoothing are not particularly sensitive to this choice.

Stophat smoothing of torsions effectively amounts to a multiplication by $\sin(n\lambda)/(n\lambda)$, where n is the periodicity of the torsional potential. When $n\lambda = \pi$, the torsion term is zero for all values of ϕ . If $n\lambda > \pi$, the torsion term is nonzero, with its sign inverted. This undesirable behavior is an artifact of the periodicity of the torsion term. To eliminate it, we set all torsion energies for which $n\lambda \geq \pi$ to zero.

III Potential Smoothing and Search Procedure

The potential smoothing and search (PSS) protocol used here is essentially that of Pappu *et al.*, so we will merely summarize it here (17). The energy of an arbitrary starting structure is minimized on a series of potential energy surfaces of increasing deformation, until a surface with few minima is reached. At that point, the deformation is gradually reduced, with minimizations at each stage, so that the minimum from the highly deformed surface can be tracked back to the original undeformed surface. The technique can be made more effective by performing local searching as the deformation is reduced. Specifically, when the smoothing parameter λ is reduced below a certain threshold, we perform a normal mode search from the current local minimum. The normal modes are calculated from the Hessian. The system is successively translated along individual modes; once a maximum is passed the system is reminimized. If any of the resulting minima are lower than the original minimum, the procedure is repeated starting from the new minimum. These calculations were run using a version of the PSS program from TINKER modified to use stophat smoothing (18).

There are several choices one must make in the use of PSS, namely the number and placement of smoothing levels (known as the smoothing schedule), the maximum value for λ , the number of normal modes searched, the value of λ at which searching commences, and the degrees of freedom to be minimized. When using gaussian smoothing, we typically followed the choices made by Pappu *et al.*(17). For stophat smoothing, we used a larger maximum value for λ , but the same smoothing schedule and search protocol.

We used simulated annealing as an alternative global optimization

method, using the ANNEAL program from the TINKER simulation package (18). Except where otherwise noted, the molecule was subjected to 100 ps of dynamics at 1000 K, after which the temperature was linearly reduced to 0 K over the course of 100 ns. The final structure from the trajectory was then energy minimized using the truncated newton method. The RATTLE algorithm was used to constrain all bonds containing hydrogen to their equilibrium lengths. The temperature was controlled using the Berendsen weak-coupling method (19).

Several different potential functions were used for the present calculations, including OPLS, OPLS-AA, CHARMM, and AMBER (20; 21; 22; 23; 24). The bond and angle terms were not smoothed. When using gaussian smoothing, the Lennard-Jones function was approximated as a sum of two gaussians. Stophat smoothing does not require this approximation, and directly uses the Lennard-Jones functional form. To improve stability when searching the OPLS, OPLS-AA, and AMBER potentials, non-zero van der Waal’s parameters ($\sigma = 0.3, \epsilon = 0.01$) were added to prevent hydrogens from fusing with other atoms during minimization from highly strained structures. Additionally, for OPLS and OPLS-AA the AMBER-style trigonometric improper dihedrals were replaced with CHARMM-style harmonic functions. These changes do not significantly perturb the structures and relative energies of the minima. For the sake of clarity, all reported energies were calculated following minimization using the standard form for the potential function.

IV Results

The effect of smoothing the potential energy surface is to reduce the number of minima and the magnitude of the barriers between them. This results in a surface far more amenable to conformational searching and global optimization. Gaussian smoothing and stophat both accomplish this task, but in very different ways. In this section, we apply the PSS protocol to several biomolecular problems, and compare the effectiveness of gaussian and stophat smoothing.

IV-A Butane

The simplest molecule containing a rotational degree of freedom is a united atom model of butane. As such, it makes a good model system for comparing

the behavior of gaussian and stophat smoothing. We constructed a butane molecule using the united atom OPLS potential. By applying a 200 kcal/deg² harmonic restraint to the torsion at a series of values and performing energy minimizations, we were able to calculate the potential energy for the molecule as a function of conformation and smoothing level. As demonstrated by Figure 4, stophat and gaussian smoothing have qualitatively different behavior. The stophat behavior is very straightforward: as λ is increased, the barriers between minima diminish smoothly, and the *gauche* and *trans* minima disappear at roughly the same rate, until the energy curve is flat. The locations of the minima are largely unchanged.

The gaussian smoothing behavior is quite different. The amplitude of the *cis* barrier increases for small values of λ , and gradually widens, shifting the *gauche* minimum until it merges into the *trans* minimum. This is a consequence of the effective increase in van der Waal’s radius induced by gaussian smoothing. However, as λ grows beyond about 1 Å², the barrier starts to decrease. This is because the gaussian approximation to the Lennard-Jones potential is finite at the origin, and thus in the limit of large λ will smooth to zero at all distances. If it were possible to instead use the true Lennard-Jones potential, with a pole at the origin, the barrier would continue to grow and spread for all values of λ .

IV-B Polyalanine Peptides

The next set of examples we considered were varying lengths of capped polyalanine. It is a well-known result that the α -helix is the most stable *in vacuo* structure, for polyalanine beyond a certain length (25; 26). We examined polyalanine peptides from 8 to 40 residues in length, applying the PSS procedure using stophat and conventional gaussian smoothing. The peptides were constructed using ideal bond lengths and angles and all backbone dihedrals set to 180°. Only the backbone ϕ and ψ angles were varied in the course of the PSS calculations. All of the PSS calculations used a pentic smoothing schedule, with 250 levels. All normal modes were searched at each level while λ was being reduced. The maximum smoothing value for stophat smoothing was 80 Å. This value was chosen by tracking the number of minima for enkephalin using the OPLS potential. Gaussian smoothing calculations were performed using λ_{\max} values of 10 Å² (the value used by Pappu *et al.*) and 25 Å² (17). The OPLS united-atom forcefield, modified as discussed above, was used for all calculations. Upon completion of the PSS

procedure, a final minimization was performed, using the standard OPLS potential, during which the backbone ϕ , ψ , and ω dihedrals were varied. For comparison purposes, each peptide was also constructed with ideal α -helical angles ($\phi = -57^\circ$, $\psi = -47^\circ$, $\omega = 180^\circ$), and minimized in the the same way. All minimizations were run until the RMS gradient per torsion was 0.0001 kcal/mol-radian.

The results are summarized in Table I. For all peptides longer than 8 residues, the α -helix appears to be the global minimum energy structure. Gaussian smoothing with $\lambda_{\max} = 10\text{\AA}^2$ finds the correct answer for peptides up to 14 residues, after which it finds a much less favorable extended sheet structure. If $\lambda_{\max} = 25\text{\AA}^2$, this success can be extended to 16 residue peptides, after which this procedure fails in a similar fashion.

On the other hand, stophat smoothing finds the global minimum structure for all but one of the peptides considered; a helix-turn-helix was found for the 30 residue peptide. Moreover, the intermediate structures in the PSS calculation are quite different from those seen using gaussian smoothing. Gaussian smoothing is driven by the “pushing out” of the van der Waal’s potential, which causes the molecule to assume a maximally extended conformation at large λ . By contrast, van der Waal’s repulsion exists only at very short distances in stophat smoothing (see Figure 2), with the result that the molecules become more compact at high smoothing levels. The structures for intermediate λ values all have $i:i+2$ and $i:i-2$ hydrogen bonds. This can be seen in Figure 5, which shows the intermediate structures formed by the 22 residue peptide. This hydrogen bonding pattern appears early in the calculation, while the molecule is still extended, and is maintained as the molecule explores a variety of compact structures. It only disappears upon the formation of α -helical structure, at $\lambda = 1\text{\AA}$. This basic pattern is retained in all of the longer peptides, although there is some variation in the λ value where helices form.

One of the major advantages of gaussian smoothing in global optimization is its determinate nature. If the energy surface is deformed sufficiently, the starting structure does not affect the final answer. The maximally deformed structure is necessarily extended, making conformational trapping unlikely. It is not immediately clear that stophat must behave in the same way, since its maximally deformed structures are compact. To demonstrate that the stophat results are not due to the highly symmetric starting structure, we generated 20 different structures for the 16 residue peptide, with the ϕ and ψ angles chosen randomly from the β -sheet region of the Ramachandran plot,

and performed PSS on each, using the same protocol. To test the ability of stophat smoothing to undo existing favorable interactions, we also started a PSS calculation from the β -sheet structure produced by the gaussian smoothing PSS run. All of these calculations produced the same helical structure and energy as the calculation starting from the perfectly extended state.

IV-C Enkephalin

For our next test case, we attempted to locate the minimum energy structure of Met-Enkephalin, an unblocked peptide with the sequence Tyr-Gly-Gly-Phe-Met. We did the conformational search using all cartesian degrees of freedom, using the PSS procedure with gaussian and stophat smoothing, and using simulated annealing. Four separate forcefields were used, one using united atoms (OPLS), and three all-atom forcefields (OPLS-AA, CHARMM, and AMBER).

The all-atom forcefields have no potential term which penalizes chiral inversion. This is not a problem for most applications, since the barrier to inversion is of order 1000 kcal/mol, due mostly to angle strain. However, the normal mode search procedure used by PSS is athermal, and can cross these barriers. We chose to prevent this by adding flat-bottomed harmonic improper dihedral terms to each chiral center (the C α s for Tyr, Phe, and Met). Two terms were used for each center, each of which contained the C α and three of the four atoms bound to it. The potential term was zero for dihedrals between 0 and 70 degrees, and climbed harmonically with a spring constant of 1 kcal/mol-degree² beyond that. Physically reasonable conformations are rarely far from 35°, so these terms have no effect on the low-energy structures with correct chirality while penalizing inversions by approximately 2500 kcal/mol. This restraint is necessary because for two of the three all-atom potentials, the global minimum energy structure has a chiral inversion, which stophat detects.

The PSS protocols used for the various calculations all used 250 cubically distributed smoothing values, and involved searching all torsional modes once $\lambda < 5$. For gaussian smoothing, $\lambda_{\max} = 30\text{\AA}^2$, while for stophat smoothing $\lambda_{\max} = 80\text{\AA}$. More expensive PSS protocols, using more smoothing levels, larger λ_{\max} , and beginning searching at higher smoothing values, produced identical results. For comparison purposes, we also attempted to find the optimum enkephalin structure for each force field using simulated annealing. Forty independent trajectories were run using OPLS, while 20 trajectories

were run using each of the all-atom force fields. Several runs were required because simulated annealing is not a deterministic process and multiple trials can produce different results.

The results of the searches are summarized in Table II. In each case, PSS with stophat smoothing finds the lowest energy structure on the potential energy surface. However, in two cases (OPLS-AA and AMBER), the structures found involved inversions of chirality of a single C α (Tyr-1 for OPLS-AA, Met-5 for AMBER). These structures, while energetically most favorable, are not accessible to enkephalin in its biologically relevant form. Accordingly, the calculations were rerun with restraints designed to disfavor the inversions, as described above. The restrained calculations did not find the global minimum of the biologically-relevant surface. For AMBER, stophat found the second-best minimum while for OPLS-AA it found the third-best structure.

PSS with gaussian smoothing found the chirally-correct minimum on two of the potential surfaces (OPLS-AA and CHARMM), but not the others. It found the second best structure for united-atom OPLS, and the fourth-best for AMBER. Gaussian smoothing never produced a structure with inverted chirality.

The 100 ns simulated annealing runs find the minimum energy biologically-relevant structure some fraction of the time, with the other trajectories producing structures with energies as much as 3.5 kcal/mol higher. For OPLS-AA, we performed an additional 1 μ s annealing run, using an exponential cooling schedule. The structure it produced had an energy 2.3 kcal/mol higher than the global minimum, indicating that even a tenfold increase in trajectory time is not sufficient to render the annealing procedure deterministic.

Table III contains the backbone dihedrals for the structures discussed in Table II. Most of the structures found fall into two clusters. The global biologically-relevant minima on the OPLS-AA, CHARMM, and AMBER surfaces are quite similar. PSS with gaussian smoothing finds this same basic fold on the OPLS surface as well. The global minimum on this surface, found by stophat smoothing and simulated annealing, is quite different, involving a physically unlikely hydrogen bond between one of the oxygens of the C-terminus and that residue's amide hydrogen.

The structures found by stophat smoothing on the OPLS-AA surface with and without chirality constraints are also very similar. This is unlikely to be a coincidence, but rather reflects the effects of averaging on the surface.

When the surface is smoothed, each interaction is replaced an average of that interaction over a range of distances; the interactions which make the chirally inverted structure most favorable are averaged into the smoothed energy for the chirally correct structure, even if inversion is forbidden. The inverted structure found using AMBER is also similar to the one found with restraints.

It is also instructive to examine the path by which the different search methods find the global minimum energy structure. When applied to CHARMM enkephalin, all three methods considered find the same structure by very different paths. The qualitative differences in the intermediate structures can be seen by considering the molecule’s radius of gyration over the course of the search. Figure 6 plots the minimized structure for each PSS level for enkephalin using the CHARMM potential. The curve for gaussian smoothing shows that the intermediate structures are very extended. By contrast, the intermediate structures for stophat smoothing are very compact, more so even than the final structure. These behaviors are both driven by the deformation of the van der Waal’s potential; gaussian smoothing increases effective van der Waal’s radii, while stophat smoothing decreases them. Simulated annealing also increases the radius of gyration, although not as dramatically. In this case, the expansion is driven by entropic factors, specifically the greater conformational space available to an extended molecule.

IV-D Glycophorin

Previous work from this laboratory successfully applied the PSS procedure to the packing of rigid transmembrane helices (16). That work demonstrated that, starting from either experimental or ideal α -helices, this procedure correctly finds the global minimum energy structure on the OPLS potential surface, and that this structure is quite similar to the experimental one (C α RMS deviation of 0.64 and 0.73 Å, starting from experimental and ideal helices, respectively). We have repeated this work, using the PSS procedure with stophat smoothing. The protocol used was very similar to that used by Pappu *et al.*: 100 smoothing levels on a cubic schedule, searching all rigid body normal modes, minimizing until the gradient was 0.0001 kcal/mol-Å per degree of freedom, with $\lambda_{\max} = 10\text{Å}$ (16). For the purposes of these calculations, electrostatics were ignored; the potential function consisted solely of interhelical van der Waal’s interactions.

This procedure correctly locates the global minimum structure independent of the starting conformation. The structure generated differs very slightly from that reported by Pappu *et al.*, because those calculations used a 2-gaussian approximation for the Lennard-Jones function as opposed to the more conventional 6-12 functional form used with stophat smoothing (16; 17). These structures are nearly identical, with a heavy-atom RMS difference of 0.02 Å. Re-minimizing the gaussian smoothing structure using the pure OPLS potential produces the structure calculated with stophat smoothing, with an energy -39.4632 kcal/mol. This structure is shown in Figure 7.

V Discussion

Many of the standard smoothing methods, such as the Diffusion Equation Method (2), the Effective Diffused Potential (8), Gaussian density annealing (6), and Gaussian Packet Annealing (11), have a built in physical interpretation: atoms are represented as gaussian probability distributions, which smooths the potential energy surface. Alternatively, one can focus on the energy landscape, and the effects these methods have upon it. Gaussian smoothing as applied using the diffusion equation is an averaging procedure, where each point on the many-dimensional potential surface is replaced with an average of that surface, weighted by a gaussian centered at that point. Because this scheme weights nearby points more heavily, the change in the energy of a conformation as the smoothing value is increased yields information about similar conformations. For example, the relative favorability of low entropy conformations (narrow minima surrounded by steep walls) will be dramatically lessened on smoothed surfaces. The analogy between this behavior and that of temperature in simulated annealing has been explored elsewhere (27).

Moreover, it may be possible to interpret the overall structure of an energy landscape by examining the way minima merge together as the surface deformation is increased. This interpretation is complicated by the fact that gaussian smoothing has no single lengthscale, in that the gaussian smoothing kernel extends over all of conformational space. While it is true that distant points on the energy surface have less effect than the immediate environment, there is no clean dividing line.

Smoothing methods exist which have a finite transformation lengthscale. For example, the tophat kernel used in the method of bad derivatives has a

finite range, unlike a gaussian. The deformed energy surface corresponds to a local average, calculated by averaging the $(3N - 6)$ -dimensional potential surface over the volume of a hypercube with sides of length λ . This formulation is advantageous, in that it is unnecessary to compute an analytic form for the deformed surface; the original, undeformed surface is evaluated at the corners of the cube, and the gradient of the smoothed surface is calculated in the form of a numerical derivative. Since only energies are involved in this process, any scoring function can be used, even if it is discontinuous or hard to differentiate. For example, the method of bad derivatives can be applied directly to the Boltzmann factor $\exp(-E/k_B T)$, which is be extremely difficult with other smoothing procedures. Unfortunately, the method of bad derivatives is computationally costly, with gradient evaluation scaling $O(N^3)$ in the number of atoms, as opposed to $O(N^2)$ for standard molecular mechanics forcefields. Moreover, the volume of integration lacks symmetry; using a cube instead of a sphere imposes orientation dependence on the averaging on the procedure.

Tophat smoothing (derived in Section II-B) can be viewed as the analytic equivalent of the method of bad derivatives. It is once again a local average, calculated by integrating the total potential over a hypersphere. The analytic functional forms for smoothing the coulomb and van der Waal’s potentials are complicated, but energy calculation scales $O(N^2)$. This makes it a very promising method for studying the properties of the potential hypersurface.

Unlike tophat and gaussian smoothing, stophat smoothing is not an average over the total potential energy hypersurface. Each individual energy term is averaged over a sphere, but the center of that sphere is shifted outward toward larger interatomic separations. As a result, different interactions are averaged over distinct regions of the $(3N - 6)$ -dimensional space. This complicates physical interpretation of the smoothed surface. However, stophat smoothing appears to have some practical benefits, as described above. Moreover, there is some precedent for smoothing methods not involving averaging. For example, the distance scaling method operates by redefining the interatomic distance r before inserting it into the standard potential functions (28). Other workers have used functional forms with softer repulsive barriers to moderately smooth the surface (29).

The behavior of a van der Waal’s interaction under gaussian smoothing is dominated by the peak at the origin. As λ increases, the peak figures significantly into the averaging at larger atomic separations, with the result that the repulsive wall “pushes outward”. During smoothing the function’s

behavior at physically relevant distances (roughly $r \geq 0.8\sigma$) is polluted by the very high peaks at short separations. As such, these short separations, which are not physically or thermodynamically relevant, affect the energetic favorability of otherwise physically reasonable structures. As a result, gaussian smoothing disfavors compact structures in favor of extended ones, since the latter have fewer close van der Waal’s contacts. This effect would be even more pronounced if it were possible to use a true Lennard-Jones function instead of the 2-gaussian approximation; the gaussian approximation becomes zero everywhere as $\lambda \rightarrow \infty$, while a Lennard-Jones function would simply continue to push out, since the pole could not be smoothed away. To a lesser extent, this effect is seen if one compares the smoothing behavior of the 2-gaussian Lennard-Jones fit to that of a 4-gaussian fit. Because the latter has a much higher value at zero separation, the degree of push out for a given value of λ is greatly increased.

Gaussian smoothing diminishes the number of realistic minima on a surface via two distinct mechanisms: 1) it flattens the surface, diminishing the impact of torsional rotation and favorable long-range interactions, and 2) it diminishes the volume of accessible conformations, by effectively expanding the van der Waal’s radii until nearly all conformations are excluded. This behavior can be inferred from Figure 4, if one considers butane to be a model for a torsion within a larger molecule. All of the minima found on a significantly deformed surface will be extended; even those minima which can be tracked to compact conformations on the undeformed surface are greatly expanded. For example, Figure 6 shows that even at moderate smoothing values, the enkephalin molecule has a very large radius of gyration, compared to dynamics at 1000K.

By contrast, stophat smoothing explicitly ignores the short-range behavior of the energy functions; *the actual interatomic separation defines the minimum distance involved in the averaging*. Instead of pushing out, the van der Waal’s barrier moves to shorter distances with increasing λ (see Figure 2). As a result, the volume of accessible conformational space is increased by stophat smoothing. As a result, stophat smoothing does not decrease the number of minima as dramatically as gaussian smoothing; even at relatively large smoothing levels, both compact and extended minima are present. As λ gets very large, the energy surface has a single, highly degenerate minimum, where all bonds and angles are at their ideal values, and the molecule is free to rotate about its torsions. The only states forbidden are those where two atoms lie directly on top of each other. If λ is reduced, there are immediately

a significant number of minima, far more than on a corresponding gaussian-deformed surface. It is not clear that this is a desirable behavior, since it appears that this would cause a PSS procedure to depend on the starting structure. We have not experienced such difficulties, perhaps because the barriers between the states are reduced at all deformations such that the normal-mode search is able to compensate.

In choosing between various algorithms for global optimization, two criteria are of primary importance: efficacy and computational cost. The above results demonstrate that the PSS procedure using stophat smoothing is a good global minimizer for a variety of molecular problems. However, to be considered a useful alternative to standard optimization techniques, it must also be at least comparably efficient computationally. We compared the computational cost for optimizing enkephalin using the CHARMM potential, the only example where all three methods successfully locate the global minimum. The results are summarized in Table IV. It is clear that the cost per trial of PSS against 100 ns of dynamics favors smoothing. However, one can run simulated annealing trajectories of any length, so the relevant quantity is the effective cost to find the global minimum. We define the effective cost as the cost per annealing trajectory multiplied by the number of independent trials required to give a 90% chance of finding the global minimum at least once, based on the success rates given in Table II. Using this metric, it becomes clear that PSS is far more efficient than simulated annealing. Moreover, the balance would be tilted further in favor of smoothing for cases like OPLS enkephalin, where the simulated annealing success rate was lower.

It could be argued that the computational cost would be reduced by using a smaller number of longer simulated annealing trials. However, the 1 μ s annealing run on enkephalin required more than 200 CPU hours and still missed the global minimum by over 2 kcal/mol. This suggests that simulated annealing trajectories long enough to consistently find the global minimum are not computationally tractable for a molecule the size of enkephalin.

It is expected that the relative efficiency of smoothing over simulated annealing will only increase with increasing problem size. The computational cost of the local minimizations in PSS scale roughly as $O(N^3)$ with the number of atoms, while the amount of searching required scales at worst linearly with the number of degrees of freedom, resulting in an $O(N^4)$ algorithm. Analysis of the computational cost of simulated annealing is more complicated. The number of conformational transitions required for equilibration at each temperature scales $O(M^2)$, where M is the number of states available to

the system (30). Since M in turn generally scales exponentially with system size for molecular examples, the algorithm as a whole scales exponentially.

VI Conclusion

The present work describes the derivation and performance of the stophat smoothing method when applied to a variety of biomolecular examples. It is clear that stophat smoothing is an effective tool for global optimization as part of a potential smoothing and search protocol. As a result, it should be a valuable tool for predicting the packing of transmembrane helices (16). Furthermore, the fact that stophat smoothing lowers all barriers between conformations suggests that it could be coupled to stochastic sampling methods such as molecular dynamics and Monte Carlo to efficiently generate thermodynamic averages.

VII Acknowledgements

This work was funded by the National Institutes of Health under grant GM58712 to J.W.P. A.G. was supported in part by an NIH NRSA fellowship (NS42975). We would like to thank Dr. Pengyu Ren for many helpful discussions and comments on the manuscript, and Michael Schneiders for comments on the manuscript.

References

- [1] Kirkpatrick, S.; Gelatt, C. D.; Vecchi, M. P., *Science* 1983, 220, 671–680.
- [2] Piela, L.; Kostrowicki, J.; Scheraga, H. A., *J Phys Chem* 1989, 93, 3339–3346.
- [3] Kostrowicki, J.; Piela, L.; Cherayil, B. J.; Scheraga, H., *J Phys Chem* 1991, 95, 4113–4119.
- [4] Kostrowicki, J.; Scheraga, H. A., *J Phys Chem* 1992, 96, 7442–7449.
- [5] Pillardy, J.; Liwo, A.; Scheraga, H. A., *J Phys Chem A* 1999, 9370–9377.
- [6] Ma, J.; Straub, J. E., *J Chem Phys* 1994, 101, 533–541.
- [7] Amara, P.; Straub, J. E., *J Phys Chem* 1995, 99, 14840–14853.
- [8] Schelstraete, S.; Verschelde, H., *J Phys Chem A* 1997, 310–315.
- [9] Verschelde, H.; Schelstraete, S.; Vandekerckhove, J.; Verschelde, J. L., *J Chem Phys* 1997, 106, 1556–1568.
- [10] Schelstraete, S.; Verschelde, H., *J Chem Phys* 1998, 108, 7152–7160.
- [11] Shalloway, D., in C. A. Floudas; P. M. Pardalos, editors, *Recent advances in global optimization*, Princeton University Press, 1992 433–477, 433–477.
- [12] Piela, L., *Collect Czech Chem Commun* 1998, 63, 1368–1380.
- [13] Schelstraete, S.; Schepens, W.; Verschelde, H., in P. B. Balbuena; J. M. Seminario, editors, *Molecular Dynamics From Classical to Quantum Methods*, Elsevier Science, 1999 129–186, 129–186.
- [14] Andricioaei, I.; Straub, J. E., *J Comput Chem* 1998, 19, 1445–1455.
- [15] Moré, J. J.; Wu, Z., in P. Pardalos; D. Shalloway; G. Xue, editors, *DI-MACS Series in Discrete Mathematics and Theoretical Computer Science: Global Minimization of Nonconvex Energy Functions: Molecular Conformation and Protein Folding*, American Mathematical Society, 1995 151–168, 151–168.

- [16] Pappu, R. V.; Marshall, G. M.; Ponder, J. W., *Nat Struct Biol* 1999, 6, 50–55.
- [17] Pappu, R. V.; Hart, R. K.; Ponder, J. W., *J Phys Chem B* 1998, 102, 9725–9742.
- [18] Ponder, J. W., *TINKER: Software Tools for Molecular Design*, v. 3.8 1998, current version available from <http://dasher.wustl.edu/tinker/>.
- [19] Berendsen, H. J. C., P.; M., J. P.; van Gunsteren, W. F.; DiNola, A.; Haak, J. R., *J Chem Phys* 1984, 81, 3684–3690.
- [20] Jorgensen, W. L.; Tirado-Rives, J., *J Am Chem Soc* 1988, 110, 1657–1666.
- [21] Jorgensen, W. L.; Maxwell, D. S.; Tirado-Rives, J., *J Am Chem Soc* 1996, 11225–11236.
- [22] MacKerell, Jr., A. D.; Bashford, D.; Bellott, M.; Dunbrack Jr., R.; Evanseck, J.; Field, M.; Fischer, S.; Gao, J.; Guo, H.; Ha, S.; Joseph-McCarthy, D.; Kuchnir, L.; Kuczera, K.; Lau, F.; Mattos, C.; Michnick, S.; Ngo, T.; Nguyen, D.; Prodhom, B.; Reiher, III, W.; Roux, B.; Schlenkrich, M.; Smith, J.; Stote, R.; Straub, J.; Watanabe, M.; Wiorkiewicz-Kuczera, J.; Yin, D.; Karplus, M., *J Phys Chem B* 1998, 102, 3586–3616.
- [23] Foloppe, N.; MacKerrell, Jr., A. D., *J Comput Chem* 2000, 21, 86–104.
- [24] Cornell, W. D.; Cieplak, P.; Bayly, C. I.; Gould, I. R.; Merz, K. M., J.; Ferguson, D. M.; Spellmeyer, D. C.; Fox, T.; Caldwell, J. W.; Kollman, P. A., *J Am Chem Soc* 1995, 117, 5179–5197.
- [25] Scott, R. A.; Scheraga, H. A., *J Chem Phys* 1966, 45, 2091–2101.
- [26] Ooi, T.; Scott, R. A.; Vanderkooi, G.; Scheraga, H. A., *J Chem Phys* 1967, 46, 4410–4426.
- [27] Hart, R. K.; Pappu, R. V.; Ponder, J. W., *J Comput Chem* 2000, 21, 531–522.
- [28] Pillardy, J.; Piela, L., *J Comput Chem* 1997, 18, 2040–2049.

- [29] Crivelli, S.; Byrd, R.; Eskow, E.; Schnabe, R.; Yu, R.; Philip, T. M.; Head-Gordon, T., *Comput Chem* 2000, 489–497.
- [30] Aarts, E. H. L.; Korst, J. H. M.; van Laarhoven, P. J. M., in E. Aarts; J. K. Lenstra, editors, *Local Search in Combinatorial Optimization*, John Wiley and Sons Ltd., 1997 .




Figure 1: The coulomb potential smoothed with stophat smoothing. Stophat smoothing shifts the potential by λ .




Figure 2: The Lennard-Jones potential smoothed with stophat smoothing. The location of the repulsive wall moves inward with increasing λ . With conventional forms of smoothing, the repulsive wall moves outward.




Figure 3: Definition of D_{14} in terms of internal coordinates.

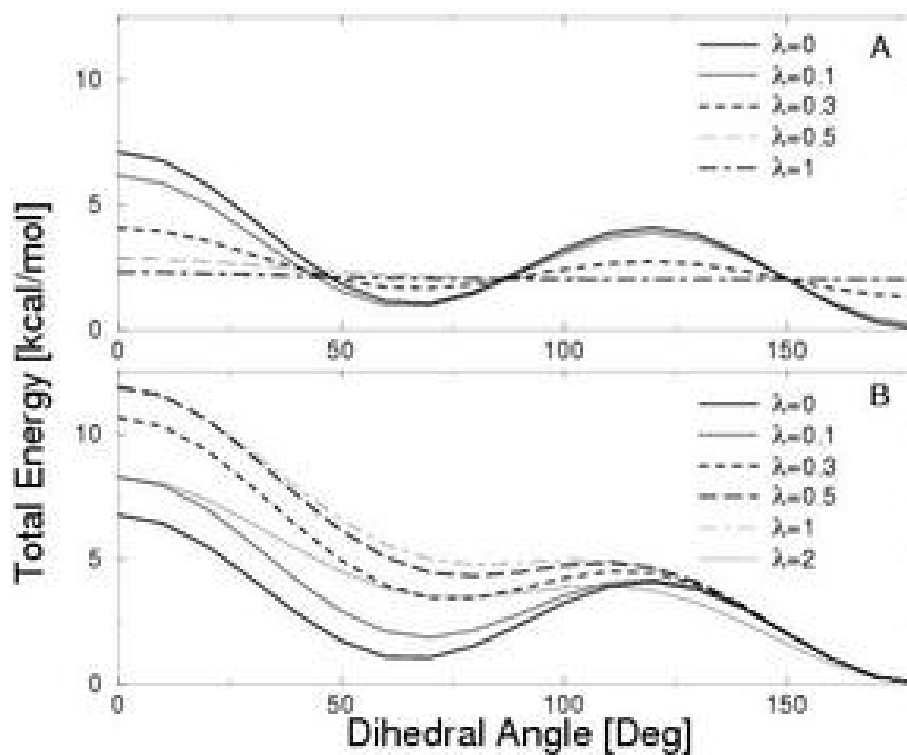


Figure 4: Potential energy of united atom butane. Part A shows the potential as deformed by stophat, Part B by gaussian smoothing. The gaussian smoothing energy curves in Part B have been shifted such that the energy at $\phi = 180^\circ$ remains constant.




Figure 5: Structures assumed by a capped alanine 22-mer during PSS smoothing. Structures A to F are the final structures from levels 148 (forward smoothing), 268, 350, 395, 396, and 450 ($\lambda = 5.81, 55.06, 6.22, 1.05, 1.00, 0.03$ Å). The first structure exhibits the $i:i+2$ and $i:i-2$ hydrogen bonding pattern. This pattern is preserved in each of the first four structures, as the molecule gradually becomes more compact, forming a knot-like structure at level 395. At level 396, the first helical structure appears. At level 450, the helix is fully formed.

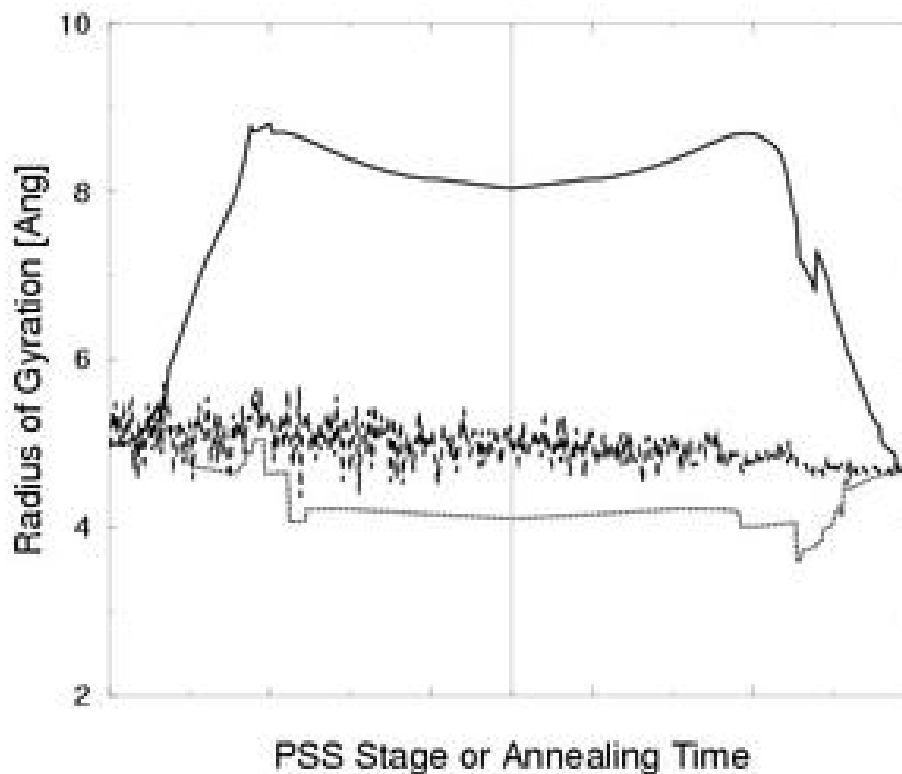


Figure 6: Radius of gyration for enkephalin over the course of sample PSS and simulated annealing trials using the CHARMM potential. The solid and dotted lines are from 500-level PSS runs using gaussian and stophat smoothing, respectively. The dashed line is from a 100 ns simulated annealing run. The midpoint of the PSS calculations, where the smoothing value λ is at its maximum, is marked by the vertical line. All three calculations located the global minimum structure.




Figure 7: Global minimum energy structure for the glycophorin A transmembrane helix dimer, found using PSS with stophat and gaussian smoothing.

Length	Helix	Gaussian ($\lambda = 10$)		Gaussian ($\lambda = 25$)		Stophat	
8	-292.6184	-293.3406	T			-293.3329	T
10	-369.8012	-369.7823	H			-369.7856	H
12	-447.3603	-447.2697	H			-447.2723	H
14	-525.1728	-525.0305	H			-525.1663	H
16	-603.1383	-582.6604	T	-602.9653	H	-603.1346	H
18	-681.2149	-652.2730	T	-652.2909	T	-681.1988	H
20	-759.3669	-727.2988	T	-727.3048	T	-759.3580	H
22	-837.5756			-837.5812	H	-837.5813	H
24	-915.8280			-871.9698	T	-915.8324	H
26	-994.1128			-941.8530	T	-994.0976	H
28	-1072.4252					-1072.4017	H
30	-1150.7579					-1115.7017	H2
40	-1542.6341					-1542.6531	H

Table I: OPLS capped polyalanine minimized using PSS with torsional minimization and searching. The first column is the number of residues. The second column is the energy for the structure built as an ideal helix, then minimized in torsion space. The next three columns report the results of PSS minimization from an extended conformation, for standard gaussian smoothing with $\lambda_{\max} = 10\text{\AA}^2$, gaussian smoothing with $\lambda_{\max} = 25\text{\AA}^2$, and stophat smoothing. “H” and “T” indicate helices and extended β -turns, respectively. “H2” denotes that the minimization found a helix-turn-helix structure.

Model	Method	Energy	Frequency
OPLS	Stophat	-304.2242	
	Gaussian	-304.0738	
	Annealing	-304.2242	0.15
OPLS-AA	Stophat*	-264.6201	
	Stophat	-261.3666	
	Gaussian	-263.3354	
	Annealing	-263.3354	0.50
CHARMM	Stophat	-66.4162	
	Gaussian	-66.4162	
	Annealing	-66.4162	0.25
AMBER	Stophat*	-167.1609	
	Stophat	-162.7474	
	Gaussian	-162.3629	
	Annealing	-164.2013	0.30

Table II: Energies for the enkephalin structures produced using different forcefields and search methods. The first column indicates the potential function used. The second column indicates the method used to generate the structure (PSS using stophat and gaussian smoothing, simulated annealing). The structures labelled “Stophat*” have C α s with inverted chirality (Tyr-1 for OPLS-AA, Met-5 for AMBER), while the rows labelled “Stophat” were run with restraints to prevent chiral inversion. The “Frequency” column indicates the fraction of the simulated annealing trials that found the lowest energy structure. The success frequency is either 1 or 0 for all PSS methods, since they are deterministic.

Model	Method	Tyr-1	Gly-1		Gly-3		Phe-4		Met-5	
		ψ	ϕ	ψ	ϕ	ψ	ϕ	ψ	ϕ	ψ
OPLS	Stophat	-70	127	44	-163	-59	-124	-11	-150	-175
	Gaussian	-47	-81	-57	-113	-114	-74	-46	-104	129
OPLS-AA	Stophat*	47	105	78	109	-54	-81	60	66	72
	Stophat	61	118	63	111	-51	-78	59	54	-114
	Gaussian	-37	-92	-50	-116	-128	-65	-48	-63	-58
CHARMM	All	-46	-87	-39	-127	-123	-69	-49	-65	-48
AMBER	Global	-60	-80	-32	-131	-124	-64	-54	-55	122
	Stophat*	89	96	19	147	6	-148	61	49	-136
	Stophat	74	137	21	117	-38	-73	62	39	-132
	Gaussian	132	-161	41	83	-54	-65	176	-50	138
ECEPP	Global	155	-160	71	80	-74	-117	12	-157	159

Table III: Backbone dihedrals for enkephalin structures produced using different forcefields and search methods. The forcefields and methods are labelled as in Table II. For AMBER, the backbone dihedrals for the global minimum are also shown. All search methods produced the global minimum structure for CHARMM. Backbone dihedrals for the global minimum for ECEPP-3 are taken from Reference (4).

Method	Time	Effective Cost
Stophat	4.5	
Gaussian	5.6	
Annealing	29.8	238.4

Table IV: Computational cost of PSS and simulated annealing on enkephalin using the CHARMM potential. The values are real elapsed times in hours for computations performed on a 950 MHz Athlon PC running Linux, using a modified version of TINKER 3.9 compiled with the Portland Group Fortran compiler. The last column is an estimate of the effective cost of running a sufficient number of simulated annealing runs to give a 90% chance of finding the minimum energy structure.