

Molecular Conformer Search with Low-Energy Latent Space

Xiaomi Guo, Lincan Fang, Yong Xu, Wenhui Duan, Patrick Rinke, Milica Todorović,* and Xi Chen*



Cite This: *J. Chem. Theory Comput.* 2022, 18, 4574–4585



Read Online

ACCESS |



Metrics & More

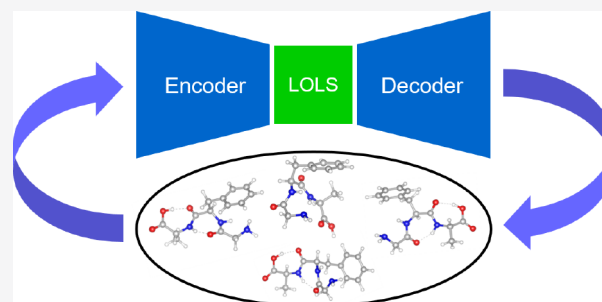


Article Recommendations



Supporting Information

ABSTRACT: Identifying low-energy conformers with quantum mechanical accuracy for molecules with many degrees of freedom is challenging. In this work, we use the molecular dihedral angles as features and explore the possibility of performing molecular conformer search in a latent space with a generative model named variational auto-encoder (VAE). We bias the VAE towards low-energy molecular configurations to generate more informative data. In this way, we can effectively build a reliable energy model for the low-energy potential energy surface. After the energy model has been built, we extract local-minimum conformations and refine them with structure optimization. We have tested and benchmarked our low-energy latent-space (LOLS) structure search method on organic molecules with 5–9 searching dimensions. Our results agree with previous studies.



INTRODUCTION

Organic molecules are typically very flexible, and any molecule with rotatable bonds can adopt multiple energetically accessible conformations, each associated with different chemical and electronic properties.^{1,2} Identifying the low-energy molecular conformers and determining their energy ranking is therefore a topic of great importance in computational chemistry,³ cheminformatics,⁴ computational drug design,⁵ and structure-based virtual screening.⁶ However, the dimension of configurational spaces and the complexity of energy landscapes increases drastically with the size of the molecule. This makes molecular conformer search one of the persistent challenges in molecular modeling.^{1,7}

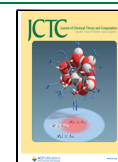
A variety of methods and tools have been developed for molecular conformer search. Systematic methods use a grid to sample all possible torsion angles in a molecule. This approach is deterministic but limited to small molecules due to its poor scaling with increasing search dimensions. Conversely, methods such as Monte Carlo annealing,⁸ minima hopping,⁹ basin hopping¹⁰ and genetic algorithms¹¹ sample configurational space stochastically. Stochastic methods can be applied to larger molecules with high-dimensional search spaces, but due to the random nature of the process, extensive sampling is required to achieve convergent results. To balance the accuracy and computational cost, hierarchical methods which first scan a large portion of configurational space, and then refine the promising candidate with more costly and accurate computations have been developed.^{12,13} Since simulation methods at different levels of accuracy may predict different potential energy surfaces (PES), a large number of structures still needs to be optimized at the higher level to avoid missing the true low-energy conformers.¹²

In recent years, machine learning techniques such as artificial neural networks,^{14,15} Gaussian process regression (GPR),^{16–19} and machine-learned force fields²⁰ have been successfully applied to accelerate structure-to-energy predictions and geometry optimization for molecules. However, most of these schemes require large data sets, usually costly to compute with ab initio methods.

In our recent work, we presented a new approach based on Bayesian Optimization and quantum chemistry methods for molecular conformer identification and ranking.²¹ We first kept all bond lengths and angles fixed, and selected the dihedral angles as the features to form the search space. Then we employed the BOSS code^{22,23} to actively learn the PES of the molecule by Bayesian Optimization iterative data sampling. After the PES converged, we analyzed the PES to extract the local minima locations and related structures, and optimized the structures with density functional theory (DFT) and other post-processings. We have tested our method on cysteine, serine, tryptophan, and aspartic acid. The method shows both high accuracy and efficiency, and can be easily automated for extensive searches. The excellent efficiency is partly due to learning the PES in the reduced conformational space of dihedral angles and only refining the local minima structures with DFT, and partly because Bayesian Optimization creates small and compact data sets. However, our method is not directly transferable to molecules with high-dimensional search

Received: March 23, 2022

Published: June 13, 2022



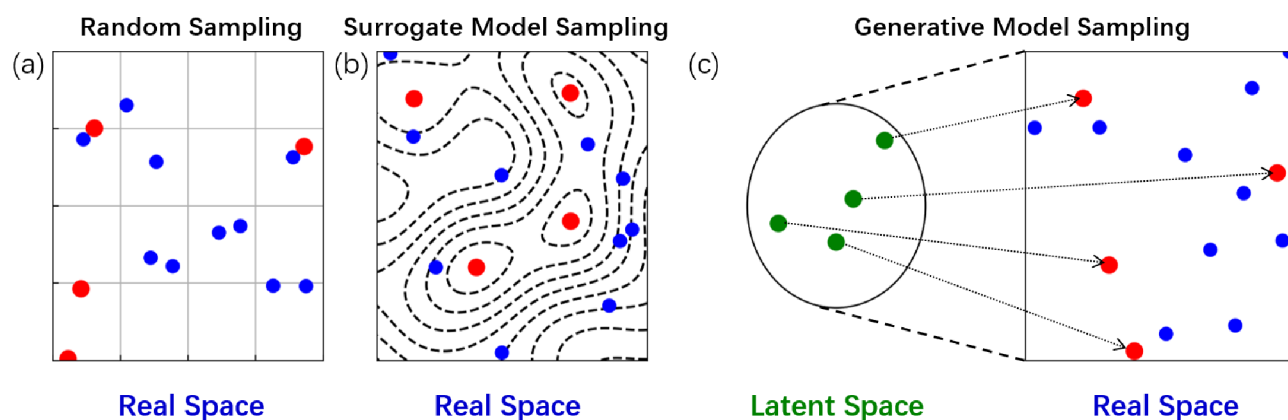


Figure 1. Schematic illustration of sampling methods. The blue and red dots represent acquired samples and candidates for the next sampling steps. In (a) the candidates are randomly picked and have no relation with already acquired samples. The dash lines in (b) represent the contour lines of the surrogate model which is fitted to the acquired samples and the local maxima or minima of the model will be the next acquisition candidates. The green dots in (c) represent samples in latent space. The generator maps them to real space.

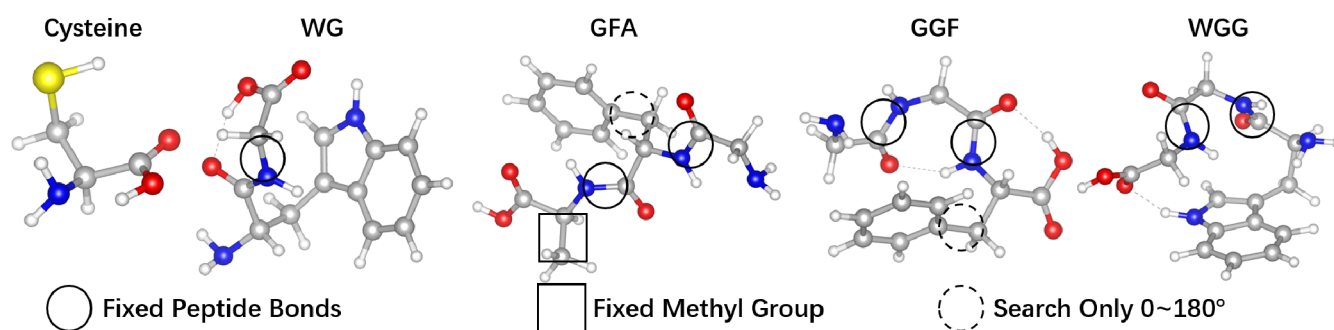


Figure 2. Ball-and-stick models of cysteine, tryptophyl-glycyl (WG), glycyl-phenylalanyl-alanyl (GFA), glycyl-glycyl-phenylalanyl (GGF) and tryptophyl-glycyl-glycyl (WGG). Red atoms denote oxygen, white hydrogen, gray carbon, blue nitrogen, and yellow sulfur. The dashed circle mark the dihedral angles that have a reduced search range of $[0^\circ, 180^\circ]$. The solid circles and squares mark peptide bonds and dihedral angles that are kept fixed during sampling. All other dihedral angles belong to our space with their full range $[0^\circ, 360^\circ]$.

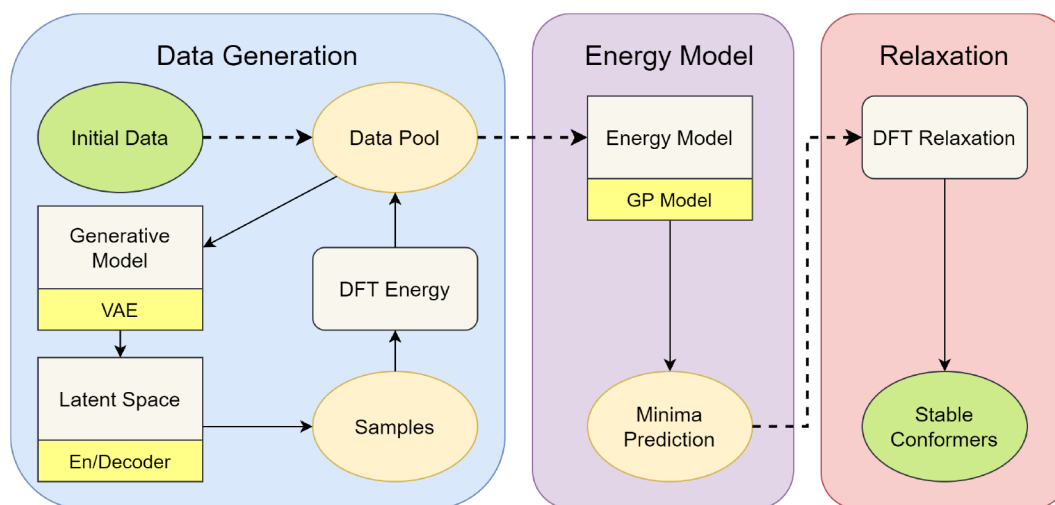


Figure 3. The LOLS workflow starts from initial data and finishes with the structures and energies of stable conformers. The eclipses represent data, the rectangles machine learning models, and the rectangles with round corners represent DFT calculations.

spaces. The data required for building reliable PESs increases rapidly with search dimensions. With increasing data set size, the cost to compute the necessary data with quantum mechanical methods and to build the surrogate model of the PES in BOSS grows and eventually becomes prohibitively expensive.

To address this challenge, we will explore the possibility of using a generative model to acquire samples in a latent space for

molecular conformer search. We decided on variational autoencoders (VAEs) as the generative model, because the neural network structure of VAEs is typically simple; and VAEs are equipped with a regularization term in the loss function to prevent over-fitting. VAEs combine an encoding neural network (encoder) with a decoding neural network (decoder). The encoder compresses data from real space (here the space of

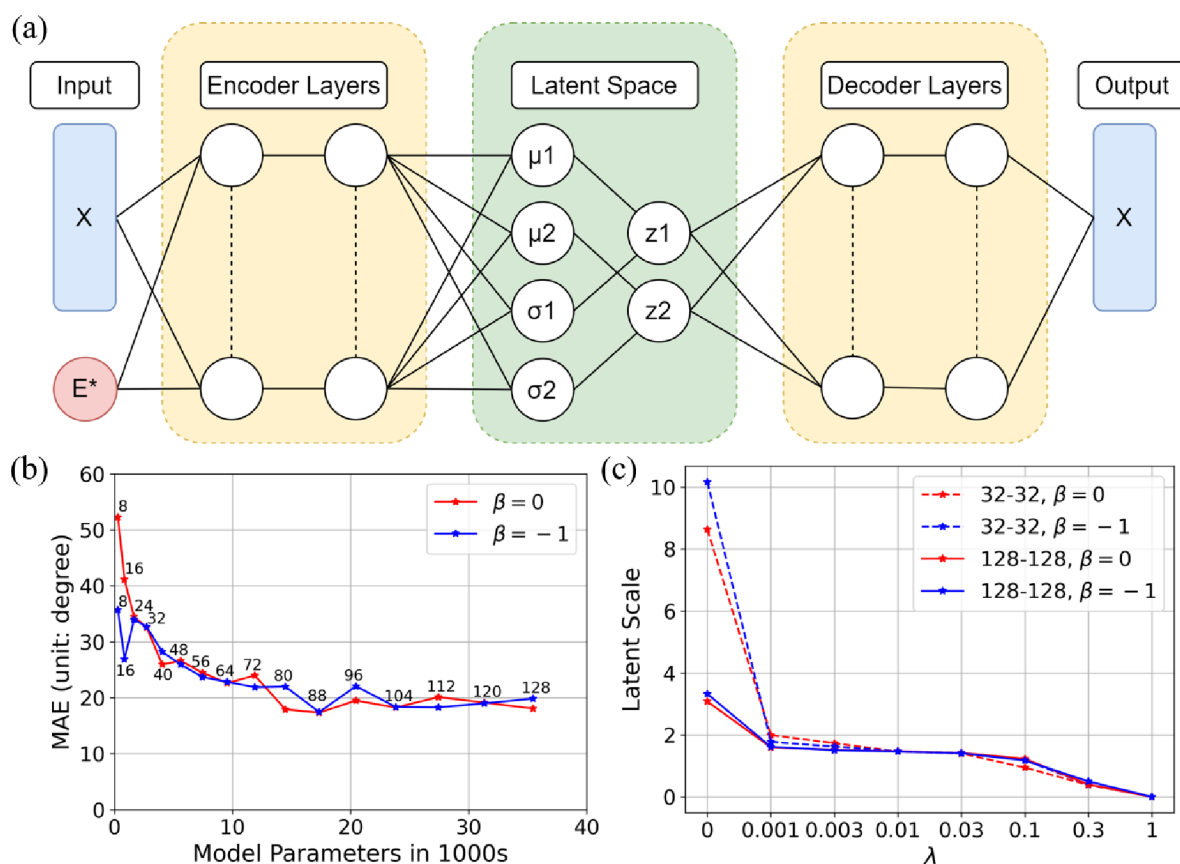


Figure 4. (a) Architecture of our variational auto-encoder. The input X consists of molecular dihedral angles, and input E^* is the scaled energy of the corresponding structure. The output X are again dihedral angles, but no energy. Both encoder and decoder have two layers with the same layer size. The mean (μ) and variance (σ^2) are the outputs of the encoder and the normal distribution is taken for trainable samples z . (b) The mean absolute difference (MAE) between input and output dihedral angles for the VAE models with different numbers of trainable parameters. The text near the data points shows the numbers of neurons of each fully connected layer in the en/decoder (*layersize*). During the test of *layersize*, λ is fixed to 0.01. (c) The relationship between the latent-space scale L and the hyperparameter λ . The test were performed with 32–32, 128–128 neural networks and $\beta = 0, -1$.

dihedral angles) into a latent space. This compression ideally retains the essential data correlations in the reduced representation. The decoder maps latent vectors back to the original representation. Figure 1 illustrates how sampling in latent space with a generative model (c) differs from conventional random sampling in real space (a) and from our previous approach of employing a surrogate model and an acquisition strategy (b).

To sample more efficiently with our generative approach, we are steering the VAE towards low-energy molecular configurations during the training. The latent space then predominantly encodes information on the relevant, low-energy region of the PES. As in previous work, we use dihedral angles to represent the different molecular conformations. We also extract local minima structures and apply structure optimization only after a meaningful PES has been learned.

In brief, in this work we designed a low-energy latent-space (LOLS) structure search method for molecular conformer search and determined appropriate settings and suitable hyperparameters for it. We tested LOLS on cysteine and four peptides tryptophyl-glycyl (WG), glycyl-phenylalanyl-alanyl (GFA), glycyl-glycyl-phenylalanyl (GGF) and tryptophyl-glycyl-glycyl (WGG) (Figure 2). The main reasons for choosing these molecules are: First, amino acids and peptides are important biomolecules. Second, peptides are very flexible and exhibit complex PESs, making them a challenging system for

conformer search. Third, previous studies provide reference data.^{21,24,25} Another objective of our work is to gain insight into the nature and properties of latent space. For this, we visualize and analyze the latent spaces of cysteine and GFA. Our method and our results will be presented in the following sections.

METHODS

Our LOLS method consists of three steps (Figure 3). In step 1, we employ an active learning approach to generate data on-the-fly. We combine two strategies to steer the generative model towards generating more low-energy data, which helps us build a compact and reliable model for the low-energy regions of the PES. Strategy one is data processing. We scale the energy of training data with a non-linear function and exclude high-energy data. Strategy two attributes more weight to lower energy data in the loss function of the generative model. Both strategies will be discussed in the following sections. In step 2, we build a Gaussian process (GP) regression model in real space. We extract the local minima from the GP and use them to initialize DFT geometry optimizations. In step 3, the candidate structures are further optimized with DFT structure relaxation. Details of our method will be explained in the following sections.

Data Generation Loop. The left part of Figure 3 shows the data generative loop we designed for sampling informative data. The “data pool” is initialized with an initial data set, in which each data point represents the dihedral angles and DFT energy

of a conformation. Then we set up an active learning approach and iteratively acquire samples from the latent space. For each new sample, the structural features are decoded by the VAE into real space, then the energy is calculated with DFT. As we add new samples to the data pool, we keep retraining the VAE.

Each time we carry out three parallel runs to average out the effects of randomization in the sampling method, and continue the data generation loop up to a preset maximum number of iterations. If the global minimum and at least 70% of the reference targets are found, we stop the data generation, otherwise we continue. The details are explained in the following sections.

VAE and Latent Space. Figure 4a shows the architecture of our VAE. The encoder layers reduce the dimension of the input data and cast the input data into a distribution in latent space with mean μ_j and variance σ_j^2 (j represents the axis number of latent space). During the training stage, the vector z in latent space is generated by $z_j = \mathcal{N}(\mu_j, \sigma_j^2)$,²⁶ where \mathcal{N} is the normal distribution. The vector z can be mapped back to real space by the decoder layers.

Table 1. General Parameters of LOLS Used for all the Molecules in this Work

	name	value
VAE	latent dimension	2
	cutoff threshold (α)	2
	energy weight (β)	0,-1,-3
	loss ratio (λ)	0.01
	training epochs	100
sampling	expansion rate	20%
	batch size	50
GP model	kernel	STDP
	fitting noise	0.001

Data Preprocessing. The raw data includes the dihedral angles of sampled molecular structures and their DFT-calculated energy E . We preprocess the data in two stages. In stage one, the dihedral angles are normalized from $[0,360]$ to $[-1,1]$, and the total DFT energy is scaled according to the following equation

$$E^* = \begin{cases} E - E_0 & E \leq E_0 \\ \ln(1 + E - E_0) & E > E_0 \end{cases} \quad (1)$$

where E_0 is a threshold energy that is used to shift the DFT energies close to zero. E_0 is system dependent but once chosen is kept constant for the same molecule (see Table 2). We adopt the logarithmic function in eq 1 to scale down high energies ($E > E_0$), because we are primarily interested in the low energy region

Table 2. Molecule-Dependent Parameters of LOLS

name	cysteine	WG	GFA	GGF	WGG
search dimensionality	5	7	9	9	9
initial data size	100	350	350	350	350
en/decoder layer size (layersize)	80	128	128	128	128
maximum iteration (M)	40	120	120	120	140
energy model interval (k)	5	20	20	20	20
threshold energy (E_0 /eV)	-19,635	-24,320	-27,467	-26,399	-29,977

and wish to avoid high energy regions that can obstruct model fitting.

In stage two, data with a scaled energy larger than $E_{\max}^* = \text{mean}(E^*) + \alpha \times \text{std}(E^*)$ is excluded from the training set of the VAE, since the corresponding structures frequently exhibit steric clashes and are therefore not relevant. In this work, we set the cutoff threshold $\alpha = 2$, which resulted in a data exclusion of 3–6% from the training set of the VAE. The excluded data is usually 5 to 25 eV above the global minimum, and was still kept in the data pool and used to build the energy model in step 2.

Loss Function. The trainable parameters of the VAE are optimized by minimizing the total loss function, which consists of two contributions

$$\delta_{\text{total}} = \delta_{\text{rec}} + \lambda \delta_{\text{reg}} \quad (2)$$

The first part is the reconstruction loss (δ_{rec}), which forces the encoder-decoder pair to minimize information loss (i.e., minimize the difference between input and output). The second part is the regularization (δ_{reg}) that confines the latent space by forcing the encoder output towards a standard normal distribution. λ is a hyperparameter that controls the ratio between the two loss terms.

To make the VAE more sensitive to low-energy structures, we weight the reconstruction loss term (δ_{rec}) with the corresponding scaled energy $\exp(\beta E^*)$,

$$\delta_{\text{rec}} = \frac{\sum_{i=1}^N \exp(\beta E_i^*) \times \text{Diff}(x_i^{\text{in}}, x_i^{\text{out}})}{\sum_{i=1}^N \exp(\beta E_i^*)} \quad (3)$$

where β is a hyperparameter which will be explored and discussed later. In this work, we varied β from 0 to -3 . A negative β assigns a smaller weight to higher energy structures in the reconstruction loss. They therefore become less important in VAE training. i refers to the i th training data and N is the size of the training data. $\text{Diff}(x_i^{\text{in}}, x_i^{\text{out}})$ returns the difference between input and output for the i th training data. Since our VAE does not output the scaled energy, we define Diff only in terms of the scaled dihedral angles

$$\text{Diff}(x_i^{\text{in}}, x_i^{\text{out}}) = \frac{1}{D} \sum_{j=1}^D (((x_{ij}^{\text{in}} - x_{ij}^{\text{out}} + 1) \bmod 2) - 1)^2 \quad (4)$$

Here D refers to the number of dihedral angles and j to the j th input and output vectors.

The regularization term (δ_{reg}) can be expressed as the Kulback-Leibler (KL) divergence (δ_{kld}) between the returned distribution and a standard Gaussian.²⁶ According to ref 26, the KL divergence is calculated by the encoder output mean μ_{ij} and variance σ_{ij}^2 , where i is the i th training data, j the axis number of latent space and d is the dimension of latent space

$$\delta_{\text{kld}} = \frac{1}{Nd} \sum_{i=1}^N \sum_{j=1}^d - \frac{1}{2} (1 + \log \sigma_{ij}^2 - \mu_{ij}^2 - \sigma_{ij}^2) \quad (5)$$

The total loss function (δ_{total}) in our work is

$$\delta_{\text{total}} = \delta_{\text{rec}} + \lambda \delta_{\text{kld}} \quad (6)$$

Next, we will select a suitable value for λ and the right neural network settings for the cysteine data set we generated in our previous work.²¹ The data set consists of 800 cysteine structures and their corresponding DFT energies from a BOSS run. We

refer to this data set as CYS800. The dihedral angle and energy distributions of this data set are shown in Figure S1.

Neural Network Configurations. We chose 2 as the latent space dimension, for the simple reason that two dimensions are convenient to visualize. Visualizing and analyzing the latent space will help us gain insight into the nature of the latent space and develop suitable sampling methods. It remains an open question if increasing the dimension of latent space would help sample more informative data and thus increase the efficiency of the approach. We will return to this question in future work.

For both encoder and decoder, we used two fully connected layers of the same size and ReLU as activation function. We varied the number of neurons in each fully connected layer in the encoder or decoder (*layersize*) from 8 to 128 and checked the mean absolute error (MAE) between inputs and outputs. The CYS800 data set was used in all the tests. Similar to eq 4, the MAE is defined as

$$\text{MAE}(x_i^{\text{in}}, x_i^{\text{out}}) = \frac{1}{D} \sum_{j=1}^D |((x_{ij}^{\text{in}} - x_{ij}^{\text{out}} + 1) \bmod 2) - 1| \quad (7)$$

In Figure 4b we show the MAE as a function of the number of neural network parameters, which is determined by the *layersize*. The MAE decreases with increasing *layersize*, but eventually converges around 20°. We believe that with a higher dimensional latent space (i.e., less information loss) we could further reduce the MAE, but we deemed 20° sufficient for our purposes. We therefore picked a *layersize* of 80 for cysteine and extended it to 128 for other molecules in this work with higher search dimensions.

The VAE was trained for 100,000 epochs to ensure the convergence of the total loss function δ_{total} (Figure S2). The value of the energy weight hyperparameter ($\beta = 0$ or $\beta = -1$) has no significant effect on the MAE for the CYS800 data set, as shown in Figure 4b. However, β will play an important role in the active learning workflow (shown in Figure 3). We will discuss its effect in the “Results and Discussion” section.

Loss Ratio λ and Latent Space. After the training is finished, the encoder maps the training data into the latent space as the latent-space data $z_{ij} = \mu_{ij}$. The encoder output variances σ_{ij}^2 are only used in the reparameterization during the training stage and ignored after training. The hyperparameter λ controls the ratio between the reconstruction loss and the KL-divergence, thus determining the shape and distribution of the latent-space data. We introduce the latent-space scale L to measure the size of the latent space

$$L = \sqrt{\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^d \mu_{ij}^2} \quad (8)$$

Figure 4c shows that L varies by one order of magnitude for λ between 0 and 1. Between $\lambda = 0.001$ and 0.03, L stabilizes around 1.47 and changes little, indicating we should pick λ from this region. In this range, L is also almost independent of the size of the neural network.

Figure S3 shows the data distribution in latent space for different λ values. The shape and size of latent spaces are highly dependent on λ . When $\lambda = 0.01$, the latent-space data distributes uniformly inside a circle (Figure 5), which may benefit sampling. Therefore, we set $\lambda = 0.01$ for all networks in the following.

Sampling Method. After generating the latent space, we can sample it. Every sample will be decoded into dihedral angles to

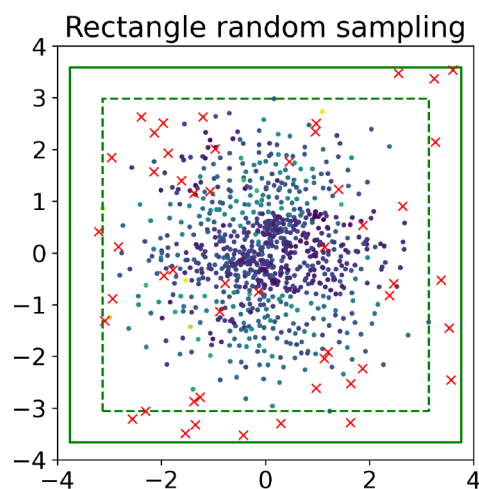


Figure 5. The dashed and solid lines show the minimal rectangle and the extended rectangle, and the red crosses show a new sampling batch inside the extended rectangle. The latent space is the same as in Figure S3.

reconstruct the atomic structure in real space. Then the DFT energy of this structure is calculated. The combination of scaled dihedral angles and DFT energy (x, E^*) is collected as new data.

We use a random sampling method to pick new structures from latent space. We had considered building a surrogate model of latent space with BOSS and sampling from its acquisition function, but the complex structure of latent space (which will be discussed in more detail in the “Results and Discussion” section) does not lend itself to more advanced sampling methods. More specifically, we use a rectangle random sampling method (Figure 5), which contains the following steps. First, we create a minimal rectangle that covers all of the latent-space data. Then we increase the width and height of the minimal rectangle with an expansion rate. The expansion rate is a hyperparameter that can be varied. We use a rate of 20% in this work, which balances sampling from known latent space areas with the need to explore unknown areas away from available latent-space data. Finally, we choose positions randomly in the extended rectangle as samples.

In LOLS, the generation loop will keep running until the number of iterations reaches the preset maximum. At each iteration, the VAE is retrained and a data batch is acquired. These newly acquired data points are added to the data pool for training the new VAE in the next iteration. In this work, we fix the batch size in each iteration to 50, which is small enough to track changes in latent space and large enough to effect a change in the VAE.

Energy Model. We fit a surrogate model in real space after every k iterations of the generation loop. We call this the “energy model” as it establishes a relation between the dihedral angles and the energy. k is the energy model interval. Here we choose $k = 5$ for cysteine and $k = 20$ for other molecules, which helped us find the relevant conformers without performing too many structure optimizations. The number of optimized structures is about 10–15% of the number of samples (See Table 3). We could use a smaller k to build more energy models and extract more local minima, but this would also require performing more DFT structure optimizations in step 3.

We use BOSS²³ to fit a GP to the energy model. The kernel is set to standard periodic (STDP) to account for the periodicity of the dihedral angles, with inverse gamma priors employed to

Table 3. Final Results for all Five Molecules^a

material	dim	β	target	achieved	new	single	relax	converged	achieved details
cysteine	5	0	11	11	16	6000	919	919	●●●●●●●●
cysteine	5	-1	11	11	18	6000	922	922	●●●●●●●●
cysteine	5	-3	11	11	20	6000	944	944	●●●●●●●●
WG	7	0	13	13	46	18,000	2314	2172	●●●●●●●●
WG	7	-1	13	12	47	18,000	2292	2177	●●●●●●●○
WG	7	-3	13	13	45	18,000	2214	2087	●●●●●●●●
GFA	9	0	16	9	15	18,000	1443	1227	●○●○●○●○●○●○
GFA	9	-1	16	13	23	18,000	1872	1588	●●●○●○●○●○●○
GFA	9	-3	16	13	27	18,000	1873	1597	●●●○●○●○●○●○
GGF	9	0	13	9	23	18,000	1870	1555	●●●○●○●○●○●○
GGF	9	-1	13	10	22	18,000	1645	1381	●●●○●○●○●○●○
GGF	9	-3	13	9	22	18,000	1553	1379	●○●○●○●○●○●○
WGG	9	0	13	7	13	21,000	2536	2004	○●○●○●○●○●○
WGG	9	-1	13	7	10	21,000	3073	2508	●○●○●○●○●○●○
WGG	9	-3	13	9	9	21,000	2844	2270	●●●○●○●○●○●○

^aResults for three parallel run are merged. "Achieved" means the number of targets we found. "New" means the number of stable structures we found but missed by the reference. (Only the ones with energy less than the maximum energy of targets are counted.) "Achieved details" enumerate the targets sorted by energy, where ● and ○ represent found and missed targets. "Single" means the total number of single-point energy calculations during the three parallel samplings. "Relax" shows the number of optimized structures. "Converged" gives the number of stable structures that are converged within 200 geometry optimization steps.

stabilise kernel hyperparameters. The noise is set to 0.001 eV, comparable to the accuracy of DFT calculations. We set an uninformative prior on the GP mean to avoid biasing the model. After the energy model in real space is built by BOSS, we take the training data as the initial positions and apply the conjugate gradient method to find local minima. Only different local minima are kept and duplicates are purged. In accordance with our previous work, we fully optimize all molecular degrees of freedom with DFT for only these unique minima structures.

DFT Method. In this work, we employed the all-electron code FHI-aims^{27–29} for all DFT calculations. We used "tight" numerical settings, "tier 2" basis sets, the PBE exchange-correlation functional³⁰ and many-body dispersion (MBD) van der Waals corrections.³¹ For a few structures, in which two or more atoms come too close to each other, the FHI-aims single-point calculations fail. We consider these structures invalid (steric clashes). For different molecules, 3–6% of samples were invalid and we omitted them.

For geometry optimization, a geometry was considered to be converged when the maximum residual force was below 0.01 eV/Å. We stopped geometry optimization after a maximum of 200 steps to reduce the calculation costs. Any structure that is not converged after 200 steps is excluded. For cysteine, all structures are converged in less than 200 relaxation steps, but for larger molecules, 5–20% of structures do not converge (see Table 3).

Complete Workflow. Algorithm 1 shows the complete workflow of LOLS. We have defined the parameters *initdata*, *layersize*, E_0 , α , β , λ , k , and the noise in the previous sections. In addition, M represents the maximum iterations.

Algorithm 1 Complete workflow

Require: *initdata*, *layersize*, E_0 , α , β , λ , *noise*, M , k

- 1: **DataPool** = *initdata*
- 2: **StableComformers** = \emptyset
- 3: **for** $i = 1 \dots M$ **do**
- 4: **Initialize**(VAE, *layersize*)
- 5: *data* = **Trim**(DataPool, E_0 , α)
- 6: **Optimize**(VAE, β , λ , *data*)
- 7: *latent* = VAE \rightarrow *Encode*(*data*)
- 8: *sample* = **TakeSamples**(*latent*)
- 9: *decoded* = VAE \rightarrow *Decode*(*sample*)
- 10: **for** *vector* \in *decoded* **do**
- 11: *atoms* = **Vec2Atoms**(*vector*)
- 12: *energy* = **DFTEnergy**(*atoms*)
- 13: DataPool \leftarrow {*vector*, *energy*}
- 14: **end for**
- 15: **if** ($i \equiv 0 \pmod k$) **then**
- 16: **Initialize**(GP, DataPool)
- 17: **Optimize**(GP, *noise*)
- 18: **for** *vector*, *energy* \in DataPool **do**
- 19: **Optimize**(*vector*, GP)
- 20: *atoms* = **Vec2Atoms**(*vector*)
- 21: **Optimize**(*atoms*)
- 22: StableComformers \leftarrow *atoms*
- 23: **end for**
- 24: **end if**
- 25: **end for**
- 26: **Return** StableComformers

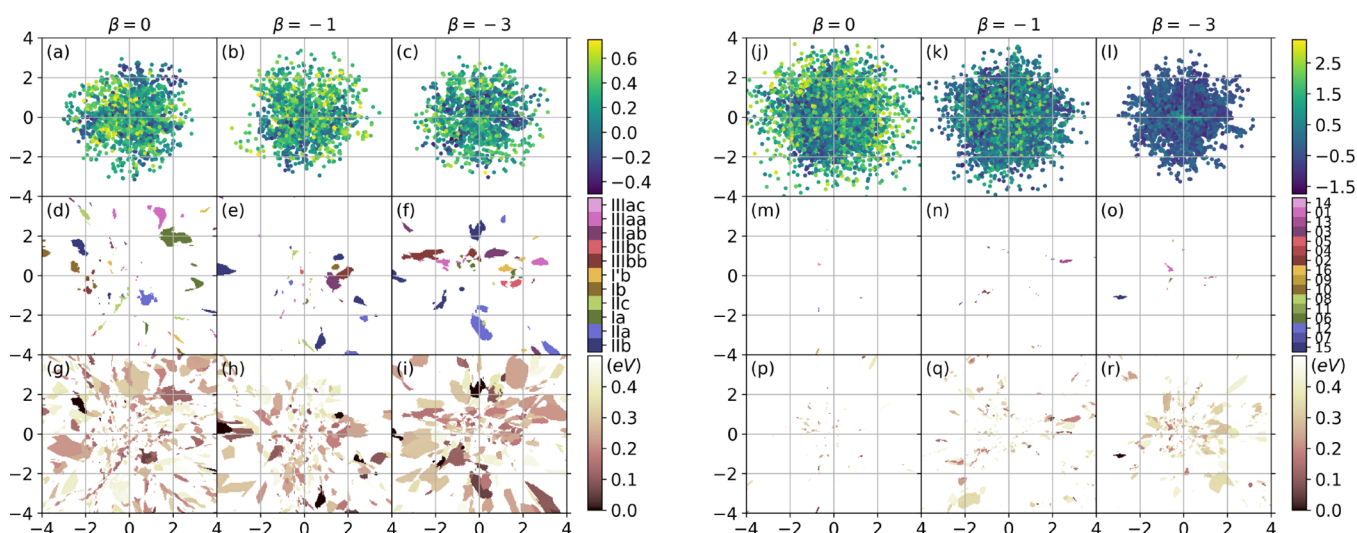


Figure 6. The latent spaces of cysteine (left, (a–i)) and GFA (right, (j–r)) for different β are visualized in three ways. In (a–c) and (j–l) the latent spaces are formed by the latent-space data generated by the encoders with different β . The color represents the scaled energy. In (d–i) and (m–r) the points in latent space are decoded into real space, and the reconstructed structures are compared with a series of targets. If the mean difference of dihedral angles (MAE) between the reconstructed structure and the nearest target is less than 30° , the points are colored. In (d–f) and (m–o) the targets are from the references,^{21,25} and the same color represents the same target. In (g–i) and (p–r) the targets are all stable conformers we found with energy less than 0.5 eV above the global minimum. In (g–i) and (p–r) the color represents the energy of the nearest conformer: the darker color, the lower energy.

We applied our LOLS method to cysteine and the peptides WG, GFA, GGF and WGG. Figure 2 shows how we chose the dihedral angles as features. The dihedral angles of the peptide bonds in WG, GFA, GGF and WGG are fixed at 180° for the *trans* conformation because they usually have lower energy than the *cis* isomers. For GFA and GGF, the dihedral angles of the benzene rotation are only searched from 0 to 180° due to symmetry. For GFA, the dihedral angle of the methyl rotation is fixed at 180° . The final dimension of features for cysteine, WG, GFA, GGF, and WGG are 5, 7, 9, 9 and 9.

The LOLS algorithm is general, but for new molecules, some input parameters may need to be modified, and the VAE retrained. In this work, the parameters in Table 1 are shared by all molecules. We do not fine-tune them for individual molecules because all the molecules in this work are small and organic. The molecule-dependent parameters are shown in Table 2.

We could initialize LOLS with random data. However, since BOSS performs active learning for optimal knowledge gain and BOSS sampling is very fast for small amounts of data, we use samples from one BOSS run as the initial data in this work. The initial data size is also shown in Table 2.

During testing on cysteine, we noticed that some targets that were correctly identified at a certain point would disappear, if we continued iterating (see Figure S4), due to statistical fluctuations of GP fitting. Because of this observation, we not only take the result from the final energy model with maximum data size but also from previous energy models.

RESULTS AND DISCUSSION

We applied LOLS to cysteine, WG, GFA, GGF and WGG. For cysteine, we mainly compared the results to our previous study,²¹ which used BOSS and quantum chemistry methods. The conformer structures in ref 21 obtained with the same DFT settings as this work were selected as targets for cysteine. For the other molecules, we compared our results to the database generated by Valders et al.²⁵ The authors first ran molecular dynamics/quenching (MD/Q) simulations with tight-binding

DFT to scan the free energy surfaces and then recalculated the low-energy structures with high-level quantum chemistry methods. We reoptimized their structures in the database with our DFT functional and settings before using them as targets. The mean difference in dihedral angles between our reoptimized and the geometries in ref 25 are generally less than 5° , except WG 03 (22.6°), GGF 05 (12.3°), GGF 13 (11.9°), and WG 11 (7.0°). Two structures are considered similar when the mean difference in the dihedral angles is less than 15° . In the series of similar structures, only the structure with the lowest energy is kept. If the maximal difference of dihedral angles between one target and one of our results is less than 15° , we state that the target has been reached. Otherwise, we consider that a new structure has been found.

Cysteine. First we analyze the VAE training process and acquired samples. The training loss, the latent-space scale, and the average energy of samples were all within reasonable values during the training, proving that the training went well for cysteine (see Figure S5 and S1). Next we analyze the latent space of cysteine. The trained VAE has two components: the encoder and the decoder. The latent-space data generated by encoders with different β are shown in Figure 6a–c. The latent-space data is distributed uniformly as a circle in the latent spaces. For $\beta = 0$, low- and high-energy data are mixed. For $\beta < 0$, low- and high-energy regions start to form that become more pronounced for more negative β s.

To understand the correspondence between the eleven target cysteine conformers²¹ and the latent space, we discretized latent space on a 400×400 grid, and mapped all points back to real space with the decoder. We assigned each corresponding structure to one of the eleven conformers, if the MAE is smaller than 30° . If it is larger, the structure remains unassigned. The result is a map of islands in latent space, shown in panels (d) to (f). The total area of islands are 6.1, 3.7 and 7.3% of the latent space for $\beta = 0, -1$ and -3 . However, the mapping is not always unique, and multiple islands may map to the same target, such as

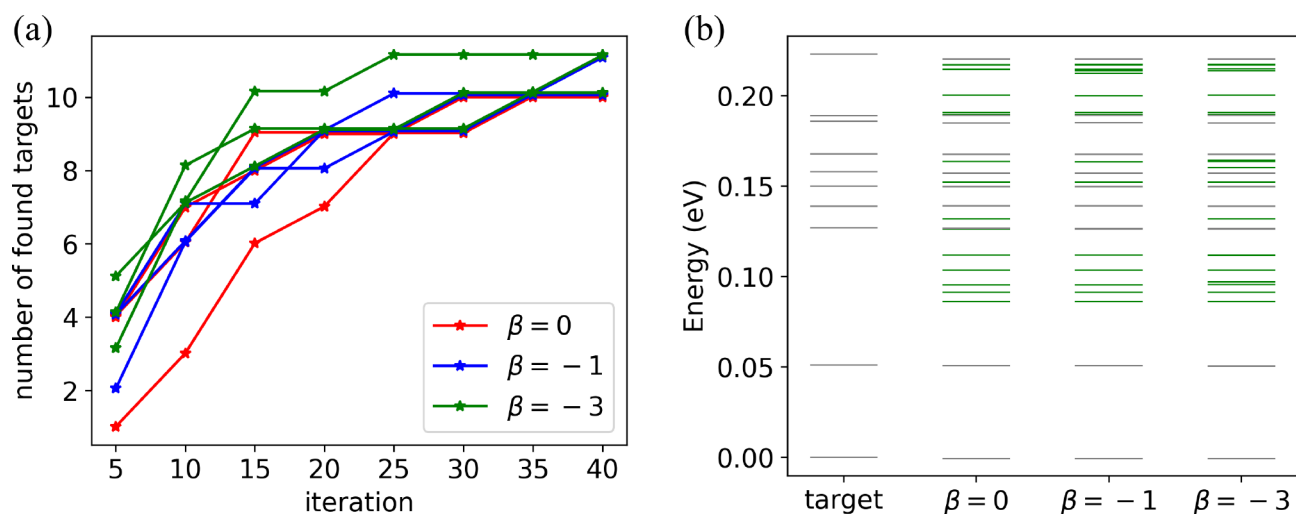


Figure 7. The results of cysteine. (a) The accumulative result for every single run. Three parallel runs are performed for each β . (b) Stable conformers found in this work. Conformers are represented by lines and sorted by energy. The gray lines represent the found structures that are the targets from the reference.²¹ The green line represents the new conformers that are not the 11 targets.

la in Figure 6d. This suggests that structures that are similar in real space are not necessarily close in latent space.

The eleven targets are distributed within $[0, 0.25 \text{ eV}]$ from the global minimum. We repeated the same procedure described in the last paragraph, but now use all the conformers we identified in the energy window $[0, 0.5 \text{ eV}]$ from the global minimum as references. We colored the latent space by the energies of these reference conformers and call the colored area low-energy areas. For $\beta = 0, -1$, and -3 , the low-energy areas cover 36.2, 26.6, and 36.8% of the latent space in Figure 6g–i.

Finally, we analyze and evaluate the performance of LOLS for cysteine conformer search. Figure 7a shows the numbers of targets found in the nine parallel runs. The same color is used for results with the same β value. The y value gives the accumulative number of correctly identified targets before that iteration. The best outcome is in one run with $\beta = -3$ (top green curve), while the worst result has $\beta = 0$ (bottom red curve). The other seven runs perform accordingly. $\beta = -3$ runs are among the best, $\beta = -1$ average and $\beta = 0$ the worst. We therefore recommend a β value smaller than zero. The results using the same β for three parallel runs are merged into one and shown in Figure 7b. The figure shows that all the eleven targeted conformers (along with some new ones) were found regardless of β . This is also shown again in Table 3.

Glycyl-phenylalanyl-alanyl. Next we applied LOLS to GFA. The training loss, the latent-space scale and the average energy of the samples in Figure S6 indicate that the training went well for GFA. In addition, we observed that more low-energy data is generated for non-zero β s. More discussions can be found in the SI. We plot the latent spaces of GFA for $\beta = 0, -1$ and -3 using the same mapping methods as for cysteine. Figure 6j–l show the latent-space data generated by the encoder. For $\beta = -3$, the latent-space data is more compact and contains more low-energy data than for $\beta = 0$ and $\beta = -1$. Figure 6m–o show the correspondence of the latent-space data to the target GFA conformers²⁵ in real space generated by the decoder. Unlike for cysteine, the colored latent space of GFA is quite empty. The total area of colored islands are 0.06, 0.28 and 0.28% for $\beta = 0, -1$ and -3 . Figure 6p–r are colored in the same way as Figure 6g–i. The low-energy areas ($[0, 0.5 \text{ eV}]$) cover 1.0, 7.6 and 11.7% of the latent space of GFA, for $\beta = 0, -1$, and -3 . The

coverages are much smaller than in cysteine. We believe that this due to the higher dimensionality of GFA (9 compared to 6). Higher-dimensional systems usually have more complex PESs, and less area can be associated with low-energy conformers, which may explain the emptiness of latent space.

For every β , the accumulative results of three parallel runs were merged into one and shown in Figure 8. We used the

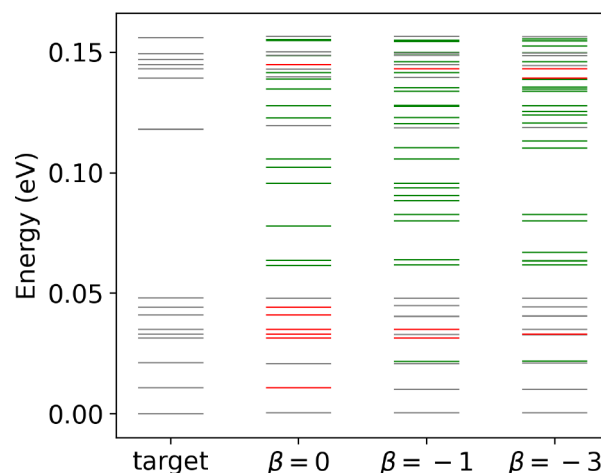


Figure 8. Stable conformers of GFA found in this work. The gray lines mean the targets from the reference²⁵ are found, while the red lines mean we missed the targets. The green line represents the new conformers that we found but missed by the reference.

sixteen GFA structures reported in ref 25 as our targets. We found nine, thirteen and thirteen out of the sixteen targets for $\beta = 0, -1$ and -3 . Among the three values of β , $\beta = -3$ performs best for GFA, $\beta = -1$ has similar performance as $\beta = -3$, but $\beta = 0$ missed six out of nine lowest energy targets. As mentioned in ref 25, these targets can be divided into six structural types according to the different hydrogen bonds. All six types are found with β of 0, -1 and -3 .

The differences between our results and the reference results²⁵ are mainly due to the flexibility of the end groups of GFA. The $-\text{CH}_2\text{NH}_2$ branch and the $-\text{C}_6\text{H}_6$ branch (benzene ring) of GFA have several stable configurations which have energy differences

within 10 meV. The two groups are at the end of the peptide, thus having little effect on the overall structures of GFA, however resulting in the different conformers. For example, GFA 06, GFA 11, and GFA 08 have very similar structures (see Figure S7). The only difference between GFA 06 and GFA 11 is the configuration of the benzene ring, which causes an 1.7 meV energy difference. And the only difference between GFA 11 and GFA 08 is the configuration of the $-\text{CH}_2\text{NH}_2$ branch, which causes a difference of 1.9 meV. We found GFA 11 but missed GFA 06 and 08 in the result with $\beta = -1$, and we missed GFA 11 but found GFA 06 and 08 with $\beta = -3$. Importantly, the global minimum (GFA 15) is always found by our method even with different β s. The reference did not find any conformers in the energy range from 0.05 to 0.12 eV above the global minimum. However, we found six, ten, and eight new structures in this energy region using $\beta = 0, -1$, and -3 . Overall, we have achieved comparable accuracy as the reference.

WG, WGG and GGF. We also tested WG, WGG and GGF, whose search dimensions are seven, nine and nine, respectively, and compared them with Ref 25. For each β of 0, -1 , and -3 , three parallel runs were carried out for WG and GGF, with maximum iteration count $M = 120$. Unfortunately, we did not find the global minimum of WGG at 120 iterations for any value of β , so we ran an additional 20 iterations for WGG. The accumulative results are shown in Figure S8. The results of all the five molecules are also summarized in Table 3.

For WG, using $\beta = 0$ or -3 , we found all the thirteen targets, but $\beta = -1$ missed the highest energy target. For GGF, the performance for different β were close but $\beta = -1$ found the most targets. For WGG, $\beta = 0$ missed the global minimum, which was found with $\beta = -1$ or $\beta = -3$. Combined with the results for cystine and GFA, we can state that non-zero β is at least beneficial for larger molecules such as GFA, GGF and WGG. Except cystine, all other molecules are peptides which are very flexible molecules. It is therefore no surprise that our structure lists are not exactly the same for different β or as the ones in ref 25. We have missed some targets but also found some new ones in the same energy region. Overall we achieved the same level of performance as the reference.

Comparison to Real Space Search. In this section, we compare our VAE approach and random sampling on the dihedral spaces while keeping all other parts in the LOLS workflow the same. We refer to random sampling as real space search. In the real space search, we took samples randomly from real space and fitted a GP surrogate model every k samples, gathered the local minima as the relaxation starting points, relaxed the geometries with DFT, removed duplicates and then compared them with targets (Algorithm S1). In other words, the real space search workflow replaces the VAE data generation loop by taking random samples directly in real space but keeps the other steps of LOLS. We tested the real space search workflow on cysteine (5-D), WG (7-D), and GFA (9-D). For each molecule, we carried out three parallel runs. The results of the parallel runs were merged and compared to LOLS with $\beta = -3$ in Figure 9. The details of the observed targets are shown in Figure S9.

Figure 9 presents the number of targets found versus the number of samples used to build the energy models. For cysteine, the real space search found all the eleven targets with 2250 samples, while LOLS ($\beta = -3$) required 3000 samples. For WG, the real space search and LOLS both took 18,000 samples to find all the thirteen targets. LOLS's performance is similar to the real space search for cysteine and WG. However, for GFA

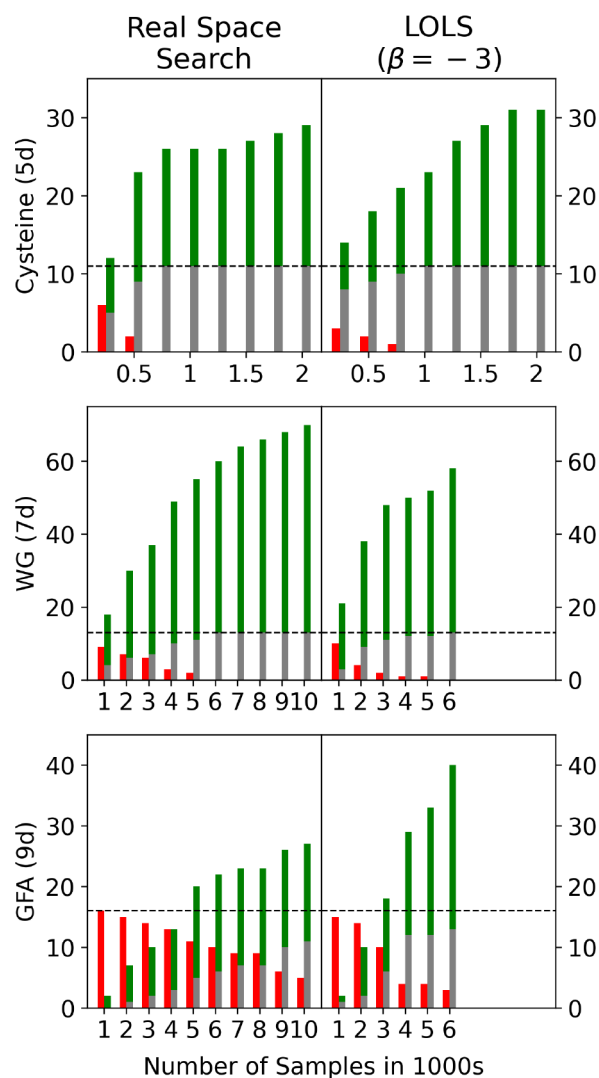


Figure 9. Results of the real space search workflow vs the LOLS (with $\beta = -3$) for cysteine, WG, and GFA. The x axis is the number of samples in each run. The results from the three parallel runs were merged. The height of bars represents number of conformers. Red means missing targets, gray means achieved targets and green means new conformers. The black dash lines represent the number of targets.

LOLS starts to provide an advantage. The real space search found eleven out of sixteen targets using 30,000 samples, while LOLS ($\beta = -3$) found twelve targets with 12,000 samples and thirteen targets with 18,000 samples. LOLS clearly outperforms the real space search.

Discussion. First, we discuss the properties of latent spaces in this work. Our analysis of the 2-D latent spaces generated by encoders revealed them to be neither smooth nor continuous (see Figure 6a–c,j–l). High- and low-energy areas appear intermixed in the latent space, and it proved difficult to fit GP models to latent-space data and extract any information on low energy regions. Moreover, casting previously known conformers into latent space demonstrated that the same conformer structure can be mapped into different locations in latent space (see Figure 6d–i,m–r). This suggests that similar structures in real space are not necessarily close in latent space. For these reasons, we did not further pursue designing acquisition functions or minima searches in latent space. Instead we use the fast, explorative and space-filling random sampling

approaches to sample latent space. Increasing the dimension of latent space may create a more smooth and continuous latent space, which potentially allows us to develop more sophisticated sampling methods. We will test high-dimensional latent spaces in future work.

We analyze the low-energy area ($[0, 0.5 \text{ eV}]$) of our latent spaces. For cysteine, all workflows with the different $\beta = 0, -1, -3$ achieved good results, which may be due to the similar coverage of low-energy area ($\sim 30\%$). However, for GFA, only 1% of latent space corresponds to low-energy structures for $\beta = 0$, which is likely to be the reason for missing most of the targets (See Figure 8). This percentage increases to 10% for $\beta = -1$ and -3 , and we achieved much better results. This analysis suggests that a non-zero β is an advantage for LOLS. More detailed discussion of how the β affects the datapool can be found in SI (Figures S10 and S11).

Next, we discuss the efficiency of LOLS. Building a high-dimensional energy model and thoroughly exploring it requires a large amount of data. For example, if we take the grid sampling method in nine-dimensional space and divide each dimension into ten equal parts, we would need 10^9 samples. Although our work does not aim to achieve the highest efficiency, we acquired enough data to build a reliable energy model for nine-dimensional peptides with 18,000–21,000 single-point energy calculations. We compared LOLS to a real space sampling algorithm, and conclude that LOLS found more conformers with fewer samples than the real space search algorithm for 9-D molecules. However, for small molecules such as cysteine (5-D) and WG (7-D), our method is unlikely to outperform this real space sampling. We conclude that LOLS is more suitable for larger molecules with more degrees of freedom, because it is challenging to sample a high-dimensional PES in real space.

A note of caution has to be added for multi-reference states, which may be intrinsically present in the molecule or arise from stretched bonds.³² Multi-reference states are notoriously difficult to treat in DFT and may adversely affect DFT-based structure search. LOLS provides an advantage in this regard, because it avoids stretched-bond related multi-reference states. LOLS starts from an equilibrium geometry free of stretched bonds and sample only the space of dihedral angles. Full geometry optimization for bond length and bond angles is then performed only for the resulting local-minima configurations and will thus not encounter stretched bonds, unless a conformer exhibits such a geometric feature.

The LOLS algorithm is flexible and ideally can be applied to any physical system with a similar search dimensionality as our molecules. We have determined suitable parameter values for peptides, but these settings may not be optimal for other, more complex systems. In our experience, the most critical parameters are the architecture of the neural networks (e.g., number of layers and nodes) in the VAE, and the hyperparameters β and λ . The architecture of the neural network and λ can be re-optimized by training the VAE on a test set and monitoring the reconstruction loss and the latent-space scale L . However, adjusting β requires the execution of the whole workflow. Also, we chose the rectangle random sampling method for simplicity. Replacing it with some more sophisticated sampling methods may further increase the efficiency. After acquiring a fixed amount of samples, we use a GP model to build the energy model and gather the local minima for post-relaxation. Here the GP model could be replaced by any continuous model, for example, a neural network.

As a new method, LOLS also has some limitations. First, latent space is an abstract entity and it is not trivial to ascertain, if it is sufficiently expressive or well sampled. For this reason, we focused on systems with existing benchmarks to determine how many independent LOLS runs and iteration steps are required. For new systems without literature references, we recommend performing several independent LOLS runs with a sufficiently large iteration step until no new structures can be found in the relevant energy window. Second, we only tested a two-dimensional latent space in this work. We suspect that two-dimensional latent spaces might become too sparse when the number of dihedral search dimensions increases, which might degrade the efficiency of LOLS. We will explore optimal latent space dimensions in future work. Lastly, we emphasize that LOLS was designed to sample the low-energy parts of the PES and to find low-energy conformers. If one is interested in the whole PES, e.g., for transition state and reaction path searches, other PES sampling methods such as Global Reaction Route Mapping (GRRM)³³ might be more appropriate.

CONCLUSIONS

In this work, we have developed the active learning workflow LOLS for molecular conformer search. LOLS is a stochastic method that contains two machine learning models: the generative model VAE for data sampling and the GP for energy model fitting. We introduced the hyperparameter β to steer the latent space towards low-energy molecular configurations for generating more informative data. We have applied LOLS to cysteine and the peptides WG, GFA, GGF, and WGG, and achieved a similar level of accuracy as the references. For small molecules such as cysteine, it is more efficient to sample data in real space; however, LOLS is more suitable for larger molecules such as peptides. LOLS is still at an early stage of development: further optimization of the generative model and energy model may increase the efficiency and facilitate applications to other systems beyond molecules.

We have also gained insight into the nature and properties of latent space both quantitatively and qualitatively. Quantitatively, we found that the distribution of latent-space data can be controlled by the hyperparameter λ that is used to balance the reconstruction loss and regulation term in the loss function of the VAE. By tuning λ , a more uniform latent space can be formed, which is beneficial for sampling. In addition, we found that the latent-space scale (L) is a good parameter to measure the size of latent space. Qualitatively, we found for cysteine and GFA that latent space is neither smooth nor continuous in the low-energy regions. Moreover, the structures are close in real space might not be close in latent space. Therefore we recommend exploratory and space-filling sampling approaches for latent space sampling.

ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.2c00290>.

(Figure S1) The distributions of dihedral angles and scaled energy of the CYS800 data set; (Figure S2) the progression of training loss with training epochs; (Figure S3) the latent-space data distributions of the CYS800 data set with different λ ; (Figure S4) the targets found at different iterations; (Figures S5 and S6) the latent-space scale and the energies of samples during the data

generation step for cysteine and GFA; (Figure S7) three very similar GFA conformers: GFA 08, GFA 11 and GFA 06.; (Figure S8) the accumulative results for WG, GGF and WGG; (Figure S9) comparison of LOLS and real space search workflow on GFA; (Figure S10) the relationship between the reconstruction error and the scaled energy in the last iteration of LOLS on cysteine; (Figure S11) the energy distribution of the data in the last LOLS iteration for cysteine, WG and GFA (PDF)

The structures and energies of the stable conformers can be found in NOMAD Repository and Archive (DOI: 10.17172/NOMAD/2022.06.08-1).

AUTHOR INFORMATION

Corresponding Authors

Milica Todorović – Department of Mechanical and Materials Engineering, University of Turku, FI-20014 Turku, Finland; orcid.org/0000-0003-0028-0105; Email: milica.todorovic@utu.fi

Xi Chen – Department of Applied Physics, Aalto University, Espoo 00076, Finland; orcid.org/0000-0001-6149-2270; Email: xi.6.chen@aalto.fi

Authors

Xiaomi Guo – State Key Laboratory of Low Dimensional Quantum Physics and Department of Physics, Tsinghua University, Beijing 100084, China; Department of Applied Physics, Aalto University, Espoo 00076, Finland

Lincan Fang – Department of Applied Physics, Aalto University, Espoo 00076, Finland

Yong Xu – State Key Laboratory of Low Dimensional Quantum Physics and Department of Physics, Tsinghua University, Beijing 100084, China; Frontier Science Center for Quantum Information, Beijing 100084, China; RIKEN Center for Emergent Matter Science (CEMS), Saitama 351-0198, Japan; orcid.org/0000-0002-4844-2460

Wenhui Duan – State Key Laboratory of Low Dimensional Quantum Physics and Department of Physics and Institute for Advanced Study, Tsinghua University, Beijing 100084, China; Frontier Science Center for Quantum Information, Beijing 100084, China

Patrick Rinke – Department of Applied Physics, Aalto University, Espoo 00076, Finland; orcid.org/0000-0003-1898-723X

Complete contact information is available at: <https://pubs.acs.org/10.1021/acs.jctc.2c00290>

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

X.G., L.F. and X.C. acknowledge the financial support from the Academy of Finland (project numbers 308647, 314298, 335571). X.G, Y.X. and W.D. acknowledge the financial support from the Basic Science Center Project of NSFC (grant no. 51788104), the National Science Fund for Distinguished Young Scholars (grant no. 12025405), the National Natural Science Foundation of China (grant no. 11874035), and the Beijing Advanced Innovation Center for Future Chip (ICFC), M.T. and P.R. have received funding from the Academy of Finland via the Artificial Intelligence for Microscopic Structure Search (AIMSS) project no. 316601 and the Flagship programme:

Finnish Center for Artificial Intelligence FCAI. X.G. and X.C. Generous computational resources were provided by CSC – IT Center for Science, Finland, and the Aalto Science-IT project. L.F. also acknowledges financial support from the Chinese Scholarship Council (grant no. [2017]3109).

REFERENCES

- (1) Hawkins, P. C. D. Conformation Generation: The State of the Art. *J. Chem. Inf. Model.* **2017**, *57*, 1747–1756.
- (2) Wales, D. J.; Bogdan, T. V. Potential Energy and Free Energy Landscapes. *J. Phys. Chem. B* **2006**, *110*, 20765–20776.
- (3) Gronert, S.; O'Hair, R. A. J. Ab Initio Studies of Amino Acid Conformations. 1. The Conformers of Alanine, Serine, and Cysteine. *J. Am. Chem. Soc.* **1995**, *117*, 2071–2081.
- (4) Ferro-Costas, D.; Mosquera-Lois, I.; Fernández-Ramos, A. TorsiFlex: an automatic generator of torsional conformers. Application to the twenty proteinogenic amino acids. *Aust. J. Chem.* **2021**, *13*, 100.
- (5) Schwab, C. H. Conformations and 3D pharmacophore searching. *Drug Discovery Today* **2010**, *7*, e245–e253.
- (6) Hawkins, P. C. D.; Skillman, A. G.; Nicholls, A. Comparison of Shape-Matching and Docking as Virtual Screening Tools. *J. Med. Chem.* **2007**, *50*, 74–82. PMID: 17201411
- (7) Friedrich, N.-O.; de Bruyn Kops, C.; Flachsenberg, F.; Sommer, K.; Rarey, M.; Kirchmair, J. Benchmarking Commercial Conformer Ensemble Generators. *J. Chem. Inf. Model.* **2017**, *57*, 2719–2728.
- (8) Chang, G.; Guida, W. C.; Still, W. C. An internal-coordinate Monte Carlo method for searching conformational space. *J. Am. Chem. Soc.* **1989**, *111*, 4379–4386.
- (9) Goedecker, S. Minima hopping: An efficient search method for the global minimum of the potential energy surface of complex molecular systems. *J. Chem. Phys.* **2004**, *120*, 9911.
- (10) Wales, D. J.; Doye, J. P. K. Global Optimization by Basin-Hopping and the Lowest Energy Structures of Lennard-Jones Clusters Containing up to 110 Atoms. *J. Phys. Chem. A* **1997**, *101*, 5111–5116.
- (11) Supady, A.; Blum, V.; Baldauf, C. First-Principles Molecular Structure Search with a Genetic Algorithm. *J. Chem. Inf. Model.* **2015**, *55*, 2338–2348.
- (12) Rossi, M.; Scheffler, M.; Blum, V. Impact of Vibrational Entropy on the Stability of Unsolvated Peptide Helices with Increasing Length. *J. Phys. Chem. B* **2013**, *117*, 5574–5584.
- (13) Chutia, S.; Rossi, M.; Blum, V. Water Adsorption at Two Unsolvated Peptides with a Protonated Lysine Residue: From Self-Solvation to Solvation. *J. Phys. Chem. B* **2012**, *116*, 14788–14804.
- (14) Chen, X.; Jørgensen, M. S.; Li, J.; Hammer, B. Atomic Energies from a Convolutional Neural Network. *J. Chem. Theory Comput.* **2018**, *14*, 3933–3942.
- (15) Smith, J. S.; Nebgen, B. T.; Zubatyuk, R.; Lubbers, N.; Devereux, C.; Barros, K.; Tretiak, S.; Isayev, O.; Roitberg, A. E. Approaching coupled cluster accuracy with a general-purpose neural network potential through transfer learning. *Nat. Commun.* **2019**, *10*, 2903–2910.
- (16) Denzel, A.; Kästner, J. Gaussian process regression for geometry optimization. *J. Chem. Phys.* **2018**, *148*, No. 094114.
- (17) Meyer, R.; Hauser, A. W. Geometry optimization using Gaussian process regression in internal coordinate systems. *J. Chem. Phys.* **2020**, *152*, No. 084112.
- (18) Kaappa, S.; del Río, E. G.; Jacobsen, K. W. Global optimization of atomic structures with gradient-enhanced Gaussian process regression. *Phys. Rev. B* **2021**, *103*, 174114.
- (19) Schmitz, G.; Godtliebsen, I. H.; Christiansen, O. Machine learning for potential energy surfaces: An extensive database and assessment of methods. *J. Chem. Phys.* **2019**, *150*, 244113.
- (20) Unke, O. T.; Chmiela, S.; Sauceda, H. E.; Gastegger, M.; Poltavsky, I.; Schütt, K. T.; Tkatchenko, A.; Müller, K.-R. Machine Learning Force Fields. *Chem. Rev.* **2021**, *121*, 10142–10186. PMID: 33705118

- (21) Fang, L.; Makkonen, E.; Todorović, M.; Rinke, P.; Chen, X. Efficient Amino Acid Conformer Search with Bayesian Optimization. *J. Chem. Theory Comput.* **2021**, *17*, 1955–1966.
- (22) Todorović, M.; Gutmann, M. U.; Corander, J.; Rinke, P. Bayesian inference of atomistic structure in functional materials. *npj Comput. Mater.* **2019**, *5*, 35.
- (23) <https://gitlab.com/cest-group/boss>.
- (24) Sanz, E. M.; Blanco, S.; López, J. C.; Alonso, J. L. Rotational Probes of Six Conformers of Neutral Cysteine. *Angew. Chem., Int. Ed.* **2008**, *47*, 6216–6220.
- (25) Valdes, H.; Pluháčková, K.; Pitonák, M.; Řezáč, J.; Hobza, P. Benchmark database on isolated small peptides containing an aromatic side chain: comparison between wave function and density functional theory methods and empirical force field. *Phys. Chem. Chem. Phys.* **2008**, *10*, 2747–2757.
- (26) Kingma, D. P.; Welling, M. *Auto-Encoding Variational Bayes*. 2014.
- (27) Blum, V.; Gehrke, R.; Hanke, F.; Havu, P.; Havu, V.; Ren, X.; Reuter, K.; Scheffler, M. Ab initio molecular simulations with numeric atom-centered orbitals. *Comput. Phys. Commun.* **2009**, *180*, 2175–2196.
- (28) Havu, V.; Blum, V.; Havu, P.; Scheffler, M. Efficient O(N) integration for all-electron electronic structure calculation using numeric basis functions. *J. Comput. Phys.* **2009**, *228*, 8367–8379.
- (29) Ren, X.; Rinke, P.; Blum, V.; Wieferink, J.; Tkatchenko, A.; Sanfilippo, A.; Reuter, K.; Scheffler, M. Resolution-of-identity approach to Hartree–Fock, hybrid density functionals, RPA, MP2 and GW with numeric atomcentered orbital basis functions. *New J. Phys.* **2012**, *14*, No. 053020.
- (30) Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Phys. Rev. Lett.* **1996**, *77*, 3865.
- (31) Tkatchenko, A.; DiStasio, R. A.; Car, R.; Scheffler, M. Accurate and Efficient Method for Many-Body van der Waals Interactions. *Phys. Rev. Lett.* **2012**, *108*, 236402.
- (32) Duan, C.; Liu, F.; Nandy, A.; Kulik, H. J. Data-Driven Approaches Can Overcome the Cost–Accuracy Trade-Off in Multi-reference Diagnostics. *J. Chem. Theory Comput.* **2020**, *16*, 4373–4387.
- (33) Ohno, K.; Maeda, S. Global Reaction Route Mapping on Potential Energy Surfaces of Formaldehyde, Formic Acid, and Their Metal-Substituted Analogues. *The Journal of Physical Chemistry A* **2006**, *110*, 8933–8941. PMID: 16836457