

# A Reinforcement-Learning-Based Approach to Enhance Exhaustive Protein Loop Sampling

Amélie Barozet<sup>1,2,\*</sup>, Kevin Molloy<sup>3</sup>, Marc Vaisset<sup>1</sup>, Thierry Siméon<sup>1</sup> and Juan Cortés<sup>1,\*</sup>

<sup>1</sup>LAAS-CNRS, Université de Toulouse, CNRS, Toulouse, France

<sup>2</sup>Sanofi recherche & développement, Integrated Drug Discovery, Molecular Design Sciences, 13 quai Jules Guesde, BP 14, 94403 Vitry-sur-Seine Cedex, France

<sup>3</sup>James Madison University, Harrisonburg, Virginia, USA

\*To whom correspondence should be addressed.

## S1 Methods: details and analysis

### S1.1 MoMA-LoopSampler in Brute Force mode

Algorithm S1.1 shows the Brute Force version of MoMA-LoopSampler. When the conformational space of the loop is restricted, it captures the full conformational ensemble that can be generated by MoMA-LoopSampler.

Since the construction method perturbs the sampled tripeptides, every brute force search in our tests is carried out three times to obtain a set of loop conformations that is representative of the conformational space reachable by MoMA-LoopSampler. The reinforcement-learning-based heuristic ( $H_{RL}$ ) can also be used to obtain the final learning tree and analyze the distribution of the solutions.

### S1.2 Tripeptide database construction

A database of tripeptide states was constructed using the structures of protein domains obtained from SCOP 2.06 (Fox *et al.*, 2014). This collection contains 244,326 domains, extracted from 77,439 PDB entries. The 95% ID filtered subset of the domains, consisting of PDB-style files for 28,011 domains, was utilized to build the structural database. DSSP (Kabsch and Sander, 1983) is employed to assign secondary structure labels to each residue in these files.

Each structure file was processed by passing a sliding window of size 3 along the amino acid sequence. Each resulting tripeptide was added to the database if all of its 3 residues have a DSSP code of T, S, B, G or no code (which corresponds to an unclassified structural type). In other words, no portion of the tripeptide participates in an alpha-helix or beta-strand. The tripeptide state, which corresponds to its 9 backbone dihedral angles (3 sets of  $\phi$ ,  $\psi$ , and  $\omega$ ) was recorded in the database and indexed by its corresponding amino acid sequence.

A slightly different treatment was applied when the provided domain structure file originated from NMR data. For each structural file that contained more than one model, a distance filter was applied to corresponding tripeptides in each model to avoid redundancy in the database. A tripeptide state was considered sufficiently distant from another tripeptide state, and was thus added to the database, if it met at least one of the two following criteria: the RMSD on  $\omega$ ,  $\phi$  and  $\psi$

---

### Algorithm S1.1: Build Loop Brute Force

---

```

1 BuildAllLoops( $C_{init}$ ,  $L_{start}$ ,  $L_{end}$ )
2   Plans  $\leftarrow$  ConstructLoopPlans( $L_{start}$ ,  $L_{end}$ )
3   Trees  $\leftarrow$  ConstructRLTrees(Plans)
4   foreach plan  $\in$  Plans do
5     tree  $\leftarrow$  SelectTree(Trees, plan)
6     BuildLoopsPosBF(plan, tree,  $C_{init}$ , 1)
7 int BuildLoopsPosBF(plan, tree, C,  $pos_{tri}$ )
8    $nb_{sols} \leftarrow 0$ 
9   Tripeptides  $\leftarrow$  GetAllTripeptidesStates(plan,  $pos_{tri}$ )
10  foreach tripeptide  $\in$  Tripeptides do
11    tripeptide  $\leftarrow$  PerturbState(tripeptide)
12    C',  $success \leftarrow$  InstallTripeptide(C, tripeptide)
13    tree  $\leftarrow$  RecordSuccessPlacement(tripeptide, tree,  $success$ )
14    if  $success$  then
15      if  $pos_{tri} = \text{plan.lastIndex}$  then
16         $success \leftarrow$  CloseLoop(plan, C')
17        tree  $\leftarrow$  RecordSuccessClosure(tripeptide, tree,  $success$ )
18      else
19         $nb_{closed} \leftarrow$  BuildLoopsPosBF(plan, C',  $pos_{tri} + 1$ )
20         $nb_{sols} \leftarrow nb_{sols} + nb_{closed}$ 
21        tree  $\leftarrow$  RecordSuccessClosureNb(tripeptide, tree,  $nb_{closed}$ )
22  return  $nb_{sols}$ 

```

---

angles is above 0.3 radians, or one of the nine dihedral angles differs by more than 1 radian.

### S1.3 Tripeptide database analysis

#### S1.3.1 Motivation

Employing a database of protein configurations to discretize the conformational space capitalizes on the prior knowledge that the backbone dihedral angles only occupy a limited range of values, which are dependent on their neighboring amino acids. The technique of sampling from databases has been utilized in many structural biology problems, including loop sampling and de novo structure prediction.

While the success of these applications may show this approach has merit, we extend this idea by validating the

resulting tripeptide states that are formed by joining two tripeptides into a structure. Note that these tripeptides may be concatenated in one direction or in another, so that both directions have to be tested. Figure S1 shows two tripeptide sequences (MVK and PGT) that have been extracted from our database and joined together to construct a larger protein structure. The red lines highlight new tripeptides that are formed from the overlaps. We refer to the states of these tripeptides as synthetic states, since they were not sampled from the database, and thus, their structural validity is unknown.

We propose validating these synthetic states to strengthen the theoretical basis for our proposed approach. This validation occurs as follows. For each pair  $(i, j)$  of tripeptides in the database, we extract 4 synthetic states (as shown in Figure S1), formed by the concatenations  $ij$  and  $ji$ . Next, for each synthetic state, we search the database using the resulting amino acid keys. For the example in Figure S1, these keys would be VKP, KPG, GTM, and TMV respectively. If we are able to locate a state in the database that is similar, we label the synthetic state as valid. Given that our database is built from a small subset of the protein universe, we can not say anything about synthetic states for which we do not find a similar neighbor.

In this work, similarity is measured as the root mean square deviation (RMSD) of the 3 sets of  $\phi$ ,  $\psi$  and  $\omega$  backbone dihedral angles that define each tripeptide state. Similarity is established when the RMSD is less than some threshold  $\epsilon$ .



**Fig. S1.** Synthetic tripeptide states creation by sampling two states from the database. The four red lines represent the 4 synthetic tripeptide states extracted.

### S1.3.2 Results

The database validation analysis was performed with a threshold of 0.5 radians for the dihedral RMSD. To speed up the process, only 1 out of 1,000 synthetic states were randomly selected and tested for a close neighbor. This represents around 9.5 billion synthetic states tested. On average, 85 % of synthetic states had a close neighbor in the database. Figure S2 gives more precise results by sequence. Unsurprisingly for a database of tripeptides involved in coils, sequences containing glycines are the most populated. However, sequences containing rarer amino acids like cysteines, histidines, methionines and tryptophans contain much fewer states. The distribution of synthetic states with a close neighbor in the database shows that sequences containing these amino acids are also the ones for which the proportion of synthetic tripeptides with a close neighbor in the database are the lowest. This may point to a lack of data for these relatively rare sequences. This fact is also supported by Figure S2(d), in which sequences with few representatives in the database are also the ones for which the average distance of synthetic states to the database is the highest.

Cases in which the absence of close neighbors for a synthetic state is not due to lack of data, but to their very low probability to exist in physiological conditions, do not constitute an

important issue if one wants to exhaustively sample the conformational space. Generating loops with such states does not prevent any acceptable conformations from being sampled. Conversely, it might be an issue if one’s goal is to preserve the distribution of structural preferences encoded in the database when generating a loop ensemble. Nevertheless, generating statistically-meaningful conformational ensembles goes beyond the scope of this paper.

This analysis contributes to the validation of the method consisting in concatenating the tripeptides to build the loop. Indeed, most synthetic states without a close neighbor in the database actually coincide with a rare tripeptide sequence for which data is insufficient. However, the results reveal a limitation of the method, which is the availability of experimental data for rare sequences. Future improvements of MoMA-LoopSampler will therefore include enriching the database with additional states, potentially sharing states across similar sequences.

### S1.4 Loop closure

In this section, we provide details for the method `CloseLoop` (Algorithm 1, line 20).

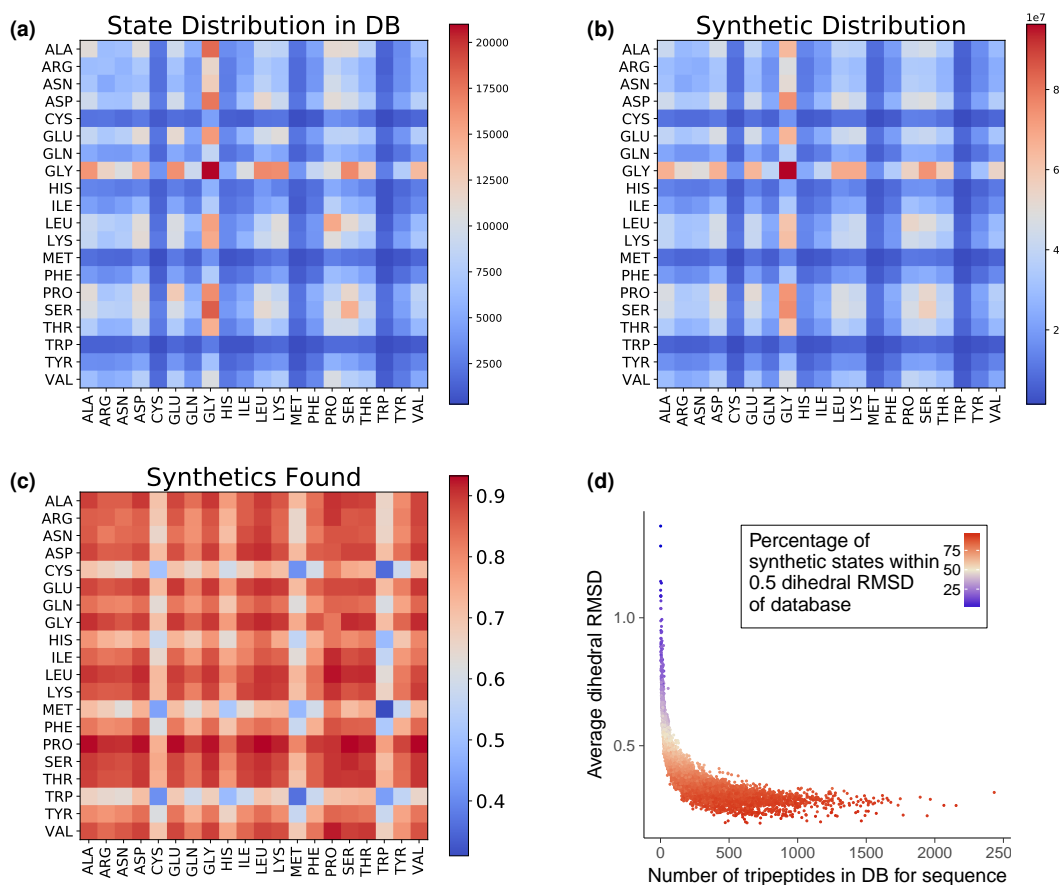
`PerturbOmegas` (line 25) samples the three  $\omega$  angles by first selecting their *cis* or *trans* configuration. The probability to generate a *cis*  $\omega$  angle is taken as the frequency of this event in the database for the corresponding residue. According to the selected configuration, the angle is then sampled around the value 0 or  $\pi$  with a Gaussian distribution whose standard deviation follows that observed in the database for this type of residue.

An in-house inverse kinematics (IK) solver (Cortés *et al.*, 2004) is used to solve the six remaining  $\phi$  and  $\psi$  angles of the backbone. However, the design of MoMA-LoopSampler allows for other inverse kinematics solvers to be employed (e.g. Manocha and Canny, 1994; Dinner, 2000; Coutsias *et al.*, 2004). Solving this problem may yield up to 16 potential solutions.

The IK solution yields tripeptide states that may not exist within the structural database. Section S1.5 describes a study showing that for the loops with a very low energy after relaxation, the IK solved tripeptide generally has a close structural neighbor within the database (Figure S3). This analysis also shows that upon loop relaxation, the distance of the final tripeptide to the database tends to lower, suggesting that the final tripeptide acts as a buffer that “absorbs” the rigidity of the other tripeptides by being more “lenient” on the distance to the database, and thus more flexible (Figure S4). Therefore, setting a threshold on the distance of the final tripeptide to the database may lead to the generation of higher quality loops, but care has to be taken in order not to be too restrictive on accepted loop conformations.

### S1.5 IK-solved tripeptide: distance to database and loop quality

The last tripeptide used to close the loop is the only tripeptide of the generated conformation that has not been sampled from the database. We define the distance between this tripeptide and another tripeptide state in the database as the angular root-mean-squared deviation of their nine backbone dihedral angles. The distance of the last tripeptide to the database is defined as the lowest distance between this tripeptide and



**Fig. S2.** Each tripeptide state in the database is projected into a 20 by 20 map based on its first two amino acid residues. (a) The top left heatmap shows the distribution of tripeptides using this projection for the SCOP 95% similarity database filtered to exclude states that participate within a secondary structure element. This database contains 2.2 million states. (b) The top right plot is the distribution of synthetic states created by concatenating tripeptide pairs in the database (approximately 9.5 billion states). (c) The bottom left plot shows for each of the synthetic states created, what percentage of these had close structural neighbors in the original database (neighbor distance  $< 0.5$  dihedral RMSD). (d) The bottom right plot shows for each tripeptide sequence the average dihedral RMSD to database obtained for the tested synthetic states depending on the number of actual states in the database.

a tripeptide in the database having the same amino-acid sequence.

For two loop systems, we performed brute force searches of loop conformations with side-chain placement. We then measured the distance to the database for the closing tripeptide of each sampled conformation and recorded the energy of each loop using Amber. The generated loops were then relaxed using Amber and their energy after relaxation was measured. The distance to database of the relaxed closing tripeptide was calculated again.

Figure S3 shows the relationship between the energy of generated loops and the distance of their closing tripeptide to the database. We see clear positive correlation between energy and the distance of the closing tripeptide to the database before relaxation. After relaxation, the profile changes a lot. Although the correlation is still present, it is mainly apparent for low energy loops. Setting a threshold on the distance to database for the closing tripeptide would thus have to be done very carefully.


Figure S4 shows the distribution of the distance to database differences after minus before relaxation for the closing tripeptide. Results show that the distribution is skewed to negative values, indicating a tendency to lower the distance

of the closing tripeptide to the database upon relaxation. This suggests that, by being unconstrained with regard to the database, the closing tripeptide acts as a buffer that absorbs the rigidity of the other tripeptides. Relaxation causes the other tripeptides to relax and the closing tripeptide to move closer to states in the database.


Setting a threshold on the distance of the closing tripeptide to the database should be done with a high tolerance since even loops with closing tripeptides far from the database may have a low energy. Moreover, looking at Figure S3(d), we can see that even after relaxation, a loop with a closing tripeptide at distance 1 radian of any tripeptide in the database may be relatively good compared to the generated loop ensemble.

### S1.6 Learning mode: leaf subdivision

A leaf in the learning structure can be split if the results obtained for the tripeptide states it contains become too heterogeneous with regard to construction success. The states contained in a leaf should be similar enough to have comparable levels of placement and subsequent loop closing success. When the results show otherwise, the leaf is split around its center. Since each placement event (successful



**Fig. S3.** Amber energy vs. distance to database for the closing tripeptide. (a) and (c) are the graphs for loop 45 (9 residues) while graphs (b) and (d) are the graphs for loop 31 (9 residues). (a) and (b) show the energy and distance before relaxation while (c) and (d) show the energy and distance to database after relaxation.




**Fig. S4.** Histogram of distance to database difference after and before relaxation. Negative values indicate that the distance to database lowered upon relaxation. (a) Loop 45. (b) Loop 31.

placement of a tripeptide, steric clash, loop closure) is recorded along with its corresponding tripeptide state, it is easy to recompute the statistics in each of the  $2^m$  newly created child leaves in case the leaf splits.

Each child leaf obtained after a split is assigned a new tree corresponding to the next position in the plan. Consequently, the statistics from the formerly pointed to tree are also

distributed among the newly created trees and their leaves. After this process, the scores of all newly created leaves obtained after the subdivision are the same as what would have been obtained if the tree had utilized this structure from the beginning of the sampling process.

Typically, a leaf split can happen when a state leads to a steric clash between 25% and 75% of the time. Splitting can



**Fig. S5.** Frames attached to the beginning and the end of a tripeptide. R1, R2 and R3 designate the first, second and third residues of the tripeptide, respectively. The x-, y- and z- axes are represented in red, green and blue, respectively. Only backbone atoms are represented for clarity. The C atom of the preceding residue and the N atom of the following residue are also visible.

help isolate the states that are responsible for steric clashes, while regrouping the states that lead to a successful placement. Another criterion for splitting the leaf is based on the frequency of forming a closed loop.

## S1.7 Learning mode: comparison of tripeptide projections

### S1.7.1 Principle

This analysis is aimed at comparing options for tripeptide state projections in the learning process. These options are compared on their ability to evenly distribute states in space and to cluster tripeptides states according to their success in forming closed loops.

### S1.7.2 Definitions

The definitions of the different tripeptide state projections are based on the tripeptide geometry corresponding to the nine dihedral angles constituting the state. We associate a reference frame to the beginning and the end of each tripeptide. The frame associated to the beginning of a tripeptide is centered on the N atom of the first residue, while the frame associated to the end is centered on the C atom of the last residue (see Figure S5). The length of a tripeptide is defined as the distance between the first and the last atom of the tripeptide backbone.

Several options were tested for tripeptide state projection. They mostly involve the relative position or orientation of the two ends of the tripeptide. For the orientation, several representations were tested. More precisely, the tested tripeptide state projections are (with the associated dimension  $m$  in parentheses):

**Position ( $m = 3$ ):** The vector of the translational part of the transformation between the beginning and the end of the tripeptide (as defined by the associated frames).

**Euler angles ( $m = 3$ ):** The vector of the Euler angles of the rotational part of the transformation between the beginning and the end of the tripeptide.

**Euler angles and length ( $m = 4$ ):** The vector containing the three Euler angles of the rotational part of the transformation between the beginning and the end of the tripeptide and the length of the tripeptide.

**Quaternion ( $m = 4$ ):** The quaternion representing the rotational part of the transformation between the beginning and the end of the tripeptide.

**Quaternion and length ( $m = 5$ ):** The vector containing the quaternion representing the rotational part of the transformation between the beginning and the end of the tripeptide, and the length of the tripeptide.

**Axis-angle ( $m = 4$ ):** The axis-angle representation of the rotational part of the transformation between the beginning and the end of the tripeptide.

**Axis-angle and length ( $m = 5$ ):** The axis-angle representation of the rotational part of the transformation between the beginning and the end of the tripeptide, and the length of the tripeptide.

### S1.7.3 Performance of the different projections

A good projection should provide a good distribution of the tripeptide states in space, and more importantly group together tripeptide states that are almost interchangeable in the loop building process. For example, if two states are very close to one another in the projection space and one is in a collision when placed, then the other should have a strong likelihood to be in collision as well.

A high value of the dimension  $m$  would result in a data structure that has very high memory requirements. This would likely result in leaves with too few tripeptides, which is not desirable. Indeed, when a tripeptide state is sampled, its results are exploited to make predictions about the success of all its neighbors in the leaf. If the neighbors are not numerous,  $H_{RL}$  will not be very effective.

We ran brute force searches for six loop systems using each of the different projections. We compare the first level of the octrees in all the runs performed, in terms of distribution of tripeptides and success probability of each cell. The first level of the octrees corresponds to the cells obtained after dividing the bounding box containing all tripeptide states once in each dimension. The number of cells at this level is thus  $2^m$  where  $m$  is the dimension of the projection. Our method employs several loop construction plans for each loop system (one per tripeptide in the loop, corresponding to a plan ending with this tripeptide). There are therefore 3 plans for the 9-residue loops, and 4 plans for 12-residue loops.

The success probability of a cell is defined as the number of successful combinations of tripeptides using a state from this cell, divided by the theoretical number of tripeptide combinations that use a state from this cell. Ideally, a good projection should distribute the tripeptides as uniformly as possible among the cells, while clearly separating states that work (i.e. that can be used to build a successful loop conformation) from those that cannot lead to a successful loop conformation. Another criterion to consider is the dimension of the projection. A projection with a higher dimension will have higher memory requirements.

Figure S17 (at the end of this document) shows the results for the first levels of all loop systems, for each of their loop construction plan employed. Looking at the different heatmaps, it is clear that no projection performs consistently better than all others. However, we will try to analyze the differences in the results. First, looking at the 3-dimensional projections *position* and *Euler angles*: tripeptide distribution

seems to be rather good in both cases, with very few empty cells. However, the distribution of solutions is more heterogeneous with *position*. Indeed, with this projection, more of the first level cells are void of working solutions. This is particularly striking for 1dim-12 (A213-A224) (loop 55), for the three last plans. *Euler angles* is not able to gather the solutions into only a few cells, whereas *position* concentrates the solutions into two or three cells. The same observation can be made for the other systems. *Position* therefore seems to be a better predictor of tripeptide success in building a loop than *Euler angles*. Comparing the 4-dimensional projections is more delicate. Indeed, *axis-angle*, *Euler angles and length* and *quaternion* behave differently between the different systems. It seems that *Euler angles and length* better gathers the solutions for 1dim-12 (A213-A224) (loop 55) than other projections, but on 153l-12 (A98-A109) (loop 61), *quaternion* is the one that better gathers solutions, while *axis-angle* gathers solutions relatively well in all the cases. For the distribution of tripeptides, the same observation can be made. The quality of the distribution differs depending on the tripeptide sequence. However, 4-dimensional projections do not seem to perform better than *position*. For most systems and most loop plans, there are more than 8 cells left with solutions. Of course those cells contain fewer tripeptide states on average, which is why comparing 3-dimensional and 4-dimensional projections is delicate. In 5-dimensional projections, the distribution of tripeptides results in an undesirable sparsity, where the solutions are unsurprisingly found where tripeptides are located. Therefore these do not stand out from all the projections either.

This comparison is only possible on first level cells with these plots. Many other levels are left to explore and each projection is likely to better separate working and non working states in lower levels. With our analysis, we decided to retain a 3-dimensional projection in order to limit the size of the learning tree in memory. In light of the results, we thus selected *position* as the tripeptide projection for our tests.

## S2 Results

### S2.1 Loop systems

The performance of MoMA-LoopSampler was tested on three benchmark sets of 9-residue loops, 12-residue loops and 15-residue loops. The 9-residue test set is a subset of the loops gathered by Jacobson and colleagues (Jacobson *et al.*, 2004). 2alp-139 and 8ruc-79 were removed because they are not included in the set by Soto and co-workers (Soto *et al.*, 2008). 1lvd-244 and 1pda-108, which were excluded from the modified Fiser set by DePristo and colleagues for either poor quality or missing side-chain atoms and gaps, were also removed (DePristo *et al.*, 2003). Finally, 4gcr-94 was removed because it only contains 3 turn residues surrounded by  $\beta$ -sheet and  $\alpha$ -helix residues, and is not strictly speaking a loop. Thus, the 9-residue test set involves 53 loops. The 12 residue test set contains the ten 12-residue loops gathered by Jacobson and colleagues (Jacobson *et al.*, 2004), and the 15 residue test set contains the 30 15-residue loops used in the analysis by Zhao *et al.*, 2011. The full lists of loop systems used in this work are provided in Tables S1-S3.

**Table S1.** 9-residue Test Set

Loop number	Corresponding PDB numbering
1	3pte-09 (A107-A115)
2	1xyz-09 (A795-A803)
3	1lkk-09 (A193-A201)
4	1mla-09 (A194-A202)
5	2ayh-09 (A169-A177)
6	1arb-09 (A90-A98)
7	1mrp-09 (A284-A292)
8	1cse-09 (E95-E103)
9	1tca-09 (A217-A225)
10	2eng-09 (A172-A180)
11	1xyz-09 (A568-A576)
12	1aba-09 (A69-A77)
13	1nif-09 (A266-A274)
14	1arp-09 (A127-A135)
15	1noa-09 (A99-A107)
16	1lkk-09 (A142-A150)
17	1xif-09 (A59-A67)
18	1rhs-09 (A216-A224)
19	1fus-09 (A91-A99)
20	1php-09 (A91-A99)
21	2cpl-09 (A24-A32)
22	1nls-09 (A131-A139)
23	1xnb-09 (A133-A141)
24	1btl-09 (A102-A110)
25	1mrj-09 (A92-A100)
26	1gpr-09 (A63-A71)
27	1csh-09 (A252-A260)
28	1sgp-09 (E109-E117)
29	3pte-09 (A78-A86)
30	1isu-09 (A30-A38)
31	1noa-09 (A9-A17)
32	2hbg-09 (A18-A26)
33	1nfp-09 (A12-A20)
34	1pgs-09 (A117-A125)
35	1tca-09 (A170-A178)
36	1ptf-09 (A10-A18)
37	1mpk-09 (A102-A110)
38	3pte-09 (A215-A223)
39	1ra9-09 (A142-A150)
40	1mrk-09 (A53-A61)
41	1wer-09 (A942-A950)
42	3tgl-09 (A56-A64)
43	2sil-09 (A183-A191)
44	1amp-09 (A57-A65)
45	1aac-09 (A58-A66)
46	1arb-09 (A168-A176)
47	1fus-09 (A31-A39)
48	1byb-09 (A246-A254)
49	1xnb-09 (A116-A124)
50	1ede-09 (A257-A265)
51	1aru-09 (A36-A44)
52	1onc-09 (A70-A78)
53	1noa-09 (A76-A84)

**Table S2.** 12-residue Test Set

Loop number	Corresponding PDB numbering
54	1arb-12 (A74-A85)
55	1dim-12 (A213-A224)
56	1xyz-12 (A813-A824)
57	1bkf-12 (A9-A20)
58	2ayh-12 (A21-A32)
59	1akz-12 (A181-A192)
60	1luc-12 (A158-A169)
61	153l-12 (A98-A109)
62	1cex-12 (A40-A51)
63	1ixh-12 (A160-A171)

### S2.2 Sampling efficiency of basic MoMA-LoopSampler

#### S2.2.1 Distance to native loop

To test the ability of MoMA-LoopSampler to sample loop configurations close to the native state, we ran a series of tests on each test dataset. The run time for each test was determined

**Table S3.** 15-residue Test Set

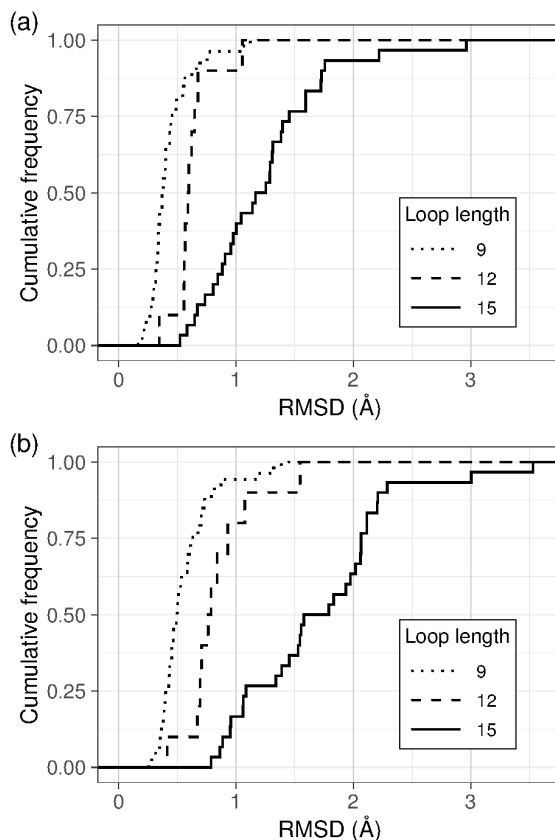
Loop number	Corresponding PDB numbering
64	2v3v-15 (A382-A396)
65	1qqf-15 (A1112-A1126)
66	1h4a-15 (X19-X33)
67	2aeb-15 (B156-B170)
68	3a3p-15 (A286-A300)
69	1wui-15 (L454-L468)
70	1qaz-15 (A298-A312)
71	3css-15 (A95-A109)
72	3a64-15 (A350-A364)
73	2o2k-15 (A1220-A1234)
74	1ju3-15 (A486-A500)
75	2h3l-15 (A1339-A1353)
76	2cjp-15 (A58-A72)
77	1ryo-15 (A172-A186)
78	1ah7-15 (A157-A171)
79	1s95-15 (A477-A491)
80	1ra0-15 (A361-A375)
81	1wb4-15 (A1033-A1047)
82	1y12-15 (A10-A24)
83	3f1l-15 (A99-A113)
84	2pkf-15 (A26-A40)
85	2oit-15 (A290-A304)
86	2dsj-15 (A354-A368)
87	1bhe-15 (A121-A135)
88	3ea1-15 (A136-A150)
89	2b0t-15 (A701-A715)
90	1ra0-15 (A283-A297)
91	1zhx-15 (A392-A406)
92	3bb7-15 (A231-A245)
93	3bf7-15 (A49-A63)

by the loop length being sampled: 2 hours for 9-residue loops, 4 hours for 12-residue loops, and 6 hours for 15-residue loops. Each experiment was repeated 4 times.<sup>1</sup>

Figure S6 gives the distribution of the lowest RMSD to the native loop for both the best and worst of the four executions. In all four tests, at least one conformation within 2 Å of native (using backbone RMSD as a distance metric) was sampled for each of the 9- and 12-residue loops, and between 23 and 26 of the 15-residue loops. Decreasing the threshold to test for a sample within 1 Å of the native state, the results are still very good for 9- and 12-residue loops: with 51 9-residue loops (out of 53) and 9 12-residue loops (out of 10). For the 15-residue loops however, this number drops to 9 loops (over all four tests, this was achieved for 11 loops). The fact that these results vary from one test to another suggests that the sampling time is not sufficient for some 15-residue loops. With such a loop length, the number of possible conformations may be very large when the loop environment is not strongly constrained, and thus, the sampling time must be adapted to the size of the conformational space.

The times observed to sample a loop within 1 or 2 Å of the native are reported in Table S4. Computational requirements increase with loop length. Given that the native loop is one of many valid conformations, the relatively high variance in the time to generate a nearby configuration for long loops is not surprising. This is due to the stochasticity of the method. The variation of  $\text{RMSD}_{\min}$  for 15-residue loops suggests that the provided sampling time is not adequate to perform a sufficient sampling of the loop’s conformational space.

<sup>1</sup> Computing times reported throughout the paper correspond to runs on a single core of a 2.30 GHz Intel® Xeon® E5-2695 v3 processor.



**Fig. S6.** Cumulative distribution of the lowest backbone RMSD to native among sampled conformations for the 93 loops in our test sets. Since sampling was performed four times, (a) shows the distribution of the best of the four results for each loop, while (b) shows the distribution of the worst of the four results.

These results show that MoMA-LoopSampler can construct loop ensembles for the 9 and 12-residue cases that include the native state with high precision, while 15-residue loops still present a formidable challenge, due to the high dimensionality of the search space possibly coupled with a less constrained environment.

### S2.2.2 Comparison of MoMA-LoopSampler with loop prediction methods

We compare the sampling performance of MoMA-LoopSampler to that of DiSGRO (Tang *et al.*, 2014) and of the updated version of RCD (Chys and Chacón, 2013; López-Blanco *et al.*, 2016). Note that, based on the results from references (Soto *et al.*, 2008) and (Tang *et al.*, 2014), DiSGRO performs better than earlier loop closure or loop prediction methods such as CCD (Canutescu and Dunbrack, 2003), Wriggling (Cahill *et al.*, 2003), PLOP-build (Jacobson *et al.*, 2004), LOOPY<sub>bb</sub> (Xiang *et al.*, 2002), Random Tweak (Shenkin *et al.*, 1987), or Direct Tweak (Xiang *et al.*, 2002; Xiang, 2006). Therefore, we do not compare directly to these older methods.

Source code for DiSGRO was obtained from <http://tanto.bioe.uic.edu/DiSGro/download.html>. The code had to be slightly modified to output exactly the required number of clash-free conformations (instead of the subset of clash-free conformations among a required number of closed ones), and

**Table S4.** Time needed to generate a conformation within 1 or 2 Å of the native loop. Statistics are calculated on all the runs for which such a distance was reached. MoMA-LoopSampler sampled a loop within 2 Å of native in at least one of the four tests for all 9- and 12-residue loops and 28 15-residue loops. It sampled a loop within 1 Å of native in at least one of the four tests for 51 9-residue loops, 9 12-residue loops and 11 15-residue loops.

Test Set	Time to reach...							
	2 Å to native				1 Å to native			
	Min	Median	Mean	Max	Min	Median	Mean	Max
9 residues	0.3 sec	4.5 sec	21.4 sec	11.8 min	1.0 sec	28.6 sec	4.1 min	1.7 h
12 residues	0.89 sec	23.4 sec	2.5 min	28.9 min	3.26 sec	5.6 min	16.5 min	2.2 h
15 residues	2.84 sec	19.6 min	1 h	5.9 h	4 min	1 h	1.6 h	5.1 h

conformations were generated without side-chains. Binaries for RCD version 1.40 were downloaded from <http://chaconlab.org/modeling/rcd/rcd-download>. MoMA-LoopSampler uses stricter constraints than DiSGRO and RCD, in particular for steric clash detection. Therefore, in order to more adequately compare running times, we also tested a variant of MoMA-LoopSampler (Soft MoMA-LoopSampler) that uses collision constraints comparable to that of DiSGRO and RCD. This variant uses a van der Waals scaling factor of 0.6 (instead of 0.7 for the other tests), does not use enlarged  $C_\beta$  atoms, and uses a lower  $max_{IK}$ . These changes are expected to lower the quality of the ensemble and its exhaustiveness, but also to considerably decrease sampling time, allowing a more straightforward time comparison with DiSGRO and RCD. The computational time and ability to generate near-native loops are compared, using the same computational resources for all four methods.

#### Difference between the sampling methods

Different sets of constraints are enforced by the three sampling methods. Concerning collisions, DiSGRO employ an energy function that makes steric clashes unlikely. The maximum allowed ratio between non-bonded atom pair distances and the sum of their van der Waals radii (called van der Waals scaling factor) was set to 0.6 in DiSGRO (value found in the source code) and to 0.5 for intra-loop backbone collisions in RCD. Collisions with the rest of the protein are handled differently in RCD: this method uses a grid and considers that there is a collision if an atom of the loop is placed in a non-empty cell. This collision detection method is much faster but also less accurate than considering the actual distance between atoms. In MoMA-LoopSampler, steric clash avoidance is a crucial component: backbone atoms and the  $C_\beta$  atoms (with enlarged volumes to account for side-chain placement) of the sampled loops are placed without major steric clash among themselves or with the rest of the protein. While the van der Waals scaling factor was set to 0.6 in the more collision-tolerant version of MoMA-LoopSampler (Soft MoMA-LoopSampler), we set this cutoff at 0.7 to test the basic and  $H_{RL}$  MoMA-LoopSampler.

Structural knowledge is included in all three methods. Although RCD can be considered an *ab initio* method, it samples dihedral angles following neighbor-dependent Ramachandran probability distributions. DiSGRO includes a stronger knowledge-based component, with a more complex dihedral angle sampling that follows distributions extracted from a structural database. Finally, MoMA-LoopSampler is

strongly dependent on structural knowledge, since it directly uses fragments from experimentally-solved protein structures.

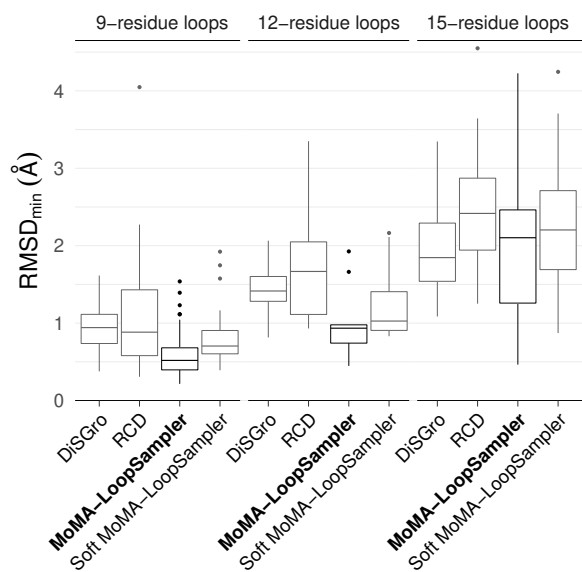
#### Results

Among 5000 sampled conformations,  $RMSD_{min}$  obtained by MoMA-LoopSampler are much lower than that obtained by the other methods for 9- and 12-residue loops (Figure S7(a)). Soft MoMA-LoopSampler also obtains  $RMSD_{min}$  lower than DiSGRO and RCD, but higher than the basic version of MoMA-LoopSampler. For the 15-residue loops,  $RMSD_{min}$  obtained on 5,000 sampled conformations are comparable for the four methods. However, looking at 100,000 sampled conformations for these longer loops, MoMA-LoopSampler obtains a much lower  $RMSD_{min}$  (Figure S7(b)). Note that the  $RMSD_{min}$ s obtained for DiSGRO on the 15-residue loop test set are lower than the ones the authors report for 100,000 sampled conformations on the same test set (Tang *et al.*, 2014).

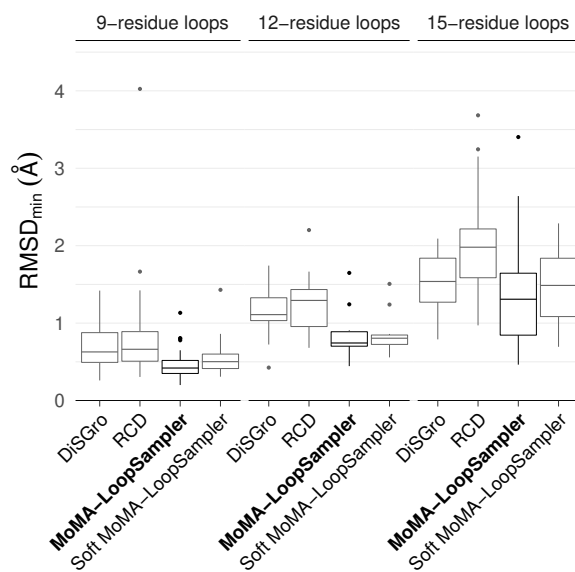
Generating 100,000 conformations instead of 5,000 lowered the  $RMSD_{min}$  for all the methods (Figure S9). However, this decrease varies from one method to another. It is very limited for MoMA-LoopSampler on the 9- and 12-residue loops, while being considerable for other methods. Note that despite this, the  $RMSD_{min}$  obtained by MoMA-LoopSampler on these loops is still much lower compared to the other tested methods. The 15-residue loops show the opposite trend: the decrease in  $RMSD_{min}$  is much larger for MoMA-LoopSampler and its “soft” version than for the other methods on 15-residue loops, resulting in a lower  $RMSD_{min}$  for MoMA-LoopSampler compared to RCD and DiSGRO.

Concerning running times, MoMA-LoopSampler and its soft version are much less sensitive to the length of the loop than the other two methods. The soft version of MoMA-LoopSampler is slightly slower than other methods on 9-residue loops but faster than DiSGRO for 12- and 15-residue loops. We note that the running times obtained for DiSGRO are higher than those reported in Tang *et al.*, 2014. In some tests, DiSGRO got blocked during the sampling process, in which case we started the run again. These failed sampling attempts are not counted in the running times we report. The basic version of MoMA-LoopSampler is unsurprisingly slower than the other methods, due to the stronger constraints it enforces. Overall, results show that the “soft” version of MoMA-LoopSampler has running times comparable to that of other methods and provides slightly lower  $RMSD_{min}$ . The basic version of MoMA-LoopSampler on the other hand, trades computational efficiency off for better filtering of sampled ensembles.






(a) 5,000 sampled conformations.

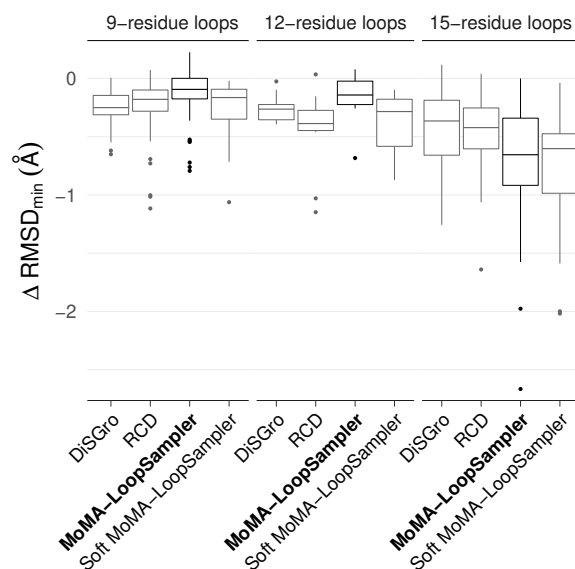


(b) 100,000 sampled conformations.

**Fig. S7.** Minimum distance to native ( $\text{RMSD}_{\min}$ ) obtained among sampled conformations, without side-chains.



**Fig. S8.** Median time per sampled conformation (estimated on 5000 sampled conformations, without side-chain placement). Computations were performed using a single core of a 2.30 GHz Intel® Xeon® E5-2695 v3 processor.



**Fig. S9.** Difference in  $\text{RMSD}_{\min}$  obtained when generating a larger number of conformations (100,000 conformations instead of 5,000).

### Discussion

The difference in the collision constraints enforced by the different methods can help to explain two major observations in the results: (1) MoMA-LoopSampler (especially the basic version) is less sensitive to loop length than RCD and DiSGro, from a running time point of view; (2) while MoMA-LoopSampler obtains a much lower  $\text{RMSD}_{\min}$  than other methods for 9- and 12-residue loops on ensembles of 5,000 conformations, a higher number of sampled conformations is necessary for 15-residue loops to observe a difference.

To explain the first point, we hypothesize that MoMA-LoopSampler is more sensitive to the environment of the loop,

and whether it is constrained or not, than to the length of the loop. First, an essential component of the method consists in checking that the distance between the two working loop ends can be covered by the tripeptides left to place. Making this verification after a tripeptide is added facilitates the closing of long loops. Second, shorter loops generally have a more constrained environment than longer loops. As previously mentioned, MoMA-LoopSampler is more intolerant to collisions compared to other methods. Therefore, generating a conformation for a loop in a more constrained environment is a problem with a difficulty comparable to that of sampling a longer unconstrained loop for MoMA-LoopSampler, which

explains why median running times per conformation vary little with loop length.

To explain the second point, a similar reasoning can be conducted. MoMA-LoopSampler only samples the accessible conformational space by carefully avoiding collisions in generated conformations. Therefore, it finds the native conformations using fewer samples than other methods. With more constrained environments, 9- and 12- residue loops have a resulting conformational space that is particularly reduced, which is why MoMA-LoopSampler finds the native conformation very early into the search. Longer loops are usually more flexible, and a larger portion of the conformational space is allowed. Therefore, methods that are overall more tolerant to collisions sample fewer bad-quality conformations in proportion. For these loops, more conformations need to be sampled to achieve a better coverage of the conformational space. The benefit in  $\text{RMSD}_{\min}$  for MoMA-LoopSampler is thus logically observable when sampling a higher number of conformations (100,000) for 15-residue loops (Figure S7(b)).

In simpler terms, this means that MoMA-LoopSampler better explores the conformational space, performing an exhaustive exploration using fewer samples than other methods. The difference in  $\text{RMSD}_{\min}$  observed after generating larger ensembles further supports this idea. Indeed, 5,000 samples from MoMA-LoopSampler are enough to explore the conformational space of 9- and 12-residue loops, explaining why the  $\text{RMSD}_{\min}$  barely decreases when generating a much larger number of conformations. For other methods, the  $\text{RMSD}_{\min}$  considerably decreases upon generating more conformations, showing that these methods keep discovering relevant conformations among the extra conformations. Conversely, for 15-residue loops, the difference in  $\text{RMSD}_{\min}$  obtained upon generating 100,000 conformations instead of 5,000 is considerable for all four methods, with MoMA-LoopSampler and Soft MoMA-LoopSampler showing a much larger decrease than RCD and DiSGRO. This confirms (1) that 5,000 conformations are not enough to cover the much larger conformational space of these longer loops, and (2) that MoMA-LoopSampler performs a more efficient exploration, discovering relevant conformations using fewer samples than RCD and DiSGRO.

Obtaining an ensemble of conformations of good quality is essential considering the costly downstream processing steps of applications involving loop sampling, in particular in the context of stable states prediction. These steps include side-chain addition, relaxation, scoring, clustering or filtering, and can be extremely time-consuming. In that regard, generating fewer conformations, but which are more representative of the ensemble overall, is perfectly satisfactory. This suggests that MoMA-LoopSampler is a good candidate for the sampling stage of many structural bioinformatics applications, including stable states prediction, since it obtains the same  $\text{RMSD}_{\min}$  as other methods (or a lower one) without needing to sample as many conformations as these methods do.

## S2.3 Application to a multi-state loop: complements

### S2.3.1 Post-processing of streptavidin loop conformations

Several experimentally-determined structures were gathered for streptavidin and sampling for the flexible loop was

performed from each of them. For all the conformations obtained after three rounds of brute force search, side-chain placement was attempted with an in-house method using continuous rotamers from BASILISK (Harder *et al.*, 2010). Only conformations for which side-chains could be placed without steric clash were retained and relaxed using energy minimization protocols from Amber 16 (Case *et al.*, 2005, 2016). This represents 5338 conformations from scaffold 2F01 chain A, 3825 from 2F01 chain B, 4042 from 3RY1 chain A, 1304 from 3RY1 chain B, 702 from 3RY1 chain C, 2820 from 3RY1 chain D, 5320 from 3RY2 chain A and 3611 from 3RY2 chain B. No constraints were applied for the relaxation, so that both backbone and side-chain movements were allowed. Bond length, bond angles and dihedral angles were all free to move. The relaxation took 133 s per loop on average. The first 250 cycles used steepest descent minimization, while the remaining cycles applied conjugate gradient. The maximum number of cycles was set to 500 and the minimization was considered to have converged when the root-mean-square of the cartesian elements of the gradient was lower than 0.1 kcal/(mol·Å). Energy was calculated using the ff14SBonlysc force field and a simple Generalized Born implicit solvent model ( $igb = 1$  and the  $mbondi$  radii sets, as recommended in the Amber manual).


### S2.3.2 Detailed analysis of the energy landscape of streptavidin’s flexible loop

The projection plots show major differences depending on the starting structure. Sometimes, both the “open” and “closed” conformations appear to be in low-energy basins (2F01-A, 2F01-B, 3RY2-A, 3RY1-B and 3RY1-C). Other projection plots show only one low-energy basin, around the crystallographic conformations (3RY2-B, 3RY1-A and 3RY1-D). The energies were calculated without a ligand, showing the strong influence of the local structure around the loop. In the case of 2F01-A, 2F01-B, and 3RY2-A, the environment surrounding the loop allows it to adopt both conformations, but the presence of the ligand probably stabilizes the loop in one of the two basins. Energy barriers of different heights separate the “closed” and “open” basins. In 3RY2-B, it seems that only the “closed” conformation is stable, suggesting that the environment of the loop also changes (due to crystal packing or ligand binding), and stabilizes this conformation. A profile similar to that of 3RY2-B is found for 3RY1-A, although this subunit is unbound in the crystallographic structure. This is an indication of the conformational changes that occur around the loop in the crystal structure. For 3RY1-D, the “open” conformation is in a large low energy basin. A few low energy conformations are found in the “closed” loop region but a large, high energy barrier separates the two regions. In the case of 3RY1-B and 3RY1-C, the conformational space appears tighter. The loop environment is probably more constrained sterically. Nevertheless, both energy basins are found, with a lower energy barrier separating them, and a much lower energy minimum for the “open” basin, explaining the “open” conformation adopted by these subunits in the crystal.


### S2.3.3 Landscapes obtained using other sampling methods

The protocol for modeling and visualizing the energy landscape of streptavidin (Section 3.2.2) was repeated using DiSGRO

**Fig. S10.** 2D projections of conformations sampled using DiSGRO for a loop in the streptavidin protein, from eight starting X-ray crystallography protein structures. The first dimension,  $d_1$  (x-axis), is the distance (Å) between an atom located in the middle of the loop and a fixed atom in the protein. The second dimension,  $d_2$  (y-axis), is the angle (degrees) formed by three atoms: an atom at approximately one quarter, one half, and three quarters of the way down of the loop. The conformations from the crystallographic structures are shown in black. For each system, the loop with the lowest energy was identified and each conformation was then colored according to the difference between its energy and this lowest energy.



**Fig. S11.** 2D projections of conformations sampled using RCD+ for a loop in the streptavidin protein, from eight starting X-ray crystallography protein structures. The first dimension,  $d_1$  (x-axis), is the distance (Å) between an atom located in the middle of the loop and a fixed atom in the protein. The second dimension,  $d_2$  (y-axis), is the angle (degrees) formed by three atoms: an atom at approximately one quarter, one half, and three quarters of the way down of the loop. The conformations from the crystallographic structures are shown in black. For each system, the loop with the lowest energy was identified and each conformation was then colored according to the difference between its energy and this lowest energy.



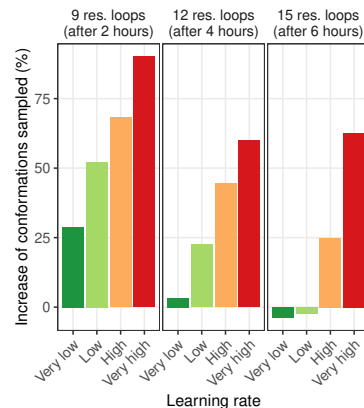
and RCD+ (López-Blanco *et al.*, 2016) as sampling methods. RCD+ is a web server that performs loop modeling using RCD as a first sampling step and then performs side-chain placement and refinement. For each scaffold, DiSGRO and RCD+ were employed to sample the same number of conformations as sampled using MoMA-LoopSampler in brute force mode. Side-chain placement was activated for DiSGRO. The sampled conformations were relaxed and projected in 2D. The resulting landscapes are shown in Figures S10 and S11.

The landscapes are very different from those obtained using MoMA-LoopSampler. Although some common features can be observed (such as the basin around the “open” conformations from scaffolds 3RY1 B, C and D), the landscapes are much rougher and harder to interpret. They are also more spread than landscapes obtained when sampling with MoMA-LoopSampler. The fact that these methods are more collision-tolerant may explain these observations. Indeed, many statistically unlikely conformations are generated, perturbing the analysis of the landscape. By creating a better filtered ensemble, MoMA-LoopSampler clarifies the analysis of energy landscapes for this loop.

## S2.4 Effects of $H_{RL}$ on the generated ensembles

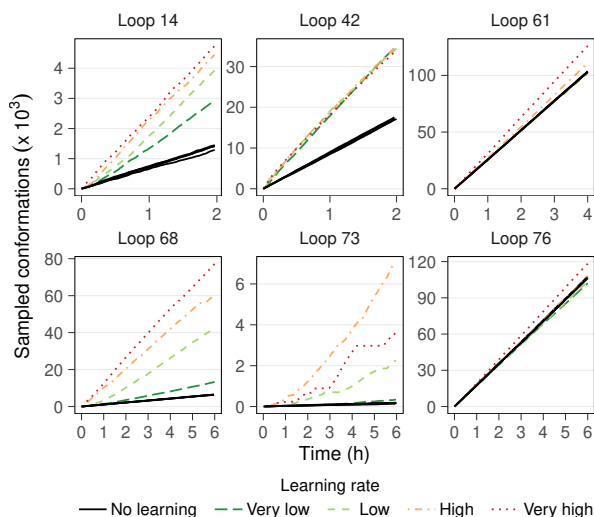
### S2.4.1 Influence of learning rate over the number of sampled conformations

Figure S12 shows that the higher the learning rate, the larger the number of conformations sampled. This is true for many



**Fig. S12.** Median percentage of increase in number of conformations sampled in  $H_{RL}$  modes compared to the basic mode across loop systems.

systems such as loops 14 and 68 (Figure S13), which show very different sampling speed depending on the learning rate. However, the effect of  $H_{RL}$  depends on both the length of the loop and the loop/protein system itself. While most 9- and 12-residue loops (and a few 15-residue loops) exhibit this expected behavior, for other loops (and for many 15-residue loops), only runs performed with *high* and *very high* learning rates are capable of generating a larger number of conformations (e.g. loop 61). As mentioned in the main paper, activating  $H_{RL}$  with a *very low* or *low* learning rate can even reduce




**Fig. S13.** Number of conformations sampled as a function of time for different learning rates for a few representative loops.

the number of sampled conformations, compared to the basic mode. This is mainly observed for 15-residue loops, such as loop 76. Loop 42 illustrates yet another case: all learning rates generate loops at comparable speed, but still much faster than MoMA-LoopSampler in basic mode does.

Loops 68 and 73 also illustrate an interesting phenomenon. The curves for *high* and *very high* learning rates show some plateaus spanning 10 minutes or more, which are due to the overhead of maintaining a very large learning tree, whose dimensionality grows exponentially in the length of the loop.

#### S2.4.2 Time needed to sample near-native conformations

The cumulative distribution of the time necessary to generate the first conformation within 1 Å RMSD to native, with or without learning, is shown Figure S14. Although these values are generally of the same order of magnitude for runs with or without learning, some observations can still be made. Looking at results loop by loop, we observe that for 9-residue loops, using  $H_{RL}$  may considerably delay the sampling of a conformation that is close to native. A possible reason is that introducing  $H_{RL}$  modifies the probability for selecting tripeptides. While in the basic mode, MoMA-LoopSampler picks tripeptide states at each step with a uniform distribution, the  $H_{RL}$  mode offers MoMA-LoopSampler the possibility to adjust the sampling of states so that a suitable distribution is obtained. Considering the tree used by  $H_{RL}$  to organize the tripeptides, MoMA-LoopSampler in the basic mode chooses each cell with a distribution that is directly proportional to the number of tripeptides it contains. Conversely, in the learning mode, MoMA-LoopSampler samples each cell according to their score. The score is currently set so as to sample effectively as many diverse conformations as possible, but other strategies may be contemplated in order to obtain a loop ensemble that follows the density of the tripeptide database for each tripeptide position. Such an ensemble could provide a more statistically accurate representation of the loop conformational space, which would be interesting to analyze entropic effects.



**Fig. S14.** Cumulative distribution of the time needed to sampled a conformation within 1 Å of the native for the 93 loops in our test sets. Four distributions were calculated: one using the shortest time of the four tests in basic mode for each loop, one using the longest time of the four tests in basic mode for each loop, one using the shortest time of the four tests in  $H_{RL}$  mode for each loop, and finally one using the longest time of the four tests in  $H_{RL}$  mode for each loop.


#### S2.4.3 Effect of $H_{RL}$ on $RMSD_{min}$

Main results and discussions regarding the effect of  $H_{RL}$  in the value of  $RMSD_{min}$  for the loop test-sets are presented in the manuscript. As a complement to these results, Table S5 gives the mean and median values obtained for  $RMSD_{min}$ .

#### S2.4.4 Evolution of sampling in time

We also show the evolution of the distribution of sampled loops with  $H_{RL}$  activated in Figure S15. These heatmaps, shown in the same two-dimensional projection, showcase the density of the sampled loops in the first and the last ten minutes of the exploration. Two effects are observed when comparing the beginning and the end of the sampling process:

(1) The first effect of learning is that the number of sampled conformations increases, and the coverage of the conformational space improves. In other words, the projection of sampled loop conformations in 2D appears to be more homogeneous and continuous. The effect is similar to that



**Fig. S15.** Evolution of the distribution of sampled loops due to  $H_{RL}$ . The plots show heatmaps of two-dimensional projections of the sampled loops for different learning rates, during the first and the last ten minutes of sampling. The x-axis gives the first projection descriptor  $d_1$  (in Å), while the y-axis gives the second projection descriptor,  $d_2$  (in degrees). (a) Loop 21, (b) Loop 26, (c) Loop 40, (d) Loop 68.

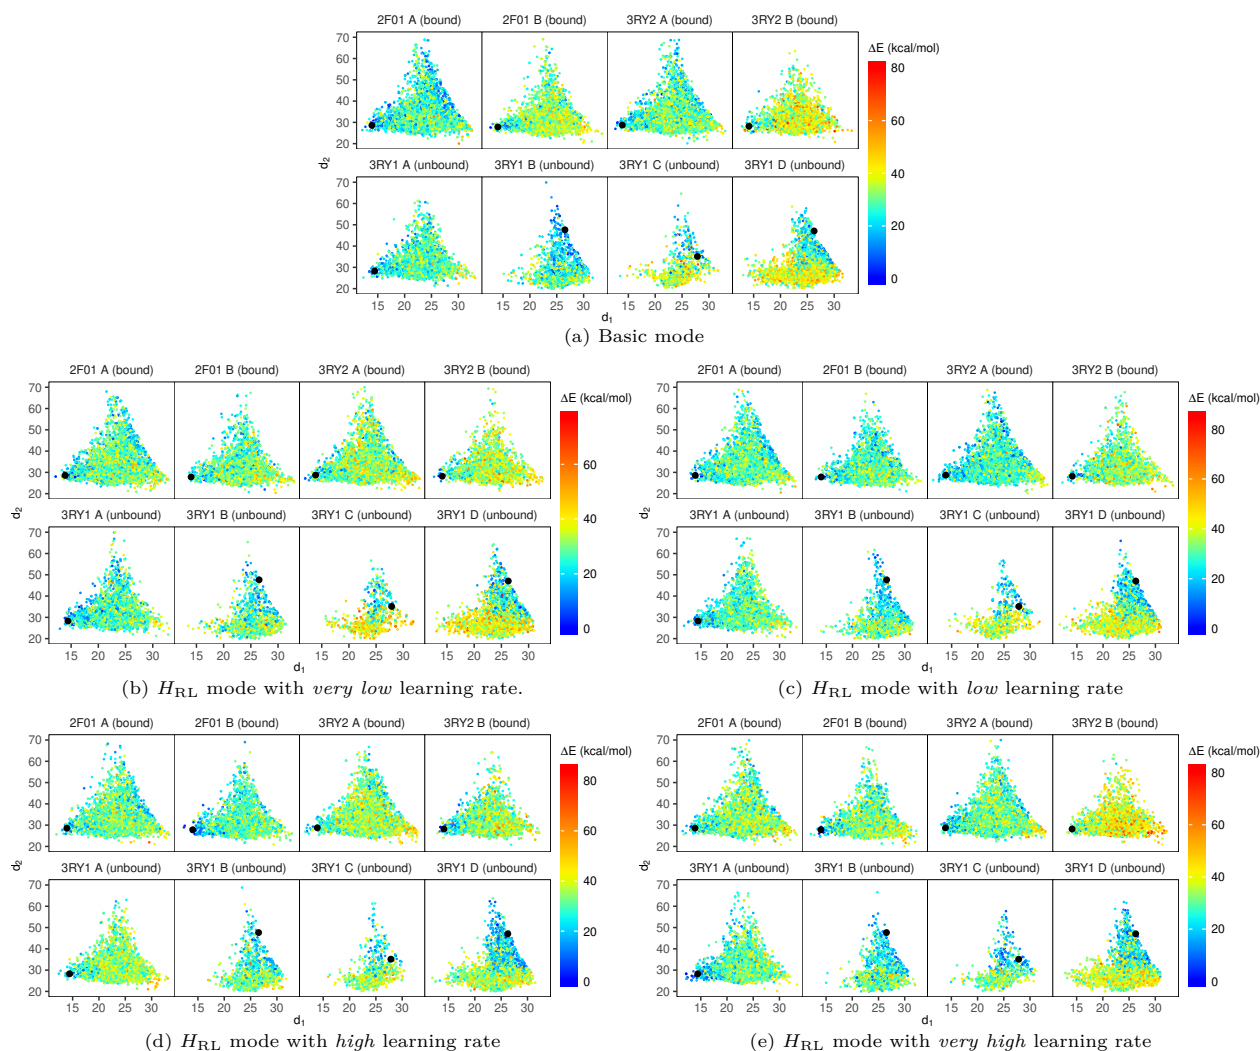
**Table S5.**  $RMSD_{min}$  obtained for different learning rates.

Length	Basic mode	$H_{RL}$ mode learning rate				
		V. low	Low	High	V. high	
9 res.	Mean (Å)	0.49	0.48	0.47	0.46	0.46
	Median (Å)	0.43	0.41	0.41	0.39	0.41
	$SD$ (Å)	0.21	0.21	0.21	0.20	0.19
12 res.	Mean (Å)	0.73	0.67	0.73	0.67	0.75
	Median (Å)	0.67	0.60	0.64	0.55	0.69
	$SD$ (Å)	0.26	0.20	0.30	0.32	0.26
15 res.	Mean (Å)	1.50	1.52	1.48	1.48	1.50
	Median (Å)	1.45	1.45	1.43	1.42	1.44
	$SD$ (Å)	0.62	0.68	0.70	0.67	0.73

observed for the different learning rates. Indeed, the algorithm progressively stops exploring the regions of space where it does not find any solution. The probability to sample a tripeptide in a cell from which all attempts have failed so far decreases with running time. Consequently, the success rate becomes

higher since MoMA-LoopSampler focuses on the vicinity of regions that are successful. This is very clear in Figure S15(d), for *very low* and *low* learning rates. It can also be observed for other systems, although to a lesser degree. The ability of  $H_{RL}$  to quickly identify areas where no solution exists depends on the positions of solutions in the conformational space and how they cluster, on the projection chosen to organize the tripeptide states, and on the learning rate.

(2) The second effect is the sudden discovery of whole regions of the conformational space. As previously mentioned, the probability to explore a region invariably found unsuccessful so far decreases based on the the number of attempts and the learning rate. If the number of attempts is too low (the learning process is too greedy), MoMA-LoopSampler can fail to explore some regions in which a few successful conformations could have been found. However, these regions may suddenly get “unlocked” after a closed loop is finally sampled. One successful conformation is necessary and sufficient to set the score of the cell leading to sampling in that region back to its maximum. This is the case for the region



**Fig. S16.** 2D projections of conformations sampled using MoMA-LoopSampler at different learning rates for a loop in the streptavidin protein, from eight starting X-ray protein structures. The first dimension,  $d_1$  (x-axis), is the distance ( $\text{\AA}$ ) between an atom located in the middle of the loop and a fixed atom in the protein. The second dimension,  $d_2$  (y-axis), is the angle (degrees) formed by three atoms: an atom at approximately one quarter, one half, and three quarters of the way down of the loop. The conformations from the crystallographic structures are shown in black. For each system, the loop with the lowest energy was identified and each conformation was then colored according to the difference between its energy and this lowest energy.

of the conformational space that is projected in the bottom right-hand corner of the heatmaps shown in Figure S15(a) at a *very high* learning rate. This phenomenon is also observed for other systems in Figure S15, albeit on smaller regions of the projection plot. This is a warning that learning is already too greedy. In theory, learning should only stop exploring regions when enough exploration has been made and the probability to find a successful loop conformation in that region is negligible. In practice, this can never be determined with certainty, as long as the full region has not been explored. For  $H_{\text{RL}}$  to have the desired behavior, a trade-off has to be found in order to enable faster exploration while not risking to lose large and/or relevant areas of the conformational space.

These results also show that the distribution of sampled conformations may be different depending on the learning rate and the running time: beyond the observation that the conformational space obtained is more continuous, the density observed in the heatmaps evolves between the different conditions. The samples appear to be more uniformly distributed when  $H_{\text{RL}}$  is involved. Another interesting observation is that, for all these four systems, some

consequences of learning become visible very early in the sampling process, within the first 10 minutes.

## S2.5 Energy landscape of a multi-state loop at different learning rates

The protocol for modeling and visualizing the energy landscape of streptavidin (Section 3.2.2) was repeated using different learning rates. For each scaffold, MoMA-LoopSampler in basic mode and in four different  $H_{\text{RL}}$  modes was employed to sample the same number of conformations as sampled using brute force. The sampled conformations were relaxed and projected in 2D. The resulting landscapes are visible in Figure S16.

The landscapes look very similar for all levels of learning, showing that  $H_{\text{RL}}$  parameterization overall preserves sampling diversity for this loop system. Only slight differences are observed when applying  $H_{\text{RL}}$  with a *very high* learning rate. In the landscape built from scaffold 2F01 A, the basin around the “open” conformation is not as clearly apparent as for less greedy  $H_{\text{RL}}$  levels.

## References

- Cahill, S., Cahill, M., and Cahill, K. (2003). On the kinematics of protein folding. *J. Comput. Chem.*, **24**(11), 1364–1370.
- Canutescu, A. A. and Dunbrack, R. L. (2003). Cyclic coordinate descent: A robotics algorithm for protein loop closure. *Protein Sci.*, **12**(5), 963–972.
- Case, D., Betz, R., Cerutti, D., T.E. Cheatham I., Darden, T., Duke, R., Giese, T., Gohlke, H., Goetz, A., Homeyer, N., Izadi, S., Janowski, P., Kaus, J., A. Kovalenko, T. L., LeGrand, S., Li, P., C.Lin, Luchko, T., Luo, R., Madej, B., Mermelstein, D., Merz, K., Monard, G., Nguyen, H., Nguyen, H., I.Omelyan, Onufriev, A., Roe, D., Roitberg, A., Sagui, C., Simmerling, C., Botello-Smith, W., Swails, J., Walker, R., Wang, J., Wolf, R., Wu, X., Xiao, L., and Kollman, P. (2016). Amber 2016, University of California, San Francisco.
- Case, D. A., Cheatham, T. E., Darden, T., Gohlke, H., Luo, R., Merz, K. M., Onufriev, A., Simmerling, C., Wang, B., and Woods, R. J. (2005). The Amber biomolecular simulation programs. *J. Comput. Chem.*, **26**(16), 1668–1688.
- Chys, P. and Chacón, P. (2013). Random Coordinate Descent with Spinor-matrices and Geometric Filters for Efficient Loop Closure. *J. Chem. Theory Comput.*, **9**(3), 1821–1829.
- Cortés, J., Siméon, T., Remaud-Siméon, M., and Tran, V. (2004). Geometric algorithms for the conformational analysis of long protein loops. *J. Comput. Chem.*, **25**(7), 956–967.
- Coutsias, E., Seok, C., Jacobson, M., and Dill, K. (2004). A kinematic view of loop closure. *J. Comput. Chem.*, **25**(4), 510–528.
- DePristo, M. A., de Bakker, P. I. W., Lovell, S. C., and Blundell, T. L. (2003). Ab initio construction of polypeptide fragments: Efficient generation of accurate, representative ensembles. *Proteins*, **51**(1), 41–55.
- Dinner, A. R. (2000). Local deformations of polymers with nonplanar rigid main-chain internal coordinates. *J. Comput. Chem.*, **21**(13), 1132–1144.
- Fox, N. K., Brenner, S. E., and Chandonia, J.-M. (2014). SCOPe: Structural Classification of Proteins—extended, integrating SCOP and ASTRAL data and classification of new structures. *Nucleic Acids Res.*, **42**(D1), D304–D309.
- Harder, T., Boomsma, W., Paluszewski, M., Frelsen, J., Johansson, K. E., and Hamelryck, T. (2010). Beyond rotamers: a generative, probabilistic model of side chains in proteins. *BMC Bioinf.*, **11**(1), 306.
- Jacobson, M. P., Pincus, D. L., Rapp, C. S., Day, T. J. F., Honig, B., Shaw, D. E., and Friesner, R. A. (2004). A hierarchical approach to all-atom protein loop prediction. *Proteins*, **55**(2), 351–367.
- Kabsch, W. and Sander, C. (1983). Dictionary of protein secondary structure: Pattern recognition of hydrogen-bonded and geometrical features. *Biopolymers*, **22**(12), 2577–2637.
- López-Blanco, J. R., Canosa-Valls, A. J., Li, Y., and Chacón, P. (2016). RCD+: Fast loop modeling server. *Nucleic Acids Res.*, **44**(W1), W395–W400.
- Manocha, D. and Canny, J. F. (1994). Efficient inverse kinematics for general 6r manipulators. *IEEE Trans. Robot. Autom.*, **10**(5), 648–657.
- Shenkin, P. S., Yarmush, D. L., Fine, R. M., Wang, H. J., and Levinthal, C. (1987). Predicting antibody hypervariable loop conformation. I. Ensembles of random conformations for ringlike structures. *Biopolymers*, **26**(12), 2053–2085.
- Soto, C. S., Fasnacht, M., Zhu, J., Forrest, L., and Honig, B. (2008). Loop modeling: Sampling, filtering, and scoring. *Proteins*, **70**(3), 834–843.
- Tang, K., Zhang, J., and Liang, J. (2014). Fast Protein Loop Sampling and Structure Prediction Using Distance-Guided Sequential Chain-Growth Monte Carlo Method. *PLoS Comput. Biol.*, **10**(4), e1003539.
- Xiang, Z. (2006). Advances in Homology Protein Structure Modeling. *Curr. Protein Pept. Sci.*, **7**(3), 217–227.
- Xiang, Z., Soto, C. S., and Honig, B. (2002). Evaluating conformational free energies: the colony energy and its application to the problem of loop prediction. *Proc. Natl. Acad. Sci. U. S. A.*, **99**(11), 7432–7437.
- Zhao, S., Zhu, K., Li, J., and Friesner, R. A. (2011). Progress in Super Long Loop Prediction. *Proteins*, **79**(10), 2920–2935.

## Appendix: heatmaps comparing tripeptide projections

Following are the figures mentioned in Section S1.7 comparing the different tripeptide state projections.

**Fig. S17.** Comparison of projections. Several different projections were tested for  $H_{RL}$ . The following heatmaps show the ability of each projection to distribute tripeptides and to regroup states that succeed in creating a closed loop together. The results are presented for 6 different systems. For each system, the heatmaps of the learning trees corresponding to the different loop plans are shown. A loop plan is designated by the position of the last tripeptide used to close the loop (IK position). The first of the two heatmaps for a given learning tree gives the distribution of tripeptides in the top level cells of the root tree. The second shows the success probability of the leaf, meaning the number of successful loop conformations that start with a state from this leaf divided by the theoretical number of tripeptide combinations starting with a state from this leaf.




Fig. S17((a)): Position






Fig. S17((b)): Euler angles




Fig. S17((c)): Euler angles and length

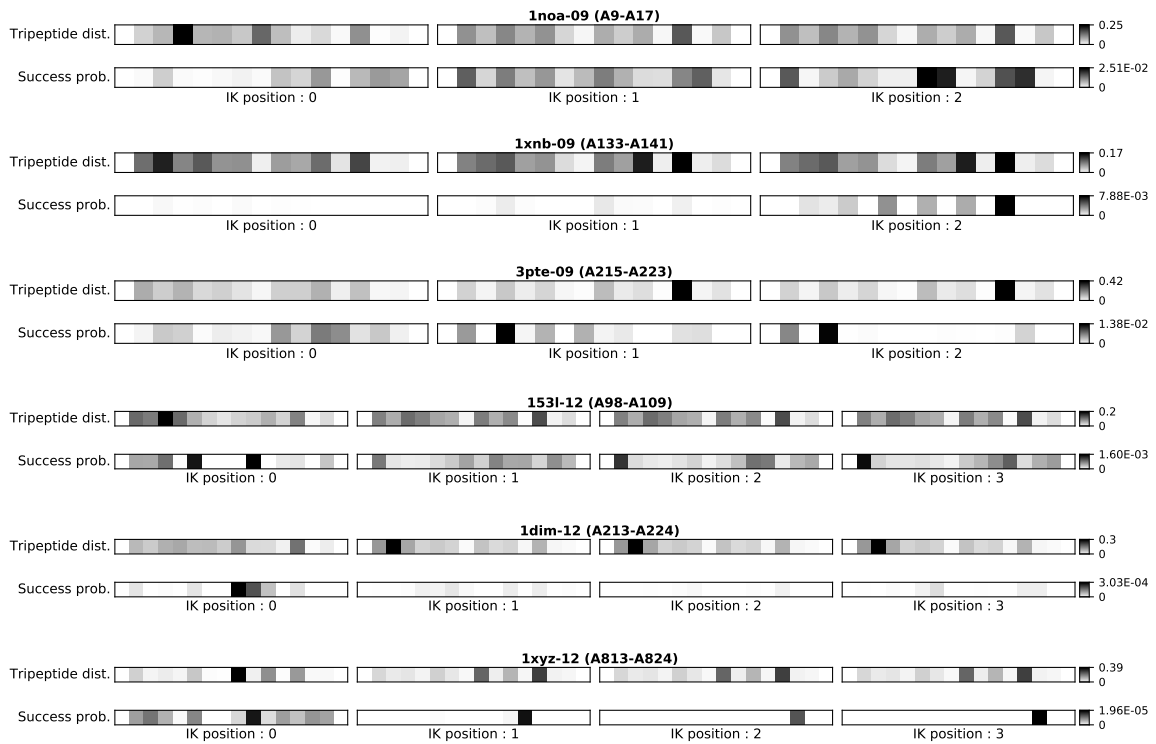


Fig. S17((d)): Quaternion




Fig. S17((e)): Quaternion and length




Fig. S17((f)): Axis-angle




Fig. S17((g)): Axis-angle and length