

Interactive Protein Manipulation

Oliver Kreylos*

Bernd Hamann*

Nelson L. Max[†]

Silvia N. Crivelli[‡]

E. Wes Bethel[‡]

1 Introduction

One of the grand challenges in computational biology is the prediction of the three-dimensional structure of a protein from its chemical makeup alone. A protein's *primary structure*, i.e., its amino acid sequence, is directly encoded in its DNA sequence; this, however, is a purely one-dimensional structure that does not directly encode a three-dimensional shape. It is commonly believed that the "natural" shape of a protein is the one corresponding to the global minimum of its internal energy; and the *protein folding problem* has been treated as an optimization problem in recent years, to some success. As with any optimization problem, it is important to start solving it from a "good" set of initial configurations that allow the optimization code to search the complete optimization space for a minimum. Our work focuses on providing an interactive, visual tool to rapidly create many initial configurations for a given amino acid sequence, which are then fed into an optimization algorithm.

2 Protein Structure Hierarchy

Proteins considered by our tool have three levels of structure [1]:

Primary Structure A protein's primary structure is its amino acid sequence. It is directly encoded in a protein's gene (each group of three bases defines one amino acid). The chemical makeup of a (simple) protein is a single chain of amino acid residues connected by peptide bonds, see Figure 1.




Figure 1: Part of the primary structure of a protein

Secondary Structure Adjacent amino acid residues inside a protein can interact with each other to form substructures: α -helices, see Figure 2, and β -strands. Inside an α -helix, each residue forms hydrogen bonds with two other residues, accounting for their rigidity. For each amino acid type, the probabilities of it forming either one of these are known, and neural networks are used successfully to predict secondary structure occurrences from amino acid sequences [2].




Figure 2: An α -helix. Hydrogen bonds are depicted as dashed yellow lines

Tertiary Structure A protein's three-dimensional structure is formed by amino acid residues from distant parts of the chain forming bonds with each other. β -strands, not very rigid by themselves, hydrogen-bond with each other to form stable β -sheets, see Figure 3, whereas α -helices cluster to each other to hide their hydrophobic amino acids from the surrounding watery solution. Prediction of tertiary structure is still an unsolved problem.




Figure 3: Two β -sheets. Hydrogen bonds are depicted as dashed yellow lines. Top row: Anti-parallel sheet (left) and cartoon rendering (right). Bottom row: Parallel sheet (left) and cartoon rendering (right)

Quaternary Structure Many proteins, e.g., hemoglobin, contain more than one amino acid chain. For those, quaternary structure describes how separate chains interact with each other to form an overall shape. Our tool, and the used optimization code, do not consider non-simple proteins.

3 Protein Folding as Optimization Problem

An optimization problem is defined by its configuration space and target function. In the case of protein folding, the configuration space is the space of all possible three-dimensional configurations of a given protein, and the target function is its internal energy. Although simple proteins are single molecules, they are of surprising flexibility. Each

*Center for Image Processing and Integrated Computing (CIPIC), Department of Computer Science, University of California, Davis; {kreylos,hamann}@cs.ucdavis.edu

[†]Lawrence Livermore National Laboratory; max2@llnl.gov

[‡]Lawrence Berkeley National Laboratory;
{sncrivelli,ewbethel}@lbl.gov

amino acid residue has two rotational degrees of freedom, the two *dihedral angles* ϕ and ψ , see Figure 4. Typical proteins contain between hundreds and thousands of residues, making the optimization space high-dimensional. Furthermore, the internal energy function has local minima in abundance. Together, these two facts make protein folding a very difficult optimization problem.




Figure 4: Rotational degrees of freedom along a residue chain. Adjacent residues are separated by dashed lines; side chains are denoted by R

4 Creating Initial Configurations

In order to exhaustively search the space of all possible configurations of a given protein for its global energy minimum, the optimization algorithm needs to be provided with dozens of initial configurations spread out over the entire space. In the past, initial configurations were designed by the computational biologists, and then created by a separate constrained optimization system forming them out of an unfolded residue chain. This process was time-consuming and non-intuitive, and many possible initial configurations were overlooked in the process. We have attempted to overcome this problem by creating a visual tool to directly manipulate protein structures. The idea is to let biologists assemble proteins as if using plastic stick-and-ball models. Our goal was to keep proteins intact during manipulation, by using their intrinsic degrees of freedom to achieve the movement intended by a user. If a user selects a secondary structure, say a β -strand, and moves it towards another β -strand to form a β -sheet, then the amorphous *coil regions* between those two structures will bend and twist to allow the motion.

5 Inverse Kinematics for Molecular Modelling

Since proteins are inherently flexible enough to allow a wide range of motion, the main problem became to translate a user's 6-degree-of-freedom motions into changes of a chain segment's dihedral angles ϕ_i and ψ_i . This problem, *inverse kinematics* (IK) [3], has been studied in the field of robotics, where it is used to translate intended motion of a robot's "hand" into changes in joint parameters along a robot's "arm." What makes this application of IK more difficult is scale: A robot assembly typically has up to a dozen joints, whereas we encounter linked assemblies of 40–80 joints in medium-sized proteins. Nevertheless, IK turned out to be the method of choice for natural interaction with large molecules.

6 Using the Visual Modelling Tool

A typical modelling session starts out by reading a prediction file created by one of many publicly available secondary

structure prediction servers [2]. These files contain the sequence of amino acid residues, and, for each residue, an indicator whether that residue is part of an α -helix, a β -strand, or an amorphous coil region. From this information, and a set of standard amino acid structure files, our program creates a "pre-configuration" consisting of fully formed secondary structures, but no tertiary structure, see Figure 5.



Figure 5: Pre-configuration for protein 1PGX

After that, a user proceeds to align adjacent β -strands to form initial β -sheets. Later, the order of β -strands is permuted to quickly create several dozen configurations. During protein assembly, users ignore the exact alignment of α -helices and coil regions since the subsequent optimization process handles them well. Figure 6 shows two initial configurations for the same protein.




Figure 6: Two initial configurations for 1PGX

7 Conclusions and Future Work

The tool described above has been deployed for use by computational biologists at the Lawrence Berkeley National Laboratory, the University of California at Berkeley, and the University of Colorado at Boulder. This group of scientists is currently competing in the international CASP5 protein structure prediction competition, and our tool is being used to create dozens of initial protein structures every day. With the new tool, the group has been able to attack proteins of sizes that were not manageable before, and the high quality of the created configurations has exposed new behaviour in the existing optimization algorithm. In the near future, we plan to integrate the tool with the optimization code, to be used as a front-end for monitoring and steering of the massively parallel optimization code itself.

References

- [1] Lehninger, A.L., et al., Principles of Biochemistry, 2nd ed. (1993), Worth, New York
- [2] McGuffin, L.J., et al., *PSIPRED: a Protein Structure Prediction Server*, <http://www.psipred.net>
- [3] Welman, C., *Inverse Kinematics and Geometric Constraints for Articulated Figure Manipulation*, Master's Thesis, Simon Fraser University, 1993