



Published in final edited form as:

Commun ACM. 2019 October ; 62(10): 76–84. doi:10.1145/3338124.

Protein Design by Provable Algorithms

MARK A. HALLEN¹, BRUCE R. DONALD²

MARK A. HALLEN: mhallen@ttic.edu; BRUCE R. DONALD: cacm19@cs.duke.edu

¹Research assistant professor at the Toyota Technological Institute at Chicago, IL, USA.

²James B. Duke Professor of Computer Science at Duke University, as well as a professor of chemistry and biochemistry in the Duke University Medical Center, Durham, NC, USA.

Graphical Abstract



PROTEINS ARE A class of large molecules that are involved in the vast majority of biological functions, from cell replication to photosynthesis to cognition. The chemical structure of proteins is very systematic⁵— they consist of a chain of atoms known as the *backbone*, which consists of three-atom (nitrogen-carbon-carbon) repeats known as *residues*, each of which features a *sidechain* of atoms emanating from the first carbon. In general, there are 20 different options for sidechains, and a residue with a particular type of sidechain is known as an *amino acid* (so there are also 20 different amino acid types). For billions of years, the process of evolution has optimized the sequence of amino acids that make up naturally occurring proteins to suit the needs of the organisms that make them. So we ask: Can we use

The authors are founders of Gavilán Biodesign, Inc.

computation to design non-naturally occurring proteins that suit our biomedical and industrial needs?

This question is a combinatorial optimization problem, because the output of a protein design computation is a sequence of amino acids. Due to the vast diversity of naturally occurring proteins, it is possible—and very useful—to begin a protein design computation with a naturally occurring protein and then to modify it to achieve the desired function. In this article, we focus on protein design algorithms that perform this optimization using detailed modeling of the 3D structure of the protein.^{5,8} Thus, they will begin with a *starting structure*, a 3D structure of a (typically naturally occurring) protein we wish to modify.

To illustrate this concept, imagine we wish to perform a simple example modification to a protein to make it more stable, so it can still function at higher temperatures. In this case, we must minimize the protein's energy with respect to its sequence of amino acids. In structure-based design, energy is typically estimated using *energy functions*, which map the 3D geometry of a molecule to its energy, so the optimization becomes slightly more complex: we minimize the energy with respect to both the *sequence* (of amino acids) and the *conformation* (the 3D geometry of the protein, that is, the locations of all its atoms in space). While the sequence is a discrete variable, the conformation is a continuous one because coordinates in \mathbb{R}^3 are continuous variables. There are some physical (for example, holonomic) constraints on how atoms can move relative to each other, and thus the conformational space can be represented most effectively using internal coordinates, resulting in the joint angle *configuration space* familiar in robotics and motion planning in computer science. Nevertheless, the full conformational space of a protein is too vast to search exhaustively, especially with a simultaneous search over sequence space.

Computational structure-based protein design arose as a response to this difficulty. Its initial goal was to overcome certain combinatorial obstructions to designing with a discretized version of the conformational space. Hence, in order to study protein design, it is first necessary to understand the structure of this simpler (but still non-trivial) discrete optimization problem. To this end, we first give a flavor for the issues that arise in discrete optimization. We examine a very special case—the case of discrete *rotamers* and a simple Markov random field (MRF)-like energy function. Next, we carefully define a mixed discrete-continuous optimization problem that gives sidechains and then backbones continuous flexibility within a conformational voxel. Then, we present algorithms that provably approximate partition functions over many states, analogously to well-known statistical inference and machine learning computations, and that exploit improved, more realistic energy functions.

It is also often useful in protein design to optimize objectives other than simply the energy of a protein. However, many useful design objectives can still often be posed in terms of the energies of multiple *biophysical states* of a protein—for example, states where it is bound to particular other molecules. Thus, the problem of *multistate design*, which we will formalize, is appropriate for tasks like optimizing the binding of one protein to another molecule, or even specific binding to a second molecule while excluding binding to a third molecule.

Together with some novel types of objective functions, multistate design is a more general tool to optimize the desired function of a protein with respect to sequence.

We will highlight provable computational techniques employed for each of these problems. These include techniques from combinatorial optimization, constraint satisfaction, machine learning, and other areas. For the relatively simple protein design problems addressed in this article, we find that algorithms with a beautiful mathematical structure suffice. This permits us to illustrate by specific examples the situation confronting practical protein designers in academic or biopharmaceutical laboratories. Throughout the article, we review algorithms of intrinsic mathematical interest and with the potential for high impact on the engineering of new molecular therapies for human disease.

In addition to this review of core algorithmic work, we will briefly discuss methods to accelerate protein design computations using GPU hardware, as well as some cases in which proteins designed using provable algorithms have performed well in experimental tests. Protein design with provable algorithms has already had success in the design of novel enzymes and proteins with therapeutic applications. As the field matures and significant errors are eliminated from more steps of the protein design process, we expect to see even more successes from this promising technique.

The Pairwise Discrete Model

Problem definition.

We will now formalize this problem of stabilizing a protein using some simplifying assumptions, which will yield the most commonly used mathematical formulation of the protein design problem. We will present several algorithms to attack this problem as well as enhancements to the formulation with more sophisticated objectives and/or modeling assumptions.

Changing the sequence of a protein—that is, *mutating* it—does not alter the chemical structure of its backbone,^a and the largest conformational changes are typically found in sidechains near the site of the mutations (we will designate these residues as *flexible*, that is, we will consider it necessary to search their conformational space). Thus, we will assume the backbone conformation (and possibly some of the sidechain conformations for residues farther from the site of mutations) is the same as in the starting structure. Moreover, analyses of sidechain conformational space have found sidechain conformations for each amino-acid type to occur in clusters known as *rotamers*. We will refer to the modal sidechain conformation in each cluster as an *ideal rotamer*. Then, for the sidechains with respect to whose amino-acid type and conformation we wish to optimize, we will assume the sidechain conformations will be ideal rotamers, meaning we need only optimize over a discrete set of (sequence, conformation) pairs in which each residue must be assigned an amino-acid type and one of the ideal rotamers for that amino-acid type.

^aActually, there is one amino acid—proline—whose sidechain bonds to the backbone in two places, but it does not alter the repeating nitrogen-carbon-carbon pattern of backbone atoms.

Let \mathbf{r} be a list of rotamers (which may be of any amino-acid type) for the residues that we are treating as flexible and/or mutable. If we use only ideal rotamers, \mathbf{r} fully defines a sequence and conformation for the protein, so our energy function gives us a well-defined energy $E(\mathbf{r})$, and our optimization problem becomes simply finding $\arg \min E(\mathbf{r})$. However, one more simplifying assumption is often applied: that we are using a *pairwise energy function*, which is a sum of terms that each depend on the amino-acid types and conformations of at most two residues. In this case, we can expand

$$E(\mathbf{r}) = \sum_i E(i_{\mathbf{r}}) + \sum_{j < i} E(i_{\mathbf{r}}, j_{\mathbf{r}}) \quad (1)$$

where i and j are residues, and $i_{\mathbf{r}}$ is the rotamer that \mathbf{r} assigns to residue i (we place the residue position in the subscript, following the convention of the field). The pairwise energy function gives us a well-defined 1-body energy $E(i_{\mathbf{r}})$ and 2-body energy $E(i_{\mathbf{r}}, j_{\mathbf{r}})$ for any rotamers $i_{\mathbf{r}}$ and $j_{\mathbf{r}}$, and indeed these energies can be precomputed (generating an *energy matrix*) before the process of optimization begins, allowing the optimization to simply operate on the energy matrix rather than calling the energy function directly. Thus, we can formalize the protein design problem in this simple pairwise discrete model as

$$\arg \min_{\mathbf{r}} \left(\sum_i E(i_{\mathbf{r}}) + \sum_{j < i} E(i_{\mathbf{r}}, j_{\mathbf{r}}) \right). \quad (2)$$

We will refer to the solution of Eq. (2) as the *global minimum-energy conformation*, or GMEC. This problem is equivalent to finding the maximum-likelihood solution for a Markov random field with only pairwise couplings.^{5,7}

Finding the GMEC is unfortunately NP-hard,²⁷ even to approximate.¹ But much algorithmic and development work has attacked it, and most biophysically relevant cases of the problem can be solved efficiently in practice with provable guarantees of accuracy. We now review some of this work.

Work on this problem using heuristic protocols such as simulated annealing, Monte Carlo simulation, and genetic algorithms is surveyed comprehensively in Donald⁵ and Gainza et al.⁸ Moreover, Monte Carlo simulation in this context is often not ergodic, rendering it less reliable than mathematical methods like Monte Carlo integration that can obtain accurate error bars based on the variance of an ergodic simulation. As a result, estimates of the GMEC even from a highly optimized Monte Carlo/simulated annealing protocol exhibit empirically significant deviations from the true optimum.³¹ Similar empirical deviations have been found in several other areas of structural biology requiring global minimizers, as reviewed in Gainza et al.⁸ For these reasons, in this article we concentrate on provable algorithms that may be of greater interest to computer scientists.

Approaches to the problem:

The classic DEE/A framework.* The first breakthrough toward solving Eq. (2) was the DEE algorithm⁴ (with refinements due to Goldstein), which eliminates rotamers that cannot be part of the GMEC. It works by comparing two rotamers i_r and i_t for the same residue. i_r can be pruned if every conformation \mathbf{r} containing i_r is higher in energy than the corresponding conformation in which i_r has been replaced by i_t , that is, if

$$\min_{\mathbf{r}} \left(E(i_r) - E(i_t) + \sum_{j \neq i} E(i_r, j_{\mathbf{r}}) - E(i_t, j_{\mathbf{r}}) \right) > 0. \quad (3)$$

Evaluating Eq. (3) is as difficult as finding the GMEC directly. But the sum of minima is always a lower bound for the minimum of a sum, so we obtain the following sufficient condition for Eq. (3), which can be evaluated in time linear in the number of residues:

$$E(i_r) - E(i_t) + \sum_{j \neq i} \min_s (E(i_r, j_s) - E(i_t, j_s)) > 0. \quad (4)$$

We call Eq. (4) the DEE criterion. By evaluating it for each residue i and each pair of rotamers i_r and i_t that are available at i , we can greatly prune the space of rotamers that may be part of the GMEC. This pruning step is polynomial-time.⁵ Thus, the combinatorial bottleneck must occur later, in the enumeration step.

DEE is an efficient algorithm, but it still may leave multiple possible rotamers for some or all of the residues. This problem has been solved by deploying the A* algorithm from artificial intelligence to find the GMEC using only the rotamers remaining, that is, using DEE/A*.²² Briefly, the A* algorithm in this context builds a priority queue of nodes that represent a partially defined conformation \mathbf{q} , which consists of rotamer assignments for only a subset $\mathcal{S}(\mathbf{q})$ of the residues. The score of a node is a lower bound on the energy of any conformation containing all the rotamers in \mathbf{q} (that is, $\min_{i_r = i_{\mathbf{q}} \forall i \in \mathcal{S}(\mathbf{q})} E(\mathbf{r})$). We repeatedly

extract the lowest-scoring node from the queue and expand it by creating nodes for which one more residue has a defined rotamer. Eventually the lowest-scoring node will be a fully defined conformation. Since all conformations in other nodes must have higher energies (based on the nodes' lower bounds), this fully defined conformation must be the GMEC.

This shows that it is possible to find the GMEC with guaranteed accuracy, and indeed to do so significantly faster (in practice) than exhaustive enumeration of conformations. We will now discuss even more sophisticated and efficient algorithms for this problem.

Algorithms from weighted constraint-satisfaction problems.

One source of such improved algorithms is from the field of weighted constraint-satisfaction problems (WCSPs), of which the pairwise discrete protein design problem (Eq. 2) can be seen as a special case. To use these techniques, the energy matrix is encoded as a cost function network (CFN), which includes the same type of 1- and 2-body terms as an energy

matrix from protein design.³³ The most efficient provably accurate algorithms for WCSPs perform a tree search like A*, but with much more refined heuristics to guide the search (including both upper and lower bounds). They also usually employ a depth-first branch-and-bound approach rather than a best-first search like A*. As a result, far less memory is required in practice. A large set of empirical benchmarks in Traoré et al.³² showed the Toulbar package for WCSPs significantly improved the state-of-the-art efficiency for protein design in the discrete pairwise model. Moreover, this increase in efficiency allowed direct comparison of the true GMEC (computed by WCSP algorithms) to estimated GMECs from the popular but non-provable simulated annealing algorithm, as implemented in the Rosetta software, for very large protein design problems. Significant discrepancies were found,³¹ and indeed the error in simulated annealing's estimates increased with protein size. This highlights the need for algorithms with provable guarantees for protein design.

A related and also provable approach is to reduce Eq. (2) to an integer linear programming problem.²¹

Algorithms making sparsity assumptions.

Although protein design as expressed in 2Eq. (2) is NP-hard even to approximate,¹ it is possible to add additional assumptions that make it solvable in polynomial time. Suppose we assume that some pairs of residues have uniformly zero interaction energies, such that the graph whose nodes are residues and whose edges denote residue pairs with nonzero 2-body energies is sparse, making it a *sparse residue interaction graph* (SPRIG, see Figure 1). The TreePack algorithm³⁶ can find the GMEC in polynomial time when the SPRIG has constant tree width. Moreover, the BWM* algorithm can find the GMEC in polynomial time and also efficiently enumerate the k best conformations in gap-free order when the SPRIG has constant branch width (where k is requested by the user).

Improved Models

The pairwise discrete model (Eq. 2) captures the most essential aspects of computational protein design, but it falls short for many practical applications. Despite the prevalence of rotameric conformations of protein sidechains, real proteins do have significant continuous flexibility in the neighborhood of each ideal rotamer. Backbone motions due to mutations are often non-negligible as well. Moreover, the energy model in Eq. (2) falls short in two ways: the most accurate energy functions are not explicitly pairwise, and the behavior of a protein is actually determined by its free energy—a quantity based on the distribution of its conformations' energies—rather than on the single minimum-energy conformation. Finally, as mentioned earlier, it is often useful to have a more sophisticated objective function than simply minimizing the energy of a single biophysical state of a protein. Here, we review algorithms to address these five shortcomings (vide supra) of the discrete pairwise model of protein design.

Continuous flexibility:

Defining the problem.—The problem of continuously flexible protein design differs from Eq. (2) in that each rotamer is no longer a single conformation of its residue. Rather, each

rotamer is a set of conformations, which we can model as a *voxel* in the form of bounds on each of several continuous internal coordinates. Sidechain flexibility in proteins occurs mainly in the form of changes in dihedral angles, and thus the conformation space of a protein can be modeled accurately as a union of voxels in dihedral angle space. For example, in Georgiev et al.,¹² each voxel is centered at an ideal rotamer, and allows up to $\pm 9^\circ$ of flexibility in each dihedral angle in either direction from the ideal rotamer's dihedral angle. The problem is then to find the list of rotamer assignments \mathbf{r} whose voxel contains the lowest-energy conformation—the *minGMEC*.

This problem has both discrete and continuous components, much like AI planning, where there are discrete steps like STRIPS or TWEAK and continuous steps like motion planning. Like robust optimization, its aim is to prevent error due to insufficiently fine sampling of conformational space—we wish to avoid eliminating a rotamer merely because its ideal rotameric conformation appears unfavorable, since a small continuous adjustment may turn out to make it optimal. Indeed, it is relatively common for ideal rotamers to be physically infeasible due to a clash (a pair of atoms too close to each other), but for a small continuous adjustment to suffice to find a favorable conformation^{9,10,12} (as illustrated in Figure 2). Moreover, the optimal sequence is often significantly different, and more biophysically realistic, when continuous flexibility is taken into account than when it is neglected.^{9,10}

Notably, no benefit in design is obtained by simply performing a discrete optimization and then continuously minimizing the energy of the discrete GMEC post hoc: such minimization does not change the optimal sequence that is selected. Rather, to obtain the full benefits of continuous flexibility, one must perform *minimization-aware* design that finds the minGMEC with guarantees of accuracy by taking continuous flexibility into account from the beginning. There are two general approaches to minimization awareness.

Adapting discrete algorithms to bound the continuous problem.—Algorithms for discrete protein design can be adapted to be minimization-aware by having them prune using bounds on conformational energies rather than using conformational energies directly. If a list of voxels \mathbf{r} represents a region in conformational space rather than a single conformation, then its energy (Eq. 1) may not be well defined per se, but a lower bound on its energy can be expressed in the form of Eq. (1)), simply by minimizing each of the 1- and 2-body energy terms over the voxel. Discrete protein design algorithms can then be used to enumerate conformations in order of lower bound. Once these conformations have been continuously minimized, additional conformations can be pruned based on their lower bounds as well, allowing provable computation of the minGMEC. This approach has been developed effectively by Gainza et al.⁹ and Georgiev,¹² who adapt the entire DEE/A* framework to be minimization-aware.

Other discrete algorithms also fit well into the framework of minimization-awareness based on bounds. For example, both belief propagation (BP) and the self-consistent mean field method (SCMF) are usually employed to estimate a GMEC, with no proofs of closeness to the optimal solution. However, SCMF can generate a provably correct lower bound on the GMEC energy, while tree-weighted belief propagation can generate a provably correct upper bound. Thus, by operating on bounds, both algorithms become provable. This contrasts with

the exact rigid energies used with methods from weighted constraint satisfaction and integer linear programming.

Reducing the continuous problem to a discrete one.—A more recent approach to minimization-aware protein design is based on machine learning and reducing the continuous protein design problem to a discrete one, without significantly compromising accuracy. Although the energy of a voxel \mathbf{r} is not explicitly in the form required for discrete protein design algorithms (Eq. 1), there is a well-defined energy $E(\mathbf{r})$ (generally the continuously minimized energy) that we want to optimize, and we can fit it to the form of Eq. (1) using machine learning. This approach is very efficient as implemented in the LUTE algorithm,¹⁷ and also accommodates other improvements in biophysical modeling^b because the user can choose the function $E(\mathbf{r})$ that is taken as input. The implementation of LUTE described in Hallen et al.¹⁷ also incorporates some elements of the bound-based approach to continuous flexibility, because it uses iMinDEE,⁹ a minimization-aware version of DEE, as a preprocessing step, resulting in a critical improvement in its training and test error.

Backbone flexibility.—Continuous sidechain flexibility handles discrepancies between ideal rotamers and the actual sidechain conformation. But an additional type of continuous flexibility—backbone flexibility—is necessary to handle discrepancies between the starting structure's backbone conformation (experimentally observed for the original sequence) and the backbone conformation that is optimal for each mutant sequence. Like continuous sidechain flexibility, backbone flexibility can be handled using voxels, which can bound the backbone's continuous internal coordinates in a neighborhood around the starting structure's backbone. The main difference is that the choice of internal coordinates is less straightforward—one must find coordinates that adequately represent the biophysically important backbone flexibility in the vicinity of the mutations without obtaining an intractably large conformational space to search. These are properties that are satisfied by sidechain dihedrals, whose locality makes them the obvious choice of internal coordinates for sidechains. But they are not satisfied by the standard backbone dihedrals ϕ and ψ , because local changes in the backbone dihedrals will propagate throughout the protein, disrupting its large-scale structure unless the changes are very small.

The DEEPer algorithm¹⁸ addresses this problem by using only backbone motions based on experimental observations, such as the backrub motion observed in crystallographic alternates. The CATS algorithm¹⁶ allows a larger degree of continuous motion by constructing a new type of backbone internal coordinates that can model the local motion of a contiguous segment of the protein backbone in all biophysically feasible directions (Figure 3). Both algorithms can be used in conjunction with continuous side-chain flexibility modeling and design.

Multistate design.

Defining the multistate problem.—Protein design software is already quite effective at stabilizing proteins, but we must pursue other objectives if it is truly to meet the full range of

^bSuch as non-pairwise energy functions, including those modeling solvation effects, quantum chemistry, and continuous entropy.

biomedical and bioengineering needs for modified proteins. Most of the important objectives involve binding—for example, binding to a protein in the human body that is involved with disease, and also not binding to other, possibly similar, proteins that are essential to normal functioning of the body. These objectives can be modeled in terms of multiple *biophysical states*—states in which the protein being designed is unbound, bound to a particular desired target, or bound to a particular undesired target, and so on. Each state a has an energy $E_a(s)$, which we can approximate as the energy of the lowest energy conformation for the state (as a function of sequence s). We want favored states to be low in energy and unfavored states to be high in energy, since this will cause the protein to adopt the favored states in preference to the unfavored ones.

Thus, following Hallen and Donald,¹⁵ we can pose the problem of multistate design as a kind of linear programming on protein state energies. We will define linear multistate energies (LMEs), which are functions of sequence s , in the form

$$c_0 + \sum_a c_a E_a(s), \quad (5)$$

where the coefficients c are chosen by the user. For example, to make an LME representing the binding energy between the protein we are designing and another molecule, we would set $c_b = 1$ and $c_u = -1$ where b is the bound state and u is the unbound state. We then wish to minimize not a single state's energy, but an LME, with respect to sequence. We may also wish to constrain other LMEs to have values above or below a user-specified threshold—for example, we may wish to keep the binding energy to an undesired target higher than the observed binding energy of the unmutated protein to that undesired target.

Algorithms for multistate design.—The formulation in the previous section comes from Hallen and Donald,¹⁵ who also present the first provable algorithm to solve this problem without exhaustive enumeration of sequences. This algorithm, COMETS, builds an A* tree with nodes representing partial sequences. Conformational search is handled with a combination of bounding techniques and construction of a “tree within a tree” for each promising sequence. The main tree is thus responsible for sequence search, while the inner trees each correspond to a single node of the main tree and perform conformational search for the sequence corresponding to that node.

DEE itself has also been adapted for multistate design. Specifically, within each sequence and biophysical state, multistate design (as defined previously) is simply computing a GMEC, and as a result it is provably accurate to perform DEE pruning within each biophysical state as long as only competitor and candidate rotamers of the same amino-acid type are considered.³⁷ This technique is known as *type-dependent DEE*. The multistate design problem has also been addressed using belief propagation.⁷

As in the case of continuous flexibility, machine learning has yielded a novel and very promising technique for multistate design. The *cluster expansion* technique¹⁴ calculates energies for a training set of sequences (for each state) and then learns an energy function

that is a sum of terms dependent only on 1 or a few residues' amino acid types. In this formulation, multistate design becomes mathematically equivalent to discrete single-state design, although combinatorially easier because there are fewer amino acid types than possible rotamers. This technique has yielded designer peptides with high selectivity for their desired target in experimental tests.¹³

Finally, other formulations of multistate design besides that discussed earlier have been used quite fruitfully. The paradigm of meta-multistate design,³ which accounts for protein dynamics, has yielded designed proteins known as DANCERS (Dynamic And Native Conformational ExchangeRs), which not only exchange between specified conformational states, but do so on the timescale of milliseconds.

Improved energy modeling.

We have so far taken the energy function as an *input* to the algorithm, and assumed that given a sequence and a biophysical state, a protein will necessarily be found in the lowest-energy conformation. However, to correctly model reality, we must dig deeper.

Free energy.—Physically, we must define the *energy* of a conformation c as a quantity proportional to $-T \ln P(c)$, where $P(c)$ is the probability of finding the molecule in conformation c and T is the temperature. Without loss of generality, we will choose a proportionality constant R (this defines units for the energy); R is the universal gas constant. Since different biophysical states are ultimately just different regions of conformational space, this notion of energy suffices to perform any single- or multistate design: we simply wish to maximize the probability of the molecule being in the state we desire. The probability of a biophysical state s is the sum (or integral) of the probabilities of each of its conformations $c \in C(s)$, and is thus proportional to the partition function q_s , where

$$q_s = \sum_{c \in C(s)} \exp\left(-\frac{E(c)}{RT}\right). \quad (6)$$

It is often useful to work not with the partition function directly, but with the *free energy* $G_s = -RT \ln q_s$ of the state. Then, we simply design to reduce the free energy of desired states and increase the free energy of undesired states. Importantly, as the temperature goes to 0, G_s becomes simply the energy of the state's lowest-energy conformation, and thus we arrive at the more approximate formulation of multistate design presented previously. But this approximation introduces error at nonzero temperature, and algorithms have been developed to actually use G_s at physiological temperatures and thus account for the distribution of energies across conformational space.

Computing the partition function is unfortunately $\#P$ -hard, analogously to similar calculations in statistics. However, the partition function can be efficiently approximated in practice for a particular sequence and biophysical state, while modeling continuous flexibility, using the K^* algorithm.^{5,23,29} The K^* algorithm builds on DEE/A* to model a thermodynamic ensemble of low-energy conformations for the bound and unbound biophysical states of a protein the user wishes to design for binding. Moreover, design based

only on GMECs has been shown not to recapitulate sequences designed with K^* that performed well empirically.²⁹

More efficient provable algorithms have also been developed for this problem. A partition function approximator similar to K^* but accelerated by WCSP techniques has achieved high efficiency,³⁴ albeit neglecting continuous flexibility, which has been shown to compromise accuracy in the K^* context.¹⁰ The BBK^* algorithm²⁵ uses an A^* tree with nodes from many sequences to compute the same top sequences as K^* , and thus provide the same guarantees of accuracy as K^* , in time sublinear in the number of sequences. Thus BBK^* achieves high efficiency while approximating free energy with continuous flexibility.

Improved energy functions.—We have not yet addressed one very important question: How do we accurately estimate $E(c)$ for a conformation c ? The most commonly used energy functions in protein design,⁵ like AMBER, EEF1, and the Rosetta energy function, make many approximations due to their prioritization of speed over accuracy. More accurate energy functions based on induced electric multipoles, quantum chemistry, and Poisson-Boltzmann solvation theory are available, but they are expensive, and they violate a key assumption of the discrete pairwise model of protein design: they are not explicitly a sum of terms depending on at most 2, or indeed on any small number of residues' conformations.

One approach to these problems is to use discrete rotamers and precompute pairwise energies by choosing a “reference” conformation, perturbing it by 1 or a few rotamers at each position, and using the differences in energy between the perturbed and reference conformations as 1-, 2-, and sometimes 3-body energies. This approach yields relatively accurate energies for many systems, using either the Poisson-Boltzmann solvation model³⁵ or the AMOEBA forcefield (featuring induced multipoles)²⁴ as the energy function.

A second approach is to *learn* a representation of the energy suitable for protein design, from a training set that can be generated with any energy function. This approach has the advantages of accommodating continuous flexibility and not requiring all the 1- through 2- or 3-body perturbed conformations from the reference conformation to be physically realizable (this can be an issue in the case of backbone flexibility). Two algorithms in the OSPREY¹⁹ protein design software exploit this approach: the EPIC algorithm learns a polynomial approximation of the continuous energy surface within a voxel, and the LUTE algorithm¹⁷ directly learns a pairwise energy matrix (possibly augmented by triples) from sampled single-voxel minimized energies. Both EPIC and LUTE have been shown to achieve small residuals, while calling the energy function just enough to obtain an accurate characterization of the energy costs of design decisions. Thus, they greatly accelerate design using quantum chemistry- and Poisson-Boltzmann-derived energies.¹⁷

“Exotic” Objective Functions

Not all protein design algorithms optimize energy with respect to sequence; we now review two other approaches.

No matter how tightly a designed protein therapeutic binds its desired target, a strong reaction by the human immune system against this new protein may prevent it from

remaining in the body for long, rendering it ineffective in the clinic. The EpiSweep algorithm²⁶ addresses this problem by finding sequences on the Pareto frontier between an OSPREY-based^{2,10,19} stability design, and an objective function based on avoiding an immune reaction.

It is also sometimes useful, even when optimizing binding, to search the space of known protein backbone conformations to find one that will place sidechains in a desired pose, as in the Rosetta-Match,³⁸ SEEDER,¹¹ and MASTER³⁹ algorithms.

Protein Design on Graphics Processing Units

In the past decade, graphics processing unit (GPU) computation has transformed nearly every area of computational science, from molecular dynamics to computer vision to quantum chemistry. For suitably structured computations, GPUs can perform approximately 1,000 times more FLOPS per dollar spent on hardware.

In the past few years, the computational tasks that are bottlenecks in protein design computation have been implemented for GPUs. For the pairwise discrete model, the bottleneck is combinatorial optimization, which the gOSPREY software⁴⁰ accelerates on GPUs. For continuously flexible protein design, continuous energy minimization within a voxel is the bottleneck. Thus, the OSPREY software, which pioneered minimization-aware protein design, allows continuous energy minimization on GPUs as of its version 3.0, achieving >10x speedups.¹⁹ This compares favorably with the previous flagship application of GPUs in computational structural biology, which is molecular dynamics (MD) simulations of proteins (temporal simulation of proteins using the classical mechanical potential defined by an energy function).

GPUs can exploit two types of parallelism in order to accelerate the biomolecular energy computations central to MD and protein design: (a) processing different conformations of a protein in parallel, and (b) processing different parts of the molecule in parallel. MD is better positioned to exploit (b) than protein design is, because MD evaluates energies for the entire molecule rather than merely the region around the mutations. On the other hand, continuously flexible protein design can minimize energies for a huge number of conformations in parallel, while MD must proceed through different conformations (such as, timesteps) in sequence. This type (a) parallelism in protein design applies both to conformations enumerated in order of lower bound, as in iMinDEE,⁹ and to conformations sampled for the purpose of learning a discrete model of the continuously minimized energy, as in LUTE.¹⁷

Thus, the success of GPUs in accelerating MD computations and the favorable parallelizability of protein design compared to MD bode well for the prospect of very efficient continuously flexible protein design on GPUs, which is already quite impressive in OSPREY 3.0.¹⁹

Successful Applications of Computational Protein Design with Provable Algorithms

Provable computational protein design algorithms have already produced many designs that perform well in experimental tests.^{5,8,10} They have engineered a shift in substrate specificity from one “molecular operand” (input molecule) to another,² and they can predict bacterial mutations in enzyme-coding genes that make the bacterial enzymes resistant to particular antibiotics (Figure 4)— predictions that have been confirmed both in vitro⁶ and in vivo.²⁸

Finally, and perhaps most importantly, proteins designed using provable algorithms have shown promise in the design of therapeutics. Using the techniques reviewed in this paper (in particular, the K^* algorithm³⁰ in OSPREY¹⁹), we collaborated with the NIH Vaccine Research Center to design a broadly neutralizing antibody against HIV with unprecedented breadth and potency (that is, stronger activity against a broader range of HIV strains) that is now in clinical trials (Clinical Trial Identifier: and six others). The OSPREY/ K^* algorithm has also produced peptides that inhibit a protein involved in cystic fibrosis.²⁹ In addition to such direct design of therapeutics, computational prediction of resistance mutations to drug candidates^{6,28} will help combat resistance against new drugs (especially antibiotics) entering the clinic.

Conclusion

Provable computational protein design algorithms have advanced significantly in the last decade. Algorithms for the pairwise discrete approximations have matured, and significant progress is being made with improved biophysical models and for the design of clinically relevant proteins and peptides. Proteins, especially antibodies, are attracting increasing attention from the pharmaceutical industry as drug candidates. These algorithms also have the potential to be transformative in the design of non-protein drugs, because unlike most drug design algorithms, they can search a large space of drug candidates in time sublinear in the size of the space and still guarantee to find the best candidates as if searching one by one.

To achieve the full potential of protein design, it is necessary to further improve the accuracy of the biophysical model. More accurate energy functions, improved modeling of protein-water interactions, and modeling of broader conformational spaces (both for search and for entropy computations) are likely to be important here. Provable guarantees are essential in this endeavor, as they ensure modeling error is the only error in protein design calculations, both allowing new models to be evaluated accurately and preventing design calculations based on accurate models from nonetheless failing due to algorithmic error. As work continues on these important problems, the future of computational protein design with provable algorithms looks bright.

Acknowledgments

Thanks to Lydia Kaviraki, Tomás Lozano-Pérez, Nate Guerin, Jeff Martin, Pablo Gainza, Cynthia Rudin, and members of the Donald Lab. We also thank Toyota Technological Institute of Chicago (M.A.H) and the NIH (grants RO1 GM-78031 and RO1 GM-118543 to B.R.D) for funding.

References

1. Chazelle B, Kingsford C and Singh M A semidefinite programming approach to side chain positioning with new rounding strategies. *INFORMS J. Computing, Computational Biology Special Issue* 16, 4 (2004), 380–392.
2. Chen C, Georgiev I, Anderson A and Donald B Computational structure-based redesign of enzyme activity. *Proc. Nat. Acad. Sci. U. S. A* 106, 10 (2009), 3764–3769.
3. Davey J, Damry A, Goto N and Chica R Rational design of proteins that exchange on functional timescales. *Nature Chemical Biology* 13, 12 (2017), 1280. [PubMed: 29058725]
4. Desmet J, Maeyer M, Hazes B and Lasters I The dead-end elimination theorem and its use in protein sidechain positioning. *Nature* 356 (1992), 539–542. [PubMed: 21488406]
5. Donald B *Algorithms in Structural Molecular Biology*. MIT Press, Cambridge, MA, 2011.
6. Frey K, Georgiev I, Donald B and Anderson A Predicting resistance mutations using protein design algorithms. *Proc. Nat. Acad. Sci. U. S. A* 107, 31 (2010), 13707–13712.
7. Fromer M, Yanover C, and Linial M Design of multispecific protein sequences using probabilistic graphical modeling. *Proteins: Structure, Function, and Bioinformatics* 78, 3 (2010), 530–547.
8. Gainza P, Nisonoff H and Donald B Algorithms for protein design. *Current Opinion in Structural Biology* 39 (2016), 16–26. [PubMed: 27086078]
9. Gainza P, Roberts K and Donald B Protein design using continuous rotamers. *PLoS Computational Biology* 8, 1 (2012), e1002335. [PubMed: 22279426]
10. Gainza P et al. OSPREY: Protein design with ensembles, flexibility, and provable algorithms. *Methods in Enzymology* 523 (2013), 87–107. [PubMed: 23422427]
11. Gainza P, Vollers S and Correia B Mining protein surfaces for binding seeds. (8 2017). RosettaCon.
12. Georgiev I, Lilien R and Donald B The minimized dead-end elimination criterion and its application to protein redesign in a hybrid scoring and search algorithm for computing partition functions over molecular ensembles. *J. Computational Chemistry* 29, 10 (2008), 1527–1542.
13. Grigoryan G, Reinke A and Keating A Design of protein-interaction specificity affords selective bZIP-binding peptides. *Nature* 458, 7240 (2009), 859–864. [PubMed: 19370028]
14. Grigoryan G, Zhou F, Lustig S, Ceder G, Morgan D, and Keating A Ultra-fast evaluation of protein energies directly from sequence. *PLoS Computational Biology* 2, 6 (2006), e63. [PubMed: 16789811]
15. Hallen M and Donald B COMETS (Constrained Optimization of Multistate Energies by Tree Search): A provable and efficient protein design algorithm to optimize binding affinity and specificity with respect to sequence. *J. Computational Biology* 23, 5 (2016), 311–321.
16. Hallen M and Donald B CATS (Coordinates of Atoms by Taylor Series): Protein design with backbone flexibility in all locally feasible directions. *Bioinformatics* 33, 14 (2017), i5–i12. [PubMed: 28882005]
17. Hallen M, Jou J and Donald B LUTE (Local Unpruned Tuple Expansion): Accurate continuously flexible protein design with general energy functions and rigid-rotamer-like efficiency In *Proceedings of the Intern. Conf. on Research in Computational Molecular Biology*. Springer, 2016, 122–136.
18. Hallen M, Keedy D and Donald B Dead-end elimination with perturbations (DEEPer): A provable protein design algorithm with continuous sidechain and backbone flexibility. *Proteins: Structure, Function and Bioinformatics* 81, 1 (2013), 18–39.
19. Hallen M et al. OSPREY 3.0: Open-source protein redesign for you, with powerful new features. *J. Computational Chemistry* 39, 30 (2018), 2494–2507.
20. Jou J, Jain S, Georgiev I, and Donald B BWM*: A Novel, Provable, Ensemble-Based Dynamic Programming Algorithm for Sparse Approximations of Computational Protein Design. *Journal of Computational Biology* 23, 6 (2016), 413–424. [PubMed: 26744898]
21. Kingsford C, Chazelle B, and Singh M Solving and analyzing sidechain positioning problems using linear and integer programming. *Bioinformatics* 21, 7 (2005), 1028–1039. [PubMed: 15546935]

22. Leach A and Lemon A Exploring the conformational space of protein side chains using dead-end elimination and the A* algorithm. *Proteins: Structure, Function, and Bioinformatics* 33, 2 (1998), 227–239.
23. Lilien R, Stevens B, Anderson A and Donald B A novel ensemble-based scoring and search algorithm for protein redesign and its application to modify the substrate specificity of the gramicidin synthetase A phenylalanine adenylation enzyme. *J. Computational Biology* 12, 6 (2005), 740–761.
24. LuCore S, Litman J, Powers K, Gao S, Lynn A, Tollefson W, Fenn T, Washington T and Schnieders M Dead-end elimination with a polarizable force field repacks PCNA structures. *Biophysical J.* 109, 4 (2015), 816–826.
25. Ojewole A, Jou J, Fowler V and Donald B BBK*(Branch and Bound Over K*): A provable and efficient ensemble-based protein design algorithm to optimize stability and binding affinity over large sequence spaces. *J. Computational Biology* (2018). Epub ahead of print.
26. Parker A, Choi Y, Griswold K and Bailey-Kellogg C Structure-guided deimmunization of therapeutic proteins. *J. Computational Biology* 20, 2 (2013), 152–165.
27. Pierce N and Winfree E Protein design is NP-hard. *Protein Engineering* 15, 10 (2002), 779–782. [PubMed: 12468711]
28. Reeve S, Gainza P, Frey K, Georgiev I, Donald B and Anderson A Protein design algorithms predict viable resistance to an experimental antifolate. In *Proceedings of the Nat. Acad. Sci. U. S. A* 112, 3 (2015), 749–754.
29. Roberts K, Cushing P, Boisguerin P, Madden D and Donald B Computational design of a PDZ domain peptide inhibitor that rescues CFTR activity. *PLoS Computational Biology* 8, 4 (2012), e1002477. [PubMed: 22532795]
30. Rudicell R et al. Enhanced potency of a broadly neutralizing HIV-1 antibody in vitro improves protection against lentiviral infection in vivo. *J. Virology* 88, 21 (2014), 12669–12682. [PubMed: 25142607]
31. Simoncini D, Allouche D, Givry S, Delmas C, Barbe S and Schiex T Guaranteed discrete energy optimization on large protein design problems. *J. Chemical Theory and Computation* 11, 12 (2015), 5980–5989.
32. Traoré S, Allouche D, André I, Givry S, Katsirelos G, Schiex T and Barbe S A new framework for computational protein design through cost function network optimization. *Bioinformatics* 29, 17 (2013), 2129–2136. [PubMed: 23842814]
33. Traoré S, Roberts K, Allouche D, Donald B, André I, Schiex T, and Barbe S Fast search algorithms for computational protein design. *J. Computational Chemistry* 37, 12 (2016), 1048–1058.
34. Viricel C, Simoncini D, Allouche D, Givry S, Barbe S and Schiex T Approximate counting with deterministic guarantees for affinity computation *Modelling, Computation and Optimization in Information Systems and Management Sciences*. Springer, 2015, 165–176.
35. Vizcarra C, Zhang N, Marshall S, Wingreen N, Zeng C, and Mayo S. An improved pairwise decomposable finite-difference Poisson-Boltzmann method for computational protein design. *J. Computational Chemistry* 29, 7 (2008), 1153–1162.
36. Xu J and Berger B Fast and accurate algorithms for protein side-chain packing. *J. ACM* 53, 4 (2006), 533–557.
37. Yanover C, Fromer M and Shifman J Deadend elimination for multistate protein design. *J. Computational Chemistry* 28, 13 (2007), 2122–2129.
38. Zanghellini A, Jiang L, Wollacott A, Cheng G, Meiler J, Althoff E, Röthlisberger D and Baker, D. New algorithms and an in silico benchmark for computational enzyme design. *Protein Science* 15, 12 (2006), 2785–2794. [PubMed: 17132862]
39. Zhou J and Grigoryan G Rapid search for tertiary fragments reveals protein sequence-structure relationships. *Protein Science* 24, 4 (2015), 508–524. [PubMed: 25420575]
40. Zhou Y, Xu W, Donald B and Zeng J An efficient parallel algorithm for accelerating computational protein design. *Bioinformatics* 30, 12 (2014), i255–i263. [PubMed: 24931991]

key insights

- Protein design algorithms optimize proteins for desired properties, such as improved stability and stronger and more specific binding to their ligands. These algorithms have many applications in the design of new therapeutics.
- Highly efficient, provably accurate algorithms are available for protein design using a simplified “pairwise discrete” model of protein chemistry.
- Recent algorithms have maintained these provable guarantees and computational tractability while introducing more physically realistic models of protein chemistry. These models account for proteins’ and ligands’ continuous flexibility, model multiple binding or structural states of proteins, and blend more complicated energetic effects.




Figure 1. Pairwise energy functions.

(a) Pairwise energy functions compute energies between pairs of mutable residues (colored) in a protein design problem, but in practice many pairs have very small interaction energies (marked with Xs).

(b) A sparse residue interaction graph (SPRIG) has mutable residues as nodes; edges with small interaction energies can be deleted, enabling highly efficient protein design computations. Figure adapted with permission from Jou et al.²⁰




Figure 2. Favorable conformation.

(a) A conformation modeled using ideal rotamers may have steric clashes—atom pairs that are unphysically close together—even when (b) continuous minimization of the conformation's energy, without changing the rotamers of any residues, results in a very favorable energy. This underscores the need to account for continuous flexibility throughout sequence and conformational search for protein design.

Figure adapted with permission from Gainza et al.⁹




Figure 3. Backbone flexibility.

(left) Mutating residue 54 of the anti-HIV antibody VRC07 to the amino-acid tryptophan (W) improves its function in experimental tests,³⁰ but rigid-backbone modeling of this mutation shows unavoidable steric clashes (purple conformation).

(right) CATS finds a non-clashing conformation (green), resolving this conundrum, while DEEPer (blue) alleviates the clashes partially. Figure adapted with permission from Hallen and Donald.¹⁶




Figure 4. Computational prediction of antibiotic resistance.

(a) the bacterial (*Staphylococcus aureus*) enzyme dihydrofolate reductase binds a drug candidate (“Cpd 1”) tightly, inhibiting the enzyme’s function, but (b) mutating position 31 of the enzyme from amino-acid type valine to leucine causes steric clashes that impeded binding, allowing the bacteria to resist the antibiotic. This predicted resistance mutation was observed experimentally after being predicted by the K^* algorithm as implemented in the OSPREY software.¹⁹ Figure adapted with permission from Reeve et al.²⁸