

LUTE (Local Unpruned Tuple Expansion): Accurate continuously flexible protein design with general energy functions and rigid-rotamer-like efficiency

Mark A. Hallen¹, Jonathan D. Jou¹, and Bruce R. Donald^{1,2,3,*}

Departments of ¹ Computer Science and ² Chemistry, Duke University, Durham, NC 27708 ³ Department of Biochemistry, Duke University Medical Center, Durham, NC 27710 *Corresponding author, brd+recomb16@cs.duke.edu

Abstract. Most protein design algorithms search over discrete conformations and an energy function that is residue-pairwise, i.e., a sum of terms that depend on the sequence and conformation of at most two residues. Although modeling of *continuous flexibility* and of *non-residue-pairwise energies* significantly increases the accuracy of protein design, previous methods to model these phenomena add a significant asymptotic cost to design calculations. We now remove this cost by modeling continuous flexibility and non-residue-pairwise energies in a form suitable for direct input to highly efficient, discrete combinatorial optimization algorithms like DEE/A* or Branch-Width Minimization. Our novel algorithm performs a local unpruned tuple expansion (LUTE), which can efficiently represent both continuous flexibility and general, possibly non-pairwise energy functions to an arbitrary level of accuracy using a discrete energy matrix. We show using 47 design calculation test cases that LUTE provides a dramatic speedup in both single-state and multi-state continuously flexible designs.

1 Introduction

Protein design algorithms compute protein sequences that will perform a desired function.⁵ They generally do this by minimizing the energy of a desired binding or structural state (or some combination thereof^{20,31}) with respect to sequence.^{4,5,7,9,13,15,26,28} Given a model of the conformational space of a protein and its energy function (which maps conformations to their energies), this is a well-defined computational problem.⁵

Previously, this minimization problem has been most efficient to solve if two restrictions are imposed on the model. First, the conformational space of the protein is modeled as discrete. Specifically, each residue takes on conformations from a discrete set (typically, experimentally observed sidechain conformations known as *rotamers*²⁴). Hence, we optimize with respect to the amino-acid type and rotamer of each residue. Second, the energy function is assumed to be residue-pairwise, i.e., it is assumed to be a sum of terms that each depend on the amino-acid types and conformations of at most two residues.

A large body of efficient algorithms has been developed for this restricted case of the protein design problem, many of which offer provable accuracy. In particular, the dead-end elimination (DEE) algorithm⁴ removes rotamers that provably cannot be part of the global minimum-energy conformation (GMEC). The A* algorithm from artificial intelligence²³ finds the optimal conformation using these unpruned rotamers.³⁰ This DEE/A* framework has been generalized to model free-energies for each sequence instead of simply GMECs^{13,33} (the K^* algorithm). It has also been generalized to optimize combinations of stability and specificity by minimizing, with respect to sequence, a linear combination of the conformationally optimized energies of several bound and unbound states of a protein, instead of just the energy of a single state²⁰ (the COMETS algorithm). Several methods in addition to DEE/A* have also been used to address the protein design problem. Some of these, such as Metropolis Monte Carlo and simulated annealing,^{28,32} lack provable guarantees of accuracy, and thus may miss the optimal conformation significantly.⁴² Other algorithms with provable accuracy are also available, largely building on techniques from integer linear programming^{27,38} and weighted constraint satisfaction.^{38,46,47} Notably, treewidth- and branch-width-based algorithms, such as TreePack⁴⁹ and BWM*,²⁵ solve this problem with provable accuracy in polynomial time for systems whose residue interaction graph has treewidth or branch-width bounded by a constant.²⁵

However, proteins are actually continuously flexible, and continuous flexibility both in the sidechains⁹ and backbone²² has been shown to result in significantly lower energies and biologically better sequences.^{9,22} Although a residue sidechain will usually be found in the *vicinity* of the modal conformation for a rotamer, its dihedral angles will often differ from this mode by 10° or more.²⁴ These continuous adjustments are often critical for determining what conformations are sterically feasible.⁹ Thus, incorporation of continuous flexibility modeling substantially increases the accuracy of designs. The minDEE and iMinDEE methods^{9,13} do this for continuous sidechain flexibility, and DEEPe^r²² for simultaneous continuous sidechain and backbone flexibility. These methods replace the traditional discrete rotamers used in DEE/A* with voxels in the conformation space of each residue, called *residue conformations (RCs)*. An RC is defined as an amino acid type together with bounds on each of the conformational degrees of freedom of the residue (e.g., sidechain dihedrals).²² The modal conformation for a rotamer is usually found at the center of this voxel. In this model, the conformation space of an entire protein is a union of voxels, each of which is constructed as the cross-product of single-residue voxels. Thus, each voxel in the conformation space of the entire protein is represented by a list of RCs, one for each residue being modeled as flexible. RCs are constructed to be small enough that we can use local minimization to find the optimal energy within the voxel. This applies to both the single-residue and entire-protein voxels.

However, the global minimum energy in this model could not previously be computed directly by DEE/A*. Instead, DEE/A* was used to enumerate *RC lists* (protein conformational voxels) in order of a *lower bound* on minimized energy.^{9,13,22} Subsequently, the optimal energy for each RC list with a sufficiently

low-energy lower bound was computed by minimization. The lower bound was computed from minimized pairwise interaction energies.¹³ This minimization was accelerated significantly by precomputing polynomials to approximate the energy landscape, using the EPIC algorithm.²¹ However, minimization was still the bottleneck in continuously flexible designs and prevented them from approaching the efficiency of designs with discrete flexibility. **In essence, these previous methods modeled continuously flexibility by modifying DEE/A* and making it do much more work. In contrast, LUTE achieves much greater efficiency by representing continuous flexibility in a form suitable for direct input into DEE/A*.**

We must also address the question of the energy function. The energy landscape of a real protein is not residue-pairwise, or otherwise exactly described solely as the sum of local terms. There is, however, ample evidence that protein interactions are local in a more general sense^{6,21,48,50}—i.e., that the cross-derivative of the energy with respect to conformational degrees of freedom of two residues will tend to zero fairly quickly as the distance between the residues increases. These properties are also observed for more realistic energy functions that return an energy for the entire protein, rather than breaking the energy into terms as molecular mechanics does. For example, the Poisson-Boltzmann model for implicit solvation⁴³ and quantum-chemical models return an energy for the entire system on which they are run. Thus, a viable approach to modeling protein energies more realistically is to infer local terms from full-protein energies. Vizcarra et al.⁴⁸ apply this approach to the Poisson-Boltzmann model, calculating pairwise energies from differences in full-protein conformational states and achieving a pairwise energy matrix that quite accurately matches the Poisson-Boltzmann energies of full conformations. However, their method can only accommodate rotamer pairs, does not support continuous flexibility, and can only be used when substituting a single rotamer into a conformation is possible while maintaining the conformation of the other residues. This is impossible when residues share conformational degrees of freedom, which is typically needed for backbone flexibility,^{12,22} and may also cause problems in the case of steric clashes. Also, DEE/A* has been generalized to accommodate higher-than-pairwise energy terms if these terms are computed explicitly for particular tuples, e.g., triples of residues.³⁴ However, most energy functions modeling higher-than-pairwise effects, including Poisson-Boltzmann, return a single energy for the entire system, rather than a sum of explicit local terms as required by algorithms such as those in Ref. 34.

Hence, today’s protein and drug designers are faced with a choice. They can neglect continuous flexibility and energy terms that aren’t explicitly local (e.g., explicitly pairwise), thus incurring significant error. Or they can pay a massive overhead to incorporate them—by enumerating many conformations (for continuous flexibility) or searching exhaustively (for non-pairwise energy functions). We now offer a way around this dilemma. We construct an energy function that is an explicit, discrete sum of local energy terms, which are associated with tuples of RCs. This function maps RC lists, which represent voxels in the conformation

space of a protein, to energies. But it will approximate, to arbitrary accuracy, the minimized voxel energy, which can be computed with *any* energy function: no need for residue-pairwiseness or any other local representation. Computing this approximation is a machine learning problem, and we attack it with a least-squares method. Our approach has some resemblance to cluster expansion methods, which have previously been used in quantum mechanics³ and to represent optimized energies for protein sequences.^{17,18} However, as discussed in Ref. 21, approximations of energy surfaces can be much more compact if unrealistically high-energy regions of conformational space are excluded from the approximation (and from the subsequent conformational search). Thus, unlike cluster expansion methods, we exclude *pruned* tuples of RCs, making our derived energy function a *local unpruned tuple expansion*, or LUTE. Because conformational and sequence search using the LUTE energy function is a discrete optimization problem of the type solved by DEE/A*, BMW*, and other very efficient algorithms, it allows designs to run quickly using these algorithms, while still approximating continuous flexibility and highly realistic energy functions to a high level of accuracy.

We have implemented LUTE in the OSPREY^{10,13,14} open-source protein design package, which has yielded many designs that performed well experimentally—*in vitro*^{2,8,11,16,37,41,44} and *in vivo*^{8,16,37,41} as well as in non-human primates.⁴¹ OSPREY contains a wide array of flexibility modeling options and provably accurate design algorithms,^{10,14} allowing LUTE to be used for many types of designs.

By presenting LUTE, this paper makes the following contributions:

1. A method to represent continuous flexibility and general energy functions to arbitrary accuracy in a local unpruned tuple expansion (LUTE) that can be used directly as input to discrete combinatorial search algorithms like DEE/A*.
2. A free implementation of LUTE in our laboratory’s open-source OSPREY protein-design software package,^{2,8,13,14} available for download¹⁴ upon publication as free software,¹⁴ supporting representation of both continuous sidechain and backbone flexibility and of molecular-mechanics and Poisson-Boltzmann energy functions.
3. Integration of LUTE with the DEE/A*,³⁰ iMinDEE,⁹ BMW*,²⁵ and COMETS²⁰ algorithms for sequence and conformational search.
4. Bounds on the time and space complexity of protein design calculations that model continuous flexibility and/or use energy functions with non-local terms. The time and space complexity are exponential merely in the branch-width w of the residue interaction graph, and thus the designs can be done in polynomial time for systems whose branch-width is bounded by a constant.
5. Experimental results for 47 computational design calculations on 36 protein structures using LUTE, which demonstrate its accuracy and efficiency in single-state designs, multistate designs and for both n -body Poisson-Boltzmann and pairwise energy functions.

2 Methods

The basic strategy of LUTE is to create a discrete, quick-to-evaluate energy matrix that tells us everything we need to know for design purposes about the

continuous energy landscape of a protein. We will now describe this energy matrix and how it works.

Our goals in protein design (both GMEC^{9,30} and binding/partition function^{13,33} calculations) can be posed in terms of a discrete function $E(\mathbf{r})$ that maps an ordered list \mathbf{r} of RCs to an energy. The list \mathbf{r} contains exactly one RC per residue and thus represents a voxel $V(\mathbf{r})$ in conformation space, where a vector \mathbf{x} of sequence and conformational degrees of freedom satisfies $\mathbf{x} \in V(\mathbf{r})$ if the degree-of-freedom bounds defined by each RC in \mathbf{r} are respected by every degree of freedom in \mathbf{x} . The conformational degrees of freedom in \mathbf{x} will generally be continuous internal coordinates, e.g., sidechain dihedrals. We let $E'(\mathbf{x})$ denote the energy of the protein system, as a function of all its degrees of freedom.

For calculation of the GMEC energy E_g , we wish to minimize $E'(\mathbf{x})$ with respect to \mathbf{x} . Letting R be the set of all possible voxels, the domain over which we minimize is a finite union of voxels $\bigcup_{\mathbf{r} \in R} V(\mathbf{r})$:

$$E_g = \min_{\substack{\mathbf{x} \in \bigcup_{\mathbf{r} \in R} V(\mathbf{r}) \\ \mathbf{r} \in R}} E'(\mathbf{x}) = \min_{\mathbf{r} \in R} \min_{\mathbf{x} \in V(\mathbf{r})} E'(\mathbf{x}), \quad (1)$$

which can be expressed in the form $\min_{\mathbf{r} \in R} E(\mathbf{r})$ where

$$E(\mathbf{r}) = \min_{\mathbf{x} \in V(\mathbf{r})} E'(\mathbf{x}). \quad (2)$$

Similarly, partition function calculations seek to calculate the partition function

$$q = \int_{\bigcup_{\mathbf{r} \in R} V(\mathbf{r})} \exp\left(-\frac{E'(\mathbf{x})}{RT}\right) d\mathbf{x} = \sum_{\mathbf{r} \in R} \int_{V(\mathbf{r})} \exp\left(-\frac{E'(\mathbf{x})}{RT}\right) d\mathbf{x} \quad (3)$$

where R is the gas constant and T is the temperature. Letting

$$E(\mathbf{r}) = -RT \ln \left(\int_{V(\mathbf{r})} \exp\left(-\frac{E'(\mathbf{x})}{RT}\right) d\mathbf{x} \right) \quad (4)$$

we have a formulation of q in terms of the discrete free energy function $E(\mathbf{r})$:

$$q = \sum_{\mathbf{r} \in R} \exp\left(-\frac{E(\mathbf{r})}{RT}\right) \quad (5)$$

Alternately, if we use the definition in Eq. (2) to define $E(\mathbf{r})$, then Eq. (5) gives us the approximation used in Refs. 33 and 13 for the partition function.

Because \mathbf{r} is a discrete variable, the energy $E(\mathbf{r})$ can be decomposed as a sum of energies associated with tuples of RCs (Fig. 1). If all the RCs in a tuple are in the list \mathbf{r} , then that tuple's energy will contribute to $E(\mathbf{r})$. Most higher-order tuples of RCs consist of residues too far apart to have higher-order interactions, and thus do not contribute significantly to the energy (see Section 1 and Ref. 21). We can reduce the number of tuples needed substantially further if we only try to represent favorable, non-clashing conformations. By eliminating high-energy conformations, this restriction of conformational space greatly reduces the range

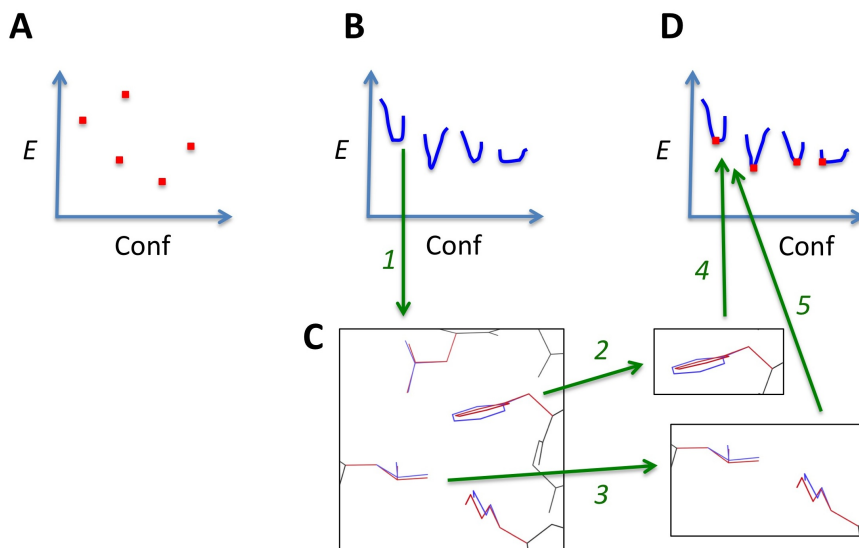


Fig. 1. LUTE makes continuously flexible design efficient by representing continuous flexibility using local, discrete energy terms. (A) Protein design with discrete flexibility searches over a discrete (albeit large) conformational space (“Conf”), looking for low-energy (“ E ”) conformations. Highly efficient algorithms like DEE/ A^* are available for this problem. (B) Protein design with continuous flexibility must search over a large space of voxels (blue) in a continuous conformational space, but we are usually interested only in the minimum-energy point of each voxel. We thus want a way to search combinatorially over these minimum-energy points. (C) The minimized energy of a voxel in protein conformational space depends on all rotamers in the voxel (arrow 1). But we can expand this minimized energy as a sum of local contributions from low-order tuples (e.g., pairs) of residues (arrows 2, 3). (Minimized conformations shown in red, ideal rotamers in blue). (D) This expansion, known as LUTE, gives us a discrete combinatorial search problem of the same form as protein design with discrete flexibility (arrows 4, 5). But this new discrete problem searches over the minimum-energy points (red) of voxels in continuous conformational space (blue). We can solve this problem very efficiently. Figure shows Leu 29, Leu 51, Phe 55, and Lys 59 of the Atx1 metallochaperone (PDB id 1CC8⁴⁰).

of energy values over which $E(\mathbf{r})$ must be accurate. To achieve this, we prune tuples that cannot be part of favorable conformations, and consider only conformations whose tuples are all unpruned. Our expansion is much more efficient to compute after provably unfavorable tuples are pruned. Hence, we are able to represent the energy $E(\mathbf{r})$ as a local unpruned tuple expansion, or LUTE.

Let us consider a conformational space with continuous and discrete degrees of freedom, consisting of RCs, and a mapping $E(\mathbf{r})$ that we can readily calculate. For example, in a typical continuously flexible design, $E(\mathbf{r})$ is defined by Eq. (2), which we assume can be calculated by local minimization. Suppose we have a set T of tuples of RCs at different residue positions. T can contain pairs but also

may contain triples, etc. We then define our local unpruned tuple expansion as a mapping $m : T \rightarrow \mathbb{R} \cup \{\perp\}$. m defines a real coefficient for each tuple $t \in T$, except for pruned tuples, for which $m(t) = \perp$. Let $T_{\mathbf{r}}$ denote the set of tuples in T that consist only of RCs in \mathbf{r} . For example, if T is the set of all possible RC pairs, then $T_{\mathbf{r}}$ will consist of all pairs of RCs in the list of RCs \mathbf{r} . Then LUTE predicts \mathbf{r} to be a pruned conformation if $m(t) = \perp$ for any $t \in T_{\mathbf{r}}$, and otherwise it predicts $E(\mathbf{r}) = \sum_{t \in T_{\mathbf{r}}} m(t)$. We refer to the data structure representing the mapping m as the *LUTE energy matrix*. We call it an energy matrix because it takes a form similar to that of traditional pairwise energy matrices,^{4,9,22,30} although it contains significantly different numerical values when computed for the same design system.

The limiting behavior of LUTE is favorable. As we expand the set T , we must eventually approach perfect accuracy, because if T is the set of all tuples of RCs at different positions, then m can represent $E(\mathbf{r})$ for each full RC list \mathbf{r} explicitly.

If we assume locality of $E(\mathbf{r})$ (see Section 1), we can expect inaccuracies to diminish fairly quickly with increasing size of T , because the component of $E(\mathbf{r})$ modeling the interactions of a residue i will depend only on the RCs assigned to residues fairly close in space to i . As a result, we expect a relatively compact LUTE expansion for any practical protein design problem. In practice, expansions in pairs and triples have worked well (see Section 3).

Most algorithms for protein design with discrete rotamers take a matrix of pairwise energies as input. By simply substituting a LUTE energy matrix for this pairwise energy matrix, we can convert any of these algorithms into an equally efficient design algorithm that searches a continuous search space instead of a discrete one, and/or that optimizes a non-pairwise energy function instead of a pairwise one. The LUTE energy matrix is computed once, before the search, which takes only polynomial time in the number of residues. For example, we need quadratic time to compute a LUTE matrix for which T is all pairs of RCs. Details of the computation by least squares of the LUTE matrix, and of the use of this matrix in search algorithms, are provided in the **Supplementary Information (SI)**¹⁹.

3 Results

We present here complexity results and computational experiments regarding the performance of LUTE. In Section 3.1, we show that the combination of LUTE with the BWM*²⁵ search algorithm is guaranteed to solve continuously flexible protein designs in polynomial time given a residue interaction graph with branch-width bounded by a constant. In Sections 3.2 and 3.3, we present 30 single-state and 17 multistate protein design calculations using LUTE. We measure the gains in efficiency provided by LUTE and its ability to accurately and efficiently perform calculations that, due to their large amount of continuous flexibility (Section 3.2) or non-pairwise energy function (Section 3.3), are inaccessible to previous algorithms. These results include designs with both continuous sidechain and backbone flexibility. Sidechain dihedrals were allowed 9° of continuous motion in either direction relative to the modal value for each

sidechain rotamer,²⁴ while backbone flexibility (when present) was modeled as in Ref. 22.

3.1 Polynomial-time protein design with continuous flexibility

Protein design in the general case is NP-hard.^{1,36} In practice, however, many designs exhibit special properties that make them more tractable. For example, the residue interaction graph—the graph whose edges encode nonnegligible interactions between pairs of residues—of practical designs often has low branch-width. It has been previously shown that protein design with discrete rotamers can be performed in asymptotic time exponential only in the branch-width²⁵ w . Furthermore, these branch-widths can be small irrespective of the number of mutable residues.²⁵ Thus, for many protein designs with discrete rotamers the corresponding GMEC can be found in polynomial time. If one substitutes the LUTE matrix for the discrete pairwise energy matrix in this complexity result, then design with continuous flexibility and a constant-bounded branch-width can be solved in polynomial time as well. We can make

this rigorous using the following theorem, whose proof is provided in SI¹⁹ Section E. In this theorem, a LUTE energy function is a function $E(\mathbf{r}) = \sum_{t \in T_r} m(t)$,

where $m : T \rightarrow \mathbb{R} \cup \{\perp\}$ maps RC tuples to real coefficients (for this purpose, the coefficient \perp of a pruned tuple is effectively ∞). Let n be the number of mutable residues, q be the maximum number of allowed RCs at any mutable residue position, and β_t and β_s be the time and space costs (respectively) to compute the branch-decomposition. This theorem establishes the complexity both of GMEC calculations and of enumeration of subsequent conformations in gap-free ascending order of energy. The latter is essential for calculation of partition functions, which can be used to account for entropy in predictions of binding.^{13,25,33}

Theorem 1. *For a LUTE energy function whose residue interaction graph has branch-width w , the GMEC can be computed in $\mathcal{O}(nw^2q^{\frac{3}{2}w} + \beta_t)$ time and $\mathcal{O}(nwq^{\frac{3}{2}w})$ space, and each additional conformation can be enumerated in order of LUTE energy in $\mathcal{O}(n \log q)$ time and $\mathcal{O}(n)$ space.*

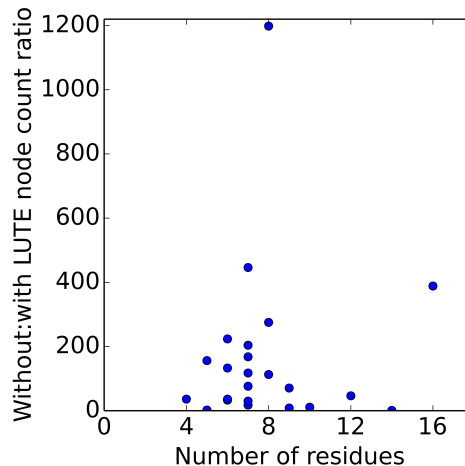


Fig. 2. LUTE markedly reduces the cost of continuously flexible conformational search. Ratios (without LUTE:with LUTE) of the number of nodes in the A* tree before enumeration of the GMEC (or of the last conformation if several conformations closely spaced in energy were calculated; see Ref. 21), versus number of flexible residues. A 20-residue sidechain placement with node ratio 2×10^5 is not shown because it would break the scale.

3.2 Continuous flexibility

LUTE single-state designs were run on 23 protein design systems from Ref. 21 with 4-16 mutable residues, as well as five larger systems (17-40 mutable residues), to measure the efficiency of LUTE and to observe the behavior of LUTE on the larger systems. Many of these larger systems are intractable by previous methods (except *post-hoc* minimization methods that do not account for continuous flexibility during search). The results show that the discrete DEE/A* search with LUTE is dramatically more efficient even compared to EPIC, which offers previously state-of-the-art efficiency for continuously flexible design²¹ (Fig. 2). They also demonstrate that LUTE can handle very large continuously flexible designs—including a 40-residue sidechain placement, which covers a large fraction of the residues in the Atx1 metallochaperone (Fig. 4, left), and a 20-residue design on the same structure with 5 amino-acid types allowed at every position. Furthermore, the LUTE energy matrix consistently represented the true energy landscape very closely (Fig. 3). Optimal sequences and conformations with LUTE differed significantly from the same designs run without continuous flexibility: the same top conformation was returned in only 2 of the 28 single-state designs. On average, 31% of the RCs in the optimal conformations differed from each other. This is consistent with previous work showing that protein design calculations with and without continuous flexibility differ significantly in their results.^{9,22}

For many systems, LUTE achieved a fit with residual under $0.01 \text{ (kcal/mol)}^2$ with only a pairwise expansion. In cases when the pairwise expansion’s residual was higher, an expansion in sparse triples was performed instead. In all but one case, the triples expansion’s residual was less than thermal energy at room temperature (0.59 kcal/mol , i.e., $0.35 \text{ (kcal/mol)}^2$), and thus deemed insignificant.

The one outlier case was a 14-residue design on ponsin (PDB id: 2O9S). It exhibited significant local minimization errors, which caused even the matrix of pairwise lower-bound energies (computed before LUTE precomputation begins) to have errors of at least $\sim 10 \text{ kcal/mol}$. These errors indicate the failure of either our local minimizer or our assumption that local minimization suffices within RCs. As a result of these errors, the LUTE residual even with triples was 1.9 kcal/mol for this system, seven times worse than the next worst residual (the 40-residue Atx1 design). Our software now detects this problem and warns the user before the LUTE computation begins.

17 multistate protein designs were also performed, using a combination of LUTE with our COMETS²⁰ multistate protein design algorithm (see SI¹⁹ Section C). The systems from these designs were taken from Ref. 20; details are provided in SI¹⁹ Section G. The same designs were run with and without continuous flexibility, with LUTE used in the continuous case. As discussed in Ref. 20, COMETS provably returns the same results as exhaustive search over sequences, but it provides a speedup compared to that exhaustive search by (a) considering only a portion of the sequences in the search space explicitly, and (b) only performing a full conformational optimization for a small portion of the sequences in (a). However, previously²⁰ (a) was only significant in designs without continuous flexibility, and (b) was much more pronounced without continuous




Fig. 4. LUTE enables very large provably accurate protein designs with continuous flexibility (left) and with Poisson-Boltzmann energy functions (right). Left: previously, protein designs with continuous flexibility only finished when performed with significantly fewer flexible residues, compared to designs with discrete rotamers. Even 20-residue designs were often intractable. But LUTE solved a sidechain placement problem with continuous flexibility in which 40 residues (purple) were made flexible in the Atx1 metallochaperone (PDB id 1CC8⁴⁰). Right: previous designs using the Poisson-Boltzmann energy function could not optimize this function directly, but only used Poisson-Boltzmann energies to rerank top hits from optimization of a simpler, pairwise energy function. But LUTE can optimize the Poisson-Boltzmann energy function directly—e.g., in a sidechain placement of 20 residues (purple) of Atx1.

flexibility. LUTE brings continuously flexible COMETS designs up to speed with discrete designs on the same system (SI¹⁹ Fig. S3).

3.3 Designs that provably optimize Poisson-Boltzmann energies

We also ran LUTE conformational optimization calculations on two proteins using the Poisson-Boltzmann energy function, which is non-pairwise. This energy was evaluated using Delphi^{35,39} in place of the pairwise EEF1²⁹ solvation energy that is used by default in OSPREY. Interestingly, triple energies did not provide significant benefit here, but LUTE was found to describe the Poisson-Boltzmann energy landscape with a high degree of accuracy. Previous work has shown that an accurate pairwise representation can be obtained for Poisson-Boltzmann energies of *discrete, rigid* rotamers,⁴⁸ but our LUTE results show that a very accurate representation of *continuously minimized Poisson-Boltzmann energies* is possible as well. With continuous flexibility, a 6-residue sidechain placement on the unliganded TIR1/IAA7 complex (PDB code 2P1Q⁴⁵) with continuous flexibility achieved a total residual of 6×10^{-4} and took about 4 days. Furthermore, a 20-residue sidechain placement without continuous flexibility on the bacterial metallochaperone protein Atx1 (PDB code 1CC8;⁴⁰ Fig. 4, right) was solved in 2.5 hours, with total residual $0.04 \text{ (kcal/mol)}^2$. Unlike previous protein design calculations that use Poisson-Boltzmann energies, our new calculations provably return the minimum of the (LUTE-approximated) Poisson-Boltzmann energy

over the entire conformational space, rather than simply over a set of top hits from an initial search that used a cheaper energy function.

4 Conclusions




Fig. 3. LUTE accurately represents continuously minimized energies. Residuals for LUTE ($(\text{kcal/mol})^2$) on the cross-validation data set, measuring the difference between the EPIC energy and a pairwise expansion (blue) or one with sparse triples (red; computed only if pairwise residual exceeded 0.01). x axis: number of flexible residues. Inset: All the same data plotted on a linear scale.

The protein design problem enjoys a wide array of powerful algorithms for conformational and sequence search. These algorithms take a discrete energy matrix and perform sequence optimizations, both in the single-state and multistate cases. At the same time, previous work in bioinformatics and quantum chemistry has made great progress toward quantitatively accurate modeling of the flexibility and energy landscapes of biomolecular systems. Uniting these fields to perform designs with highly realistic modeling would result in great biomedical impact, both in protein and drug design. However, because state-of-the-art flexibility and energy modeling methods do not produce a discrete matrix, there is a gap between these fields. LUTE offers a strategy to bridge this gap. By representing continuous flexibility and general energy functions in a discrete matrix, it greatly increases the realism of the modeling that discrete combinatorial optimization algorithms like

DEE/A* can directly accommodate. We thus believe that LUTE can serve as a foundation for greatly improved biomolecular design protocols.

Acknowledgments We would like to thank Drs. Kyle Roberts and Pablo Gainza for providing PDB files and scripts for testing; all members of the Donald lab for helpful comments; and the PhRMA and Dolores Zohrab Liebmann foundations (MAH) and NIH (grant R01-GM-78031 to BRD) for funding.

References

1. Bernard Chazelle, Carl Kingsford, and Mona Singh. A semidefinite programming approach to side chain positioning with new rounding strategies. *INFORMS Journal on Computing, Computational Biology Special Issue*, 16(4):380–392, 2004.
2. Cheng-Yu Chen, Ivelin Georgiev, Amy C. Anderson, and Bruce R. Donald. Computational structure-based redesign of enzyme activity. *Proceedings of the National Academy of Sciences of the USA*, 106(10):3764–3769, 2009.
3. Jiří Čížek. On the use of the cluster expansion and the technique of diagrams in calculations of correlation effects in atoms and molecules. In *Correlation Effects in*

- Atoms and Molecules*, volume 14 of *Advances in Chemical Physics*, pages 35–90. John Wiley and Sons, 2009.
- Johan Desmet, Marc de Maeyer, Bart Hazes, and Ignace Lasters. The dead-end elimination theorem and its use in protein side-chain positioning. *Nature*, 356:539–542, 1992.
 - Bruce R. Donald. *Algorithms in Structural Molecular Biology*. MIT Press, Cambridge, MA, 2011.
 - N. Flocke and Rodney J. Bartlett. A natural linear-scaling coupled-cluster method. *Journal of Chemical Physics*, 121(22):10935–10944, 2004.
 - C. A. Floudas, J. L. Klepeis, and P. M. Pardalos. Global optimization approaches in protein folding and peptide docking. In *Mathematical Support for Molecular Biology*, volume 47 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 141–172. American Mathematical Society, 1999.
 - Kathleen M. Frey, Ivelin Georgiev, Bruce R. Donald, and Amy C. Anderson. Predicting resistance mutations using protein design algorithms. *Proceedings of the National Academy of Sciences of the USA*, 107(31):13707–13712, 2010.
 - Pablo Gainza, Kyle Roberts, and Bruce R. Donald. Protein design using continuous rotamers. *PLoS Computational Biology*, 8(1):e1002335, 2012.
 - Pablo Gainza, Kyle E. Roberts, Ivelin Georgiev, Ryan H. Lilien, Daniel A. Keedy, Cheng-Yu Chen, Faisal Reza, Amy C. Anderson, David C. Richardson, Jane S. Richardson, and Bruce R. Donald. OSPREY: Protein design with ensembles, flexibility, and provable algorithms. *Methods in Enzymology*, 523:87–107, 2013.
 - I. Georgiev, P. Acharya, S. Schmidt, Y. Li, D. Wycuff, G. Ofek, N. Doria-Rose, T. Luongo, Y. Yang, T. Zhou, B. R. Donald, J. Mascola, and P. Kwong. Design of epitope-specific probes for sera analysis and antibody isolation. *Retrovirology*, 9(Suppl. 2):P50, 2012.
 - Ivelin Georgiev and Bruce R. Donald. Dead-end elimination with backbone flexibility. *Bioinformatics*, 23(13):i185–i194, 2007.
 - Ivelin Georgiev, Ryan H. Lilien, and Bruce R. Donald. 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. *Journal of Computational Chemistry*, 29(10):1527–1542, 2008.
 - Ivelin Georgiev, Kyle E. Roberts, Pablo Gainza, Mark A. Hallen, and Bruce R. Donald. OSPREY (Open Source Protein Redesign for You) user manual. Available online: www.cs.duke.edu/donaldlab/software.php (2009). Updated, 2015. 94 pages.
 - Ivelin S. Georgiev, Rebecca S. Rudicell, Kevin O. Saunders, Wei Shi, Tatsiana Kirys, Krisha McKee, Sijy O’Dell, Gwo-Yu Chuang, Zhi-Yong Yang, Gilad Ofek, Mark Connors, John R. Mascola, Gary J. Nabel, and Peter D. Kwong. Antibodies VRC01 and 10E8 neutralize HIV-1 with high breadth and potency even with Ig-framework regions substantially reverted to germline. *The Journal of Immunology*, 192(3):1100–1106, 2014.
 - Michael J. Gorczynski, Jolanta Grembecka, Yunpeng Zhou, Yali Kong, Liya Roudaia, Michael G. Douvas, Miki Newman, Izabela Bielnicka, Gwen Baber, Takeshi Corpora, Jianxia Shi, Mohini Sridharan, Ryan Lilien, Bruce R. Donald, Nancy A. Speck, Milton L. Brown, and John H. Bushweller. Allosteric inhibition of the protein-protein interaction between the leukemia-associated proteins Runx1 and CBF β . *Chemistry and Biology*, 14:1186–1197, 2007.
 - Gevorg Grigoryan, Aaron W. Reinke, and Amy E. Keating. Design of protein-interaction specificity affords selective bZIP-binding peptides. *Nature*, 458(7240):859–864, 2009.

18. Gevorg Grigoryan, Fei Zhou, Steve R. Lustig, Gerbrand Ceder, Dane Morgan, and Amy E. Keating. Ultra-fast evaluation of protein energies directly from sequence. *PLoS Computational Biology*, 2(6):e63, 2006.
19. Supplementary material available online:
<http://www.cs.duke.edu/donaldlab/Supplementary/recomb16/lute/>
20. Mark A. Hallen and Bruce R. Donald. COMETS (Constrained Optimization of Multistate Energies by Tree Search): A provable and efficient algorithm to optimize binding affinity and specificity with respect to sequence. In *Research in Computational Molecular Biology*, volume 9029 of *Lecture Notes in Computer Science*, pages 122–135. Springer International Publishing, 2015.
21. Mark A. Hallen, Pablo Gainza, and Bruce R. Donald. A compact representation of continuous energy surfaces for more efficient protein design. *Journal of Chemical Theory and Computation*, 11(5):2292–2306, 2015.
22. Mark A. Hallen, Daniel A. Keedy, and Bruce R. Donald. Dead-end elimination with perturbations (DEEPer): A provable protein design algorithm with continuous sidechain and backbone flexibility. *Proteins: Structure, Function and Bioinformatics*, 81(1):18–39, 2013.
23. Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.
24. Joël Janin, Shoshanna Wodak, Michael Levitt, and Bernard Maignet. Conformation of amino acid side-chains in proteins. *Journal of Molecular Biology*, 125(3):357–386, 1978.
25. Jonathan D. Jou, Swati Jain, Ivelin Georgiev, and Bruce R. Donald. BWM*: A novel, provable, ensemble-based dynamic programming algorithm for sparse approximations of computational protein design. *Journal of Computational Biology*, 2015. In press.
26. John Karanicolas and Brian Kuhlman. Computational design of affinity and specificity at protein-protein interfaces. *Current Opinion in Structural Biology*, 19(4):458–463, 2009.
27. Carleton L. Kingsford, Bernard Chazelle, and Mona Singh. Solving and analyzing side-chain positioning problems using linear and integer programming. *Bioinformatics*, 21(7):1028–1039, 2005.
28. Brian Kuhlman and David Baker. Native protein sequences are close to optimal for their structures. *Proceedings of the National Academy of Sciences of the USA*, 97(19):10383–10388, 2000.
29. Themis Lazaridis and Martin Karplus. Effective energy function for proteins in solution. *Proteins: Structure, Function, and Bioinformatics*, 35(2):133–152, 1999.
30. Andrew R. Leach and Andrew P. Lemon. Exploring the conformational space of protein side chains using dead-end elimination and the A* algorithm. *Proteins: Structure, Function, and Bioinformatics*, 33(2):227–239, 1998.
31. Andrew Leaver-Fay, Ron Jacak, P. Benjamin Stranges, and Brian Kuhlman. A generic program for multistate protein design. *PLoS One*, 6(7):e20937, 2011.
32. Christopher Lee and Michael Levitt. Accurate prediction of the stability and activity effects of site-directed mutagenesis on a protein core. *Nature*, 352:448–451, 1991.
33. Ryan H. Lilien, Brian W. Stevens, Amy C. Anderson, and Bruce R. Donald. 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. *Journal of Computational Biology*, 12(6):740–761, 2005.

34. Stephen D. LuCore, Jacob M. Litman, Kyle T. Powers, Shibo Gao, Ava M. Lynn, William T. A. Tollefson, Timothy D. Fenn, M. Todd Washington, and Michael J. Schnieders. Dead-end elimination with a polarizable force field repacks PCNA structures. *Biophysical Journal*, 109(4):816–826, 2015.
35. Anthony Nicholls and Barry Honig. A rapid finite difference algorithm, utilizing successive over-relaxation to solve the Poisson-Boltzmann equation. *Journal of Computational Chemistry*, 12(4):435–445, 1991.
36. Niles A. Pierce and Erik Winfree. Protein design is NP-hard. *Protein Engineering*, 15(10):779–782, 2002.
37. Kyle E. Roberts, Patrick R. Cushing, Prisca Boisguerin, Dean R. Madden, and Bruce R. Donald. Computational design of a PDZ domain peptide inhibitor that rescues CFTR activity. *PLoS Computational Biology*, 8(4):e1002477, 2012.
38. Kyle E. Roberts, Pablo Gainza, Mark A. Hallen, and Bruce R. Donald. Fast gap-free enumeration of conformations and sequences for protein design. *Proteins: Structure, Function, and Bioinformatics*, 83(10):1859–1877, 2015.
39. Walter Rochia, Sundaram Sridharan, Anthony Nicholls, Emil Alexov, Alessandro Chiabrera, and Barry Honig. Rapid grid-based construction of the molecular surface and the use of induced surface charge to calculate reaction field energies: Applications to the molecular systems and geometric objects. *Journal of Computational Chemistry*, 23(1):128–137, 2002.
40. A. C. Rosenzweig, D. L. Huffman, M. Y. Hou, A. K. Wernimont, R. A. Pufahl, and T. V. O’Halloran. Crystal structure of the Atx1 metallochaperone protein at 1.02 Å resolution. *Structure*, 7(6):605–617, 1999.
41. Rebecca S. Rudicell, Young Do Kwon, Sung-Youl Ko, Amarendra Pegu, Mark K. Louder, Ivelin S. Georgiev, Xueling Wu, Jiang Zhu, Jeffrey C. Boyington, Xuejun Chen, Wei Shi, Zhi-Yong Yang, Nicole A. Doria-Rose, Krisha McKee, Sijy O’Dell, Stephen D. Schmidt, Gwo-Yu Chuang, Aliaksandr Druz, Cinque Soto, Yongping Yang, Baoshan Zhang, Tongqing Zhou, John-Paul Todd, Krissey E. Lloyd, Joshua Eudailey, Kyle E. Roberts, Bruce R. Donald, Robert T. Bailer, Julie Ledgerwood, NISC Comparative Sequencing Program, James C. Mullikin, Lawrence Shapiro, Richard A. Koup, Barney S. Graham, Martha C. Nason, Mark Connors, Barton F. Haynes, Srinivas S. Rao, Mario Roederer, Peter D. Kwong, John R. Mascola, and Gary J. Nabel. Enhanced potency of a broadly neutralizing HIV-1 antibody *in vitro* improves protection against lentiviral infection *in vivo*. *Journal of Virology*, 88(21):12669–12682, 2014.
42. David Simoncini, David Allouche, Simon de Givry, Celine Delmas, Sophie Barbe, and Thomas Schiex. Guaranteed discrete energy optimization on large protein design problems. *Journal of Chemical Theory and Computation*, 11(12):5980–5989, 2015.
43. Doree Sitkoff, Kim A. Sharp, and Barry Honig. Accurate calculation of hydration free energies using macroscopic solvent models. *Journal of Physical Chemistry*, 98:1978–1988, 1994.
44. Brian W. Stevens, Ryan H. Lilien, Ivelin Georgiev, Bruce R. Donald, and Amy C. Anderson. Redesigning the PheA domain of gramicidin synthetase leads to a new understanding of the enzyme’s mechanism and selectivity. *Biochemistry*, 45(51):15495–15504, 2006.
45. Xu Tan, Luz Irina A. Calderón-Villalobos, Michal Sharon, Changxue Zheng, Carol V. Robinson, Mark Estelle, and Ning Zheng. Mechanism of auxin perception by the TIR1 ubiquitin ligase. *Nature*, 446:640–645, 2007.

46. Seydou Traoré, David Allouche, Isabelle André, Simon de Givry, George Katsirelos, Thomas Schiex, and Sophie Barbe. A new framework for computational protein design through cost function network optimization. *Bioinformatics*, 29(17):2129–2136, 2013.
47. Seydou Traoré, Kyle E. Roberts, David Allouche, Bruce R. Donald, Isabelle André, Thomas Schiex, and Sophie Barbe. Fast search algorithms for computational protein design. *Journal of Computational Chemistry*, 2016.
48. Christina L. Vizcarra, Naigong Zhang, Shannon A. Marshall, Ned S. Wingreen, Chen Zeng, and Stephen L. Mayo. An improved pairwise decomposable finite-difference Poisson-Boltzmann method for computational protein design. *Journal of Computational Chemistry*, 29(7):1153–1162, 2008.
49. Jinbo Xu and Bonnie Berger. Fast and accurate algorithms for protein side-chain packing. *Journal of the ACM*, 53(4):533–557, 2006.
50. Da W. Zhang and J. Z. H. Zhang. Molecular fractionation with conjugate caps for full quantum mechanical calculation of protein-molecule interaction energy. *Journal of Chemical Physics*, 119(7):3599–3605, 2003.