Structure-Guided Deimmunization of Therapeutic Proteins

ANDREW S. PARKER¹, YOONJOO CHOI¹, KARL E. GRISWOLD², and CHRIS BAILEY-KELLOGG¹

ABSTRACT

Therapeutic proteins continue to yield revolutionary new treatments for a growing spectrum of human disease, but the development of these powerful drugs requires solving a unique set of challenges. For instance, it is increasingly apparent that mitigating potential antitherapeutic immune responses, driven by molecular recognition of a therapeutic protein's peptide fragments, may be best accomplished early in the drug development process. One may eliminate immunogenic peptide fragments by mutating the cognate amino acid sequences, but deimmunizing mutations are constrained by the need for a folded, stable, and functional protein structure. These two concerns may be competing, as the mutations that are best at reducing immunogenicity often involve amino acids that are substantially different physicochemically. We develop a novel approach, called EpiSweep, that simultaneously optimizes both concerns. Our algorithm identifies sets of mutations making such Pareto optimal trade-offs between structure and immunogenicity, embodied by a molecular mechanics energy function and a T-cell epitope predictor, respectively. EpiSweep integrates structure-based protein design, sequence-based protein deimmunization, and algorithms for finding the Pareto frontier of a design space. While structure-based protein design is NPhard, we employ integer programming techniques that are efficient in practice. Furthermore, EpiSweep only invokes the optimizer once per identified Pareto optimal design. We show that EpiSweep designs of regions of the therapeutics erythropoietin and staphylokinase are predicted to outperform previous experimental efforts. We also demonstrate EpiSweep's capacity for deimmunization of the entire proteins, case analyses involving dozens of predicted epitopes, and tens of thousands of unique side-chain interactions. Ultimately, Epi-Sweep is a powerful protein design tool that guides the protein engineer toward the most promising immunotolerant biotherapeutic candidates.

Key words: deimmunization, experiment planning, multi-objective optimization, sequence-based protein design, structure-based protein design, therapeutic proteins.

1. INTRODUCTION

THE EVER-EXPANDING TOOLBOX of approved therapeutic proteins is one of the crowning achievements of modern biotechnology. Biotherapeutic agents are providing new and efficacious treatment options for common diseases such as cancer, diabetes, rheumatoid arthritis, anemia, heart attacks, strokes, and more. Moreover, they are breaking new ground in helping to treat a broad spectrum of previously intractable

¹Department of Computer Science, ²Thayer School of Engineering, Dartmouth College, Hanover, New Hampshire.

illnesses such as multiple sclerosis, cystic fibrosis, congenital lipid and carbohydrate storage disorders, and HIV/AIDS to name a few (Leader et al., 2008). While revolutionizing the treatment of numerous diseases, the rapid growth of the biotherapeutics market has also exposed a range of new challenges in drug design and development. One of the limitations inherent to proteinaceous drugs is the fact that, unlike small molecules, they are subject to surveillance by the human immune system. The recognition of a biotherapeutic as "nonself" can result in the patient's body mounting a concerted anti-biotherapeutic immune response (aBIR). Although in some instances this immune response has no clinical significance, in other cases it manifests a variety of detrimental outcomes ranging from loss of drug efficacy to severe anaphylactic shock (Schellekens, 2002). Therapeutic proteins and peptides of nonhuman origin are especially prone to such an immune response, but the complexity of human immune surveillance can result in immune reactions even against exogenously administered human proteins (Kessler et al., 2006). Regardless of a protein's origin, the aBIR is fundamentally driven by molecular recognition of antigenic peptide sequences embedded within the full-length protein, and removing or manipulating these sequences represents an effective strategy for biotherapeutic deimmunization.

Grafting-based "humanization" strategies simply swap segments of a biotherapeutic candidate for comparable segments of a homologous human protein. While particularly effective for deimmunizing therapeutic antibodies (Jones et al., 1986, Hwang and Foote, 2005), these methodologies are predicated on the availability of a homologous human protein as well as detailed knowledge of underlying structure–function relationships. Non-immunoglobulin proteins, which represent a rich but largely untapped reservoir of prospective therapeutic agents, often fail to meet one or both of these criteria. As a result, there exists a growing need for more broadly applicable protein deimmunization methodologies, and the development of such methods will undoubtedly spur further innovations in disease treatment.

As noted above, the aBIR is driven by molecular recognition of immunogenic peptides, hereafter epitopes, which are found within the biotherapeutic's primary sequence. Immune surveillance is initiated when antigen-presenting cells internalize the therapeutic protein, which is then hydrolyzed into smaller peptide fragments. Fragments that represent potential antigenic epitopes are loaded into the groove of cognate type II major histocompatibility complex (MHC II) proteins (Fig. 1A), and the complexes are trafficked to the cell surface for display to the extracellular milieu. There, the peptide–MHC II complexes are free to interact with T-cell receptor proteins on the surface of T-cell lymphocytes, and true immunogenic epitopes are recognized upon formation of ternary peptide–MHC II–T-cell receptor complexes. The subsequent signaling cascade leads to maturation and proliferation of B-cell lymphocytes that ultimately secrete immunoglobulin molecules able to bind the original biotherapeutic agent.

This well-defined immunological pathway suggests that therapeutic proteins might be engineered such that their peptide fragments would evade MHC II/T-cell receptor binding and thereby block the undesirable aBIR (De Groot et al., 2005, De Groot and Martin, 2009). Indeed, the mutation of key anchor residues within known epitopes of therapeutic proteins has yielded partially deimmunized versions of



FIG. 1. EpiSweep. (**A**) An immune response to a therapeutic protein is initiated by MHC II (red/white) recognition of an immunogenic peptide epitope (blue) digested from the protein. Our goal is to mutate the protein so that no such recognition will occur. (**B**) MHC/T-cell epitopes are pervasive: exposed on the surface, buried in the core, and covering active sites. Shown is staphylokinase (PDB ID 2sak), with black backbone for no epitopes starting at the residue and highlighted sausage for the predicted binding by eight common MHC II alleles, ranging from thin pale yellow (bind just one) to thick bright red (bind all eight). Putative active sites are denoted with green arrows. (**C**) EpiSweep explores the Pareto frontier simultaneously optimizing structure (molecular mechanics energy function, *y*-axis) and immunogenicity (epitope score, *x*-axis). The left circled plan has fewer predicted epitopes than the right one, but the right one has better predicted energy. MHC II, type II major histocompatibility complex; PDB, Protein Data Bank.

staphylokinase (Warmerdam et al., 2002) and erythropoietin (Tangri et al., 2005), among others. It is important to note that unlike antibody epitopes, which localize exclusively to the solvent-exposed protein surface, T-cell epitopes result from proteolytic processing and may therefore be found anywhere within a protein's primary sequence (Fig. 1B). As a result, experimental determination of T-cell epitopes requires technically challenging immunoassays on large pools of overlapping component peptides that span the entire protein of interest. Upon identifying the immunodominant peptides, the protein engineer then performs scanning alanine or other systematic mutagenesis on each one so as to identify critical binding residues whose substitution ameliorates the undesired immune response. Efficacious mutations must then be engineered back into the full-length protein, where they often affect structural stability or therapeutic function. For example, hydrophobic amino acids typically anchor peptide epitopes within the MHC II binding groove, and these amino acids therefore represent attractive targets for mutagenic epitope deletion. Those same residues, however, can also play a central role in stabilizing the close-packed core of the fulllength protein, and substitution at these positions may compromise protein folding. Intuitively, deimmunizing mutations must not undermine the target protein's native structure or function, and consequently protein deimmunization is inherently a multi-objective optimization problem.

To efficiently direct experimental resources toward the most promising sets of mutations, we have developed EpiSweep, a novel approach that integrates validated immunoinformatics and structural modeling methods within a framework for identifying Pareto optimal designs (Fig. 1C). These designs (sets of mutations) make the best trade-offs between the two objectives of stability and immunogenicity in that no design is simultaneously better for both objectives. (Stability) We compute protein stability using a highly successful, structure-based protein design strategy that seeks to optimize side-chain packing (Dahiyat and Mayo, 1997; Lilien et al., 2004; Chen et al., 2009). In this approach, the protein backbone is fixed, and the best side-chain conformations (allowing for amino acid subsitutions) are chosen from a discrete set of common, low-energy rotamers. Individual rotamers are selected so as to minimize the total protein energy, calculated with a molecular mechanics energy function. The side-chain packing approach assumes that a design with low energy for the fixed-target backbone will in fact adopt that target backbone. While this assumption has been borne out by the experimental demonstration of stable, active proteins, it may be advantageous to iterate fixed-backbone design with structure prediction, as is done in RosettaDesign (Kuhlman and Baker, 2000), in order to assess whether the designed sequence is likely to adopt the desired backbone conformation. (Immunogenicity) To assess immunogenicity, we leverage the well-established development of T-cell epitope predictors that encapsulate the underlying specific recognition of an epitope by an MHC II protein (De Groot and Moise, 2007). MHC II proteins from the predominant human leukocyte antigen DR isotype (HLA-DR) have a recognition groove whose pockets form energetically favorable interactions with specific side-chains of peptides approximately nine residues in length (Fig. 1A). Numerous computational methods are available for identifying peptide epitopes, and studies have shown these methods to be predictive of immunogenicity (Wang et al., 2008, De Groot and Martin, 2009). Here we assess each constituent peptide of our protein and optimize the total.

EpiSweep is the first protein design tool that simultaneously optimizes primary sequence, reducing immunogenicity and tertiary structure, maintaining stability and function. It significantly extends structurebased protein design by accounting for the complementary goal of immunogenicity. It likewise significantly extends our previous work on Pareto optimization for protein engineering in general (Zheng et al., 2009, He et al., 2012) and for deimmunization in particular, which assessed effects on structure and function only according to a sequence potential (Parker et al., 2010; Parker et al., 2011a, Parker et al., 2011b). Inspired by an approach for optimization of stability and specificity of interacting proteins (Grigoryan et al., 2009), we employ a sweep algorithm that minimizes the energy of the design target at decreasing predicted epitope scores. The sweep reveals an energy–epitope landscape of Pareto-optimal plans (Fig. 1C) and can also produce near-optimal plans. Although, beyond the scope of this article, EpiSweep promises to inform protein engineering experiments [as our sequence-based algorithms have done (Osipovitch et al., 2012)] seeking sets of effective deimmunizing mutations for the development of enhanced biotherapeutics.

2. METHODS

We seek to make mutations to a target protein so as to reduce its immunogenicity, as evaluated by a sequence-based epitope score, while maintaining its stability and activity, as evaluated by a structure-based

effective energy function. We now formalize this as a Pareto optimization problem that extends the standard side-chain packing formulation of structure-based protein design (Dahiyat and Mayo, 1997) with the complementary/competing epitope score. In general, structure-based protein design problems have been shown to be NP-hard (Pierce and Winfree, 2002).

Problem 1 (Structurally Guided Deimmunization) We are given a protein sequence A of n amino acids, along with a 3D structure that includes a backbone as well as a rotamer sequence R paralleling A. We are also given a set M of mutable positions (amino acid types allowed to vary), and a set F of flexible positions (side-chains allowed to vary) with $M \subset F$; by default $M = F = \{1, ..., n\}$. Finally, we are given a mutational load m. Let a design be an m-mutation sequence A', with mutations only at positions in M, along with a set R' of selected rotamers, differing from R only at positions in F. Our goal is to determine all Pareto optimal designs, minimizing the two objectives

$$f_{\epsilon}(A') = \sum_{i=1}^{n-8} \epsilon(A'[i \dots i+8])$$
(1)

$$f_{\phi}(R') = \sum_{i \in F} \phi_i(R'[i]) + \sum_{i \in F} \sum_{j \in F, \, j \neq i} \phi_{i,j}(R'[i], R'[j])$$
(2)

based on the following contributions:

- $\epsilon: \mathcal{A}^9 \to \mathbb{N}$ gives the epitope score for a peptide (we assume a 9-mer; see below).
- $\phi_i : \mathcal{R} \to \mathbb{R}$ gives the singleton energy capturing the internal energy of a rotamer at position i plus the energy between the rotamer and the backbone structure and side-chains at nonflexible positions.
- $\phi_{i,i}: \mathcal{R} \times \mathcal{R} \to \mathbb{R}$ gives the pairwise energy between a pair of rotamers at a pair of positions i, j.

We use $\mathcal{A} = \{A, C, ..., Y\}$ for the set of amino acids and \mathcal{R} for the set of rotamers (from a rotamer library and the input structure). We subscript sequences with brackets and use the notation X[i..j] to indicate the substring of X from position *i* to *j*, inclusive. We will use function $a : \mathcal{R} \to \mathcal{A}$ to obtain the amino acid type of a given rotamer. While our implementation is modular and can support a variety of epitope scores and energy functions, we now summarize those we use for the results.

Epitope score. Our epitope score is the number of the eight most common HLA-DR alleles [representing a majority of human populations worldwide (Southwood et al., 1998)] predicted by ProPred (Singh and Raghava, 2001) to recognize a given 9-mer peptide at a 5% threshold. If desired, a "personalized" therapeutic could be designed by focusing on a particular set of alleles. ProPred is an extension of the TEPITOPE "pocket profile" method (Sturniolo et al., 1999), learning position-specific amino acid scores from experimentally measured binding affinities for individual amino acid types at individual pockets of the MHC II binding groove (Fig. 1A). Given a 9-mer, the position-specific weights for its amino acids are summed, and the total is compared to a threshold to predict whether or not the given peptide is in a given percentile of the best-recognized peptides. The amino acid in the first pocket is termed the P1 anchor for its significant contribution to binding. ProPred has been successfully employed in a number of experimental studies (e.g., Mustafa and Shaban, 2006; Dinglasan et al., 2007; Klyushnenkova et al., 2007) and was one of the best predictors in a recent independent assessment (Wang et al., 2008), achieving an average 0.73 area under the curve in epitope prediction. In our earlier work, we found its predictions to have a strikingly good correspondence with a previously published ELISPOT assay (Parker et al., 2011b), and used it, along with conservation analysis, to optimize a set of mutations for a β lactamase (Osipovitch et al., 2012).

Molecular mechanics energy function. We assess singleton and pairwise energies according to the AMBER force field as implemented in the Osprey software package for protein redesign (Chen et al., 2009). In particular, we use the van der Waals and electrotatics forces as well as an implicit solvation factor and residue reference energies as intended for standard protein redesign (as opposed to active site redesign for change of substrate specificity). We discard rotamers, and thus all conformations containing them, if they have significant overlap of van der Waals radii or otherwise exceptionally high intra- or inter-rotamer energies.

2.1. Sweep algorithm

We now develop an algorithm to identify *all* Pareto optimal designs. We rely on the fact that epitope scores are discrete at a prespecified significance level—either the peptide is deemed capable of recognition

by an MHC II allele or it is not. Thus, we can "sweep" the epitope score from the wild-type score to successively smaller values for each epitope score identifying the design with the best energy (Fig. 1C). We note that a design with a particular epitope score might actually have a worse overall energy than a design with a somewhat smaller epitope score. Thus, at each step we find the design with the best energy such that the epitope score is at most the current sweep value. A similar sweep approach was employed in developing bZIP partners optimized for stability and specificity, optimizing primarily for stability with a required specificity that is incremented by a fixed amount at each step (Grigoryan et al., 2009). Since our epitope scores are discrete, we can safely step by a value of 1 and provide a guarantee to find all and only the Pareto optimal designs.

More formally, we initialize $E = f_{\epsilon}(A)$, the epitope score of the wild-type protein. Then we repeat the following steps. Let (A', R') be the design minimizing the energy $f_{\phi}(R')$ using the optimization approach below, with the epitope score $f_{\epsilon}(A')$ constrained to be $\leq E - 1$. Output (A', R'), update E to $f_{\epsilon}(A')$, and repeat until there is no solution meeting the smaller epitope constraint. Clearly, each identified design is Pareto optimal, and all Pareto-optimal designs are identified. In addition, the approach is efficient in an output-sensitive fashion, only requiring D invocations of the optimizer to identify D Pareto-optimal designs.

At each step in the sweep, we employ an integer programming (IP) approach to find the global minimum energy conformation with epitope score at most *E*. Like previous work on side-chain packing, we ensure a consistent rotamer selection that minimizes the singleton and pairwise energy scores (Kingsford et al., 2005). However, we also add a network of constraints for the epitope sweep. Define singleton binary variable $s_{i,r}$ to indicate whether or not the rotamer at position *i* is rotamer *r*. Define pairwise binary variable $p_{i,j,r,t}$, derived from $s_{i,r}$ and $s_{j,t}$, to indicate whether or not the rotamers at positions *i* and *j* are rotamers *r* and *t*, respectively. Finally, define window binary variable $w_{i,X}$, derived from $s_{i,r}$ through $s_{i+8,t}$, to indicate whether or not the amino acids for the rotamers in the window from position *i* to position *i* + 8 corresponds to peptide $X \subset A^9$.

We rewrite f_{ϕ} in terms of these binary variables and use it as the objective function for the IP:

$$\Phi = \sum_{i,r} s_{i,r} \cdot \phi_i(r) + \sum_{i,j,r,t} p_{i,j,r,t} \cdot \phi_{i,j}(r,t)$$
(3)

We then constrain the epitope score according to the current sweep value:

$$\sum_{i,X} w_{i,X} \cdot f_{\epsilon}(X) \le E - 1 \tag{4}$$

To guarantee that the variable assignments yield a valid set of rotamers, we impose the following constraints:

$$\forall i : \sum_{r} s_{i,r} = 1 \tag{5}$$

$$\forall i, r, j > i : \sum_{t} p_{i,j,r,t} = s_{i,r} \tag{6}$$

$$\forall j, t, i < j : \sum_{r} p_{i,j,r,t} = s_{j,t} \tag{7}$$

$$\forall i, r \ \forall h \in 1..9 : \sum_{X:X[h]=a(r)} w_{i,X} = s_{i+h-1,r}$$

$$\tag{8}$$

Equation 5 ensures that only one rotamer is assigned to a given position. Equations 6 and 7 maintain consistency between singleton and pairwise variables, while Equation 8 maintains consistency between singleton and window variables.

Finally, we enforce the desired mutational load:

$$\sum_{i, r:A[i] \neq a(r)} s_{i, r} = m \tag{9}$$

We implement this IP using the Java API to the IBM ILOG optimization suite. It provides solutions that are guaranteed to be optimal, at the price of having no guarantee on the time required. However, we find that in practice (see Results section), the globally optimal solutions can in fact be found in at most 8 hours for the size of problems we and others are considering.

Extension: peptide-focused design. Since the "unit" of immunogenicity here (due to MHC binding) is a peptide, initial experimental studies are often peptide-centered. That is, individual immunodominant peptides will be identified by one round of experiments and targeted for epitope deletion in subsequent experiments (Warmerdam et al., 2002; Jones et al., 2005; Tangri et al., 2005). The set M of mutable positions allows us to focus our design effort on such a peptide. However, since only an individual peptide will be assessed for immunogenicity, we modify our epitope scoring and constraints accordingly: We sweep for the epitope score within the peptide while avoiding the introduction of additional epitopes in the "flanking regions" up- and downstream from that peptide (i.e., 9-mers spanning part of the peptide and part of the protein N- or C-terminal to it).

Let the peptide of interest span from residue p_l to residue $p_h \ge p_l + 9$ (so $M \subset \{p_l, \ldots, p_h\}$). Then there are three sets of starting positions for the 9-mers to be assessed: N-terminal flanking positions $P_l = \{\max\{1, p_l - 8\}, \ldots, \max\{1, p_l - 1\}\}$, core positions $P_c = \{p_l, \ldots, p_h - 8\}$, and C-terminal flanking positions $P_h = \{p_h - 7, \ldots, p_h\}$. We replace Equation 4 with a score for just the 9-mers in the peptide of interest (i.e., starting at the core positions):

$$\sum_{i \in P_c, X} w_{i, X} \cdot f_{\epsilon}(X) \le E$$
(10)

and initialize E with a wild-type total that is likewise restricted. We also add flanking region constraints to prevent introducing new epitopes:

$$\forall i \in P_l \cup P_h \ \forall X : \ w_{i,X} \cdot f_\epsilon(X) \le f_\epsilon(A[i..i+8]) \tag{11}$$

2.2. Preprocessing filters

In practice, it is often helpful to prune the search space in order to focus the combinatorial space of designs, employing additional background knowledge to limit the designs to be considered, as well as eliminating regions of the space guaranteed to be suboptimal. Our implementation incorporates two filters.

Homology filters. Natural variation provides insights into mutations likely to yield stable, active structures. We can restrict the allowed mutations to be considered at each position, based either on a general substitution matrix such as BLOSUM, or on a position-specific assessment of conservation within a multiple sequence alignment (MSA) of homologs to the target. For BLOSUM, we eliminate an amino acid (and thereby all its rotamers) at a position if its BLOSUM62 score vs. the wild-type amino acid exceeds a specified threshold. For conservation, we likewise eliminate an amino acid and its rotamers at a position if its frequency in the MSA is below that in a specified background distribution (McCaldon and Argos, 1988).

Dead-end elimination. The dead-end elimination (DEE) criterion provides a pruning technique to eliminate rotamers that are provably not part of the global minimum energy conformation (GMEC) (Desmet et al., 1992). DEE has been generalized and extended in a number of powerful ways (e.g., Lilien et al., 2004; Georgiev et al., 2006; Chen et al., 2009); we find the simple Goldstein DEE criterion to be sufficient for our current purposes (Goldstein, 1994). However, DEE is correct only for pruning with respect to the overall GMEC; it does not account for our epitope score constraint, which might require making more drastic mutations and thus employ rotamers that would be pruned if relying exclusively on the GMEC criterion. To ensure that we do not eliminate rotamers contributing to Pareto optimal designs, we employ the common technique of performing DEE with respect to a "buffer" energy δ , such that rotamers are pruned only if they provably do not participate in any conformation whose energy is within δ of the GMEC. We do this DEE pruning before initiating the epitope sweep. If during the sweep we find that the solution for a particular epitope value has energy exceeding the GMEC+ δ , we must decide either to terminate the sweep (we do not care to produce designs with energy so high) or to re-perform DEE with an increased δ and then re-optimize for that epitope value and continue the sweep.

2.3. Postprocessing energy minimization

The molecular mechanics energy function employed during the sweep assesses energies with respect to the fixed backbone and rigid rotamers. Therefore it is standard practice in structure-based protein design to energy-minimize each resulting design, performing a limited relaxation of the conformation in order to improve atomic interactions. We employ TINKER (Ponder, 2011) to minimize the EpiSweep designs,

using the AMBER (AMBER-f99sb) force field (Hornak et al., 2006) and the generalized Born solvent model. The backbone atoms (N, C α , C, and O) are held fixed, as in the design target, and the side-chains are allowed to relax from their discrete rotameric states.

Since a suboptimal design may relax to a conformation with a better energy than the one that was originally deemed optimal with rigid rotamers, we also consider near-optimal designs. To do so, we generate successive Pareto near-optimal curves, each one Pareto optimal with respect to all designs except those on the previously identified curves. This is done by adding constraints in successive rounds to ensure that a new design differs from each old one by at least one substitution.

In order to ensure that we do not miss designs that are suboptimal with rigid rotamers but become optimal upon energy minimization, we would need to extend the generation of designs to account for the subsequent effects of energy minimization. For example, the minDEE (Georgiev et al., 2006) and iMinDEE (Gainza et al., 2012) variants of DEE ensure that in the pre-processing stage, rotamers are eliminated only if they cannot participate in a design that would be optimal upon minimization. Likewise, a sufficient number of near-optimal Pareto curves could be produced until a bound on the improvement in energy guarantees that remaining designs cannot "leapfrog" to become optimal. We leave such extensions for future work.

2.4. Design selection

After generating and energy minimizing a number of alternative designs, the protein engineer is left with a decision as to which to experimentally evaluate. In practice, we select a number of designs along the (minimized) Pareto frontier, since the appropriate balance between these concerns is unknown ahead of time and is the reason for performing the sweep. Thus, we assess some variants that are more conservative and some that are less conservative from an energetic standpoint, expecting, respectively, less and more impact on immunogenicity. We also consider the overall diversity of the designs, in terms of mutations that are commonly employed among the variants, as well as in their relative locations in the sequence and structure.

3. RESULTS AND DISCUSSION

We demonstrate EpiSweep with case-study applications to two therapeutic proteins, SakSTAR (Protein Data Bank [PDB] code: 2SAK, chain A) and Epo (1EER, chain A), previously targeted for deimmunization by experimental methods. It should be noted that the side-chain atoms of the last residue (R166) in the Epo structure are missing except for the C β atom. The missing side-chain atom coordinates were added using SCWRL (Side-Chains With a Rotamer Library) (Canutescu et al., 2003).

Since the previous experimental work focused on identified immunogenic regions, we first perform analogous peptide-focused design, using the extension discussed in the methods. For these studies, we focus on the overall trends in the energy–epitope landscape, using EpiSweep to explore the trade-offs between maintaining stability and reducing immunogenicity. We do not bother performing energy minimization to better evaluate these designs, as that would not provide significant additional insights into these limited (peptide-focused) designs.

We then demonstrate that EpiSweep can optimize entire proteins, selecting optimal sets of mutations to hit scattered epitopes. Since the side-chains modeled by EpiSweep are from discrete sets of rigid rotamers, we better model the resulting structures by generating not only the Pareto optimal curve but also the four best near-optimal curves, and energy-minimizing all these designs.

The initial energy calculations for each peptide case study required 3–6 hours of wall-clock time on 60 nodes of a shared cluster, while the full-protein design required several days (OSPREY evaluates all pairs of rotamers; this could be reduced by restricting based on a contact map). Then the actual EpiSweep algorithm took less than an hour on a dedicated eight-core machine for each peptide design problem (target and mutational load), and less than 8 hours for the full-protein design. The energy minimization takes approximately 15 minutes of wall-clock time per structure.

3.1. Staphylokinase (SakSTAR) peptides

Staphylokinase is a fibrin-selective thrombolytic agent with potential therapeutic use in treating heart attacks and strokes. Warmerdam et al. (2002) sought to deimmunize a variant called SakSTAR, derived



FIG. 2. Pareto-optimal plans in the energy–epitope landscapes for two SakSTAR peptides. In addition to the wild-type (magenta star for Beta, not shown for C3), the plots show three different mutational loads: 1, solid blue and diamond; 2, dash red and circle; 3, dot black and square. The tables detail the mutations from the wild-type. The numbers above the wild-type sequence are the epitope scores for 9-mers starting at those positions.

from a lysogenic *S. aureus* strain. Based on T-cell profileration assays, they focused on the highly immunogenic C3 region, spanning residues 71–87. They employed alanine-scanning mutagenesis to identify mutations that reduced immunogenicity. Subsequent proliferation assays then showed that the response was indeed reduced for various designs incorporating 2–4 alanine substitutions. They did not engineer the redesigned peptides back into the whole protein to test stability and activity.

We applied peptide-focused EpiSweep to identify mutations in the C3 region as well as in an additional region, which we name Beta, spanning residues 24–38. We evaluated energies according to the deposited SakSTAR structure (2SAK) and applied homology filtering according to the family (Pfam PF02821).

The C-terminal end of SakSTAR folds over perpendicular to the C3 peptide, approximately bisecting it and leaving some residues buried, some exposed, and some half-exposed. Warmerdam et al. (2002) chose substitutions at the positions of C3 underlined in the table in Figure 2 (bottom). They chose not to mutate the large hydrophobics because of potential structural implications, but these same residues tend to anchor MHC II binding. In particular, F76 is a P1 anchor (see Methods, Epitope score) for 4 epitopes. Our energy evaluation predicts that an F76K substitution in fact improves the energy, possibly because it is solvent exposed. Thus, our plans (Fig. 2, bottom) uniformly choose it and obtain a better energy than the wild-type. Over the sweep at a fixed mutational load, we see trends from more to less energetically favorable substitutions with decreasing epitope score (left-right in the figure and bottom-up in the table). We are able to delete all the epitopes with either two or three substitutions.

We chose also to study the Beta peptide because it was predicted to be highly immunogenic in our initial T-cell epitope prediction analysis, and it is also structurally quite different from C3. Beta sits in an antiparallel beta strand with a pattern of surface-exposed and buried side-chains. The predicted epitope and energy range (Fig. 2, top) is larger than that for C3, and the number of undominated solutions greater except for the two-mutation curve. In the plans, we see that Y24H is commonly taken, as is T30K at the higher mutational loads. We again see trade-offs between stability-preserving and epitope-deleting selections and note that some of the mutations are predicted to be stability enhancing (though again, this is before energy minimization).

3.2. Erythropoietin (Epo) peptides

Erythropoietin (Epo) has therapeutic use in treating anemia but unfortunately induces an immune response in some patients (Indiveri and Murdaca, 2002). Based on an intensive analysis of peptides spanning the entire protein, Tangri et al. (2005) identified two highly immunogenic regions, spanning 101–115 and 136–150. They engineered four variants targeting the anchor residues of identified T-cell epitopes in these regions: L102P/S146D (named G2), T107D/S146D (G3), L102G/T107D/S146D (G4), and L102S/T107D/



FIG. 3. Pareto-optimal plans in the energy-epitope landscapes for two Epo peptides. See Figure 2 for description.

S146D (G5). Variants G3 and G4 reduced response in an ELISPOT assay. However, variants G2 and G5 were not bioactive, possibly because of destabilizing substitutions at L102. We target the same two solvent-exposed peptides, and develop EpiSweep designs based on the structure with PDB ID 1buy and Pfam entry PF00758.

In peptide 101–115, EpiSweep avoids the L102 position found to be a problem by Tangri et al. (2005) L102 is completely buried and tightly packed in the model structure. Instead, EpiSweep makes an A, G, or L mutation at R103 in every plan, except the energy-optimal plans with two of three mutations (Fig. 3). While these might not seem conservative substitutions from a purely sequence-based perspective, we find that the energies are favorable. Both EpiSweep and Tangri et al. make substitutions at T107. Small residues are common substitutions here, as they tend not to make good interactions with MHC II pockets, but yet are predicted to be energetically benign. The number of glycines suggests that additional modeling might be appropriate to more fully assess the impact on overall flexbility. We see that the three-mutation plan has a much nicer deimmunizing trend than the two-mutation plan, deleting all epitopes without having to take an energetically less favorable choice.

In peptide 136–150, Tangri et al. (2005) make a substitution at S146, while we concentrate instead at or around the P1 anchor at F138. As the trends show, mutations are made around, but not at 138 until the epitope sweep reaches a point where this substitution is required to eliminate as many epitopes as possible. Mutation at 138 is not the most energy favorable, and so we see many plans with nearby R139S and K140E instead. Like the other peptides, as the mutation load increases, the sweep has more options to choose substitutions that minimize energy and epitope scores simultaneously.

3.3. SakSTAR full-protein design

We next subjected the entire SakSTAR protein to EpiSweep. Warmerdam et al. (2002) verified T-cell epitopes exist outside C3 and that the SakSTAR protein requires extensive deimmunization work in other regions. Figure 4 illustrates some Pareto-optimal designs at mutational loads of 4, 6, and 8. Full lists are provided in the Supplementary Material (available online at www.liebertonline.com/cmb).¹ We see the same kinds of trade-offs as in the peptide-focused designs, moving from stability-preserving (or -enhancing) mutations that nominally reduce immunogenicity, to epitope-deleting mutations that might destabilize the protein. However, there is more freedom in the full-protein design, so we see more effective epitope-deleting mutations throughout each curve. For example Y24K and V112A replace P1 anchors for relatively good energy trade-off.

¹Provided is a spreadsheet with lists of mutations and their epitope scores and initial and minimized energies.



FIG. 4. Pareto-optimal plans in the energy–epitope landscape for the full SakSTAR protein. The plots show three different mutational loads: 4, solid blue and diamond; 6, dash red and circle; 8, dot black and square. Details are provided for only sampled sets of plans along the curves; all are detailed in the Supplementary Material.

Previously, we had redesigned SakSTAR based on a one- and two-body sequence potential, sampling linear combinations of this sequence potential and the epitope score (Parker et al., 2011b). It is interesting to compare and contrast the sequence-guided plans with these structure-guided (and Pareto optimal) plans. For example, both tend to choose Y24K and V112A, along with some substitution at M26 (M26D for sequence and M26K for structure). These choices are natural, as the wild-type residues are P1 anchors for MHC II binding, and mutation to a charged residue precludes effective binding. However, the sequence-guided approach completely avoided the C3 region; we conjectured that this was to a high degree of covariation among the amino acids in the sequence record. In contrast, many of the structure-guided plans incorporate substitutions in the C3 region at 76, 78, and 79. Specifically, F76 and V79 are solvent exposed and replacement with a hydrophillic residue improves the energy score. These two residues are replaced often. On the other hand, V78 side chain is buried and few plans show the more conservative V78L mutation. By directly modeling the energetic interactions, rather than relying just on evolutionary history, we may discover new designs that are favorable for our design goals rather than natural pressures.

Energy minimization of the rigid rotamer-based designs to relax the side-chain conformations produced structures with lower energies, as illustrated in Figure 5 and detailed in the Supplementary Material. The general trend of the energy vs. epitope trade-off is still maintained: For each mutational load, once the number of epitopes is forced to go sufficiently low, the energy increases. We note, though, that increasing the mutational load produces lower energy designs for the same epitope scores by employing less drastic substitutions. Thus, we see clear boundaries of the Pareto frontiers at different numbers of mutations.

3.4. Epo full-protein design

Our final result redesigns the entire Epo protein with EpiSweep. Tangri et al. (2005) quantified each Epo peptide binding to MHC-II protein through extensive experimental effort. They then engineered limited





• 4 Mutations • 6 Mutations • 8 Mutations



FIG. 6. Pareto-optimal plans in the energy-epitope landscape for the full Epo protein. See Figure 4 for description.

segments of Epo at some regions of high immunogenicity. However, as shown by their immunoassays, other immunogenic regions exist. Figure 6 illustrates some Pareto-optimal designs for the full protein at mutational loads of 4, 6, and 8. Several residues in the 136-150 peptide are targeted, including F138, F142, and L149. Our methodology largely avoids the 101-115 peptide (with the exception of R103G) possibly because the buried, contacting, and smaller side-chains do not allow room for substitution. Instead EpiSweep chooses P1 anchors at Y49, L75, and W88. EpiSweep also mutates R103G and K140E which, although not P1 anchors, favorably change the binding profile of the composite peptides while making a relatively good energy tradeoff for the full design.

Similar to SakSTAR, previously we had redesigned Epo based on a one- and two-body sequence potential, sampling linear combinations of this sequence potential and the epitope score (Parker et al., 2011b). And it is again interesting to compare and contrast the sequence-guided plans with these structure-guided (and Pareto optimal) plans. We notice that our previous sequence plans, like the new structure plans, avoid the 101-115 peptide but do mutate in the 136-150 peptide; in particular, both approaches make an F138A substitution. Both approaches also target L75, but the sequence-based approach employs L75S while structure-guided design prefers L75N. Strikingly, both sequence and structure plans utilize V82E, which we describe in the following paragraph.

It should be noted that for Epo, the mutants are predicted to have lower energies (median -4422 kcal/mol; the energy values of all designs are depicted in Fig. 7) than the wild-type (-3859 kcal/mol). This was not the case for SakSTAR (wild-type: -2154 kcal/mol, median mutant -2160 kcal/mol). We investigated the cause of these lower energies. Of the 293 designs, 237 of them have a V82E substitution. (See the Supplementary Material for the other mutation frequencies.) In addition to reducing the epitope score by 6, this mutation introduces a negatively charged residue that makes a strong charge–charge interaction with the positively charged adjacent residue K82. We used TINKER to model this single mutation (Fig. 8) and found that it provides a -146 kcal/mol contribution to the energy. We note that such stabilization has been



• 4 Mutations • 6 Mutations • 8 Mutations



FIG. 8. Side-chain rearrangement due to the V82E mutation that appears in 237 of the 293 Epo designs. The wild-type Epo structure is colored red, while the mutant is colored yellow and the mutated V82E is in gray. The valine residue (V82) of the wild-type is omitted for clarity. The lysine residue (K83) changes its conformation due to the charge–charge interaction (dashed line) between K83 and the mutated V82E, yielding a 146 kcal/mol reduction in energy. This figure was created using PyMol.

demonstrated experimentally, with redesigning of surface charge-charge interactions yielding enzymes that are stabilized relative to the wild-types but with their activities retained (Gribenko et al., 2009). However, this large an energy change (comparable values were produced under the CHARMM and OPLS force fields, in addition to AMBER) may overestimate the actual stabilizing effects of the substitution (de Prat Gay et al., 1994; Vijayakumar and Zhou, 2001). But we would still speculate, based on the modeling, that the variant would, in fact, exhibit improved stability relative to the wild-type Epo.

4. CONCLUSION

In order to simultaneously optimize stability/activity and immunogenicity of therapeutic proteins, we have developed a novel Pareto optimization approach that integrates methods from structure-based protein design with immunoinformatics predictors. Our EpiSweep algorithm elucidates the energy–epitope land-scape of a protein, identifying all Pareto-optimal plans, along with near-optimal plans as desired. While the underlying design problem is NP-hard, our methods are efficient in practice, requiring only hours to characterize an entire Pareto frontier even for a redesign problem considering an entire protein. We recognize that this speed is due to our reliance on computational models of both stability and immunogenicity that, while extensively validated in numerous retrospective and prospective studies, are imperfect and may yield designs that are unstable or immunogenic in practice. Due to our use of provably correct algorithms, however, this outcome reflects only on the models and offers an opportunity to improve them. Furthermore, the ability of EpiSweep to characterize the beneficial region of the energy–epitope landscape enables engineers to better identify high-confidence designs worthy of experimental evaluation.

ACKNOWLEDGMENTS

This work is supported in part by NSF grant IIS-1017231 to CBK and NIH grant R01-GM-098977 to CBK and KEG.

DISCLOSURE STATEMENT

The authors declare that no competing financial interests exist.

REFERENCES

- Canutescu, A.A., Shelenkov, A.A., and Dunbrack Jr., R.L. 2003. A graph-theory algorithm for rapid protein side-chain prediction. *Protein Science* 12, 2001–2014.
- Chen, C.-Y., Georgiev, I., Anderson, A.C., and Donald, B. R. 2009. Computational structure-based redesign of enzyme activity. *PNAS* 106, 3764–3769.
- Dahiyat, B., and Mayo, S. 1997. De novo protein design: fully automated sequence selection. Science 278, 82-87.
- De Groot, A.S., Knopp, P.M., and Martin, W. 2005. De-immunization of therapeutic proteins by T-cell epitope modification. *Dev. Biol. (Basel)* 122, 171–94.
- De Groot, A.S., and Martin, W. 2009. Reducing risk, improving outcomes: Bioengineering less immunogenic protein therapeutics. *Clinical Immunology* 131, 189–201.
- De Groot, A.S. and Moise, L. 2007. Prediction of immunogenicity for therapeutic proteins: State of the art. *Curr. Opin. Drug Discov. Devel.* 10, 332–340.
- de Prat Gay, G., Johnson, C., and Fersht, A. 1994. Contribution of a proline residue and a salt bridge to the stability of a type I reverse turn in chymotrypsin inhibitor-2. *Protein Eng.* 7, 103–108.
- Desmet, J., De Maeyer, M., Hazes, B., and Lasters, I. 1992. The dead-end elimination theorem and its use in protein side-chain positioning. *Nature* 356, 539–542.
- Dinglasan, R.R., Kalume, D.E., Kanzok, S.M., et al. 2007. Disruption of Plasmodium falciparum development by antibodies against a conserved mosquito midgut antigen. *PNAS* 104, 13461–13466.
- Gainza, P., Roberts, K., and Donald, B. 2012. Protein design using continuous rotamers. PLoS Comput. Biol. 8, e1002335.
- Georgiev, I., Lilien, R., and Donald, B. 2006. A novel 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. *Proc. RECOMB*, 530–45.
- Goldstein, R.F. 1994. Efficient rotamer elimination applied to protein side-chains and related spin glasses. *Biophysical J.* 66, 1335–1340.
- Gribenko, A.V., Patel1, M.M., Liu, J., et al. 2009. Rational stabilization of enzymes by computational redesign of surface charge interactions. *PNAS* 106, 26012606.
- Grigoryan, G., Reinke, A., and Keating, A. 2009. Design of protein-interaction specificity gives selective bZIP-binding peptides. *Nature* 458, 859–864.
- He, L., Friedman, A., and Bailey-Kellogg, C. 2012. A divide and conquer approach to determine the pareto frontier for optimization of protein engineering experiments proteins. *Proteins* 80, 790–806.
- Hornak, V., Abel, R., Okur, A., et al. 2006. Comparison of multiple amber force fields and development of improved protein backbone parameters. *Proteins* 65, 715–725.
- Hwang, W.Y.K., and Foote, J. 2005. Immunogenicity of engineered antibodies. Methods 36, 3-10.
- Indiveri, F., and Murdaca, G. 2002. Immunogenicity of erythropoietin and other growth factors. *Rev Clin Exp Hematol.* 1, 7–11.
- Jones, P.T., Dear, P.H., Foote, J., et al. 1986. Replacing the complementarity-determining regions in a human antibody with those from a mouse. *Nature* 321, 522–525.
- Jones, T.D., Phillips, W.J., Smith, B.J., et al. 2005. Identification and removal of a promiscuous CD4+ T cell epitope from the C1 domain of factor VIII. *J. Thromb. Haemost.* 3, 991–1000.
- Kessler, M., Goldsmith, D., and Schellekens, H. 2006. Immunogenicity of biopharmaceuticals. *Nephrology, Dialysis, Transplantation* 21, v9–12.
- Kingsford, C., Chazelle, B., and Singh, M. 2005. Solving and analyzing side-chain positioning problems using linear and integer programming. *Bioinf.* 21, 1028–1036.
- Klyushnenkova, E.N., Kouiavskaia, D.V., Kodak, J. A., et al. 2007. Identification of HLA-DRB1*1501-restricted T-cell epitopes from human prostatic acid phosphatase. *Prostate* 67, 1019–1028.
- Kuhlman, B., and Baker, D. 2000. Native protein sequences are close to optimal for their structures. PNAS 97, 10383– 10388.
- Leader, B., Baca, Q.J., and Golan, D.E. 2008. Protein therapeutics: a summary and pharmacological classification. *Nat. Rev. Drug Disc.* 7, 21–39.
- Lilien, R., Stevens, B., Anderson, A., and Donald, B. 2004. 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 adenlytaion enzyme. *Proc. RECOMB*, 46–57.
- McCaldon, P., and Argos, P. 1988. Oligopeptide biases in protein sequences and their use in predicting protein coding regions in nucleotide sequences. *Proteins: Structure, Function and Genetics* 4, 99–122.
- Mustafa, A.S., and Shaban, F.A. 2006. Propred analysis and experimental evaluation of promiscuous T-cell epitopes of three major secreted antigens of Mycobacterium tuberculosis. *Tuberculosis* 86, 115–124.
- Osipovitch, D., Parker, A., Makokha, C., et al. 2012. Design and analysis of immune-evading enzymes for ADEPT therapy. *Protein Eng. Des. Sel.* 25, 613–624.

- Parker, A.S., Griswold, K.E., and Bailey-Kellogg, C. 2011a. Optimization of combinatorial mutagenesis. J. Comput Biol 18, 1743–56.
- Parker, A.S., Griswold, K.E., and Bailey-Kellogg, C. 2011b. Optimization of therapeutic proteins to delete T-cell epitopes while maintaining beneficial residue interactions. *J Bioinf Comput Biol* 9, 207–229.
- Parker, A.S., Zheng, W., Griswold, K.E., and Bailey-Kellogg, C. 2010. Optimization algorithms for functional deimmunization of therapeutic proteins. *BMC Bioinf.* 11, 180.

Pierce, N., and Winfree, E. 2002. Protein design is NP-hard. Protein Eng. 15, 779-782.

- Ponder, J. 2011. TINKER: Software tools for molecular design (version 6.0). Jay Ponder Lab, Washington University, Saint Louis, MO.
- Schellekens, H. 2002. Bioequivalence and the immunogenicity of biopharmaceuticals. *Nature Reviews Drug Discovery* 1, 457–462.
- Singh, H., and Raghava, G. 2001. ProPred: prediction of HLA-DR binding sites. Bioinformatics 17, 1236–1237.
- Southwood, S., Sidney, J., Kondo, A., et al. 1998. Several common HLA-DR types share largely overlapping peptide binding repertoires. J. Immunol. 160, 3363–3373.
- Sturniolo, T., Bono, E., Ding, J., et al. 1999. Generation of tissue-specific and promiscuous HLA ligand database using DNA microarrays and virtual HLA class II matrices. *Nature Biotechnol.* 17, 555–561.
- Tangri, S., Mothe, B.R., Eisenbraun, J., et al. 2005. Rationally engineered therapeutic proteins with reduced immunogenicity. J. Immunol. 174, 3187–3196.
- Vijayakumar, M., and Zhou, H.-X. 2001. Salt bridges stabilize the folded structure of barnase. J. Phys. Chem. B 105, 7334–7340.
- Wang, P., Sidney, J., Dow, C., et al. 2008. A systematic assessment of MHC class II peptide binding predictions and evaluation of a consensus approach. *PLoS Comp. Biol.* 4, e1000048.
- Warmerdam, P.A.M., Plaisance, S., Vanderlick, K., et al. 2002. Elimination of a human T-cell region in staphylokinase by T-cell screening and computer modeling. J. Thromb. Haemost. 87, 666–673.
- Zheng, W., Friedman, A.M., and Bailey-Kellogg, C. 2009. Algorithms for joint optimization of stability and diversity in planning combinatorial libraries of chimeric proteins. *J Comput Biol* 16, 1151–1168.

Address correspondence to: Dr. Chris Bailey-Kellogg Department of Computer Science Dartmouth College 6211 Sudikoff Laboratory Hanover, NH 03755

E-mail: cbk@cs.dartmouth.edu