
Machine learning in protein engineering

Kevin K. Yang¹ Zachary Wu¹ Frances H. Arnold¹

Abstract

Machine learning-guided protein engineering is a new paradigm that enables the optimization of complex protein functions. Machine-learning methods use data to predict protein function without requiring a detailed model of the underlying physics or biological pathways. They accelerate protein engineering by learning from information contained in all measured variants and using it to select variants that are likely to be improved. In this review, we introduce the steps required to collect protein data, train machine-learning models, and use trained models to guide engineering. We make recommendations at each stage and look to future opportunities for machine learning to enable the discovery of new protein functions and uncover the relationship between protein sequence and function.

1. Introduction

Protein engineering seeks to design or discover proteins whose properties, useful for technological, scientific, or medical applications, have not been needed or optimized in nature. We can envision the mapping of protein sequence to a desired function or functions as a “fitness landscape” over the high-dimensional space of possible protein sequences (Romero & Arnold, 2009). The fitness represents a protein’s performance: expression level, catalytic activity, or other properties of interest to the protein engineer. The landscape determines the range of properties available to different sequences as well as the ease with which they can be optimized. In one limit, convex landscapes that climb smoothly to a global maximum are easy to search. In the other, rugged landscapes with many local maxima are much more difficult to traverse. Protein engineering seeks to identify sequences corresponding to high fitnesses on these landscapes.

Identifying optimal locations on a fitness landscape is extremely challenging. The space of possible protein sequences is too large to be searched exhaustively naturally, in the laboratory, or computationally (Mandecki, 1998). The problem of finding optimal sequences is NP-hard, meaning that there is no known polynomial-time method for searching this space (Pierce & Winfree, 2002). Functional proteins are also extremely scarce in this vast space of sequences. Moreover, as the threshold level of fitness increases, the number of sequences having that fitness decreases exponentially (Smith, 1970; Orr, 2006). As a result, highly fit sequences are vanishingly rare and overwhelmed by nonfunctional and mediocre sequences.

Until recently, the two main approaches for finding high-fitness protein sequences have been directed evolution and rational design. Rational design uses physics-based models to guide the search for improved sequences. These models typically contain an atomic structural representation of a protein and energy-based scoring functions to quantify the target function (Dahiyat & Mayo, 1997; Siegel et al., 2010). Rational design has been successful in identifying sequences that fold into desired static structures (Dou et al., 2018). This is an important advance, and useful when a single stable structure dictates function (Dou et al., 2018). However, because proteins are marginally stable, even small inaccuracies in energy-based scoring functions can lead to very poor performance (Dou et al., 2017).

Many protein functions such as binding, catalysis, allostery, and signalling, are mediated through recessed cavities, mobility, or multiple low-energy states (Huang et al., 2016). For example, computational enzyme design currently proceeds by designing an idealized active site for the desired reaction, matching the active site residues to stable backbones, and then using molecular dynamics (MD) simulations to winnow designs with flaws not apparent from static evaluations. MD simulations require enormous computational resources (100s of CPU hours for each variant) and are not appropriate for

¹Division of Chemistry and Chemical Engineering, California Institute of Technology, Pasadena, CA, USA. Correspondence to: Frances H. Arnold <frances@chem.caltech.edu>.

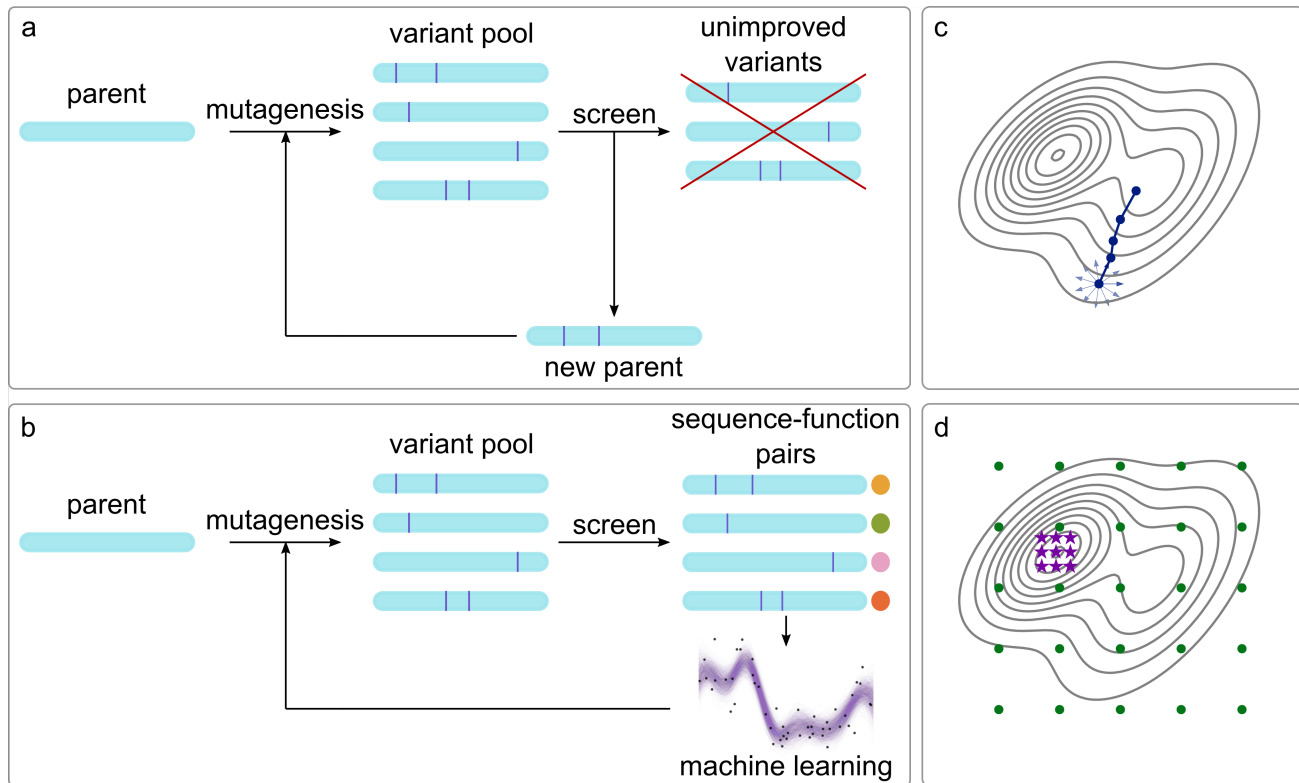


Figure 1. Directed evolution with and without machine learning. (a) Directed evolution uses iterative cycles of diversity generation and screening to find improved variants. Information from unimproved variants is discarded. (b) Machine-learning methods use the data collected in each round of directed evolution to choose the mutations to test in the next round. Careful choice of mutations to test decreases the screening burden and improves outcomes. (c) Directed evolution is a series of local searches on the fitness landscape. (d) Machine learning-guided directed evolution often rationally chooses the initial points (green circles) to maximize the information learned from the fitness landscape, allowing future iterations to quickly converge to improved sequences (violet stars).

testing many variants. This process generally yields sequences with modest activity that are finally improved with directed evolution (Garcia-Borràs et al., 2017).

Inspired by natural evolution, directed evolution climbs a fitness landscape by accumulating beneficial mutations in an iterative protocol of mutation and selection, as illustrated in Figure 1a. The first step is sequence diversification using techniques such as random mutagenesis, site-saturation mutagenesis, or recombination to generate a library of modified sequences starting from the parent sequence(s). The second step is screening or selection to identify variants with improved properties for the next round of diversification. The steps are repeated until fitness goals are achieved.

Directed evolution is limited by the fact that even the most high-throughput screening or selection methods only sample an insignificant fraction of the sequences that can be made using most diversification methods, and developing efficient screens is nontrivial. Moreover, directed evolution requires at least one minimally-functional parent and a locally-smooth fitness landscape for stepwise optimization (Romero & Arnold, 2009). Recombination methods may allow for bigger jumps in sequence space while retaining function (Drummond et al., 2005), but sequences designed using recombination are by definition restricted to exploring combinations of previously-explored mutations. No matter the diversification technique, directed evolution is energy-, time-, and material-intensive, and multiple generations of evolution may be required to achieve meaningful performance improvements.

More recently, researchers have begun using statistical, or machine-learning, methods to approximate sequence-function relationships. Machine-learning methods learn functional relationships from data (Hastie & Tibshirani, 2008; Murphy, 2012). Machine-learning models of protein function can be predictive even when the underlying mechanisms are not well-understood. As shown in Figure 1b, machine-learning models can guide directed evolution by learning from the

information contained in all measured variants. This new paradigm for protein engineering enables engineering with fewer measurements and fewer generations of evolution. Here, we cover the basics of machine learning with a focus on applications to protein engineering and protein sequence-function datasets, discuss how machine learning can be integrated with directed evolution to accelerate protein engineering, and consider the developments that are required to enable wider applications.

2. A brief introduction to machine learning on proteins

In most computational methods, the user provides a hard-coded algorithm and inputs, and the computer executes the steps provided by the human expert. In contrast, machine-learning models infer patterns from data, which can then be used to make predictions on unobserved data. The user must collect the training data, represent it in a form amenable to machine learning, decide what type of machine-learning algorithm to use, and train and interpret the model.

2.1. Protein function databases

The broadest protein datasets are UniProt, which aims to catalog all known protein sequences (UniProt Consortium, 2017), the Protein Data Bank (Rose et al., 2016), which catalogs all known protein structures, and BRENDA, which catalogs natural enzymes and their functions (Placzek et al., 2016). These databases do not specifically store sequence-function relationships.

Protein engineering experiments have generated a growing collection of well-characterized libraries. Databases that collect and organize specific categories of sequence-function data include ProTherm and ProNit for protein stability and protein-nucleic acid interactions, respectively (Kumar et al., 2006), and SKEMPI (Moal & Fernández-Recio, 2012), AB-Bind (Sirin et al., 2016), and PROXIMATE (Jemimah et al., 2017) for protein-protein complexes. The Protein Mutant Database (Kawabata et al., 1999), an early attempt to catalog the effects of mutations from protein engineering studies across different proteins and functions, has not been updated in over a decade. Currently, Protobank is an actively-maintained and updated database for general protein design and engineering data (Wang et al., 2018).

ProTherm is the oldest and most mature of these databases, but has not been updated since 2013. There have been many efforts to use the sequence-function information in ProTherm to train machine-learning models to predict the effects of mutations on stability (Tian et al., 2010; Li & Fang, 2012; Jia et al., 2015; Capriotti et al., 2005b;a; Cheng et al., 2006; Buske et al., 2009; Liu & Kang, 2012; Pires et al., 2013; Jokinen et al., 2018; Dehouck et al., 2009; Giollo et al., 2014). However, Yang *et al.* found that ProTherm contains many errors and cautions against using entries as training sets without proper validation (Yang et al., 2018b).

Datasets derived from protein engineering experiments tend to be small (10^2 - 10^3 variants) and focused on high-performing variants derived from a small number of sequences. Furthermore, the variants sampled in a protein engineering study are limited by the bias and intent of the study (Wang et al., 2018). This may make it more difficult to generalize models trained on one dataset to other variants of even the same protein generated in different ways or with a different objective. In contrast, datasets of natural variants can be quite large, with examples from many families of proteins. However, the variant distribution in natural datasets is biased by evolution and the fact that not all organisms are equally likely to be studied and sequenced. This can lead to difficulty generalizing to the non-natural sequences often encountered in protein engineering.

2.2. Vector representations of proteins

For machine-learning models to learn about protein sequences, protein sequences must be represented as vectors or matrices of numbers. How each protein sequence is represented determines what can be learned (Domingos, 2012; Bengio et al., 2013). Even the most powerful models produce poor results if an inappropriate representation is used. An ideal encoding would perform well across different proteins and functions, not require alignments, structures, or feature selection, and transfer the information contained in unlabeled sequences to specific prediction tasks. Unfortunately, such an encoding does not exist.

In general, a protein sequence is a string of length L where each residue is chosen from an alphabet of size A . The simplest way to encode such a string is to represent each of the A amino acids as a number. However, the assignment of each residue to a number enforces an ordering on the amino acids that has no physical or biological basis. Instead of representing each position as a single number, a one-hot encoding represents each of the L positions as $A - 1$ zeros and one 1, with the position of the 1 within the series denoting the identity of the amino acid at that position. Given structural information, the identity of

pairs of amino acids within a certain distance in the structure can also be one-hot encoded (Romero et al., 2013b; Bedbrook et al., 2017). Single mutations can also be encoded as a 20-dimensional vector where the original amino acid is denoted by -1, the new amino acid by 1, and all others by zero (Capriotti et al., 2004). One-hot encodings are inherently sparse, memory-inefficient, and high-dimensional. In a one-hot encoding, there is no notion of similarity between sequence or structural elements: they are either identical, or not. Furthermore, one-hot encodings of the primary sequence require that all sequence variants of interest are aligned. This alignment must be updated as sequences are added to the model. Nevertheless, one-hot encodings offer good performance for little complexity and can be considered a good baseline encoding.

A protein can also be encoded by its physical properties by representing each individual amino acid with a collection of physical properties, such as its charge or hydrophobicity, and each protein with a combination of those properties. Properties such as predicted secondary structures can also be used to represent proteins. The challenge is that there are a large number of physical properties that could be used to describe each amino acid or protein. Furthermore, the molecular properties that dictate functional properties are unknown, highly constrained, and dependent on the specific function considered. Therefore, selecting informative properties is challenging because it is difficult to know *a priori* what properties will be predictive for a particular task. Most representations using physical properties also require alignments of the input sequences.

AAIndex (Kawashima et al., 2007) attempts to systematically collect descriptors of protein sequences. AAIndex comprises three sections: AAIndex1 contains 554 amino acid features, AAIndex2 contains 94 amino acid substitution matrices, and AAIndex3 contains 47 amino acid contact potential matrices. ProfFET is another encoding system that considers physical properties of the bulk protein, the individual residues, and sub-sequences of residues (Ofer & Linial, 2015). There have also been attempts to describe each amino acid using two reduced dimensions based on volume and hydrophobicity (Barley et al., 2018) and to combine physical properties with structural information (Qiu et al., 2007; Buske et al., 2009) by encoding each position in the sequence as a combination of the properties of amino acids in its geometric neighborhood.

While there are a vast number of known protein sequences, only a tiny fraction are labeled with measured properties relevant to any specific prediction task. The number of unlabeled sequences will continue to rise as more sequences are deposited into public databases. These unlabeled sequences contain information about the frequency and patterns of amino acids selected by evolution to compose proteins, information that may be helpful across prediction tasks. The simplest examples are BLOSUM (Henikoff & Henikoff, 1992) or AAIndex2-style substitution matrices based on relative amino-acid frequencies. However, more sophisticated continuous vector encodings of sequences can be learned from patterns in unlabeled sequences (Asgari & Mofrad, 2015; Mazzaferro, 2017; Ng, 2017; Kimothi et al., 2016; Yang et al., 2018a; Schwartz et al., 2018). These representations are known as embedded representations. Conceptually, through sequence context, these representations learn a mapping from a space with one indicator (one-hot) dimension per k-mer or protein sequence to a continuous vector space with a much lower dimension such that similar sequences are close together in the continuous space. When modeling large ($> 10^4$ examples) datasets with neural networks, embeddings for individual amino acids or k-mers can be learned simultaneously with the model weights. Learned encodings are low-dimensional, do not require alignments, and may improve performance by transferring information in unlabeled sequences to specific prediction tasks. However, it is difficult to predict which learned encoding will perform well for any given task.

Just as no model will be optimal across all machine learning tasks, there is no universally optimal vectorization method (Wolpert, 1996). Because computational resources are finite, researchers must use a combination of domain expertise and heuristics to select a set of encodings to compare. For small datasets, one-hot encodings offer superior performance to general sets of protein properties (Yang et al., 2018a), although careful feature selection informed by domain knowledge may yield more accurate predictions. If accuracy is insufficient, learned encodings may be able to improve performance. The encoding should ultimately be chosen empirically to maximize predictive performance.

2.3. Models for protein data

A wide range of machine-learning algorithms exist, and no single algorithm is optimal across all learning tasks (Wolpert, 1996). Machine-learning methods can be broadly classified as supervised, unsupervised, or semi-supervised. In supervised learning, the training data consist of inputs and their associated output values (labels). Supervised methods learn a mapping from input space to output space that enables them to accurately predict outputs from new inputs. Supervised learning can be further divided into regression, which aims to predict real-valued outputs, and classification, which aims to predict class membership. In unsupervised learning, the training data consist only of input values. Unsupervised methods learn to find patterns, such as trends or clustering, in the inputs. In semi-supervised learning, the training data consist of inputs, of which a subset has associated output values. Semi-supervised methods leverage information in the unlabeled training

inputs to improve their ability to predict outputs from inputs. Supervised learning is the most developed of these approaches, and is used in the majority of applications to protein engineering. We outline some common supervised machine-learning algorithms below.

The simplest machine-learning models apply a linear transformation of the input features, such as the amino acid at each position, the presence or absence of a mutation (Fox et al., 2007), or blocks of sequence in a library of chimeric proteins made by recombination (Li et al., 2007). Linear models are simple and the learned parameters are easily interpreted by the user. Linear models are commonly used as baseline predictors before more powerful models are tried.

Classification and regression trees (Breiman, 2017) use a decision tree to go from input features (represented as branches) to labels (represented as leaves). Decision tree models are often used in ensemble methods, such as random forests (Breiman, 2001) or boosted trees (Friedman, 2002), which combine multiple models into a more accurate meta-predictor. For small biological datasets ($< 10^4$ training examples), including those often encountered in protein engineering experiments, random forests are a strong and computationally efficient baseline, and have been used to predict thermostability (Tian et al., 2010; Li & Fang, 2012; Jia et al., 2015).

Kernel methods, such as support vector machines (Cortes & Vapnik, 1995) and kernel ridge regression (Nadaraya, 1964), employ a kernel function, which calculates similarities between pairs of inputs, to implicitly project the input features into a high-dimensional feature space without explicitly calculating the coordinates in this new space. The choice of kernel profoundly affects the accuracy of these methods. While general-purpose kernels can be applied to protein inputs, there are also kernels designed for use on proteins, including spectrum and mismatch string kernels (Leslie et al., 2002; 2004), which count the number of shared sub-sequences between two proteins, and weighted decomposition kernels (Jokinen et al., 2018) and other graph kernels, which account for three-dimensional protein structure. Support vector machines have been used to predict protein thermostability (Capriotti et al., 2005b;a; Cheng et al., 2006; Buske et al., 2009; Tian et al., 2010; Jia et al., 2015; Li & Fang, 2012; Liu & Kang, 2012), enantioselectivity (Zaugg et al., 2017), and membrane protein expression (Saladi et al., 2018).

Gaussian process regression and classification combine kernel methods with Bayesian learning to produce probabilistic predictions (Rasmussen & Williams, 2006). These models rigorously capture uncertainty, which, combined with methods from Bayesian optimization, can provide principled ways to guide experimental design in optimizing protein properties. The run-time for exact GP regression scales as the number of training examples cubed, making it unsuitable for large ($> 10^3$) datasets, but there are now fast and accurate approximations (Wilson & Nickisch, 2015). Gaussian processes have been used to predict thermostability (Pires et al., 2013; Romero et al., 2013b; Jokinen et al., 2018), enzyme substrates (Mellor et al., 2016), fluorescence (Saito et al., 2018), membrane localization (Bedbrook et al., 2017), and channelrhodopsin photo-properties (Bedbrook et al., 2018).

Deep learning models, also known as neural networks, stack multiple linear layers connected by non-linear activation functions, allowing them to extract high-level features from structured inputs. Neural networks are well-suited for tasks with large labeled datasets with examples from many protein families, such as protein-nucleic acid binding (Zhang et al., 2015; Alipanahi et al., 2015; Zeng et al., 2016), protein-MHC binding (Hu & Liu, 2017), binding site prediction (Jiménez et al., 2017), protein-ligand binding (Gomes et al., 2017; Mazzaferro, 2017), solubility (Khurana et al., 2018), thermostability (Dehouck et al., 2009; Giollo et al., 2014), subcellular localization (Almagro Armenteros et al., 2017), secondary structure (Sønderby & Winther, 2014), functional class (Szalkai & Grolmusz, 2018; Cao et al., 2017), and even 3D structure (Hopf et al., 2012). Deep learning networks are also particularly useful in metabolic pathway optimization (Oyetunde et al., 2018) and genome annotation (Yip et al., 2013; Jurtz et al., 2017; Yue & Wang, 2018), which have been reviewed elsewhere.

k -nearest-neighbor (Shakhnarovich et al., 2006) methods make predictions on new data points by taking the majority (for classification) or mean (for regression) labels for the k nearest training points. The quality of the predictions can be affected by setting the neighborhood size k as well as the distance metric used to identify the nearest neighbors. Alternatively, predictions can be made as a linear combination of the training labels weighted by their closeness to the test point. Because calculating distances between protein sequences can be non-intuitive (have little meaning in the problem context) or computationally expensive, k -nearest-neighbor methods are not commonly applied to protein datasets.

2.4. Model training and evaluation

All model families have hyperparameters that determine the form of the model. Unlike model parameters, hyperparameters cannot be learned directly from the data. These may be set manually by the practitioner or determined using a procedure such as grid search, simulated annealing, random search, or Bayesian optimization. Hyperparameters may be discrete or continuous. For example, in support vector machines, the type of kernel is a hyperparameter, as are the number of layers and learning rate in a deep neural network. The vectorization method is also a hyperparameter. Even modest changes in the values of hyperparameters can improve or diminish accuracy considerably, and the selection of optimal values is often challenging, as each set of hyperparameters considered may require training a new version of the model.

The key test for a machine-learning model is the ability to accurately predict labels for inputs it has not been trained with. Therefore, when training models by learning their parameters and selecting model hyperparameters, it is necessary to estimate the model's performance on data *not* in the training set. Thus it is essential to initially remove a set of data, called the test set, to be set aside until the absolute end for model evaluation. Typically, the test set comprises approximately 20% of the data. In order to compare models and select hyperparameters during a study, the remaining data should be split into a training set and a validation set. The training set is used to learn model parameters, while the validation set is used to choose between models with different hyperparameters. If the training set is small, cross-validation may be used instead of a constant validation set. In n -fold cross-validation, the training set is partitioned into n complementary subsets. Each subset is then predicted using a model trained on the remaining subsets. Averaging accuracy across the withheld subsets provides an estimate of predictive accuracy over the entire training set.

Care must be taken when selecting the training/validation/test sets that the splits allow an accurate estimate of model performance under the conditions where it will be used. Datasets from mutagenesis studies tend to be small and focused. In this case, the best practice is to train on variants characterized in earlier rounds of mutagenesis and to evaluate model performance on later rounds in order to recapitulate the iterative engineering process. When dealing with large, diverse datasets containing examples from different protein families, the best practice is to ensure that all examples in the validation and test sets are some minimum distance away from all the training examples in order to test the model's ability to generalize to unrelated sequences. In the machine learning context, generalization refers to a model's ability to accurately predict a test set drawn from the same distribution as the training set after training, in contrast to its colloquial meaning of transferring knowledge from one distribution to another, such as from an experimental setting to a real application (Lipton & Steinhardt, 2018).

2.5. Model interpretation

Once a machine-learning model has been built for a certain protein function, the model itself can be a source of knowledge about the underlying physical or biological processes. Model interpretation is the process of determining why or how a model makes its predictions, and allows practitioners to draw biological insights from the knowledge captured by the model. Furthermore, model interpretation can lead to greater user confidence in a model's predictions or a better understanding of when and how a model can fail.

Some machine-learning algorithms are inherently easier to interpret. For example, the learned weights in a linear model indicate which mutations or sequence blocks are beneficial or detrimental for a function of interest (Fox et al., 2007; Liao et al., 2007; Li et al., 2007), and the splits in a decision tree naturally map to human-interpretable information about the features used to make predictions. However, interpretation is less straightforward for complex models with many parameters (such as neural networks) or non-parametric models (such as Gaussian processes). Unfortunately, this complexity is also what gives these models the capacity to make accurate predictions on complex systems. One way to interpret these complex models is to build local or global approximations. Local approximations, such as LIME (local interpretable model-agnostic explanations) (Ribeiro et al., 2016) build a simple, interpretable approximation to the original model in the neighborhood of a single example, as illustrated in Figure 2a. In contrast, a global approximation attempts to simplify the complex model over all examples; for example, a sparse global linear approximation to a Gaussian process regression model can be used to determine which contacting amino-acid pairs are important for channelrhodopsin membrane localization, as shown in Figure 2b (Bedbrook et al., 2017). Alternatively, the layer activations and weights within a neural network can be directly examined for inputs of interest. Convolution weights indicate the relative importance of sequence motifs to the property predicted (a convolution layer scans across a sequence looking for the presence of a learned motif). Activations within attention layers indicate the sections of each input sequence that were most important to the prediction. Figure 2c and d visualize a convolution layer and attention activations, respectively, for predictions of protein subcellular

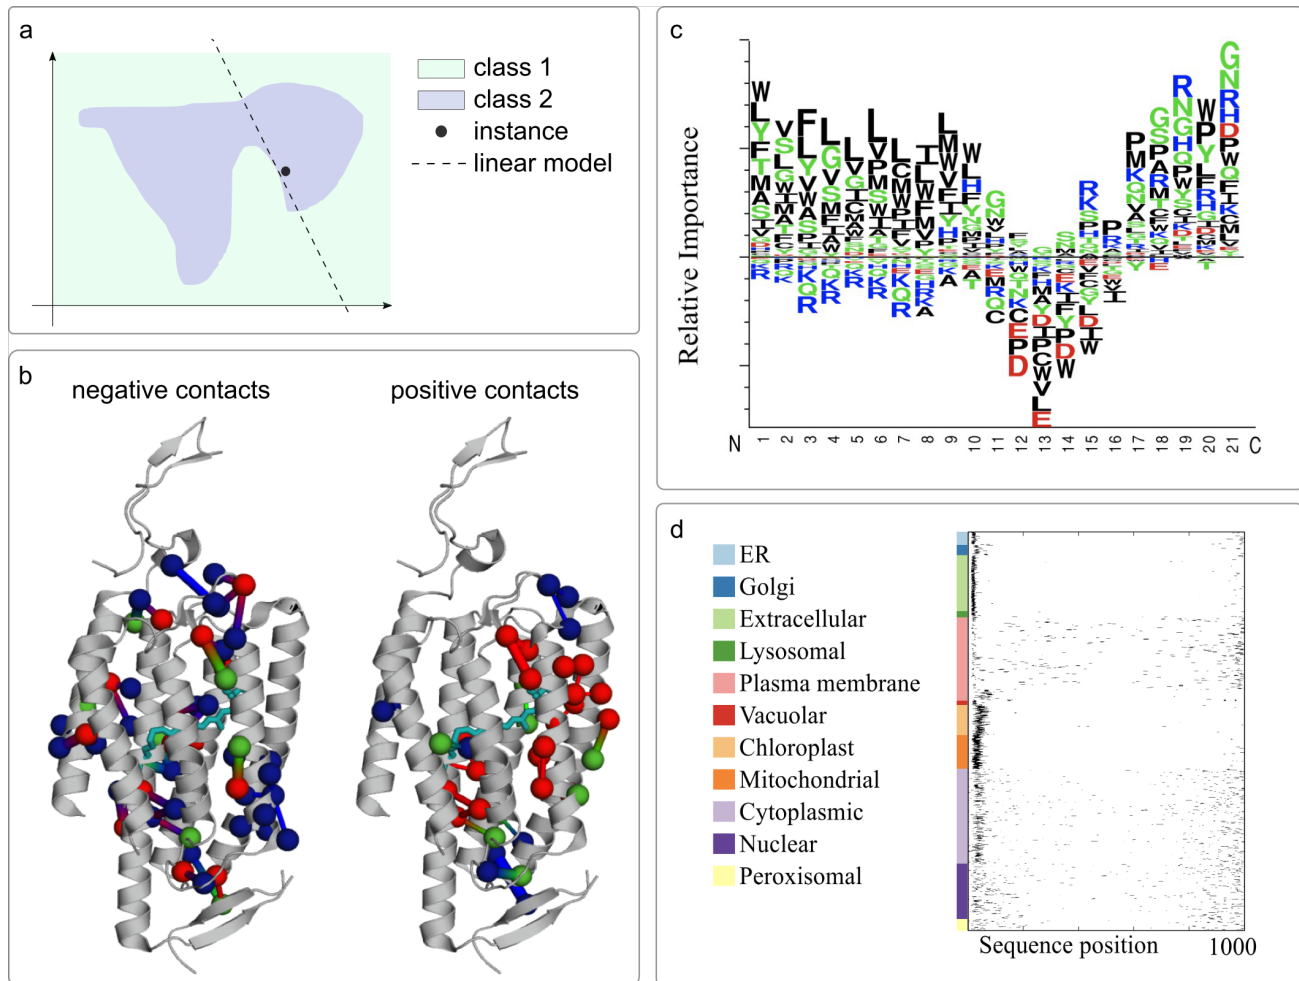


Figure 2. Examples of model interpretation. (a) Local approximation. The original model’s decision function is represented by the violet/green background, and is clearly nonlinear. The black dot is the instance being explained, and a linear model (dashed line) that approximates the original model well in the vicinity of the instance is learned. Note that the explanation is not accurate globally, but is accurate locally around the instance. (b) Most important contacts for predicting channelrhodopsin localization to the plasma membrane. Contacts with the largest positive and negative weights in a Bayesian ridge regression approximation to a Gaussian process regression model are depicted on the channelrhodopsin crystal structure (Bedbrook et al., 2017). Each set of two balls and a stick represents two contacting amino acids. The amino acids are colored according to the source parent. The bound retinal cofactor is shown in cyan. (c) Convolution weights. Visualization of one convolutional filter from a model trained to predict the intracellular location of proteins (Sønderby et al., 2015). The size of each amino acid represents its importance at that position in the k-mer. (d) Attention activations. Importance weights assigned to different regions of proteins in a model trained to predict intracellular location (Sønderby et al., 2015).

localization (Almagro Armenteros et al., 2017).

3. Machine learning as a guide to directed evolution

Directed evolution accumulates beneficial mutations in iterations of mutagenesis and selection or screening. There are an enormous number of ways to mutate any given protein: for a 300-amino acid protein there are 5700 possible single amino acid substitutions and 32,381,700 ways to make just two substitutions with the 20 canonical amino acids. Additionally, screening is expensive, time-consuming, and often limited in throughput. While directed evolution discards information about unimproved variants, machine-learning methods can use this information to expedite evolution and expand the properties that can be optimized by intelligently selecting new variants to screen. The only added costs are in computation and DNA sequencing, both of which are decreasing rapidly. Figure 1 compares directed evolution with and without machine learning as a guide, and Table 1 summarizes some studies that have used machine learning to guide directed evolution.

However, machine learning is not beneficial in all applications. In cases where the fitness landscape is smooth enough (i.e. essentially additive), machine learning may not significantly decrease screening burden or find better variants. In these cases, the added cost of sequencing DNA to form sequence-function relationships is unnecessary. Because one major benefit of machine learning is in reducing the quantity of sequences to test, machine learning will be especially useful in cases where the lack of a high-throughput screen limits or precludes directed evolution.

A machine learning-guided evolution strategy requires a method for generating diversity, a screen to evaluate diversity, a machine-learning model that learns the relationship between sequence and function, and a method to use the model to choose mutations for the next round of evolution. Examples of each are discussed here, with the exception of choosing a machine-learning model, which was outlined above.

3.1. Generating diversity

A straightforward method of generating diversity is to make random mutations throughout the length of the protein by error-prone polymerase chain reaction (PCR). In most directed evolution strategies, beneficial mutations are discovered by screening and accumulated until a satisfactory level of performance is reached. However, multiple mutations with error-prone PCR can occur, in which case mutations that are generally beneficial may be masked by co-occurrence with (more prevalent) detrimental mutations. Fortunately, even a simple linear model can be sufficient to recover accurate classifications of mutations as beneficial, deleterious, or neutral, allowing more beneficial mutations to be fixed more quickly than by directed evolution alone (Fox et al., 2003; Fox, 2005; Fox et al., 2007).

Site-saturation mutagenesis randomizes selected locations within the sequence determined to be most responsible for function or most likely to tolerate mutation. These sites are identified through previous studies or by knowledge of the protein's structure and mechanism. If only one or two pre-identified sites are considered, then machine learning is not necessary, as the entire library can be screened to identify optimal variants. If limited sets of amino acids are tested at each position, more positions can be explored simultaneously (Reetz et al., 2006).

Recombination methods can make larger jumps in sequence space while preserving a large fraction of functional sequences by only considering diversity from within a set of related proteins (Zhao et al., 1998). The sequence elements to be swapped may be chosen randomly or rationally. Structure-guided recombination uses a 3D structure to choose the boundaries of sequence blocks to swap in order to balance maximizing diversity, minimizing disruption to the structure, and having evenly-sized blocks (Endelman et al., 2004).

3.2. Using machine learning to select new variants

An initial set of variants to screen can be selected at random from the library (Fox et al., 2007), to maximize information about the mutations considered (Liao et al., 2007; Govindarajan et al., 2014; Musdal et al., 2017), or to maximize information about the remainder of the library (Romero et al., 2013a; Bedbrook et al., 2017; 2018). Further rounds of learning and selection can be done either by collecting mutations believed to be beneficial or by directly optimizing over sequences in the library. To label mutations as beneficial, linear models of the mutational effects can be learned and the parameters can be directly interpreted to classify mutations as beneficial, neutral, or deleterious. The most beneficial mutations can then be fixed, deleterious mutations can be eliminated from the pool of considered mutations, and new mutations can be added to the pool of considered mutations (Fox et al., 2007).

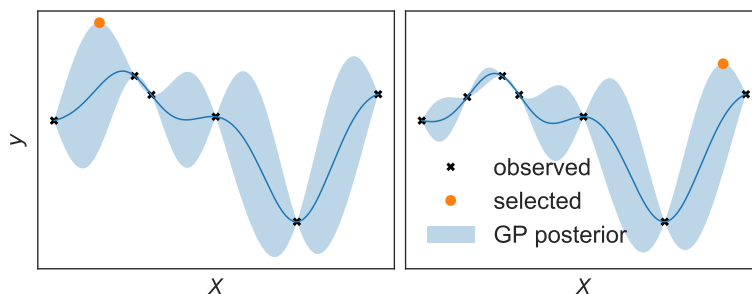


Figure 3. Gaussian Process Upper Confidence Bound algorithm. At each iteration, the next point to be sampled is chosen by maximizing the weighted sum of the posterior mean and standard deviation. This balances exploration and exploitation by exploring points that are both uncertain and have a high posterior mean.

Alternatively, learning and selection can be performed directly over sequences in the library. Typically, this is done using a non-parametric model. Instead of choosing the form of a function and learning parameters that best fit the data, non-parametric models directly learn a function that explains the data well without assuming its form. This makes non-parametric models useful for problems such as predicting protein properties from protein sequence, where the form of the function is not obvious (and indeed may not be possible to write down and parameterize). Non-parametric models used in protein engineering include adaptive substituent reordering, or more commonly, Gaussian process models. Adaptive substituent reordering algorithms (ASRAs) account for epistasis by constructing a function-free model of the underlying fitness landscape. In ASRA, a protein of length L resides in an L -dimensional space and can be described by a vector of substituents x_1, x_2, \dots, x_L . Each substituent x_i is an integer between 1 and 20 indicating the amino acid at that position. Given sequence-function measurements, properties of unmeasured sequences can be estimated by interpolation within this space. The ordering of the substituents, however, defines the smoothness of the space and therefore also the accuracy of the interpolations. In ASRA, an ordering at each site is learned that balances smoothness and training set accuracy (Liang et al., 2005; Feng et al., 2012).

In addition to being non-parametric, Gaussian processes are also probabilistic: they provide an estimate of uncertainty. This allows a principled trade-off between exploiting the information learned from previous iterations and exploring the remainder of the library at each iteration. For example, the Gaussian Process Upper Confidence Bound (GP-UCB) algorithm balances exploration and exploitation by selecting variants that maximize the weighted sum of the predictive mean and standard deviation (Srinivas et al., 2009), and is guaranteed to asymptotically minimize the cumulative regret (difference between sampled variants and best variant) over infinite iterations. Figure 3 demonstrates two iterations of the GP-UCB algorithm. Alternatively, the model and data can be fully exploited using the Gaussian Process Lower Confidence Bound algorithm, which selects variants that maximize the weighted difference between the predictive mean and standard deviation. These approaches have been combined with a structure-guided recombination library to optimize cytochrome P450 thermostability (Romero et al., 2013b), channelrhodopsin localization to mammalian cell membranes (Bedbrook et al., 2017), and channelrhodopsin light-activated conductance (Bedbrook et al., 2018). Because there is no high-throughput screen for the channelrhodopsin properties, they would have been difficult or impossible to optimize with directed evolution alone. Alternately, the GP model can be used to select combinations of mutations in a multi-site site saturation library that has a high probability of containing improved variants (Saito et al., 2018).

4. Conclusions and future directions

Supervised machine-learning methods have already demonstrated their utility in directed protein evolution. The biggest obstacle to future applications of machine learning to protein engineering is a lack of high-quality data. Protein mutation datasets are heavily influenced by experimental conditions such as buffer components, temperature, expression system, and baseline thresholds. While ProtaBank (Wang et al., 2018) is spearheading the development of a centralized collection of these variants with their experimental conditions, there is currently no organized way to collect protein mutation data from various experiments to use as benchmarks in machine learning experiments. The collections that do exist are plagued by inconsistencies and errors (Yang et al., 2018b), and significant resources must be dedicated to maintaining the quality of these databases. As demonstrated by Jokinen and coworkers, one way to circumvent limited labeled data is to augment with

Table 1. Comparison of selected studies using machine learning to guide directed evolution

Protein	Property	<i>n</i> screened	Model
halohydrin dehydrogenase	volumetric productivity	60,000	linear (Fox et al., 2003; Fox, 2005; Fox et al., 2007)
epoxide hydrolase	enantioselectivity	20,000	ASRA (Feng et al., 2012)
proteinase K	activity, heat tolerance	95	linear (Liao et al., 2007)
glutathione transferase	catalytic activity	95	linear (Govindarajan et al., 2014)
glutathione transferase	catalytic activity	95	linear (Musdal et al., 2017)
cytochrome P450	thermostability	~200	GP (Romero et al., 2013b)
channelrhodopsin	membrane localization	~200	GP (Bedbrook et al., 2017)
green fluorescent protein	fluorescence	296	GP (Saito et al., 2018)
channelrhodopsin	spectral properties	119	GP (Bedbrook et al., 2018)

computationally-generated examples (Jokinen et al., 2018). However, accurate physics-based predictors for properties more complicated than stability and binding do not yet exist. Collections of robust protein sequence-function data would allow researchers to benchmark machine-learning methods for protein functions against a variety of proteins and functions.

Deep mutational scanning (Fowler & Fields, 2014) combines a high-throughput screen with next-generation sequencing to generate large sequence-function datasets. In deep mutational scanning, variants are sorted by a selection criterion, such as fluorescence or binding affinity. The sequences are sorted into bins, and the frequency of each mutation is compared before and after selection to infer relative fitness values. There is thus no direct measurement of the property of interest for each variant, and if the gene is longer than the sequencing read length, deconvoluting the interactions of multiple distant mutations requires a DNA barcoding scheme. Nevertheless, deep mutational scanning provides large datasets that map a significant fraction of the single and some double mutants to a fitness measure. Alternatively, a deep mutational scanning dataset may map a complete multi-site site-saturation library. There are now deep mutational scanning datasets for a variety of proteins and properties, including green fluorescent protein (Sarkisyan et al., 2016), amidase (Wrenbeck et al., 2017), β -lactamase (Klesmith et al., 2017), β -glucosidase (Romero et al., 2015), HIV envelope protein (Haddox et al., 2016), influenza nucleoprotein (Ashenberg et al., 2017), influenza hemagglutinin (Doud & Bloom, 2016), PhoQ (Podgoraia & Laub, 2015), GB1 (Wu et al., 2016), the DNA-binding domain of a steroid hormone receptor (Starr et al., 2017), and Gal4 transcription factor (Kitzman et al., 2015). These datasets provide test beds for machine-learning methods that learn to predict the effects of small numbers of mutations or aim to improve on directed evolution for protein optimization (Riesselman et al., 2018).

Large quantities of unlabeled sequence data may also enable machine-learning models to generate artificial protein diversity leading to novel protein functions. Only a tiny fraction of the amino acid landscape encodes a functional protein, and the complete landscape is littered with cliffs and holes, where small changes in sequence result in complete loss of function. Natural and designed proteins are samples from the distribution of functional proteins. Selectively sampling from the distribution of possible proteins would enable large jumps to previously unexplored sections of sequence space that may contain novel functions. Generative models of the distribution of functional proteins would enable these large jumps and provide an attractive alternative to *de novo* design methods (Baker, 2010).

Unlike discriminative models that learn $p(y|x)$ in order to predict labels y given inputs x , generative models learn to generate examples that are similar to those in the training set but are not found in the training set: they learn the generating distribution $p(x)$ for the training data. Tantalizingly, generative models in other fields have been trained to generate new faces (Radford et al., 2015), sketches (Ha & Eck, 2017), and even music (Roberts et al., 2018). Variational autoencoders additionally allow interpolation between examples or the ability to mix and match properties (Kingma & Welling, 2013).

Instead of using neural network models to directly learn the mapping from protein sequence to function, Sinai et al. and Riesselman et al. trained variational autoencoders to learn the distribution of allowed mutations within functional protein families (Sinai et al., 2017; Riesselman et al., 2018). An autoencoder is a neural network that learns to encode an input as a vector (encoding) and then reconstruct the input from the vector (decoding) (Figure 4). By learning an encoding with smaller dimensionality than the original input, the model extracts the most important information from the input. The encoding can then be used as an information-dense input to other learning algorithms. In a variational autoencoder (Kingma & Welling, 2013), the learned encoding is further constrained to encourage the encodings to be densely embedded in the encoding space, allowing smooth interpolations in that space. Applied to aligned families of protein sequences, variational autoencoders can

- Baker, David. An exciting but challenging road ahead for computational enzyme design. *Protein Sci*, 19(10):1817–1819, 2010. doi: 10.1002/pro.481.
- Barley, Mark H, Turner, Nicholas J, and Goodacre, Royston. Improved descriptors for the quantitative structure–activity relationship modeling of peptides and proteins. *J Chem Inf Model*, 58(2):234–243, 2018. doi: 10.1021/acs.jcim.7b00488.
- Bedbrook, Claire N, Yang, Kevin K, Rice, Austin J, Gradinaru, Viviana, and Arnold, Frances H. Machine learning to design integral membrane channelrhodopsins for efficient eukaryotic expression and plasma membrane localization. *PLoS Comput Biol*, 13(10):e1005786, 2017. doi: 10.1371/journal.pcbi.1005786.
- Bedbrook, Claire N, Yang, Kevin K, J, Robinson Elliott, Viviana, Gradinaru, and Arnold, Frances H. Machine learning-guided channelrhodopsin engineering enables minimally-invasive optogenetics. *In preparation*, 2018.
- Bengio, Yoshua, Courville, Aaron, and Vincent, Pascal. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- Breiman, Leo. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- Breiman, Leo. *Classification and Regression Trees*. Routledge, 2017.
- Brookes, David H and Listgarten, Jennifer. Design by adaptive sampling. 2018.
- Buske, Fabian A, Their, Ricarda, Gillam, Elizabeth MJ, and Bodén, Mikael. In silico characterization of protein chimeras: relating sequence and function within the same fold. *Proteins*, 77(1):111–120, 2009. doi: 10.1002/prot.22422.
- Cao, Renzhi, Freitas, Colton, Chan, Leong, Sun, Miao, Jiang, Haiqing, and Chen, Zhangxin. ProLanGO: Protein function prediction using neural machine translation based on a recurrent neural network. *Molecules*, 22(10):1732, 2017. doi: 10.3390/molecules22101732.
- Capriotti, Emidio, Fariselli, Piero, and Casadio, Rita. A neural-network-based method for predicting protein stability changes upon single point mutations. *Bioinformatics*, 20(suppl_1):i63–i68, 2004.
- Capriotti, Emidio, Fariselli, Piero, Calabrese, Remo, and Casadio, Rita. Predicting protein stability changes from sequences using support vector machines. *Bioinformatics*, 21(suppl_2):ii54–ii58, 2005a.
- Capriotti, Emidio, Fariselli, Piero, and Casadio, Rita. I-Mutant2.0: predicting stability changes upon mutation from the protein sequence or structure. *Nucleic Acids Res*, 33(suppl_2):W306–W310, 2005b.
- Cheng, Jianlin, Randall, Arlo, and Baldi, Pierre. Prediction of protein stability changes for single-site mutations using support vector machines. *Proteins*, 62(4):1125–1132, 2006.
- Cortes, Corinna and Vapnik, Vladimir. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.
- Dahiyat, Bassil I and Mayo, Stephen L. De novo protein design: fully automated sequence selection. *Science*, 278(5335): 82–87, 1997.
- Dehouck, Yves, Grosfils, Aline, Folch, Benjamin, Gilis, Dimitri, Bogaerts, Philippe, and Rooman, Marianne. Fast and accurate predictions of protein stability changes upon mutations using statistical potentials and neural networks: PoPMuSiC-2.0. *Bioinformatics*, 25(19):2537–2543, 2009. doi: 10.1093/bioinformatics/btp445.
- Domingos, Pedro. A few useful things to know about machine learning. *Communications of the ACM*, 55(10):78–87, 2012.
- Dou, Jiayi, Doyle, Lindsey, Jr. Greisen, Per, Schena, Alberto, Park, Hahnbeom, Johnsson, Kai, Stoddard, Barry L., and Baker, David. Sampling and energy evaluation challenges in ligand binding protein design. *Protein Sci*, 26(12):2426–2437, 2017. ISSN 09618368. doi: 10.1002/pro.3317.
- Dou, Jiayi, Vorobieva, Anastassia A, Sheffler, William, Doyle, Lindsey A, Park, Hahnbeom, Bick, Matthew J, Mao, Binchen, Foight, Glenna W, Lee, Min Yen, Gagnon, Lauren A, et al. De novo design of a fluorescence-activating β -barrel. *Nature*, pp. 1, 2018. doi: 10.1038/s41586-018-0509-0.
- Doud, Michael B and Bloom, Jesse D. Accurate measurement of the effects of all amino-acid mutations on influenza hemagglutinin. *Viruses*, 8(6):155, 2016. doi: 10.3390/v8060155.

- Drummond, D Allan, Silberg, Jonathan J, Meyer, Michelle M, Wilke, Claus O, and Arnold, Frances H. On the conservative nature of intragenic recombination. *Proc Natl Acad Sci USA*, 102(15):5380–5385, 2005.
- Endelman, Jeffrey B, Silberg, Jonathan J, Wang, Zhen-Gang, and Arnold, Frances H. Site-directed protein recombination as a shortest-path problem. *Protein Eng Des Sel*, 17(7):589–594, 2004.
- Feng, Xiaojiang, Sanchis, Joaquin, Reetz, Manfred T, and Rabitz, Herschel. Enhancing the efficiency of directed evolution in focused enzyme libraries by the adaptive substituent reordering algorithm. *Chemistry*, 18(18):5646–5654, 2012. doi: 10.1002/chem.201103811.
- Fowler, Douglas M and Fields, Stanley. Deep mutational scanning: a new style of protein science. *Nat Methods*, 11(8):801, 2014. doi: 10.1038/nmeth.3027.
Deep mutational scanning can quickly generate large sequence-function datasets.
- Fox, Richard. Directed molecular evolution by machine learning and the influence of nonlinear interactions. *J Theor Biol*, 234(2):187–199, 2005.
- Fox, Richard, Roy, Ajoy, Govindarajan, Sridhar, Minshull, Jeremy, Gustafsson, Claes, Jones, Jennifer T, and Emig, Robin. Optimizing the search algorithm for protein engineering by directed evolution. *Protein Eng*, 16(8):589–597, 2003.
This study is the first to use machine learning to guide directed evolution.
- Fox, Richard J, Davis, S Christopher, Mundorff, Emily C, Newman, Lisa M, Gavrilovic, Vesna, Ma, Steven K, Chung, Loleta M, Ching, Charlene, Tam, Sarena, Muley, Sheela, Grate, John, Gruber, John, Whitman, John C, Sheldon, Roger A, and Huisman, Gjal W. Improving catalytic function by ProSAR-driven enzyme evolution. *Nat Biotechnol*, 25(3):338, 2007.
- Friedman, Jerome H. Stochastic gradient boosting. *Comput Stat Data An*, 38(4):367–378, 2002. doi: 10.1016/S0167-9473(01)00065-2.
- Garcia-Borràs, Marc, Houk, Kendall N, and Jiménez-Osés, Gonzalo. Computational design of protein function. *Computational Tools for Chemical Biology*, 3:87, 2017. doi: 10.1039/9781788010139-00087.
- Giollo, Manuel, Martin, Alberto JM, Walsh, Ian, Ferrari, Carlo, and Tosatto, Silvio CE. NeEMO: a method using residue interaction networks to improve prediction of protein stability upon mutation. *BMC Genomics*, 15(4):S7, 2014. doi: 10.1186/1471-2164-15-S4-S7.
- Gomes, Joseph, Ramsundar, Bharath, Feinberg, Evan N, and Pande, Vijay S. Atomic convolutional networks for predicting protein-ligand binding affinity. 2017.
- Goodfellow, Ian, Pouget-Abadie, Jean, Mirza, Mehdi, Xu, Bing, Warde-Farley, David, Ozair, Sherjil, Courville, Aaron, and Bengio, Yoshua. Generative adversarial nets. In *Adv Neur In*, pp. 2672–2680, 2014.
- Govindarajan, Sridhar, Mannervik, Bengt, Silverman, Joshua A, Wright, Kathy, Regitsky, Drew, Hegazy, Usama, Purcell, Thomas J, Welch, Mark, Minshull, Jeremy, and Gustafsson, Claes. Mapping of amino acid substitutions conferring herbicide resistance in wheat glutathione transferase. *ACS Synth Biol*, 4(3):221–227, 2014. doi: 10.1021/sb500242x.
- Gupta, Anvita and Zou, James. Feedback GAN (FBGAN) for DNA: a novel feedback-loop architecture for optimizing protein functions. 2018.
- Ha, David and Eck, Douglas. A neural representation of sketch drawings. 2017.
- Haddox, Hugh K, Dingens, Adam S, and Bloom, Jesse D. Experimental estimation of the effects of all amino-acid mutations to HIV’s envelope protein on viral replication in cell culture. *PLoS Pathog*, 12(12):e1006114, 2016. doi: 10.1371/journal.ppat.1006114.
- Hastie, Tibshirani and Tibshirani, R. *The Elements of Statistical Learning; Data Mining, Inference and Prediction*. Springer, New York, 2008.
- Henikoff, Steven and Henikoff, Jorja G. Amino acid substitution matrices from protein blocks. *Proc Natl Acad Sci USA*, 89(22):10915–10919, 1992.

- Hopf, Thomas A., Colwell, Lucy J., Sheridan, Robert, Rost, Burkhard, Sander, Chris, and Marks, Debora S. Three-dimensional structures of membrane proteins from genomic sequencing. *Cell*, 149(7):1607–1621, 2012. ISSN 00928674. doi: 10.1016/j.cell.2012.04.012.
- Hu, Jianjun and Liu, Zhonghao. DeepMHC: Deep convolutional neural networks for high-performance peptide-MHC binding affinity prediction. pp. 239236, 2017. doi: 10.1101/239236.
- Huang, Possu, Boyken, Scott E, and Baker, David. The coming of age of de novo protein design. *Nature*, 537(7620):320, 2016. doi: 10.1038/nature1994.
- Jemimah, Sherlyn, Yugandhar, K, and Michael Gromiha, M. PROXiMATE: a database of mutant protein–protein complex thermodynamics and kinetics. *Bioinformatics*, 33(17):2787–2788, 2017. doi: 10.1093/bioinformatics/btx312.
- Jia, Lei, Yarlagadda, Ramya, and Reed, Charles C. Structure based thermostability prediction models for protein single point mutations with machine learning tools. *PLoS One*, 10(9):e0138022, 2015. doi: 10.1371/journal.pone.0138022.
- Jiménez, J, Doerr, S, Martínez-Rosell, G, Rose, AS, and De Fabritiis, G. DeepSite: protein-binding site predictor using 3D-convolutional neural networks. *Bioinformatics*, 33(19):3036–3042, 2017. doi: 10.1093/bioinformatics/btx350.
- Jokinen, Emmi, Heinonen, Markus, and Lähdesmäki, Harri. mGPFusion: Predicting protein stability changes with Gaussian process kernel learning and data fusion. *Bioinformatics*, 34(13):i274–i283, 2018. doi: 10.1093/bioinformatics/bty238.
- Jurtz, Vanessa Isabell, Johansen, Alexander Rosenberg, Nielsen, Morten, Almagro Armenteros, Jose Juan, Nielsen, Henrik, Sønderby, Casper Kaae, Winther, Ole, and Sønderby, Søren Kaae. An introduction to deep learning on biological sequence data: examples and solutions. *Bioinformatics*, 33(22):3685–3690, 2017. doi: 10.1093/bioinformatics/btx531.
- Kawabata, Takeshi, Ota, Motonori, and Nishikawa, Ken. The protein mutant database. *Nucleic Acids Res*, 27(1):355–357, 1999.
- Kawashima, Shuichi, Pokarowski, Piotr, Pokarowska, Maria, Kolinski, Andrzej, Katayama, Toshiaki, and Kanehisa, Minoru. AAindex: amino acid index database, progress report 2008. *Nucleic Acids Res*, 36(suppl_1):D202–D205, 2007.
- Khurana, Sameer, Rawi, Reda, Kunji, Khalid, Chuang, Gwo-Yu, Bensmail, Halima, Mall, Raghvendra, and Valencia, Alfonso. DeepSol: A deep learning framework for sequence-based protein solubility prediction. *Bioinformatics*, 1:9, 2018. doi: 10.1093/bioinformatics/bty166.
- Kimothi, Dhananjay, Soni, Akshay, Biyani, Pravesh, and Hogan, James M. Distributed representations for biological sequence analysis. 2016.
- Kingma, Diederik P and Welling, Max. Auto-encoding variational Bayes. 2013.
- Kitzman, Jacob O, Starita, Lea M, Lo, Russell S, Fields, Stanley, and Shendure, Jay. Massively parallel single-amino-acid mutagenesis. *Nat Methods*, 12(3):203, 2015. doi: 10.1038/nmeth.3223.
- Klesmith, Justin R, Bacik, John-Paul, Wrenbeck, Emily E, Michalczyk, Ryszard, and Whitehead, Timothy A. Trade-offs between enzyme fitness and solubility illuminated by deep mutational scanning. *Proc Natl Acad Sci USA*, 114(9):2265–2270, 2017. doi: 10.1073/pnas.1614437114.
- Kumar, MD Shaji, Bava, K Abdulla, Gromiha, M Michael, Prabakaran, Ponraj, Kitajima, Koji, Uedaira, Hatsuho, and Sarai, Akinori. ProTherm and ProNIT: thermodynamic databases for proteins and protein–nucleic acid interactions. *Nucleic Acids Res*, 34(suppl_1):D204–D206, 2006.
- Leslie, Christina, Eskin, Eleazar, and Noble, William Stafford. The spectrum kernel: A string kernel for SVM protein classification. pp. 564–575, 2002.
- Leslie, Christina S, Eskin, Eleazar, Cohen, Adiel, Weston, Jason, and Noble, William Stafford. Mismatch string kernels for discriminative protein classification. *Bioinformatics*, 20(4):467–476, 2004.
- Li, Yougen, Drummond, D Allan, Sawayama, Andrew M, Snow, Christopher D, Bloom, Jesse D, and Arnold, Frances H. A diverse family of thermostable cytochrome P450s created by recombination of stabilizing fragments. *Nat Biotechnol*, 25(9):1051, 2007.

- Li, Yunqi and Fang, Jianwen. PROTS-RF: a robust model for predicting mutation-induced protein stability changes. *PLoS One*, 7(10):e47247, 2012. doi: 10.1371/journal.pone.0047247.
- Liang, Fan, Feng, Xiao-Jiang, Lowry, Michael, and Rabitz, Herschel. Maximal use of minimal libraries through the adaptive substituent reordering algorithm. *J Phys Chem B*, 109(12):5842–5854, 2005.
- Liao, Jun, Warmuth, Manfred K, Govindarajan, Sridhar, Ness, Jon E, Wang, Rebecca P, Gustafsson, Claes, and Minshull, Jeremy. Engineering proteinase K using machine learning and synthetic genes. *BMC Biotechnol*, 7(1):16, 2007.
- Lipton, Zachary C and Steinhardt, Jacob. Troubling Trends in Machine Learning Scholarship. pp. 1–15, 2018.
- Liu, Jianguo and Kang, Xianjiang. Grading amino acid properties increased accuracies of single point mutation on protein stability prediction. *BMC Bioinformatics*, 13(1):44, 2012. doi: 10.1186/1471-2105-13-44.
- Mandecki, Wlodek. The game of chess and searches in protein sequence space. *Trends Biotech*, 16(5):200–202, 1998.
- Mazzaferro, Carlo. Predicting protein binding affinity with word embeddings and recurrent neural networks. pp. 128223, 2017. doi: 10.1101/128223.
- Mellor, Joseph, Grigoras, Ioana, Carbonell, Pablo, and Faulon, Jean-Loup. Semisupervised Gaussian process for automated enzyme search. *ACS Synth Bio*, 5(6):518–528, 2016. doi: 10.1021/acssynbio.5b00294.
- Moal, Iain H and Fernández-Recio, Juan. SKEMPI: a Structural Kinetic and Energetic database of mutant protein interactions and its use in empirical models. *Bioinformatics*, 28(20):2600–2607, 2012. doi: 10.1093/bioinformatics/bts489.
- Müller, Alex T, Hiss, Jan A, and Schneider, Gisbert. Recurrent neural network model for constructive peptide design. *J Chem Inf Model*, 2018. doi: 10.1021/acs.jcim.7b00414.
- Murphy, Kevin. *Machine learning, a probabilistic perspective*. MIT Press, 2012.
Murphy’s textbook provides a thorough introduction to modern machine learning.
- Musdal, Yaman, Govindarajan, Sridhar, and Mannervik, Bengt. Exploring sequence-function space of a poplar glutathione transferase using designed information-rich gene variants. *Protein Eng Des Sel*, 30(8):543–549, 2017. doi: 10.1093/protein/gzx045.
- Nadaraya, E. On estimating regression. *Theor Probab Appl*, 9(1):141–142, 1964. doi: 10.1137/1109020.
- Ng, Patrick. dna2vec: Consistent vector representations of variable-length k-mers. 2017.
- Ofer, Dan and Linial, Michal. PROFET: Feature engineering captures high-level protein functions. *Bioinformatics*, 31(21): 3429–3436, 2015. doi: 10.1093/bioinformatics/btv345.
- Orr, H Allen. The distribution of fitness effects among beneficial mutations in Fisher’s geometric model of adaptation. *J Theor Biol*, 238(2):279–285, 2006.
- Oyetunde, Tolutola, Bao, Forrest Sheng, Chen, Jiung-Wen, Martin, Hector Garcia, and Tang, Yinjie J. Leveraging knowledge engineering and machine learning for microbial bio-manufacturing. *Biotechnol Adv*, 36(4):1308 – 1315, 2018. ISSN 0734-9750. doi: 10.1016/j.biotechadv.2018.04.008.
- Pierce, Niles A and Winfree, Erik. Protein design is NP-hard. *Protein Eng*, 15(10):779–782, 2002.
- Pires, Douglas EV, Ascher, David B, and Blundell, Tom L. mCSM: predicting the effects of mutations in proteins using graph-based signatures. *Bioinformatics*, 30(3):335–342, 2013. doi: 10.1093/bioinformatics/btt691.
- Placzek, Sandra, Schomburg, Ida, Chang, Antje, Jeske, Lisa, Ulbrich, Marcus, Tillack, Jana, and Schomburg, Dietmar. BRENDA in 2017: new perspectives and new tools in BRENDA. *Nucleic Acids Res*, pp. D380–D388, 2016.
- Podgornaia, Anna I and Laub, Michael T. Pervasive degeneracy and epistasis in a protein-protein interface. *Science*, 347 (6222):673–677, 2015. doi: 10.1126/science.1257360.
- Qiu, Jian, Hue, Martial, Ben-Hur, Asa, Vert, Jean-Philippe, and Noble, William Stafford. A structural alignment kernel for protein structures. *Bioinformatics*, 23(9):1090–1098, 2007.

- Radford, Alec, Metz, Luke, and Chintala, Soumith. Unsupervised representation learning with deep convolutional generative adversarial networks. 2015.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Reetz, Manfred T, Wang, Li-Wen, and Bocola, Marco. Directed evolution of enantioselective enzymes: iterative cycles of CASTing for probing protein-sequence space. *Angew Chem Int Ed Engl*, 118(8):1258–1263, 2006.
- Ribeiro, Marco Tulio, Singh, Sameer, and Guestrin, Carlos. Why should I trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1135–1144. ACM, 2016. doi: 10.1145/2939672.2939778.
- Riesselman, Adam J, Ingraham, John B, and Marks, Debora S. Deep generative models of genetic variation capture mutation effects. *Nat Methods*, 15:816–822, 2018. doi: 10.1038/s41592-018-0138-4.
This study predicts the effects of mutations without using any labeled data.
- Roberts, Adam, Engel, Jesse, Raffel, Colin, Hawthorne, Curtis, and Eck, Douglas. A hierarchical latent vector model for learning long-term structure in music. 2018.
- Romero, Philip A and Arnold, Frances H. Exploring protein fitness landscapes by directed evolution. *Nat Rev Mol Cell Biol*, 10(12):866, 2009. doi: 10.1038/nrm2805.
- Romero, Philip A, Krause, Andreas, and Arnold, Frances H. Navigating the protein fitness landscape with Gaussian processes. *Proc Natl Acad Sci USA*, 110(3):E193–E201, 2013a. doi: 10.1073/pnas.1215251110.
- Romero, Philip A, Krause, Andreas, and Arnold, Frances H. Navigating the protein fitness landscape with Gaussian processes. *Proc Natl Acad Sci USA*, 110(3):E193–E201, 2013b. doi: 10.1073/pnas.1215251110.
This is the first study to combine SCHEMA recombination with the GP-UCB algorithm to optimize a protein property.
- Romero, Philip A, Tran, Tuan M, and Abate, Adam R. Dissecting enzyme function with microfluidic-based deep mutational scanning. *Proc Natl Acad Sci USA*, 112(23):7159–7164, 2015. doi: 10.1073/pnas.1422285112.
- Rose, Peter W, Prlić, Andreas, Altunkaya, Ali, Bi, Chunxiao, Bradley, Anthony R, Christie, Cole H, Costanzo, Luigi Di, Duarte, Jose M, Dutta, Shuchismita, Feng, Zukang, Green, Rachel K, Goodsell, David S, Hudson, Brian, Kalro, Tara, Lowe, Robert, Peisach, Ezra, Randle, Christopher, Rose, Alexander S, Shao, Chenghua, Tao, Yi-Ping, Valasatava, Yana, Voigt, Maria, Westbrook, John D, Woo, Jesse, Yang, Huangwang, Young, Jasmine Y, Zardecki, Christine, Berman, Helen M, and Burley, Stephen K. The RCSB protein data bank: integrative view of protein, gene and 3D structural information. *Nucleic Acids Res*, 2016. doi: 10.1093/nar/gkw952.
- Saito, Yutaka, Oikawa, Misaki, Nakazawa, Hikaru, Niide, Teppei, Kameda, Tomoshi, Tsuda, Koji, and Umetsu, Mitsuo. Machine-learning-guided mutagenesis for directed evolution of fluorescent proteins. *ACS Synth Bio*, 2018. doi: 10.1021/acssynbio.8b00155.
- Saladi, Shyam M, Javed, Nauman, Müller, Axel, and Clemons, William M. A statistical model for improved membrane protein expression using sequence-derived features. *J Biol Chem*, 293(13):4913–4927, 2018. doi: 10.1074/jbc.RA117.001052.
- Sarkisyan, Karen S, Bolotin, Dmitry A, Meer, Margarita V, Usmanova, Dinara R, Mishin, Alexander S, Sharonov, George V, Ivankov, Dmitry N, Bozhanova, Nina G, Baranov, Mikhail S, Soylemez, Onuralp, Bogatyreva, Natalya S, Vlasov, Peter K, Egorov, Evgeny S, Logacheva, Maria D, Kondrashov, Alexey S, Chudakov, Dmitry M, Putintseva, Ekaterina V, Mamedov, Ilgar Z, Tawfik, Dan S, Lukyanov, Konstantin A, and Kondrashov, Fyodor A. Local fitness landscape of the green fluorescent protein. *Nature*, 533(7603):397, 2016. doi: 10.1038/nature17995.
- Schwartz, Ariel S, Hannum, Gregory J, Dwiell, Zach R, Smoot, Michael E, Grant, Ana R, Knight, Jason M, Becker, Scott A, Eads, Jonathan R, LaFave, Matthew C, Eavani, Harini, Liu, Yinyin, Bansal, Arjun K, and Richardson, Toby H. Deep semantic protein representation for annotation, discovery, and engineering. *bioRxiv*, pp. 365965, 2018. doi: 10.1101/365965.
- Shakhnarovich, Gregory, Darrell, Trevor, and Indyk, Piotr. *Nearest-neighbor methods in learning and vision: theory and practice*. The MIT press, 2006.

- Siegel, Justin B., Zanghellini, Alexandre, Lovick, Helena M., Kiss, Gert, Lambert, Abigail R., St.Clair, Jennifer L., Gallaher, Jasmine L., Hilvert, Donald, Gelb, Michael H., Stoddard, Barry L., Houk, Kendall N., Michael, Forrest E., and Baker, David. Computational design of an enzyme catalyst for a stereoselective bimolecular diels-alder reaction. *Science*, 329 (5989):309–313, 2010. ISSN 00368075. doi: 10.1126/science.1190239.
- Sinai, Sam, Kelsic, Eric, Church, George M, and Nowak, Martin A. Variational auto-encoding of protein sequences. 2017.
- Sirin, Sarah, Apgar, James R, Bennett, Eric M, and Keating, Amy E. AB-Bind: antibody binding mutational database for computational affinity predictions. *Protein Sci*, 25(2):393–409, 2016. doi: 10.1002/pro.2829.
- Smith, John Maynard. Natural selection and the concept of a protein space. *Nature*, 225(5232):563, 1970.
- Sønderby, Søren Kaae and Winther, Ole. Protein secondary structure prediction with long short term memory networks. 2014.
- Sønderby, Søren Kaae, Sønderby, Casper Kaae, Nielsen, Henrik, and Winther, Ole. Convolutional LSTM networks for subcellular localization of proteins. In *International Conference on Algorithms for Computational Biology*, pp. 68–80. Springer, 2015. doi: 10.1007/978-3-319-21233-3_6.
- Srinivas, Niranjan, Krause, Andreas, Kakade, Sham M, and Seeger, Matthias. Gaussian process optimization in the bandit setting: No regret and experimental design. 2009.
- Starr, Tyler N, Picton, Lora K, and Thornton, Joseph W. Alternative evolutionary histories in the sequence space of an ancient protein. *Nature*, 549(7672):409, 2017. doi: 10.1038/nature23902.
- Szalkai, Balázs and Grolmusz, Vince. Near perfect protein multi-label classification with deep neural networks. *Methods*, 132:50–56, 2018. doi: 10.1016/j.jymeth.2017.06.034.
- Tian, Jian, Wu, Ningfeng, Chu, Xiaoyu, and Fan, Yunliu. Predicting changes in protein thermostability brought about by single- or multi-site mutations. *BMC Bioinformatics*, 11(1):370, 2010. doi: 10.1186/1471-2105-11-37.
- UniProt Consortium. UniProt: the universal protein knowledgebase. *Nucleic Acids Res*, 45(D1):D158–D169, 2017. doi: 10.1093/nar/gkw1099.
- Wang, Connie Y, Chang, Paul M, Ary, Marie L, Allen, Benjamin D, Chica, Roberto A, Mayo, Stephen L, and Olafson, Barry D. ProtaBank: A repository for protein design and engineering data. *Protein Sci*, 27(6):1113–1124, 2018. doi: 10.1002/pro.3406.
- Wilson, Andrew Gordon and Nickisch, Hannes. Kernel interpolation for scalable structured Gaussian processes (KISS-GP). In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pp. 1775–1784, 2015.
- Wolpert, David H. The lack of a priori distinctions between learning algorithms. *Neural Comput*, 8(7):1341–1390, 1996.
- Wrenbeck, Emily E, Azouz, Laura R, and Whitehead, Timothy A. Single-mutation fitness landscapes for an enzyme on multiple substrates reveal specificity is globally encoded. *Nat Commun*, 8:15695, 2017. doi: 10.1038/ncomms15695.
- Wu, Nicholas C, Dai, Lei, Olson, C Anders, Lloyd-Smith, James O, and Sun, Ren. Adaptation in protein fitness landscapes is facilitated by indirect paths. *Elife*, 5:e16965, 2016. doi: 10.7554/eLife.16965.
- Yang, Kevin, Wu, Zachary, Bedbrook, Claire, and Arnold, Frances. Learned protein embeddings for machine learning. *Bioinformatics*, 2018a. doi: 10.1093/bioinformatics/bty178.
- Yang, Yang, Urolagin, Siddhaling, Niroula, Abhishek, Ding, Xuesong, Shen, Bairong, and Vihinen, Mauno. PON-tstab: Protein Variant Stability Predictor. importance of training data quality. *Int J Mol Sci*, 19(4):1009, 2018b. doi: 10.3390/ijms19041009.
- Yip, Kevin Y., Cheng, Chao, and Gerstein, Mark. Machine learning and genome annotation: a match meant to be? *Genome Biol.*, 14(5):205, 5 2013. ISSN 1474-760X. doi: 10.1186/gb-2013-14-5-205.
- Yue, Tianwei and Wang, Haohan. Deep learning for genomics: A concise overview. 2018.

- Zaugg, Julian, Gumulya, Yosephine, Malde, Alpeshkumar K, and Bodén, Mikael. Learning epistatic interactions from sequence-activity data to predict enantioselectivity. *J Comput Aided Mol Des*, 31(12):1085–1096, 2017. doi: 10.1007/s10822-017-0090-x.
- Zeng, Haoyang, Edwards, Matthew D, Liu, Ge, and Gifford, David K. Convolutional neural network architectures for predicting DNA-protein binding. *Bioinformatics*, 32(12):i121–i127, 2016. doi: 10.1093/bioinformatics/btw255.
- Zhang, Sai, Zhou, Jingtian, Hu, Hailin, Gong, Haipeng, Chen, Ligong, Cheng, Chao, and Zeng, Jianyang. A deep learning framework for modeling structural features of RNA-binding protein targets. *Nucleic Acids Res*, 44(4):e32–e32, 2015. doi: 10.1093/nar/gkv1025.
- Zhao, Huimin, Giver, Lori, Shao, Zhixin, Affholter, Joseph A., and Arnold, Frances H. Molecular evolution by staggered extension process (StEP) in vitro recombination. *Nat Biotechnol*, 16(3):258–261, 1998. ISSN 10870156. doi: 10.1038/nbt0398-258.