



# Machine learning to navigate fitness landscapes for protein engineering

Chase R Freschlin<sup>1</sup>, Sarah A Fahlberg<sup>1</sup> and Philip A Romero<sup>1,2</sup>

Machine learning (ML) is revolutionizing our ability to understand and predict the complex relationships between protein sequence, structure, and function. Predictive sequence–function models are enabling protein engineers to efficiently search the sequence space for useful proteins with broad applications in biotechnology. In this review, we highlight the recent advances in applying ML to protein engineering. We discuss supervised learning methods that infer the sequence–function mapping from experimental data and new sequence representation strategies for data-efficient modeling. We then describe the various ways in which ML can be incorporated into protein engineering workflows, including purely *in silico* searches, ML-assisted directed evolution, and generative models that can learn the underlying distribution of the protein function in a sequence space. ML-driven protein engineering will become increasingly powerful with continued advances in high-throughput data generation, data science, and deep learning.

## Addresses

<sup>1</sup> Department of Biochemistry, University of Wisconsin–Madison, Madison, WI, USA

<sup>2</sup> Department of Chemical & Biological Engineering, University of Wisconsin–Madison, Madison, WI, USA

Corresponding author: Philip A Romero ([promero2@wisc.edu](mailto:promero2@wisc.edu))

Current Opinion in Biotechnology 2022, 75:102713

This review comes from a themed issue on **Analytical Biotechnology**

Edited by **Peng Xu** and **Kang Zhou**

<https://doi.org/10.1016/j.copbio.2022.102713>

0958-1669/© 2022 Elsevier Ltd. All rights reserved.

## Introduction

Evolution has shaped proteins to perform complex chemical and biological functions with exceptional proficiency, accuracy, and specificity. These naturally occurring proteins represent a huge potential to solve some of the world's most pressing problems. Protein engineering aims to realize this potential by altering protein functions for applications in biotechnology, industry, and medicine [1–3]. Amino acid sequences encode the overall function and properties of proteins.

The relationship between a protein sequence and function can be imagined by the sequence–function 'fitness landscape', wherein protein sequences are mapped to a corresponding fitness value that represents a measurable protein property like catalytic activity or thermostability [4]. Protein engineering aims to search this landscape for high-fitness sequences.

Directed evolution navigates this landscape through iterative rounds of mutation, high-throughput functional assays, and selection to 'walk' uphill toward more fit sequences. Despite its widespread utility and success, directed evolution is resource and time intensive, and the search process is largely ignorant of the underlying sequence–function landscape. This local search is inefficient and can ultimately limit the fitness achieved in a protein engineering experiment. Machine learning (ML) has recently emerged as a highly effective method for inferring the sequence–function landscape [5–7]. ML-guided protein engineering can discover highly fit sequences more efficiently and with less experimental screening than traditional directed evolution [8]. In this review, we highlight recent approaches that leverage ML to explore, exploit, and expand the functional sequence space to engineer proteins for improved function.

## Supervised learning for modeling the protein–fitness landscape

Rational protein engineering requires a detailed, quantitative understanding of the complex biophysical interactions underlying protein sequence and function. ML can infer these interactions from experimental data to generate predictive models for protein design. While a diverse range of ML models are suitable for this task, here we focus our discussion on the recent advances in deep neural networks owing to their state-of-the-art performance and flexibility. We also discuss the importance of model evaluation and how this influences our ability to compare different modeling approaches.

Supervised deep neural networks combine multiple layers of nonlinear functions to learn highly complex input–output relationships. There are several commonly used architectures capable of modeling the sequence–function landscape, and each one carries inductive biases that make different assumptions about the relationship. Common architectures include multilayer perceptrons, which learn the nonlinear

interactions between inputs, convolutional neural networks (CNNs), which consider information within a local 'receptive field' in the input space, recurrent neural networks (RNNs), which process information in a sequential manner, and transformers, which can learn specific long-range interactions between inputs using attention mechanisms. Both CNNs and RNNs outperform the baseline linear models for protein function prediction [9–12]. CNNs also show promise for predicting the effects of mutations that were not experimentally tested. This capacity may be useful when characterizing mutations is costly or when expanding the set of beneficial mutations that can be used to design a protein [9,13].

Supervised deep learning relies on large sequence–function datasets for model training and evaluation, but many proteins have limited (10–100s) functional data, owing to the lack of high-throughput functional assays. One method to reduce the data requirement for training is to apply more informative protein representations to supervised learning tasks [14]. Rather than simply representing a protein as a sequence of amino acids, which is very high dimensional, it is possible to devise other protein representations that may capture key physiochemical properties such as total charge, evolutionary features such as site conservation, and structural characteristics such as solvent-accessible surface areas. A carefully designed set of features can distill the essential aspects of protein function and enable learning with less data [15,16]. A more recent and data-driven approach uses unsupervised ML to learn informative protein representations from the large amounts of unlabeled sequence data available in databases such as UniProt [17–22]. These models effectively learn condensed representations of protein sequences by identifying correlations and patterns within protein families. These low-dimensional protein representations can be used as inputs for supervised models to improve predictive accuracy and enable learning from smaller sequence–function datasets [8,23]. Variational autoencoders (VAEs), long short-term memory RNNs, and transformers trained on natural sequence information consistently show top performance in learning effective protein representations for supervised ML [19,23–26].

The potential for ML in protein engineering has resulted in rapid development of new ML models and strategies, which demand rigorous performance benchmarks and metrics for model comparison. The differences in datasets, train-test splits, and evaluation metrics make it nontrivial to compare ML modeling approaches, as all of these factors impact the model performance. Recent efforts to standardize the evaluation of learned representations aim to address these issues by outlining the standard datasets, splits, and tasks for assessing representation learning models [19], but evaluation of

supervised models has yet to be standardized. An important factor in standardizing evaluation of supervised models is assessing model performance. Models are often evaluated by their average prediction error, but there are more relevant metrics for evaluating the models for protein engineering applications. To address this, Mater et al. [13] devised five metrics to evaluate a model's ability to handle extrapolation, epistasis, sparse data, and variability in sequence length. Incorporating these kinds of tasks into a standard benchmark for supervised ML models will allow protein engineers to better evaluate ML models and find a model that best suits their needs. Finally, a common goal of the recent supervised learning papers has been to create models and training protocols that are more accessible for comparing ML approaches and for use by non-ML experts [9,11,27–29]. As the field of ML-based protein engineering expands, it is important to make these tools readily accessible to all.

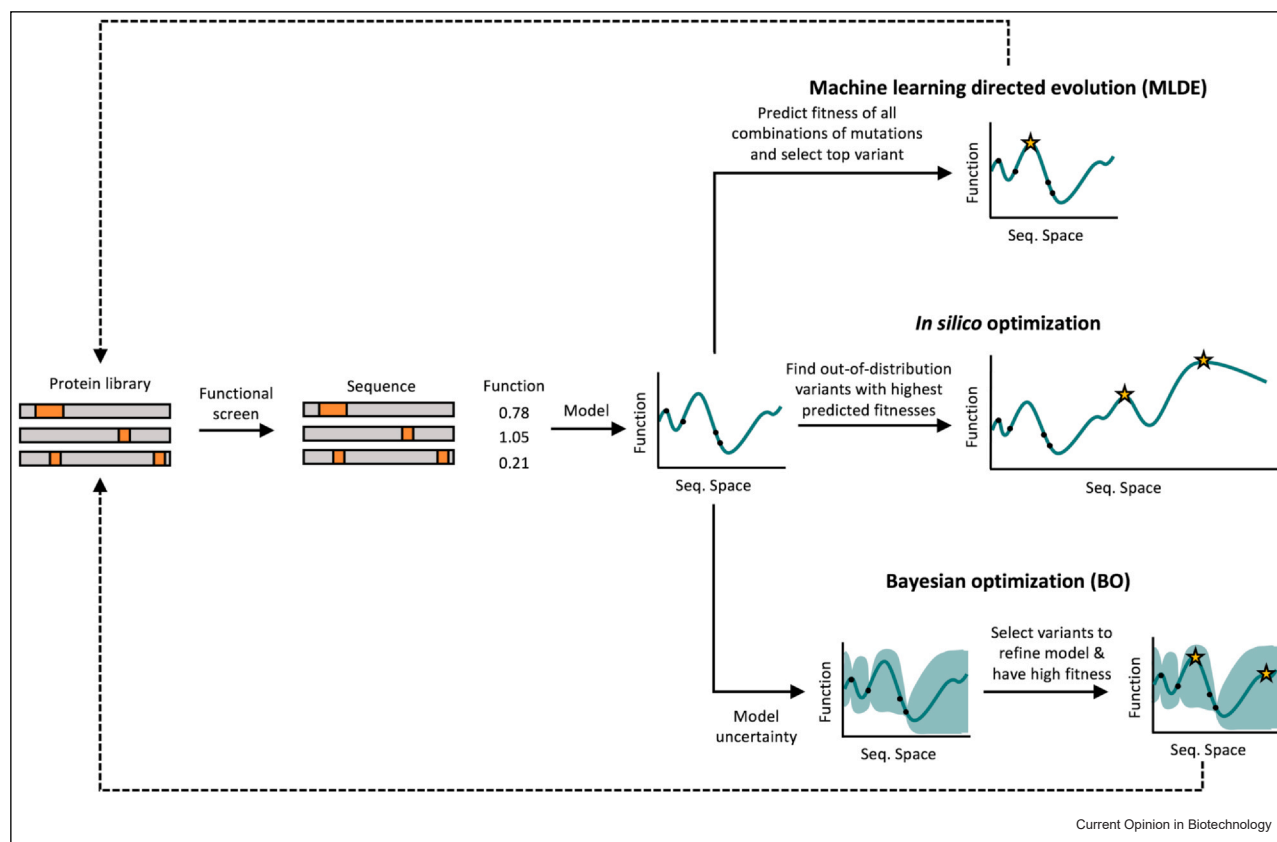
### ***In silico* optimization of the protein–fitness landscape**

An accurate supervised sequence–function model can be used to guide the search through sequence space for new and improved proteins (Figure 1). These models can extrapolate the learned relationship beyond the training set and identify highly fit sequences. We refer to this approach as '*in silico* optimization' because candidate proteins are identified by optimizing the sequence–function model.

*In silico* optimization strategies often involve landscape search heuristics such as hill climbing, simulated annealing, and genetic algorithms to identify sequences that are predicted to have high fitness values [9,12,30]. We recently applied a hill-climbing approach to design GB1 sequences based on neural network models trained on deep mutational scanning data. We found that the top designed variant with 10 mutations from wild-type GB1 was stable when expressed in *E. coli* and exhibited substantially increased binding affinity for IgG. It is also important to consider the diversity of candidate sequences in order to increase the independence of designs. Diversification strategies attempt to maximize the predicted fitness of designs while ensuring that they occupy distinct regions of the landscape [30,31].

*In silico* optimization can also be applied in small data settings using learned low-dimensional protein sequence representations [23]. The learned representations distill information pertinent to function beyond just sequence information alone and enable ML-based protein engineering in remarkably small data regimes. eUniRep is a representation model trained on 24 million unlabeled protein sequences and was used in a supervised setting to design improved green fluorescent protein (GFP)

Figure 1



ML-driven protein optimization strategies require an initial protein sequence library, typically created using error-prone PCR, site-saturated mutagenesis, or chimeragenesis. This protein sequence library is screened experimentally to determine a fitness value for each protein variant. Protein sequence–function data are modeled using supervised ML. The model can then be used to design improved variants in multiple ways. In MLDE, the fitness for other combinations of mutations in a combinatorial library is predicted, and the best variant is selected for further rounds of mutagenesis, screening, and modeling. *In silico* optimization, however, uses optimization strategies to find highly fit protein variants far from the initial training library in the larger protein sequence space. BO involves iterative rounds of mutagenesis, screening, and learning to maximize the protein fitness over fewer rounds of experimental characterization by proposing protein variants that will help refine the model and have high fitness.

variants with fewer than 100 labeled sequence–function examples. The eUniRep representation appears to guide design proposals away from nonfunctional sequence space. Embeddings may therefore be useful in scenarios where ML algorithms are attempting significant extrapolation beyond the training set. Despite these recent successes, there are still many open questions in the field related to the best sequence representations, supervised models, and the limits of model extrapolation.

### Active learning to iteratively search protein sequence space

The *in silico* optimization approach described above learns from a sequence–function dataset to design improved proteins in a one-step process. Active ML takes a different approach, where an iterative design–test–learn cycle is implemented to make the search through the sequence space more efficient. While this approach requires multiple rounds of sequence proposal and refining

the model with experimental data, the total screening burden is drastically reduced in comparison to both traditional directed evolution and most *in silico* optimization methods.

Machine learning-assisted directed evolution (MLDE) is a new approach that supplements traditional directed evolution with an underlying sequence–function model to effectively screen larger regions of sequence space [8,32]. MLDE begins with a combinatorial site-saturation mutagenesis library from which a small number of variants are screened for a function of interest. The resulting sequence–function data are used to train an ML model that predicts the function for all remaining variants in the combinatorial space. The top variant’s mutations are identified and fixed to provide the parent sequence for the next round of MLDE. By repeating this process, it is possible to efficiently traverse large swathes of sequence space to find optimal proteins.

Recent work has found that using biophysical, structural, and evolutionary information to filter the training data for diverse or highly fit variants helped achieve the maximum fitness more often [8,33]. MLDE augments the power of traditional directed evolution by providing a guided walk through sequence space while reducing the overall screening burden and accounting for epistasis.

Bayesian optimization (BO) is another approach that enables the engineering of proteins with relatively few experimental measurements. BO employs an iterative design–test–learn cycle where a model is trained on sequence–function data and used to propose new protein sequences to test that are simultaneously informative for the model and are predicted to have high fitness. Experimentally testing these sequences refines the model and also identifies new fitness peaks [34]. Gaussian process models have long been the standard for BO of the protein–fitness landscape owing to their ability to explicitly model uncertainty [35,36]. We recently applied BO to engineer acyl-ACP reductases for improved *in vivo* fatty alcohol production [37]. Over the course of ten design–test–learn cycles and fewer than 100 experimental measurements, we were able to engineer enzymes that produce over twofold more fatty alcohols than the starting natural sequences. However, Gaussian processes are relatively simple models and are limited in their ability to learn highly complex relationships compared with deep learning models. Ensembles of deep learning models, such as CNNs and RNNs, can be used to estimate the model uncertainty for BO and improve the maximum fitness achieved relative to Gaussian processes-based optimization [38]. Thus, combining active learning techniques with state-of-the-art supervised ML models is a promising approach for engineering proteins that cannot be screened in high throughput.

### Targeted exploration of protein sequence space with generative modeling

Supervised methods are powerful for inferring the protein landscape from labeled sequence–function data; however, there are millions of sequences in databases such as UniProt that lack functional data [22]. These sequences are referred to as 'unlabeled' because they have unknown functional characteristics, but they still provide valuable clues regarding the underlying landscape. Generative models learn the distribution of unlabeled protein sequence data with the assumption that any sequence falling within this distribution will exhibit similar properties to the training data. By proposing sequences from this distribution, generative models can sample novel sequences that 'look' and function like natural sequences [39,40]. However, it is difficult for these generative models to propose designs with improved properties without explicit fitness or function

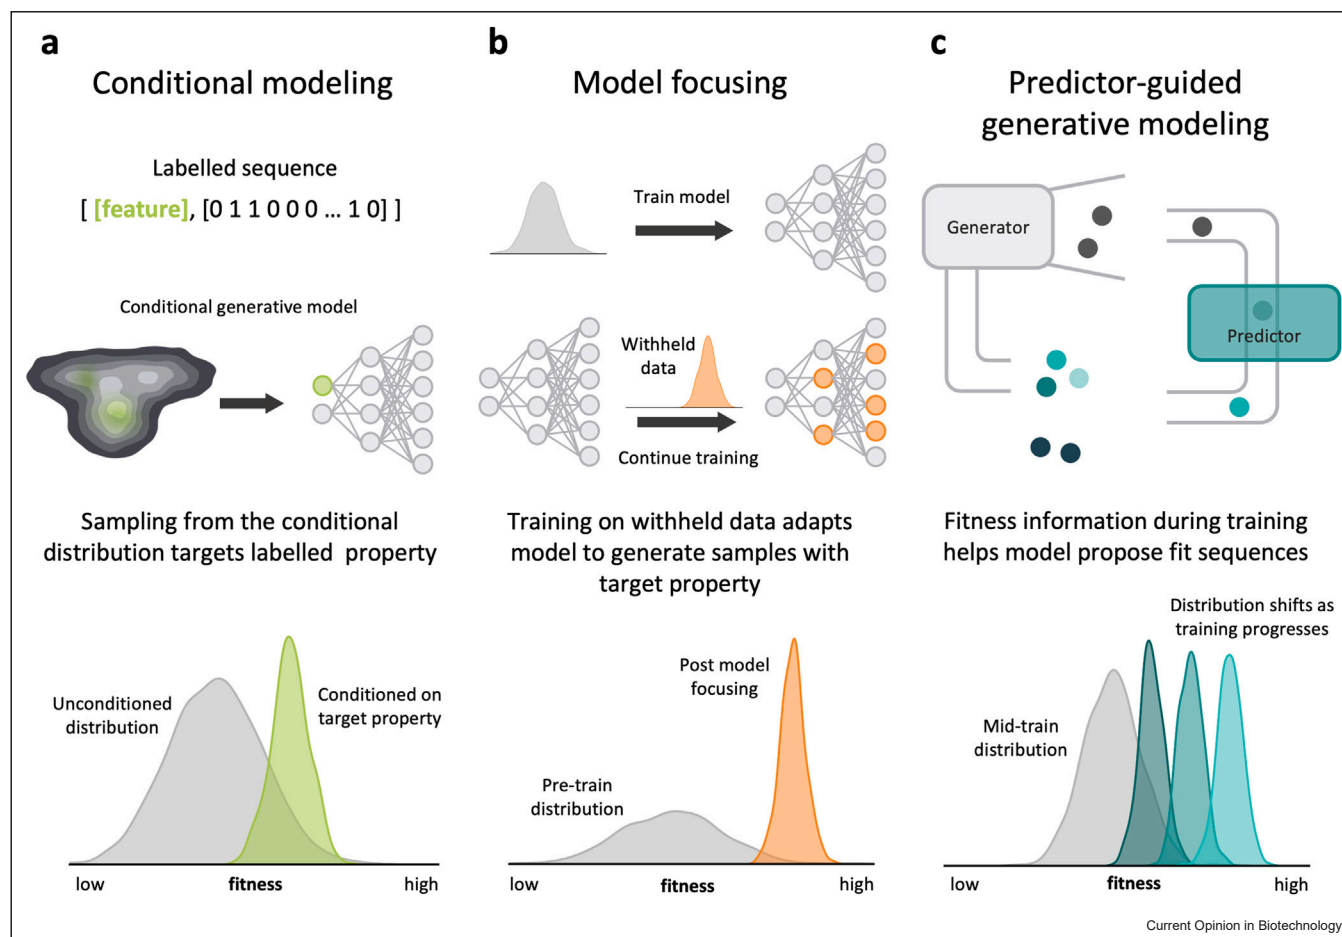
information. Here, we discuss modeling strategies to direct generative models toward proteins with improved function (Figure 2). For a comprehensive review on the broad applications of generative modeling in protein engineering, we refer the readers to the recent review by Wu et al. [5].

Conditional generative models can be used to bias the sampling process to engineer proteins with specific properties. These models learn a probability distribution over sequences and a set of defined protein attributes; this distribution can be conditioned on a desired attribute to enable controlled generation of proteins with the target attribute (Figure 2a) [41–44]. This strategy has been used on a diverse set of protein engineering tasks, including engineering luciferases with specified solubility levels [42]. Hawkins-Hooker et al. trained a conditional VAE on luciferase sequences labeled with a low, medium, or high solubility level. Conditioning on a particular solubility resulted in novel luciferases that displayed the conditioned solubility while retaining native luciferase activity. While this study targets one attribute, multiple attributes can be conditioned on simultaneously to generate proteins with many desired attributes, providing a flexible modeling framework suited to diverse applications [45,46].

Another strategy to bias generative models toward improved sequences is model focusing, where a general model's parameters are fine-tuned on a subset of data with a desired property (Fig. 2b). The resulting model is biased to sample proteins that closely match the characteristics of the tuning data. For example, Amimeur et al. [47] trained a generative adversarial network (GAN) to produce human-like antibodies. They then fine-tuned this model's parameters by focusing the training on a subset of sequences with desirable properties, such as decreased immunogenicity, that biased the model to design sequences with the target property.

Other generative protein engineering strategies explicitly incorporate predicted fitness into a generator's training process, such that over time the model is taught to propose more highly fit sequences (Fig. 2c). These predictions can be from supervised learning models or any other biophysical or structure-based model. Feedback GAN (fbGAN) couples a GAN architecture with a predictive model that predicts the fitness of the generated sequences during the training process [48]. Fit sequences are identified and fed back into the model's training data. While fbGAN modifies the training data, the Deep Exploration Network (DEN) developed by Linder et al. [30] ties the predicted function of the generated sequences directly into the generator's cost function. The DEN performed better than other state-of-the-art methods, including fbGAN, across a variety of applications, including the design of highly fit and

Figure 2



Strategies for protein engineering with generative models. **(a)** Conditional modeling labels training examples with a feature vector to identify sequences with a target attribute. Conditioning on a specific attribute generates a posterior distribution that preferentially samples proteins with the target attribute. **(b)** Model focusing adapts a general model to a specialized task with additional training on a targeted subset of data. Model focusing biases the model to generate proteins with properties that resemble the targeted training set. **(c)** Predictor-guided generative modeling incorporates fitness information into generator training, which can be provided by any model that predicts a property. As generator training progresses, the model proposes more highly fit examples.

diverse avGFP variants. The DEN also incorporates an explicit diversity and likelihood metric to efficiently search multiple regions of sequence space for fit variants with high confidence for success. Generative models can find functional protein sequences far removed from the training data — coupling this explorative power with functional information can guide this exploration toward fit proteins for effective protein engineering.

## Conclusion

Deep learning is revolutionizing our ability to decipher the complex relationships between protein sequence, structure, and function [49–51]. Predictive models of the sequence–function landscape are enabling protein engineers to efficiently and reliably search sequence space for new and useful proteins. Approaches that minimize experimental data requirements, including information-rich

protein representations and integrated design–test–learn cycles, will push the boundaries of protein engineering to functions that rival the performance and complexity of natural systems. This new frontier of biomolecular engineering will address global grand challenges in human health, agriculture, the environment, and energy.

## Conflict of interest statement

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgements

This work was funded by the United States National Institutes of Health (R35GM119854).



## References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest
- of outstanding interest.

1. Maheshri N, Koerber JT, Kaspar BK, Schaffer DV: **Directed evolution of adeno-associated virus yields enhanced gene delivery vectors.** *Nat Biotechnol* 2006, **24**:198-204.
  2. Wu Z, Yang KK, Liszka MJ, Lee A, Batzilla A, Wernick D, Weiner DP, Arnold FH: **Signal peptides generated by attention-based neural networks.** *ACS Synth Biol* 2020, **9**:2154-2161.
  3. Chevalier A, Silva D-A, Rocklin GJ, Hicks DR, Vergara R, Murapa P, Bernard SM, Zhang L, Lam K-H, Yao G, et al.: **Massively parallel de novo protein design for targeted therapeutics.** *Nature* 2017, **550**:74-79.
  4. Romero PA, Arnold FH: **Exploring protein fitness landscapes by directed evolution.** *Nat Rev Mol Cell Biol* 2009, **10**:866-876.
  5. Wu Z, Johnston KE, Arnold FH, Yang KK: **Protein sequence design with deep generative models.** *Curr Opin Chem Biol* 2021, **65**:18-27.
  6. Wittmann BJ, Johnston KE, Wu Z, Arnold FH: **Advances in machine learning for directed evolution.** *Curr Opin Struct Biol* 2021, **69**:11-18.
  7. Yang KK, Wu Z, Arnold FH: **Machine-learning-guided directed evolution for protein engineering.** *Nat Methods* 2019, **16**:687-694.
  8. Wittmann BJ, Yue Y, Arnold FH: **Informed training set design enables efficient machine learning-assisted directed protein evolution.** *Cell Syst* 2021, **12**:1026-1045 e7.
  9. Gelman S, Fahlberg SA, Heinzelman P, Romero PA, Gitter A: **Neural networks to learn protein sequence-function relationships from deep mutational scanning data.** *Proc Natl Acad Sci USA* 2021, **118**:e2104878118.
- This paper demonstrates how extrapolation of the protein fitness landscape through *in silico* optimization can be used to find drastically improved proteins from diverse regions of sequence space.
10. Xu Y, Verma D, Sheridan RP, Liaw A, Ma J, Marshall NM, McIntosh J, Sherer EC, Svetnik V, Johnston JM: **Deep dive into machine learning models for protein engineering.** *J Chem Inf Model* 2020, **60**:2773-2790.
- This paper provides a systematic evaluation of many embeddings and supervised learning models and demonstrates the promise of CNNs for modeling the protein-fitness relationship.
11. Griffith D, Holehouse AS: **PARROT is a flexible recurrent neural network framework for analysis of large protein datasets.** *eLife* 2021, **10**:e70576.
  12. Bryant DH, Bashir A, Sinai S, Jain NK, Ogden PJ, Riley PF, Church GM, Colwell LJ, Kelsic ED: **Deep diversification of an AAV capsid protein by machine learning.** *Nat Biotechnol* 2021, **39**:691-696.
  13. Mater AC, Sandhu M, Jackson C: **The NK landscape as a versatile benchmark for machine learning driven protein engineering.** *bioRxiv* 2020, <https://doi.org/10.1101/2020.09.30.319780>
  14. Bepler T, Berger B: **Learning the protein language: evolution, structure, and function.** *Cell Syst* 2021, **12**:654-669 e3.
  15. Carlin DA, Caster RW, Wang X, Betzenderfer SA, Chen CX, Duong VM, Ryklansky CV, Alpekin A, Beaumont N, Kapoor H, et al.: **Kinetic characterization of 100 glycoside hydrolase mutants enables the discovery of structural features correlated with kinetic constants.** *PLOS One* 2016, **11**:1-14.
  16. Rocklin GJ, Chidyausiku TM, Goresnik I, Ford A, Houlston S, Lemak A, Carter L, Ravichandran R, Mulligan VK, Chevalier A, et al.: **Global analysis of protein folding using massively parallel design, synthesis, and testing.** *Science* 2017, **357**:168-175.
  17. Alley EC, Khimulya G, Biswas S, AlQuraishi M, Church GM: **Unified rational protein engineering with sequence-based deep representation learning.** *Nat Methods* 2019, **16**:1315-1322.
  18. Heinzinger M, Elnaggar A, Wang Y, Dallago C, Nechaev D, Matthes F, Rost B: **Modeling aspects of the language of life through transfer-learning protein sequences.** *BMC Bioinformatics* 2019, **20**:723.
  19. Rao R, Bhattacharya N, Thomas N, Duan Y, Chen X, Canny J, Abbeel P, Song YS: **Evaluating protein transfer learning with TAPE.** *Adv Neural Inf Process Syst* 2019, **32**:9689-9701.
  20. Luo Y, Jiang G, Yu T, Liu Y, Vo L, Ding H, Su Y, Qian WW, Zhao H, Peng J: **ECNet is an evolutionary context-integrated deep learning framework for protein engineering.** *Nat Commun* 2021, **12**:5743.
  21. Yang KK, Wu Z, Bedbrook CN, Arnold FH: **Learned protein embeddings for machine learning.** *Bioinformatics* 2018, **34**:2642-2648.
  22. The UniProt Consortium: **The universal protein resource (UniProt).** *Nucl Acids Res* 2007, **36**:D190-D195.
  23. Biswas S, Khimulya G, Alley EC, Esvelt KM, Church GM: **Low-N protein engineering with data-efficient deep learning.** *Nat Methods* 2021, **18**:389-396.
- Using protein representation models, the authors demonstrate the capacity for semisupervised learning models to engineer proteins with minimal functional data. They use eUniRep protein embeddings to train a ridge regression model and engineer avGFP and TEM-1  $\beta$ -lactamase activity from a small functional dataset (< 100).
24. Rives A, Meier J, Sercu T, Goyal S, Lin Z, Liu J, Guo D, Ott M, Zitnick CL, Ma J, et al.: **Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences.** *Proc Natl Acad Sci USA* 2021, **118**:e2016239118.
  25. Hsu C, Nisonoff H, Fannjiang C, Listgarten J: **Learning protein fitness models from evolutionary and assay-labeled data.** *Nat Biotechnol* 2022.
- This work provides an excellent systematic evaluation of many current representation/supervised learning methods and provides a framework for future evaluation of models that use evolutionary and assay-labeled data.
26. Ding X, Zou Z, Brooks CL III: **Deciphering protein evolution and fitness landscapes with latent space models.** *Nat Commun* 2019, **10**:5644.
  27. Siedhoff NE, Illig A-M, Schwaneberg U, Davari MD: **PyPEF—an integrated framework for data-driven protein engineering.** *J Chem Inf Model* 2021, **61**:3463-3476.
  28. Chen KM, Cofer EM, Zhou J, Troyanskaya OG: **Selene: a PyTorch-based deep learning library for sequence data.** *Nat Methods* 2019, **16**:315-318.
  29. Favor A, Jayapurna I: **Evaluating eUniRep and other protein feature representations for in silico directed evolution.** *Authorea* 2020,1-14, <https://doi.org/10.22541/au.159683529.96283070>
  30. Linder J, Bogard N, Rosenberg AB, Seelig G: **A generative neural network for maximizing fitness and diversity of synthetic DNA and protein sequences.** *Cell Syst* 2020, **11**:49-62 e16.
- The authors design a DEN architecture that can be adapted to a variety of engineering tasks. They incorporate a similarity metric and predicted fitness into the generator's cost function to design avGFP with maximized diversity and function. The DEN architecture proposes sequences with high fitness and diversity than competing methods. They also incorporate a separately trained VAE to constrain sequence exploration to natural-like sequences.
31. Zhu D, Brookes DH, Busia A, Carneiro A, Fannjiang C, Popova G, Shin D, Chang EF, Nowakowski TJ, Listgarten J, et al.: **Machine learning-based library design improves packaging and diversity of adeno-associated virus (AAV) libraries.** *bioRxiv* 2021, <https://doi.org/10.1101/2021.11.02.467003>
  32. Wu Z, Kan SBJ, Lewis RD, Wittmann BJ, Arnold FH: **Machine learning-assisted directed protein evolution with combinatorial libraries.** *Proc Natl Acad Sci USA* 2019, **116**:8852-8858.
  33. Qiu Y, Hu J, Wei G-W: **Cluster learning-assisted directed evolution.** *Nat Comput Sci* 2021, **1**:809-818.
  34. Hie BL, Yang KK: **Adaptive machine learning for protein engineering.** *Curr Opin Struct Biol* 2022, **72**:145-152.

35. Romero PA, Krause A, Arnold FH: **Navigating the protein fitness landscape with Gaussian processes.** *Proc Nat Acad Sci* 2013, **110**:E193-E201.
36. Bedbrook CN, Yang KK, Rice AJ, Gradinaru V, Arnold FH: **Machine learning to design integral membrane channelrhodopsins for efficient eukaryotic expression and plasma membrane localization.** *PLoS Comput Biol* 2017, **13**:e1005786.
37. Greenhalgh JC, Fahlberg SA, Pflieger BF, Romero PA: **Machine learning-guided acyl-ACP reductase engineering for improved in vivo fatty alcohol production.** *Nat Commun* 2021, **12**:5825.
38. Gruver N, Stanton S, Kirichenko P, Finzi M, Maffettone P, Myers V, Delaney E, Greenside P, Wilson AG: **Effective Surrogate models for protein design with Bayesian optimization.** In *ICML 2021 Workshop on Computational Biology*; 2021.

Here, the authors combine traditional Bayesian optimization with deep neural networks and show that this strategy helps optimize the protein fitness landscape more quickly than Gaussian processes models. The authors also experimentally validate this adapted method by engineering GFP.

39. Repecka D, Jauniskis V, Karpus L, Rembeza E, Rokaitis I, Zrimec J, Poviloniene S, Laurynešas A, Viknander S, Abuajwa W, et al.: **Expanding functional protein sequence spaces using generative adversarial networks.** *Nat Mach Intell* 2021, **3**:324-333.

The authors trained a GAN on bacterial malonate dehydrogenases and design new MDH sequences with a median 64.4% sequence identity to the nearest natural sequences. They find that 13 of 19 tested MDH designs exhibit catalytic activity, including an MDH mutant with 106 mutations. This is an excellent demonstration of how generative modeling can be used to expand functional sequence space.

40. Madani A, Krause B, Greene ER, Subramanian S, Mohr BP, Holton JM, Olmos JL, Xiong C, Sun ZZ, Socher R, et al.: **Deep neural language modeling enables functional protein generation across families.** *bioRxiv* 2021, <https://doi.org/10.1101/2021.07.18.452833>
41. Sohn K, Lee H, Yan X: **Learning structured output representation using deep conditional generative models.** In *Advances in Neural Information Processing Systems*. Edited by Cortes C, Lawrence N, Lee D, Sugiyama M, Garnett R. Curran Associates, Inc.; 2015.
42. Hawkins-Hooker A, Depardieu F, Baur S, Couairon G, Chen A, Bikard D: **Generating functional protein variants with variational autoencoders.** *PLoS Comput Biol* 2021, **17**:e1008736.

The authors demonstrate the application of conditional generative modeling with a conditional VAE trained on luciferases. They engineer functional luciferases with specific solubilities and verify their results with *in vitro* solubility assays.

43. Brookes DH, Park H, Listgarten J: **Conditioning by adaptive sampling for robust design.** *arXiv* 2021. [arXiv:1901.10060](https://arxiv.org/abs/1901.10060).
44. Chan A, Madani A, Krause B, Naik N: **Deep extrapolation for attribute-enhanced generation.** In *Advances in Neural Information Processing Systems* 2021, **35**:1-13.
45. Karimi M, Zhu S, Cao Y, Shen Y: **De Novo protein design for novel folds using guided conditional Wasserstein generative adversarial networks.** *J Chem Inf Model* 2020, **60**:5667-5681.
46. Das P, Sercu T, Wadhawan K, Padhi I, Gehrmann S, Cipicigan F, Chenthamarakshan V, Strobel H, dos Santos C, Chen P-Y, et al.: **Accelerated antimicrobial discovery via deep generative models and molecular dynamics simulations.** *Nat Biomed Eng* 2021, **5**:613-623.
47. Amieur T, Shaver JM, Ketchum RR, Taylor JA, Clark RH, Smith J, Van Citters D, Siska CC, Smidt P, Sprague M, et al.: **Designing feature-controlled humanoid antibody discovery libraries using generative adversarial networks.** *bioRxiv* 2020, <https://doi.org/10.1101/2020.04.12.024844>.

These authors use transfer learning to generate a library of human-like antibodies biased toward specific features. They start with a pretrained GAN that generates human-like antibodies confirmed to fold via phage display. They continue training the model on a subset of withheld antibody sequences with the target property. After training, they confirmed that their model is biased to produce sequences with the target property through computational analyses.

48. Gupta A, Zou J: **Feedback GAN for DNA optimizes protein functions.** *Nat Mach Intell* 2019, **1**:105-111.

These authors use supervised learning to incorporate predicted fitness information to guide training of a GAN. By feeding sequences with high predicted fitnesses back into the training data during model training, the model learns to propose highly fit sequences. The authors designed antimicrobial peptide (AMP) sequences with improved AMP properties and a higher composition of alpha helices.

49. Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, Tunyasuvunakool K, Bates R, Židek A, Potapenko A, et al.: **Highly accurate protein structure prediction with AlphaFold.** *Nature* 2021, **596**:583-589.
50. Eguchi RR, Anand N, Choe CA, Huang P-S: **IG-VAE: generative modeling of immunoglobulin proteins by direct 3D coordinate generation.** *bioRxiv* 2020, <https://doi.org/10.1101/2020.08.07.242347>
51. Norn C, Wicky BIM, Juergens D, Liu S, Kim D, Tischer D, Koepnick B, Anishchenko I, Foldit Players, Baker D, et al.: **Protein sequence design by conformational landscape optimization.** *Proc Natl Acad Sci USA* 2021, **118**:e2017228118.