Assessing the performance of protein regression models

Richard Michael*[1], Jacob Kæstel-Hansen*[2], Peter Mørch Groth[1,3], Simon Bartels[1], Jesper Salomon[3], Pengfei Tian[3], Nikos S. Hatzakis[2], Wouter K. Boomsma[1]

[1]Department of Computer Science, University of Copenhagen

[2]Department of Chemistry, University of Copenhagen

[3]Bioinformatics & Design, Enzyme Research, Novozymes A/S

correspondence: wb@di.ku.dk

* These authors contributed equally to this work.

Abstract

To optimize proteins for particular traits holds great promise for industrial and pharmaceutical purposes. Machine Learning is increasingly applied in this field to *predict* properties of proteins, thereby guiding the experimental optimization process. A natural question is: How much progress are we making with such predictions, and how important is the choice of regressor and representation? In this paper, we demonstrate that different assessment criteria for regressor performance can lead to dramatically different conclusions, depending on the choice of metric, and how one defines generalization. We highlight the fundamental issues of sample bias in typical regression scenarios and how this can lead to misleading conclusions about regressor performance. Finally, we make the case for the importance of calibrated uncertainty in this domain.

Introduction

Accurate prediction of protein traits related to function and stability remains an important challenge, both for *in silico* protein engineering and for assessing the phenotypic consequences of genetic diseases[1;2;3]. Recent years have seen progress in *unsupervised* models for predicting the fitness effect arising from protein variants. For single protein families, statistical models of aligned sequences have been used successfully to characterize the effect of mutations through a change in (approximate) likelihood[4;5;6;7;8;9;10] and attention-based language models have since extended this idea beyond single protein families[11;12;13;14].

While unsupervised models are very useful descriptors of variants with respect to a natural baseline, protein engineering often requires us to go beyond the unsupervised setting. One reason is that the function optimized for will often differ from the biological function selected for by evolution[15]. Another is that we generally wish to improve our predictive capabilities as we make more experimental observations of a system, which requires that we train models in a *supervised* fashion.

Supervised learning is associated with a number of challenges. First of all, it requires us to carefully consider how the available data are applied, since they are now used both for training and evaluation. The data can be split into training and test sets in different ways, and the similarities we permit between these sets will influence our assessment. The goal should be to choose a test set that reflects the data distribution that we expect to see when the regressor is ultimately applied in practice. But this matter is complicated by the fact that in an optimization setting, regressors are frequently used partially out-of-domain, i.e. to make predictions on sequences further from the wildtype than those in the training set, or distributed differently. A representative test set might not be available in this case. Even the notion of *in-domain* and *out-of-domain* is often not clear cut, and can depend on the choice of how a protein is encoded. The embeddings computed by protein language models (PLMs) are for instance likely to have a biologically richer notion of similarity than a distance based merely on amino acid identity. The extrapolative setting also highlights the requirement for regression algorithms to provide reliable uncertainty estimates. This is important to ensure that downstream decision (by human practitioners of optimization algorithms) can be based on an informed trade-off between the predicted fitness of

proposed candidates and the uncertainty of these estimates. Finally, supervised learning on protein sequences is made difficult by the fact that data sets for a particular protein system are often very limited in size, and can have substantial sample biases.

In this paper, we systematically describe the challenges involved in supervised learning for protein engineering. We describe the constituents of a protein regression pipeline (Fig. 1), placing particular focus on the assessment stage. We show empirically that different choices at this stage can have a dramatic impact on the conclusions drawn, and discuss how we can design train/test splits to best reflect the generalization capabilities we desire for a given task. We also investigate the role of embeddings of the input protein sequences, comparing several recent language models with simpler alternatives. Finally, we consider the metrics used to measure performance, and highlight the importance of including calibration of predicted uncertainties when comparing regression procedures.

Results

The fundamental objective in protein engineering is to find modifications (variants) of naturally occurring amino acid sequences that optimize one or multiple properties of interest, *e.g.* the thermostability of an enzyme. Since there are twenty naturally occurring amino acids, the space of possible amino acid sequences of a given length L is 20^L, an astronomically large number even for short proteins. In practice, engineering R&D pipelines often restrict themselves to sequences that deviate by only a few mutations from a naturally occurring sequence - a *wildtype* (WT). Exhaustive search through this restricted space is, still intractable. In practice we are limited in time and resources, and the search strategies employed are often based on combinations of random sampling and the intuition of domain experts. Developing better optimization strategies has therefore become an active area of research [16;17;18;19].

As part of such optimization strategies, the need arises to *predict* fitness values for unseen protein sequences. One example is the *surrogate function* in a Bayesian Optimization setting, which predicts the expected fitness value (and uncertainty) of any candidate protein. Together with an *acquisition function* which specifies a trade-off between exploration and exploitation, this allows us to select new candidates optimally, according to a well-defined protocol.

Our success in optimizing proteins thus depends critically on the ability to regress the fitness of a protein against its sequence. Rather than working with the raw amino acid sequence, the input can be encoded in various ways to facilitate modelling. A simple choice is ONE-HOT-encoding, where each amino acid in a protein sequence or alignment is encoded in a d-dimensional vector, where d is the number of different amino acid labels (typically 20 plus gap symbols and nonstandard labels). In recent years, it has become clear that regression performance can often be enhanced further by regressing against learned protein representations [13;20;21;22]. This can potentially simplify the search space and make it possible to extrapolate further away from the wildtype sequence.

The next step is to choose a suitable regression algorithm. This choice will depend on factors such as 1) appropriate inductive bias, 2) expressivity relative to the amount of available data, 3) attributes of a model such as the ability to predict uncertainties. Ultimately, a dominant reason for the choice will be that a particular model is perceived to perform better on the specific task of interest. As we discuss the pros and cons of different regression algorithms and representations, it is thus of critical importance that we can reliably assess this performance. In contrast to unsupervised fitness prediction, assessment in the supervised setting will require us to split the data into train, validation, and test sets. Depending on the application task of interest, different choices are meaningful, and this choice can have a substantial impact on the reported performance. The performance will depend on the nature of the fitness landscape, and the amount of available data. Since experimental assays typically produce only tens up to a few thousands of measurements at a time, protein regression typically resides in the low-data regime, thus restricting the capacity of regression algorithms we can meaningfully apply. The experimental assays can also be of different quality, have different levels of coverage, and display more or less correlation to the functional trait of interest.

These are thus the key components in a protein regression pipeline: the nature of the data, its encoding, the regression algorithm, and the assessment strategy (Fig.1). We will now consider them in turn.

Data

On the input side, protein fitness data can have different scope, ranging from sampling broadly from naturally occurring sequences (sometimes referred to as wildtype exploration), to sampling sequences locally around a single or a few naturally occurring sequences. Since our study fo-

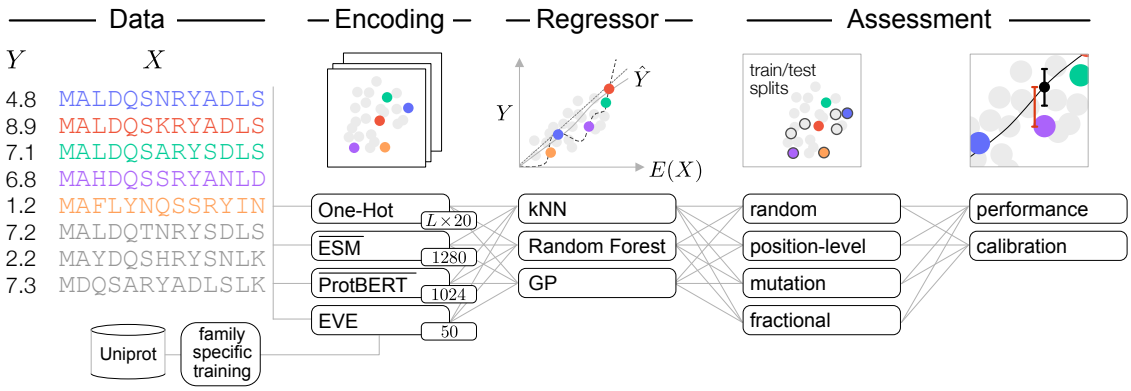

Figure 1: A general overview for supervised learning tasks on protein systems.: For each protein exists labelled experimental assay *Data*, multiple sequence alignments (MSA), and sequence representations (*Encoding* annotated with the respective dimensionality). The language model (ESM-1B and ProtBert) representations are the positional means of the internal representation (indicated by the bar). We fit and assess regressors, and select training and validation data through different protocols (*Assessment*). From the performance and calibration results we assess an in-silico optimization task for variant selection.

cuses on protein optimization, we consider the latter category here, but refer to related work for considerations about wild type exploration [22;23].

Experimental techniques for characterizing protein properties, particularly many high-throughput assays, are often imperfect proxies of the desired functional trait and exhibit considerable noise. For instance, a common first step in a protein engineering pipeline is a site-saturation mutagenesis experiment, which in a single experiment characterizes the majority of variants deviating by a single mutation from the wildtype protein. Apart from relatively high levels of noise, these initial high-throughput experiments often have only limited correlation to the functional trait of interest, and are therefore sometimes augmented with lower-throughput, more costly assays in later stages of the optimization process. After an exploration of single mutants, protocols will often proceed with variants with greater separation from the wildtype. In these later stages, there will typically be considerable selection bias in the experimentally probed sequences, as a result of new variants being selected based on results from previous experiments. These effects have consequences for the performance we can expect from our regression algorithms, and in the design of our split strategies, as we will see below.

In addition to the sequences with observed experimental values (*i.e. labelled* sequences), there are typically many sequences available for similar proteins from the databases of naturally occurring proteins. If the functional trait we are optimizing for is related to the biological function of the protein, we can hope to gain additional information from these unlabelled sequences, for instance by learning improved representations or by unsupervised predictions.

To ground our discussion on modelling choices in regressor design, we conduct an empirical evaluation on a representative set of protein variant effect data sets. The data sets were selected to vary in data quantity, in how far the observed variants are from their closest wildtype, and to be diverse respective to their biological organismal functions and structures. The experimental data we consider are from mutational scans, which recorded growth under different stresses depending on the protein systems (see Tab.2 cf. [8;14]) For each synthesizeable candidate sequence we observe a measurement $y \in Y \subset \mathbb{R}$, as a label.

For our initial experiments, we will use two commonly employed train/validation split scenarios: a random cross validation (Random CV) strategy where the data is split at random (uniformly), and a position-level cross validation (Position CV), where splits are made by positions in the chain (non-overlapping segments are assigned to either the training, validation or test group). We will return to the choice of splitting procedures in our treatment of assessment techniques later in the paper.

Protein Representations

Advances in *representation learning* have shown that different approaches to unsupervised learning from biological sequences lead to differently structured representations [8;12;13;24;25;20]. *A priori*, it is not clear how such differences in structure will translate into downstream performance. Therefore, the usefulness of a representation can only be assessed in terms of the performance of a

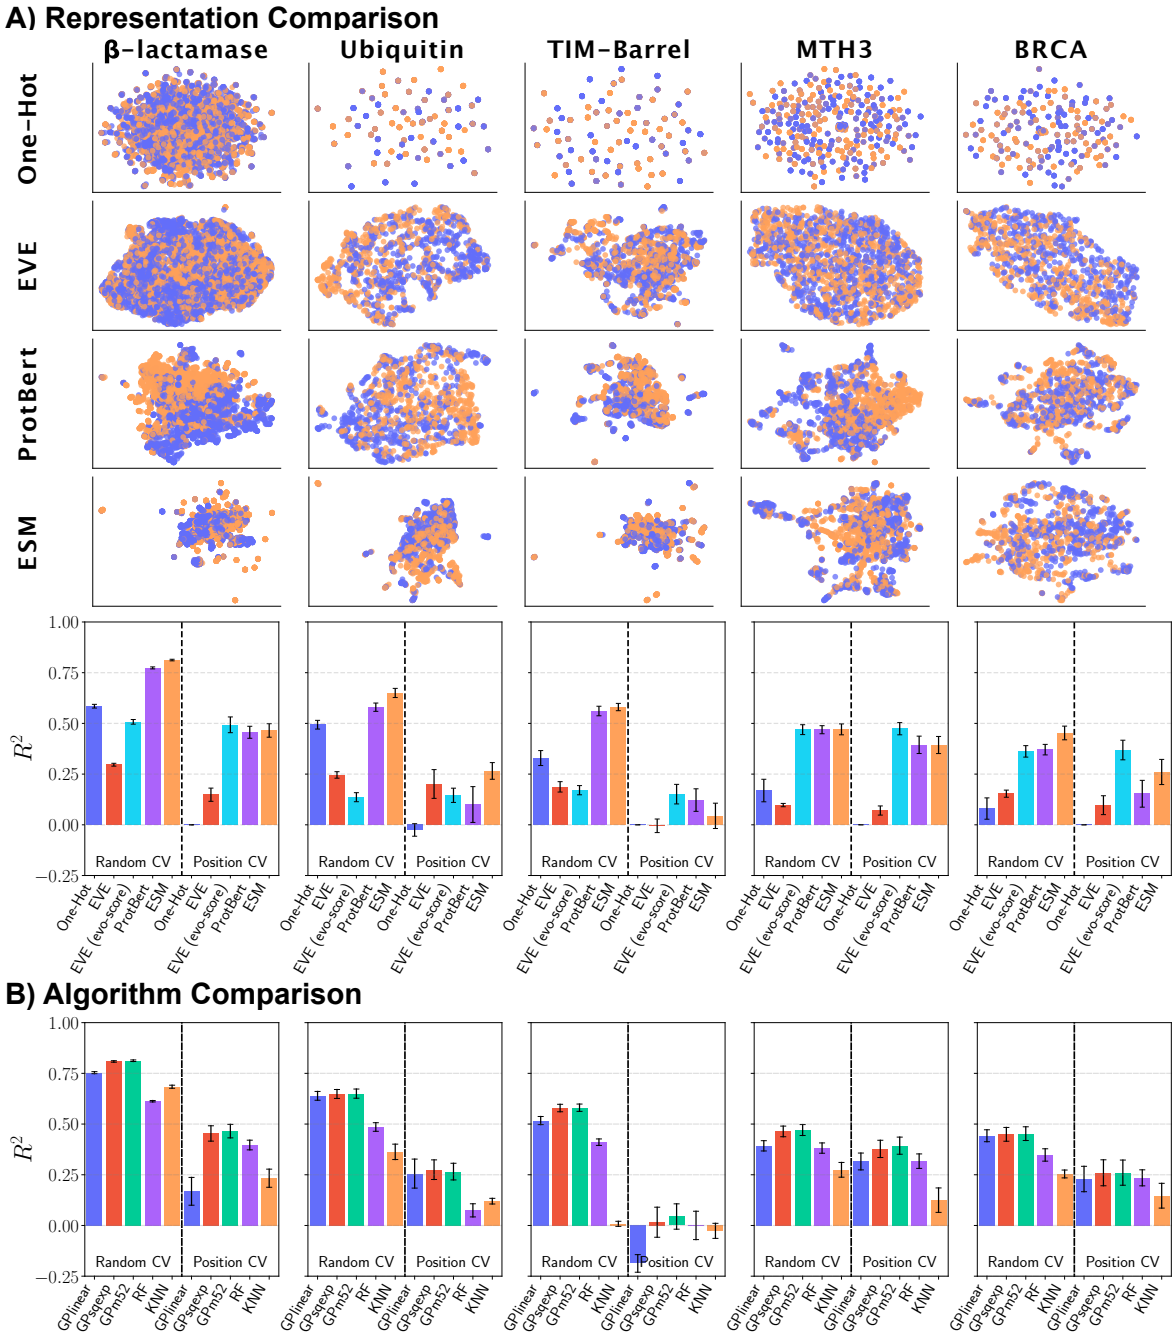

Figure 2: Representation and regressor comparison. We assess five distinct protein data sets by: A) two-dimensional spatial separation of UMAP reduced representations color-coded in functional (orange) and non-functional (purple) variants. Performance of a GP-Matern52 regressor on the representation as input reporting accuracy (adjusted R2) on the test set of the CV protocol with bars as standard-error across splits. B) All supervised regressors on the ESM-1B representation input.

downstream prediction algorithm. For the particular protein systems investigated here, we show such an analysis in Figure 2A, comparing the ONE-HOT representation to two recent large language model representations, ESM-1B[11] and PROTBERT[24], and a family-specific representation extracted from the EVE variational autoencoder model[10]. We will return to the exact choice of metric and cross validation scheme later, and focus first on the fact that regression performance is clearly heavily impacted by the choice of representation (Figure 2A, bottom). As suggested by the striking results reported for fitness prediction with large language models and family-specific density models in the zero-shot (unsupervised) setting[24;14], we might expect increased performance when using the corresponding representations as input for supervised training. Indeed, we observe that representations are capable of increasing regression performance beyond using the raw amino acid sequence. Although the strength of this effect depends on the protein, choice of regressor, and cross validation scheme, the representations extracted from language models seem to generally be a robust choice (Fig. 2A). Interestingly, we note that this does not hold true for all choices of learned representations. In particular, we observe consistently that the latent space of a protein family-specific variational autoencoder (here EVE) seems to underperform in downstream regression tasks. This is somewhat surprising, since the *evo-score*, which is the evidence lower bound (*i.e.* a proxy for the likelihood against the *WT*) extracted from the same model seems to fair better on average, despite being a one-dimensional value.

Overall, we see a large variance between protein systems on the usefulness of the different representations. This is likely to partially reflect the different nature of the assays. For instance, the cases where the evo-score is predictive of the assay value are presumably those where the assay is a good proxy for the biological function. However, as has been shown in earlier work[20], the optimal choice of representation depends on the fitness landscape one wishes to model, and it is advisable to experiment with different choices on a given problem.

One final remark on the choice of representation relates to their dimensionality. The high dimensional per-position embeddings produced by language models are not well suited for direct input to a regression model. Often, the representations are aggregated over the length of the protein by simple averaging[1]. Even when employing such an averaging strategy, one is left with a representation of (often more than) hundreds of dimensions. For certain regression algorithms, like Gaussian Processes using kernels based on Euclidean distance, such high dimensional spaces pose a challenge. In an ablation study we consider whether dimensionality-reduction pre-processing of the averaged representations could have a positive effect on regression performance, but generally found this to be detrimental (see Supplementary Material Fig. 20). However, several recent protein optimization studies have found it beneficial to include in their optimization protocol a supervised dimensionality reduction step trained specifically on the task[26;27].

Regression Algorithms

Choosing the right regression algorithm is often considered as the most critical of the model design choices in the pipeline. That most system-specific protein regression tasks exist within the low-data regime places constraints on the available model types and their capacities. In our case study, we select a set of representative methods, that are expected to show reasonable performance in this regime: (1) a kNN regressor, (2) a *Random Forest* (RF) with predictive ensemble uncertainties, and (3) three *Gaussian Process* (GP) regressors (linear kernel, squared exponential kernel, and Matérn $\frac{5}{2}$ kernel)[2]. Each regressor is implemented to provide not only a predicted value $\hat{y}_i$, but also a predictive variance estimate as an associated uncertainty $\hat{\sigma}_i^2$ (see Methods).

When comparing the performance of these regressors across the five data sets we find the GP regressors perform well despite the high dimensional input when hyperparameters are optimized (Fig. 2B). The fact that a GP with linear kernel in many cases provides a reasonable performance suggests that much of the signal in this limited data regime can be captured with a linear model, but this will depend on the system in question, the amount of available data, and the learned representation. A more general observation is that the choice of hyperparameters is at least as important as the choice of regressor; if chosen poorly, it can have a significant detrimental impact on performance (SM Fig. 18,19). Note that we focused here on regression performance using a language model representation, but find similar effects with other choices of representations (SM Fig. 19).

As in the previous section, we again observe a dramatic decrease in performance when we

[1]For simplicity, this is also the strategy employed here, although we stress that better choices exist, and can have considerable impact on performance[20].

[2]For method details (parameters, and regressor optimization) see Methods

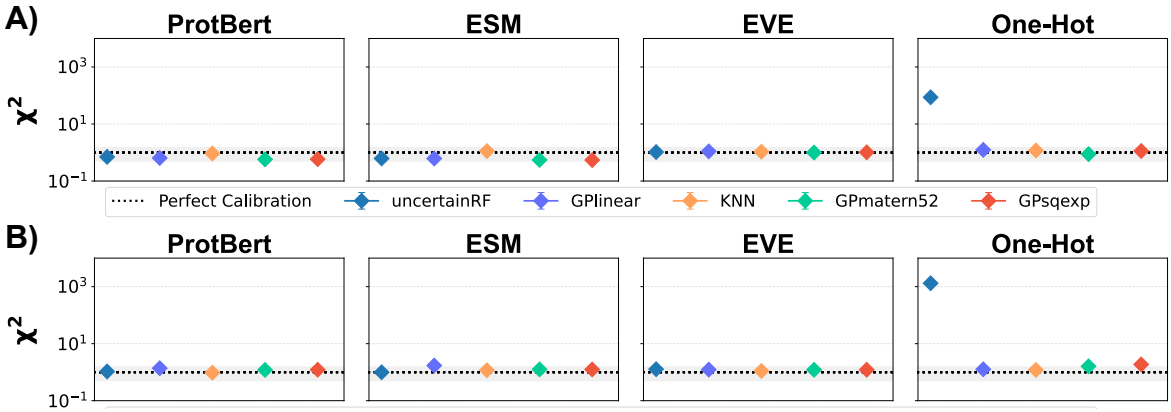

Figure 3: Calibration of regressors. Measured on β-LACTAMASEby the (reduced) χ^2 statistic when training the regressors by splitting at random (A) compared to by position (B). Generally, the regressors are well-calibrated, slightly more so for the Position CV (B). Outliers are characterized by very low predictive variance values (see ONE-HOT).

change splitting strategy from random sampling to position-level splitting (*i.e.* predicting in sequence regions unseen during training). The size of this effect is system specific. In cases like the TIM-BARREL, we have no predictive ability in the position-level split, although the random split performance is on par with other systems. We will explore this dependency in detail below.

Assessment: metrics, tasks, domains, and generalization

Metrics

An elementary part for any learning algorithm assessment is a means to assess predictive performance. To measure performance comparatively across data sets and systems we use the adjusted ratio between residuals and total residuals.: $R2 = 1 - \frac{\frac{1}{N}\sum_i (y_i - \hat{y}_i)^2}{\frac{1}{N}\sum_i (y_i - \mu_y)^2}$, which indicates an improvement of our predictions $\hat{y}$ over the training-set mean as a baseline μ_y (see Methods for details).

Since the predictive uncertainties that the regressors provide are used actively in downstream applications (*e.g.* in the acquisition function of a Bayesian Optimization procedure), we should assess our methods also on the quality of their uncertainty estimates. We can use the predictive uncertainties to assess *calibration* and *confidence* of a model[28;29;30]. The calibration implies that in expectation the average size of the prediction error should correspond to the magnitude of the grouped predictive uncertainties[31]. A method would thus be over-confident if the empirical error is larger than the uncertainties it predicts. We do this by computing a *standardized* (reduced) χ^2 statistic $\chi^2 = \frac{1}{N-1}\sum_i (\frac{y_i - \hat{y}_i}{\hat{\sigma}_i})^2$ as the squared mean of the residuals normalized by the predictive variance, which provides us with one estimate for the regressor's calibratedness. We also investigate calibration curves[28;29] by discretizing the predictions into q quantiles of the predictive variances $\hat{\sigma}_i^2$ across all observations $i \in [1, n]$, indexed by $j \in [1, q]$, so that each $\sigma_j^2 \in [0, \frac{1}{q}, ..., 1]$ describes a Gaussian $\mathcal{N}_j(0, \hat{\sigma}_j^2)$. From this we compute *calibration*, *expected calibration error* (ECE), and sharpness[3].

Plotting calibration assessments across regressors and representations (Fig. 3) for our test systems (SM Figures 8-11), we find the methods to be fairly well calibrated - though with a tendency towards overconfidence when employing poor representations.

Domains and generalization

When we create a test set we implicitly define a notion of generalization: we assume that the test set is *as representative of the training set as we expect any future application to be*. The appropriate choice of test set will thus depend on the task we intend to solve with the model subsequently.

To make the discussion more precise, we will briefly formalize the concepts of *domain* and *task*. We define the *domain* as a finite set of embedded sequences $X \subset \mathbb{R}^d$ with the distribution $p(X)$. A test set is said to be out-of-domain if there is a shift in distribution between inputs in the training and test sets - sometimes referred to as *covariate shift*. A *task* involves a set of observations $Y \subset \mathbb{R}$, and is defined by the joint distribution $p(Y, X)$. The joint distribution factorizes as $p(Y, X) = p(Y|X)p(X)$ under the usual assumptions.[4] With this definition, a task can change

[3]See Methods for definitions.

[4]For related task- and transfer-learning definitions

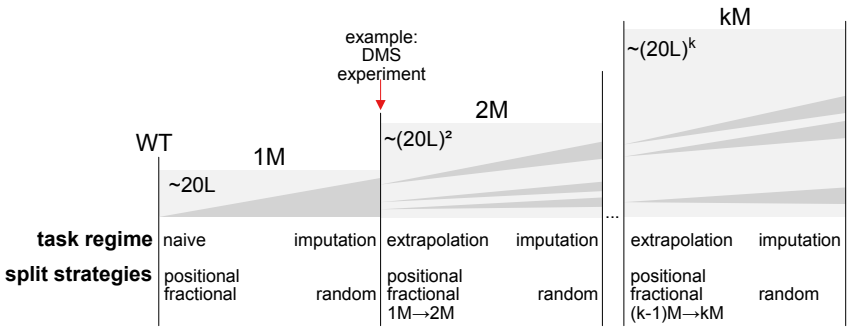

Figure 4: Schematic data over domains of variants by number of mutations (1...kM) relative to sequence wildtype (WT). Areas indicate observed variant samples (dark) against unexplored (light). Not all possible sequences can be explored for DMS experiments, due to natural constraints. The *task-regime* presents tasks applicable relative to the domain and the *split strategies* are potentially suitable protocols to assess performance on the domain by constructing training and test sets.

either through a covariate shift, or through a change in the likelihood $p(Y|X)$, *i.e.* a change in the fundamental relationship between X and Y - sometimes referred to as *concept shift*.

We fit a supervised model on a sample from a domain with the intent to make predictions related to an intended task, $f(\cdot) : \mathbb{R}^d \to \mathbb{R}$. The sample of size n for the training of the supervised model is denoted $S|_{\text{tr}} = \{(X_i, y_i)\}_{i=1..n}$. The *task* respective distribution should correspond to a downstream application, and is represented through a test set $\tilde{X} \subset \mathbb{R}^d$, $\tilde{y} \subset \mathbb{R}$, $S|_{\text{te}} = \{(\tilde{X}_j, \tilde{y}_j)\}_{j=1..n_{\text{test}}}$ with distributions $p(\tilde{X})$, $p(\tilde{Y}|\tilde{X})$, respectively. The generalization capabilities of the supervised model are assessed by residuals between the true observations and the predictions, as well as the calibration of predictive uncertainties on the test data.

The concept of domain is important in a protein engineering setting, because there is an inherent need for extrapolation. A typical engineering pipeline will gradually move further away from one or more initial (WT) sequences by introducing an increasing number of mutations. As an example of how these concepts can come into play, Fig. 4 contains a schematic of the available data in a typical protein engineering setting. The progression from left to right illustrates an exploration of variant space, organized by the number of mutations away from the wildtype (1M, 2M, ..., kM). Within each block, the dark shade denotes an increasing percentage of cumulative variants explored. If we assume that the variants within each class are sampled according to some distribution, we can consider each block as representing a domain, for which we sample an increasing number of data points (left to right). The *DMS* arrow demonstrates the typical scenario explored in (Fig. 2) where a Deep Mutational Scan (DMS) experiment provides us with a fairly complete set of single-variants. If we wish to predict missing values within the 1M class (and assume that these values are missing at random), then our task would be an imputation task. For this task, a random splitting strategy would be appropriate. In contrast, if we wish to make predictions in the 2M class, having only observed variants from the 1M class, we would be in an extrapolation setting, which is inherently a more difficult modelling task. A test set constructed on the 1M data using random splitting will not give a reliable estimate of our performance in this 2M prediction setting. The position-level splitter, which tests the ability of the model to make predictions at one position given only information at other positions, is likely to be a more useful proxy for the performance on this task.

When observing experimental data for 2M and higher order variants, we typically also face the problem that the data we encounter are biased. While the initial round of obtained variants is often sampled relatively uniformly (e.g. in a site saturation experiment), the subsequent double- and higher-order variants will generally be selected based on the best performing variants observed in earlier rounds (Fig. 4, dark vs. light grey regions). Therefore, even if we define our domain to be all variants up to a maximum mutation order k, the sample bias in the data makes it difficult to reason about expected generalization. Another way of seeing this is that if we define our domain as a uniform distribution of variants up to order k, a site-saturation experiment of all 1M variants is a heavily biased sample.

see [32;33;34].

Assessing Performance: Train/Test Splits

In our experiments comparing representations and regressors (Fig. 2A and B), we saw a clear tendency towards lower performance when splitting by *positions* of a sequence compared to splits *at random*. We now discuss these two splitting strategies, in addition to several alternatives, in the light of the discussion above.

Random CV. This random splitting protocol is a standard k-fold cross-validation protocol that shuffles training data uniformly at random, separates (X, y) into k folds, and conducts k times testing on one of the folds having trained on the remainder. When dealing with naturally occurring biological sequences, random splitting is often considered inappropriate due to the evolutionary structure in the data, and homology reduction schemes are therefore used for testing to better reflect how the model will generalize to sequences from species not included in the training set. For protein engineering where data consists of artificially constructed sequences, we do not have this concern, and there is therefore *a priori* no reason for disregarding random splitting. However, as argued above, it is appropriate only for imputing values that are missing at random, and will generally produce overly optimistic estimates of our performance beyond the imputation task.

Positional CV. This protocol partitions the sequence into segments of size p, such that we create $\frac{L}{p}$ number of splits. All variants with mutations that are in the range of positions in a segment are held-out for testing, while training is done on the remaining data. More specifically, let a sequence $X = \{x_0, ..., x_k, ..., x_L\}$ of length L have a mutation at position k. The training set is then composed of the sequences that do not have a mutation at the given position range: $S|_{\text{tr}} = \{(X_i, y_i)\}_{i=1..N, k|k \notin [j, j+p]}$ (with partition starting index j) and $S|_{\text{te}}$ is the complement of that set.[5] This protocol therefore allows us to test generalization *beyond* the sequence positions observed during training. Rather than relying on a site-specific signal (the equivalent of conservation in natural sequences), the model must exploit pairwise or higher order signals. Since this assesses the ability of the model to predict *in the context* of the other amino acid present in the sequence, we might expect it to better reflect the performance we can expect in an extrapolation setting.

[5]For computational feasibility, we consider positions in blocks of $p = 15$.

Mutational CV. This splitting strategy probes how well we generalize from a lower number of mutations to a higher number of mutations in the sequence. The protocol tests on a sample of degree-k variants in the domain kM containing m variants, $S|_{\text{te}} = \{(X_j^{(kM)}, y_j^{(kM)})\}_{j=1..m}$, after having trained on variants of degree-k' where $k' \leq k$, $S|_{\text{tr}} = \{(X_i^{(k'M)}, y_i^{(k'M)})\}_{i=1..n}$. In addition to an inherent shift in $p(X)$ there will often be an associated shift in $p(Y|X)$ when switching domains, as discussed above. The usefulness of the Mutational CV approach will thus depend on whether the variants predicted in expected application of the model are distributed similarly as the successive mutation degrees observed so far, i.e. if a similar selection protocol is used to select substitutions. Obviously, this splitting strategy is only feasible if the available data contains multiple substitutions, which is not the case for any of the data discussed so far. We will analyze such a case in detail below.

Fractional CV. We introduce this protocol to assess regressor performance as more data becomes available - from a naive, few observation setting, to the nearly full-information, imputation setting. The protocol is a k-fold CV protocol that sub-samples train and test data from the total available sequences uniformly at random. We sample each fraction $q \in [0, 1]$ uniformly at random
$S|_{\text{tr}} \sim \{(X_i^{(j)}, y_i^{(j)})\}_{i=1..\lceil q \times N \rceil}^{j=1..k}$ with k train, test iterations. This splitting strategy assesses the expected regressor performance when faced with different amounts of data within the domain. This setting we expect to occur during the initial steps of an optimization protocol, given that the selected observations are sufficiently close to the initial sequence.

Optimization. Our final assessment strategy consists of an actual optimization protocol, where we use a Bayesian optimization strategy to actively select samples based on previous observations via the *Expected Improvement*[35] acquisition function. The protocol is similar to the Fractional-CV splitter, but replaces the simplifying assumption of uniform sampling of candidates (which is guaranteed to stay in-domain) with an active selection of candidates, thus becoming biased toward higher performance. Note that this protocol still optimizes only considering the known data points, and the task is to find the optimal value in fewest iterations.

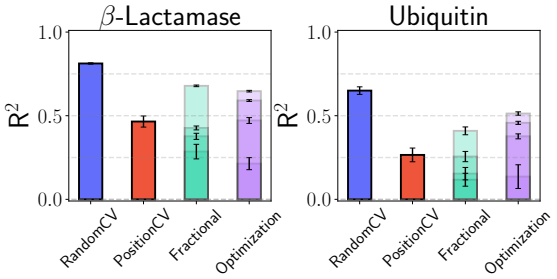

Figure 5: Performance comparison across tasks. GPmatern52 regressor performance (R2) for the ESM-1B representation of β-LACTAMASE and UBIQUITIN, when evaluated splitting at random (blue), by sequence positions (red), across fractions of the data (green), and optimization (purple). Fractional and optimization results are relative to available training data (four partitions each), with mean and std.err. reported respectively (from first partition (lowest expected value) to all available data (last partition, high expected value).

Case 1: In-domain optimization

We illustrate the Fractional-CV and Optimization splitter on the same datasets as before, comparing the performance in a single-mutation data setting, and increasing levels of data completeness (Fig. 5). We see that the Fractional-CV split and Optimization protocol give even more conservative assessments of the total performance (especially with only few available observations) than the earlier assessment criteria. Among the selection of regressors that we tested, the GPs again perform most reliably on this task, finding the optimal candidate after fewest iterations (SM Fig. 16). This might be explained by the fact that the acquisition function of the optimization protocol relies actively on the uncertainty estimates provided by the regressors. We also note that random selection and sorting by EVE-score are competitive reference baselines.

Case 2: Extrapolation

To illustrate some of the pitfalls that can occur when extrapolating to different domains, we conduct an analysis on a DMS experiment on PARDANTITOXIN[36]. This data set contains variants with multiple mutations, and is an example of the bias scenario introduced earlier, where all the triple and quadruple variants share the same initial mutation sites. Furthermore, as we will see, the different mutation degrees also have different label distributions.

In our analysis, we will use the *mutational CV* strategy introduced above, with the goal of estimating how well we can predict variants at increasing distance from the wild type. Fig. 6 demonstrates the performance obtained when training regressor models using data up to (or up to and including) the kth mutation degree, and testing on the kth. If we consider the mean square error (MSE) scores (Fig. 6a) (see Eq. 4), the results behave as expected: when predicting on the 2M (or 3M) data, our performance increases if we include a randomly selected set of 2M (or 3M) variants in the training set as well. However, if we instead consider Spearman rank correlation (see Eq. 5) as our metric, we see more surprising behavior: in the 3M case, our performance *declines* if we include 3M samples in the training set. This inconsistent behavior is caused by the fact that the label distribution $p(Y)$ changes dramatically between the 1M, 2M, 3M and 4M settings. In the 1M case, we see a bimodal distribution between non-functional and functional with the latter being dominant. In 2M, 3M and 4M cases this switches more and more dramatically to the non-functional case, increasingly centering around a single value. As we move from 2M to 4M, the MSE therefore increasingly reports on the residuals around an almost constant prediction (in this dataset, none of the 4M mutations are functional). Since the empirical variance of the distributions drops, so do these residuals, giving a misleading impression of improved performance. Note that in this discussion we used MSE rather than R2 to keep the discussion simple, since we for the R2 would see shifts both in the numerator and the denominator (*i.e.*, the performance of the reference mean predictor also changes). See SM Fig. 12 for the equivalent R2 plot.

In the Spearman correlation results we include as baseline the results from an additive model (dashed line), such that an observation of a double variant is predicted by the sum of the observations of its constituents. This simple procedure is often used in practice in engineering pipelines, in particular when optimizing for stability, where a sum of independent site contributions can be a reasonable approximation. We see that for this example, the additive model is in many cases a reasonable baseline for ranking candidates, although the additive values do not constitute competitive predictions in terms of MSE (SM Fig. 13). Since ranking candidates is often a primary concern, this illustrates the importance of including such site-independent baselines when assessing regressor performance in the multi-variant setting.

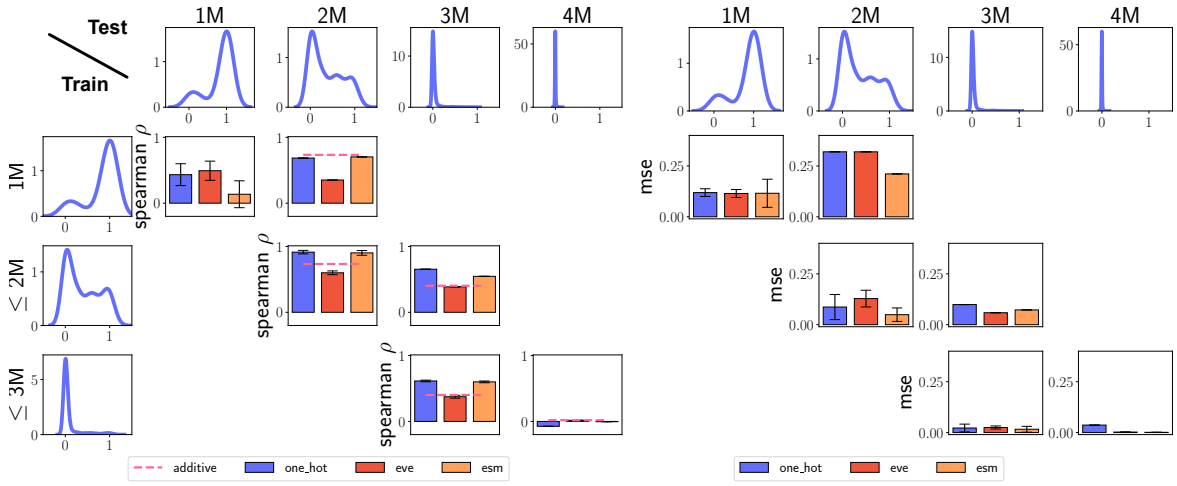

Figure 6: *mutation*-degree protocol on GP (Matern52) from first-degree (1M) to fourth-degree (4M) of ParD-antitoxin represented by ESM-1b. Functional observations of ParD-antitoxin are density curves in first row and first columns respectively. We assess rank correlation (spearman ρ) across in and out of domain (left test columns 1M to 4M). The reference is an additive benchmark baseline (dashed line), which is the addition for all variants constituting the target variants. The baseline for second-degree variants: y(Var1,Var2) = y(Var1)+y(Var2), and for triple variants the set of all combinations of the constituents. We assess accuracy (MSE) across in-domain and out of domain (right test columns 1M to 4M). The diagonal shows in-domain performance, such that we learn on sequences with less than or equal number of mutations, with a standard 5-fold CV protocol (given equal number of mutations randomly selected 20%). The off-diagonal for each source data domain (x-axis) shows what we predict in the next domain.

Case 3: Free optimization

So far, we have considered our test set as a perfect representation of a downstream task: our assessment measures how well we will do if the data in a downstream task is distributed identically to our test set. In reality, however, in a protein engineering pipeline, we conduct *free* (*i.e.*not bound to a predefined list of candidates) optimization where we *propose* candidates, and evaluate our regression algorithm on these proposals. It is therefore our responsibility to choose proposals such that they do not deviate too much from the domain on which our regressor has been trained and validated. We can address this in various ways, either: 1) ensuring that the procedure used to select new candidates is restricted to points close to the training set, 2) assessing the proximity to the training domain by modelling the input distribution directly (*i.e.* with a density model $p(X)$), or 3) employing a regression model which will associate out-of-domain predictions with high degrees of uncertainty. Since the topic of our paper is regression, we will focus on the last point here, but note that examples of the first two options exist in recent protein optimization methods, where lists of candidates are generated to be close to wild type proteins, for instance using a generative model of $p(X)$ [26;27].

The question is thus if the uncertainties predicted by our regression algorithms are accurate enough to distinguish between useful and useless predictions. We saw that within the 1M domain, our different regression methods were all reasonably well calibrated (Fig. 3), but since we only had access to 1M data, that analysis did not probe the out-of-distribution behavior. As a simple sanity check, we would expect that the uncertainties produced by our selection of regressors generally increases as we predict on sequences at larger distances from the wildtype. Unfortunately, we observe that this is only partially the case (SM Fig. 17), as most of the methods produce constant uncertainty values at increasing distances. We had expected the GPs to shine in this area, but found that only the linear GP had a relative increase in predictive variance with respect to added mutations. Furthermore, the predicted variance values differ in orders of magnitude depending on the type of regressor and underlying representation.[6]

The ParD-antitoxin data allows us to investigate this effect in greater detail on real data. Here, we observe better calibration with more data (SM Fig. 15), with generally a

[6]This observation is partially explained by a property of polynomial kernel functions (such as the linear kernel): the prior variance grows with the norm of the input[37 p. 90] which is not the case for stationary kernel functions (such as squared exponential and Matérn$\frac{5}{2}$).

larger deviation from a perfect calibration compared to the previous assessment from data such as β-Lactamase, Ubiquitin, *etc.* (SM Fig. 14). However, the Gaussian Process models seem to most robustly quantify this effect for the multi-variant ParD-antitoxin observations; such that they are better calibrated in this extrapolation setting.

An important consideration when discussing uncertainty quantification in protein optimization is that regression arises in two different contexts: 1) as surrogate models, for instance internally in a Bayesian Optimization protocol, and 2) as oracles, which sometimes serve as the objective of optimization when experiments cannot be conducted, frequently used for method development of improved optimization schemes. Uncertainty plays different roles in these two settings. In the surrogate model setting, the uncertainty of predictions can reflect both epistemic (model) uncertainty and aleatoric (data) uncertainty. Our focus in Bayesian optimization is the epistemic uncertainty, which we wish to reduce by making additional observations. In contrast, in the oracle setting, the regressor is fixed during optimization, and serves as artificial, (*in silico*) experimental data. The uncertainty of a prediction made by the oracle should therefore be considered an aleatoric uncertainty, because it is irreducible once the oracle is trained. In the Bayesian optimization setting, we expect evaluations of the surrogate function regressor exactly in regions where the uncertainty is high (depending on the acquisition function), while in an oracle setting, it is appropriate to focus on regions where the regressor produces predictions with low uncertainty. A similar situation could arise in a real experimental setting if the experimental setup is known to provide very noisy outcomes in certain parts of the domain. In such cases, it would make sense in a Bayesian optimization protocol to incorporate the aleatoric uncertainty into the objective and acquisition function.

When using a regression algorithm as an oracle, an additional step of assessment is necessary to assess the candidates found during optimization, since we have no guarantees about the quality of an oracle evaluation. In particular, if the oracle was trained only on a limited size, system-specific dataset, we should be concerned that the optimization procedure has optimized for an extrapolation artefact of the oracle, rather than a signal from the data. In general, the free optimization setting implies that we have no data for the candidate sequence selected, so standard train/test splitting is not applicable. Ideally, experiments should be conducted at this stage to validate the candidates. However, in computational labs such experiments might not be accessible, which has led research towards finding computational proxies. One approach that has been suggested is to train two oracles on subsets of the data, optimize against one, and validate against the other - averaging over both permutations[27]. Ideally, the model assumptions should be different between the two oracles, to ensure that extrapolation artefacts are different between them. This approach to quantify the epistemic uncertainty of the oracle is similar to those built into the GP and RF regressors, but incorporates uncertainty about the model class as well, and could therefore provide more useful uncertainty quantification in the extrapolation setting.

Discussion

In this study, we have identified the impact of some central components in a protein regression protocol: the data, the representations chosen to encode the data, the choice of model, and the assessment criteria. Our results confirm that some choices of representation can have positive impact on downstream regression performance, in particular those of large language models.

The choice of regression algorithm itself seems to be less critical for performance, at least in the data-limited regime that we study here, and under the assumption that hyper-parameters of these models have been carefully tuned. This choice might therefore be guided by other considerations, such as whether the model can provide meaningful uncertainty estimates. In our experiments, Gaussian Process regression turned out to be particularly well suited for this task, although this is by no means the only option.

Of all the components in a regression protocol, assessment is perhaps the most critical. We demonstrate that performance can change dramatically depending on how data is split into train/validation/test sets. While this result is perhaps unsurprising, it warrants our attention, since performance values are routinely reported when new methods are published. If splits are not defined in the same way, comparing methods is meaningless. Equally important, if splits are defined too simplistically, we risk overestimating our abilities. A choice of splitting strategy implies an expected downstream task for which we are interested in quantifying expected performance. We give several examples of splitting strategies that attempt to reflect the regression scenario faced in a protein engineering protocol. The choice of assessment strategy is, however, complicated by selection bias in the data, and the inherent drive towards extrapolation when conducting protein engineering. We argue that

for this reason, it is essential that we have the ability to associate predictions with reliable uncertainty estimates, and that we should evaluate regression methods on their ability to produce calibrated uncertainties.

Although we have stressed the importance of reliable uncertainty estimates in our analysis, the methods we considered were often lacking - especially when applied in an extrapolation setting. For Gaussian Process regression, this might be due to the high dimensionality of the inputs, and interesting future work could be to investigate whether lower dimensional approximations to learned embeddings might behave better in this setting. An alternative technique for uncertainty quantification that could be worth exploring is *conformal prediction*. In particular, recent work in this area presents a solution to the the bias that arises during optimization when iteratively adding proposals from a regressor to the available training data by an acquisition function [38].

Our previous discussion of domains and generalization indicate that it might not be fruitful to think about domains in terms of raw amino acid sequences and the number of mutations that a sequence is separated from the *WT*. If a method generalizes from 2M to 3M variants, it is presumably because it relies on *features* of the raw input that occur in both the 2M and 3M dataset, *e.g.* specific sequence patterns, charge distribution, or a combination of amino acids forming a stabilizing bond. Such features can easily be preserved between sequences deviating by tens of mutations. In a limited data regime, it can be difficult to reliably learn biochemically relevant features from the available data. Embeddings from *pre-trained* language models provide a potential shortcut for extracting such features, which explains why using them as input to a regression algorithm might improve performance. It is likely that these pre-learned representations also provide a richer notion of locality than that provided by raw input space. By making local perturbations in the representation of a protein, and decoding back to input space, one can hope to generate variants that are functionally or structurally close, but potentially more divergent in sequence space. This strategy has been employed by recent protein engineering pipelines [26]. A similar approach might prove useful in the future as a more robust basis for the definition of splits.

Given the large differences in performance arising from different split strategies and task definitions, we encourage the community to work towards standardized benchmarks that include splits of data that are well-motivated by specific biological tasks, including the extrapolative setting. Encouragingly, recent contributions seem to be moving in this direction. Apart from a standard cross-validation approach [17;39;25;19;40;41;21], previous work includes different hold-outs [11], homology splitting [13], mutational or blocks of positions splitting [42;22], and task-inspired splitting [22;23]. We stress that such benchmarks should include meaningful baselines (such as the additive model described above), and also rank models both by their ability to quantify prediction performance and uncertainty quantification. This will ensure that we can meaningfully compare performance of regression models, and provide practitioners with a realistic expectation of the performance of the state-of-the-art.

References

[1] Peter Bross, Thomas J Corydon, Brage S Andresen, Malene M Jørgensen, Lars Bolund, and Niels Gregersen. Protein misfolding and degradation in genetic diseases. *Human mutation*, 14(3):186–198, 1999.

[2] Rosamonde E Banks, Michael J Dunn, Denis F Hochstrasser, Jean-Charles Sanchez, Walter Blackstock, Darryl J Pappin, and Peter J Selby. Proteomics: new perspectives, new biomedical opportunities. *The Lancet*, 356(9243):1749–1756, 2000.

[3] James A Brannigan and Anthony J Wilkinson. Protein engineering 20 years on. *Nature Reviews Molecular Cell Biology*, 3(12):964–970, 2002.

[4] Faruck Morcos, Nicholas P Schafer, Ryan R Cheng, José N Onuchic, and Peter G Wolynes. Coevolutionary information, protein folding landscapes, and the thermodynamics of natural selection. *Proceedings of the National Academy of Sciences*, 111(34):12408–12413, 2014.

[5] Pengfei Tian and Robert B Best. How many protein sequences fold to a given structure? a coevolutionary analysis. *Biophysical journal*, 113(8):1719–1730, 2017.

[6] Thomas A Hopf, John B Ingraham, Frank J Poelwijk, Charlotta PI Schärfe, Michael Springer, Chris Sander, and Debora S Marks. Mutation effects predicted from sequence co-variation. *Nature biotechnology*, 35(2):128–135, 2017.

[7] Pengfei Tian, John M Louis, James L Baber, Annie Aniana, and Robert B Best. Coevolutionary fitness landscapes for sequence design. *Angewandte Chemie International Edition*, 57(20):5674–5678, 2018.

[8] Adam J Riesselman, John B Ingraham, and Debora S Marks. Deep generative models of genetic variation capture the effects of mutations. *Nature Methods*, 15(10):816–822, 2018.

[9] Pengfei Tian and Robert B Best. Exploring the sequence fitness landscape of a bridge between protein folds. *PLoS computational biology*, 16(10):e1008285, 2020.

[10] Jonathan Frazer, Pascal Notin, Mafalda Dias, Aidan Gomez, Joseph K Min, Kelly Brock, Yarin Gal, and Debora S Marks. Disease variant prediction with deep generative models of evolutionary data. *Nature*, 599(7883):91–95, 2021.

[11] Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Jason Liu, Demi Guo, Myle Ott, C Lawrence Zitnick, Jerry Ma, et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15), 2021.

[12] Ethan C Alley, Grigory Khimulya, Surojit Biswas, Mohammed AlQuraishi, and George M Church. Unified rational protein engineering with sequence-based deep representation learning. *Nature methods*, 16(12):1315–1322, 2019.

[13] Roshan Rao, Nicholas Bhattacharya, Neil Thomas, Yan Duan, Peter Chen, John Canny, Pieter Abbeel, and Yun Song. Evaluating protein transfer learning with tape. *Advances in neural information processing systems*, 32, 2019.

[14] Pascal Notin, Mafalda Dias, Jonathan Frazer, Javier Marchena Hurtado, Aidan N Gomez, Debora Marks, and Yarin Gal. Tranception: protein fitness prediction with autoregressive transformers and inference-time retrieval. In *International Conference on Machine Learning*, pages 16990–17017. PMLR, 2022.

[15] Zachary Wu, SB Jennifer Kan, Russell D Lewis, Bruce J Wittmann, and Frances H Arnold. Machine learning-assisted directed protein evolution with combinatorial libraries. *Proceedings of the National Academy of Sciences*, 116(18):8852–8858, 2019.

[16] Christopher A Voigt, Stephen L Mayo, Frances H Arnold, and Zhen-Gang Wang. Computational method to reduce the search space for directed protein evolution. *Proceedings of the National Academy of Sciences*, 98(7):3778–3783, 2001.

[17] Philip A. Romero, Andreas Krause, and Frances H. Arnold. Navigating the protein fitness landscape with gaussian processes. *Proceedings of the National Academy of Sciences*, 110(3):E193–E201, 2013.

[18] Kevin K Yang, Zachary Wu, and Frances H Arnold. Machine-learning-guided directed evolution for protein engineering. *Nature methods*, 16(8):687–694, 2019.

[19] Bruce J Wittmann, Kadina E Johnston, Zachary Wu, and Frances H Arnold. Advances in machine learning for directed evolution. *Current opinion in structural biology*, 69:11–18, 2021.

[20] Nicki Skafte Detlefsen, Søren Hauberg, and Wouter Boomsma. Learning meaningful representations of protein sequences. *Nature communications*, 13(1):1–12, 2022.

[21] Chloe Hsu, Hunter Nisonoff, Clara Fannjiang, and Jennifer Listgarten. Learning protein fitness models from evolutionary and assay-labeled data. *Nature Biotechnology*, 2022.

[22] Christian Dallago, Jody Mou, Kadina E Johnston, Bruce J Wittmann, Nicholas Bhattacharya, Samuel Goldman, Ali Madani, and Kevin K Yang. Flip: Benchmark tasks in fitness landscape inference for proteins. *bioRxiv*, pages 2021–11, 2022.

[23] Peter Mørch Groth, Richard Michael, Pengfei Tian, Jesper Salomon, and Wouter Boomsma. Flop: Tasks for fitness landscapes of protein families using sequence-and structure-based representations. In *OpenReview*, 2023.

[24] Ahmed Elnaggar, Michael Heinzinger, Christian Dallago, Ghalia Rihawi, Yu Wang, Llion Jones, Tom Gibbs, Tamas Feher, Christoph Angerer, Martin Steinegger, et al. Prottrans: towards cracking the language of life's code through self-supervised deep learning and high performance computing. *arXiv preprint arXiv:2007.06225*, 2020.

[25] Surojit Biswas, Grigory Khimulya, Ethan C Alley, Kevin M Esvelt, and George M Church. Low-n protein engineering with data-efficient deep learning. *Nature Methods*, 18(4):389–396, 2021.

[26] Samuel Stanton, Wesley Maddox, Nate Gruver, Phillip Maffettone, Emily Delaney, Peyton Greenside, and Andrew Gordon Wilson. Accelerating bayesian optimization for biological sequence design with denoising autoencoders. In *International Conference on Machine Learning*, pages 20459–20478. PMLR, 2022.

[27] Minji Lee, Luiz Felipe Vecchietti, Hyunkyu Jung, Hyunjoo Ro, Meeyoung Cha, and Ho Min Kim. Protein sequence design in a latent space via model-based reinforcement learning. In *MLSB Workshop*, 2023.

[28] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR, 2017.

[29] Gabriele Scalia, Colin A Grambow, Barbara Pernici, Yi-Pei Li, and William H Green. Evaluating scalable uncertainty estimation methods for deep learning-based molecular property prediction. *Journal of chemical information and modeling*, 60(6):2697–2717, 2020.

[30] Brian Hie, Bryan D Bryson, and Bonnie Berger. Leveraging uncertainty in machine learning accelerates biological discovery and design. *Cell systems*, 11(5):461–477, 2020.

[31] Marc C Kennedy and Anthony O'Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001.

[32] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.

[33] Masashi Sugiyama and Motoaki Kawanabe. *Machine learning in non-stationary environments: Introduction to covariate shift adaptation*. MIT press, 2012.

[34] Jochen Garcke and Thomas Vanck. Importance weighted inductive transfer learning for regression. In *Joint European conference on machine learning and knowledge discovery in databases*, pages 466–481. Springer, 2014.

[35] D.R. Jones, M. Schonlau, and W.J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4):455–492, 1998.

[36] Christopher D Aakre, Julien Herrou, Tuyen N Phung, Barrett S Perchuk, Sean Crosson, and Michael T Laub. Evolving new protein-protein interaction specificity through promiscuous intermediates. *Cell*, 163(3):594–606, 2015.

[37] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian Processes for Machine Learning*. MIT Press Cambridge, MA, 2006.

[38] Clara Fannjiang, Stephen Bates, Anastasios N Angelopoulos, Jennifer Listgarten, and Michael I Jordan. Conformal prediction under feedback covariate shift for biomolecular design. *Proceedings of the National Academy of Sciences*, 119(43):e2204569119, 2022.

[39] Stanislav Mazurenko, Zbynek Prokop, and Jiri Damborsky. Machine learning in enzyme engineering. *ACS Catalysis*, 10(2):1210–1223, 2019.

[40] Tristan Bepler and Bonnie Berger. Learning the protein language: Evolution, structure, and function. *Cell Systems*, 12(6):654–669, 2021.

[41] Céline Marquet, Michael Heinzinger, Tobias Olenyi, Christian Dallago, Kyra Erckert, Michael Bernhofer, Dmitrii Nechaev, and Burkhard Rost. Embeddings from protein language models predict conservation and variant effects. *Human genetics*, pages 1–19, 2021.

[42] Emmi Jokinen, Markus Heinonen, and Harri Lähdesmäki. mgpfusion: predicting protein stability changes with gaussian process kernel learning and data fusion. *Bioinformatics*, 34(13):i274–i283, 2018.

[43] Michael A Stiffler, Doeke R Hekstra, and Rama Ranganathan. Evolvability as a function of purifying selection in tem-1 β-lactamase. *Cell*, 160(5):882–892, 2015.

[44] David Mavor, Kyle Barlow, Samuel Thompson, Benjamin A Barad, Alain R Bonny, Clinton L Cario, Garrett Gaskins, Zairan Liu, Laura Deming, Seth D Axen, et al. Determination of ubiquitin fitness landscapes under different chemical stresses in a classroom setting. *Elife*, 5:e15802, 2016.

[45] Jochen Weile, Song Sun, Atina G Cote, Jennifer Knapp, Marta Verby, Joseph C Mellor, Yingzhou Wu, Carles Pons, Cassandra Wong, Natascha van Lieshout, et al. A

framework for exhaustively mapping functional missense variants. *Molecular systems biology*, 13(12):957, 2017.

[46] Yvonne H Chan, Sergey V Venev, Konstantin B Zeldovich, and C Robert Matthews. Correlation of fitness landscapes from three orthologous tim barrels originates from sequence and structure constraints. *Nature communications*, 8(1):1–12, 2017.

[47] Gregory M Findlay, R Daza, BK Martin, Melissa D Zhang, Anh P Leith, Molly Gasperini, Joseph D Janizek, Xingfan Huang, Lea M Starita, and Jay Shendure. Accurate functional classification of thousands of brca1 variants with saturation genome editing. *bioRxiv*, 2019.

[48] Liat Rockah-Shmuel, Ágnes Tóth-Petróczy, and Dan S Tawfik. Systematic mapping of protein mutational space by prolonged drift reveals the deleterious effects of seemingly neutral mutations. *PLoS computational biology*, 11(8):e1004421, 2015.

[49] Belur V Dasarathy. Nearest neighbor (nn) norms: Nn pattern classification techniques. *IEEE Computer Society Tutorial*, 1991.

[50] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[51] Leo Breiman. Random forests. *Machine learning*, 45:5–32, 2001.

[52] Thomas Minka. Bayesian linear regression. Technical report, Citeseer, 2000.

[53] Alexander G de G Matthews, Mark Van Der Wilk, Tom Nickson, Keisuke Fujii, Alexis Boukouvalas, Pablo León-Villagrá, Zoubin Ghahramani, and James Hensman. Gpflow: A gaussian process library using tensorflow. *Journal of Machine Learning Research*, 18(40):1–6, 2017.

[54] Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature Methods*, 17(3):261–272, 2020.

[55] Jonathan Wenger, Nicholas Krämer, Marvin Pförtner, Jonathan Schmidt, Nathanael Bosch, Nina Effenberger, Johannes Zenn, Alexandra Gessner, Toni Karvonen, François-Xavier Briol, et al. Probnum: Probabilistic numerics in python. *arXiv preprint arXiv:2112.02100*, 2021.

[56] David T Jones. Setting the standards for machine learning in biology. *Nature Reviews Molecular Cell Biology*, 20(11):659–660, 2019.

[57] Ryan J Tibshirani. Degrees of freedom and model search. *Statistica Sinica*, pages 1265–1296, 2015.

Methods

Computing Representations Seven distinct proteins were extracted from the ProteinGym data-set (see[14]).

The individual data is published in the following references.: β-LACTAMASE[43], UBIQUITIN[44], CALMODULIN[45], TIM-BARREL[46], BRCA1[47], T2-MTH[48], PARD-ANTITOXIN[36].

The representations are obtained by different means.: EVE is constructed from a single protein family per protein (the UNLABELLED sequences in Table 1,2), for which we compute the latent space via the encoder of the EVE ([10] from github.com/OATML-Markslab/EVE) model. We include the log-likelihood (ELBO) of the model, which has proven highly informative in an unsupervised setting (referred to as evo-score ($d = 1$)). To compute PROTBERT[24] and ESM-1B[11] representations, we reduce the high dimensionality of the language model by taking the mean over the sequence positions from the last attention layer of the transformer model as our representation (as is frequently done in practice). We compare these learned representations with a simple ONE-HOT encoding of the raw input sequence.

Optimizing Regressors We make the fairly common choice to standardize the observations in the training sample as $\bar{y} = \frac{y-\mu_y}{\text{std}(y)}$, where μ_y and $\text{std}(y)$ are the mean and standard deviation.

As regressors we fit:

1. a kNN regressor[49], as an example of a non-parametric, deterministic baseline, for which we optimize the number of neighbors[50],

2. a Random Forest (RF) as an example of a parametric, ensemble-based regressor[51], we implement the sklearn-implementation[50], which we extend with a predictive uncertainty estimate as $\hat{\sigma}^2 = \hat{y}^2 - \bar{\hat{y}}^2$;

3. a Gaussian Process (GP) regressor with: a linear kernel (equivalent to standard linear Bayesian regression[52]): $k_{\text{l}}(x, x') = \sum_{d=1}^{D} \sigma_d^2 x_d x'_d$;

4. a GP with a non-linear squared-exponential: $k_{\text{SE}}(x, x') = \sigma^2 \exp(-\frac{(x-x')^2}{2l^2})$;

5. a GP with a Matérn-52 kernel: $k_{\frac{5}{2}}(x, x') = \sigma^2(1 + \frac{\sqrt{5}(x-x')}{l} + \frac{5(x-x')^2}{3l^2}) \exp(-\frac{\sqrt{5}(x-x')}{l})$.

We compute exact GP regression[37] (using GPFlow[53]) we use a zero-mean prior $f(x) \sim \mathcal{GP}(0, k(x, x'))$. We assume Gaussian noise for each experimental observation, such that

$$y_n = f(x_n) + \epsilon, \epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2 \mathbf{I}). \tag{1}$$

The standardized observations do not affect RF or kNN predictions, and allows us to use a zero-mean prior for our GP models.

Given the GP setting, we can compute the marginal likelihood[37] for our training input X of size n, as

$$\begin{aligned} p(\bar{y}|X, \theta) = &-\frac{1}{2}\bar{y}^T [K_{XX} + \sigma_\epsilon^2 I]^{-1} \bar{y} \\ &- \frac{1}{2} \log |K_{XX} + \sigma_\epsilon^2 I| - \frac{n}{2} \log(2\pi). \end{aligned} \tag{2}$$

The hyperparameters of the GP covariance functions were chosen as $\theta = \{\sigma^2, l, \sigma_\epsilon^2\}$ [7], such that $\sigma^2 \sim \Gamma^{-1}(3, 3), l \sim \Gamma^{-1}(3, 3), \sigma_\epsilon^2 \sim \mathcal{U}[0.01, 1.0]$ are loosely constrained. Given the standardized observation this is the least imposing prior that results in a more robust optimization, as it punishes unlikely extreme values for length-scale and variance. For each model the log-marginal likelihood is optimized using an L-BFGS optimizer (max=500 iterations)[54].

We use skopt (scikit-optimize/stable) Bayesian optimization routine with an internal three-fold cross-validation, minimizing the negative mean absolute error to find the number of neighbors for the kNN regressor and number of estimators for the Random Forest[50]. The kNN regressor has $k \in [1, \lfloor 0.95\frac{N}{3} \rfloor]$ with an optimization budget of 75 using gradient boosted trees as internal surrogates (GBT). In the unoptimized setting we set number of neighbors $k = \lceil \frac{1}{3} N \rceil$. For the Random Forest regressor the number of estimators is $\in [2, N]$ with a budget of 15 (using GBT). A delta against Random Forest default parameters, and KNN fixed neighbors is computed (see SI Fig. 18).

Performance Metrics and Calibration We compute the accuracy by an adjusted R2 score:

$$R2 = 1 - \frac{\frac{1}{N}\sum_i (y_i - \hat{y}_i)^2}{\frac{1}{N}\sum_i (y_i - \mu_y)^2}. \tag{3}$$

When not normalizing by training signal we resort to the elementary error metric as mean squared residuals:

$$MSE = \frac{1}{N}\sum_i^N (y_i - \hat{y}_i)^2. \tag{4}$$

[7] Note that for the linear kernel $\theta_l := \theta \backslash \{l\}$

Table 1: Overview of the datasets.

DATASET	REF.	UNLABELLED	LABELLED N	LENGTH L
β-LACTAMASE	(43)	13191	4788	263
UBIQUITIN	(44)	34281	1255	76
CALMODULIN	(45)	37953	1729	149
TIM-BARREL	(46)	25264	1519	252
BRCA1	(47)	9575	1184	237
T2-MTH	(48)	28232	1719	330
PARD-ANTITOXIN	(36)	14672	9193	93

To compute correlations we use the Spearman rank correlation (based on the ranks r of the inputs):

$$\rho = \frac{\mathrm{cov}(\mathrm{r}(y)\mathrm{r}(\hat{y}))}{\sigma_{\mathrm{r}(y)}\sigma_{\mathrm{r}(\hat{y})}}. \tag{5}$$

We compute calibration (goodness of fit) as a reduced χ^2 statisticv:

$$\chi^2 = \frac{1}{N-1}\sum_i (\frac{y_i - \hat{y}_i}{\hat{\sigma}_i})^2. \tag{6}$$

For an implementation of the χ^2 statistic see the *evaluation* metrics in the *ProbNum* package[55]. As further uncertainty quantification assessments we compute confidence curves, as:

$$f_{\mathrm{cc}}(\mathbf{Y}) = \sum_{j=1}^{q}\frac{1}{N}\sum_{i=1}^{N} l(\hat{y}_i, y_i)\mathbf{1}\{\hat{\sigma}_i \leq h_j\}. \tag{7}$$

, the absolute deviation from calibration (ECE) and sharpness:

$$\mathrm{ece}(f_{\mathrm{cal}}) = \frac{1}{q}\sum_{j=1}^{q}|f_{\mathrm{cal}} - j|, \tag{8}$$

$$\mathrm{sharp}(\hat{\sigma}^2_{1\ldots n}) = \frac{\mathrm{std}(\hat{\sigma}^2_{1\ldots n})}{\frac{1}{n}\sum_i^n \hat{\sigma}_i^2}. \tag{9}$$

Bayesian Optimization and the unsupervised baseline We use the previously described regressors as surrogate functions and compute the expected improvement (EI)[56] as acquisition function, for the list of available labelled candidates from the dataset. The budget is 500 iterations, the regressor is optimized with the previously specified hyperparameters at each iteration. Ten different random seeds have been run for initial ordering of the available sequences. As reference baselines the EVE model evo-score was used to rank the available sequences (optimization results can be found in SM Fig. 16).

Acknowledgements

This work was in part funded by the Danish Data Science Academy, which is funded by the Novo Nordisk Foundation (NNF21SA0069429) and VILLUM FONDEN (40516). Further funding includes the Innovation Fund Denmark (1044-00158A), the MLSS Center (Basic Machine Learning Research in Life Science NNF20OC0062606), the Danish Ministry of Education and Science, Digital Pilot Hub, Skylab Digital, and the Pioneer Centre for AI (DRNF grant number P1).

Code and Data Availability

All code required to replicate the results can be found at github.com/MachineLearningLifeScience/protein_regression. All data, including computed representations are archived and publically available at ERDA .