VespaG: Expert-guided protein Language Models enable accurate and blazingly fast fitness prediction

Céline Marquet[1,✧,*], Julius Schlensok[1,✧], Marina Abakarova[2,5], Burkhard Rost[1,3,4] & Elodie Laine [2,6,*]

1 TUM (Technical University of Munich), Germany; TUM School of Computation, Information and Technology; Department of Informatics, Bioinformatics & Computational Biology - i12, Boltzmannstr. 3, 85748 Garching/Munich, Germany

2 Sorbonne Université, CNRS, IBPS, Laboratory of Computational and Quantitative Biology (LCQB), UMR 7238, 75005 Paris, France

3 Institute for Advanced Study (TUM-IAS), Lichtenbergstr. 2a, 85748 Garching/Munich, Germany

4 TUM School of Life Sciences Weihenstephan (TUM-WZW), Alte Akademie 8, Freising, Germany

5 Universite Paris Cité, INSERM UMR U1284, 75004 Paris, France

6 Institut universitaire de France (IUF)

* Corresponding authors: firstname.lastname@tum.de, firstname.lastname@sorbonne-universite.fr

✧ Céline Marquet and Julius Schlensok contributed equally to this work

Abstract

Exhaustive experimental annotation of the effect of all known protein variants remains daunting and expensive, stressing the need for scalable effect predictions. We introduce VespaG, a blazingly fast single amino acid variant effect predictor, leveraging embeddings of protein Language Models as input to a minimal deep learning model. To overcome the sparsity of experimental training data, we created a dataset of 39 million single amino acid variants from the human proteome applying the multiple sequence alignment-based effect predictor GEMME as a pseudo standard-of-truth. Assessed against the ProteinGym Substitution Benchmark (217 multiplex assays of variant effect with 2.5 million variants), VespaG achieved a mean Spearman correlation of 0.48±0.01, matching state-of-the-art methods such as *GEMME*, *TranceptEVE*, *PoET*, *AlphaMissense, and VESPA*. VespaG reached its top-level performance several orders of magnitude faster, predicting all mutational landscapes of the human proteome in 30 minutes on a consumer laptop (12-core CPU, 16 GB RAM).

Availability: VespaG is available freely at https://github.com/JSchlensok/VespaG

Main

Proteins are the essential building blocks of life, fulfilling a wide range of vital roles within cells and organisms. Hence, understanding the effect of variations such as point mutations (or single amino acid variants, SAVs) on protein stability and function is crucial for comprehending disease mechanisms (Murray, Laurieri, and Delgoda 2017) and modulating their activities through engineering. Multiplexed assays of variant effect (MAVEs), in particular deep mutational scans (DMS) (Fowler and Fields 2014), have enabled the quantification of mutational outcomes on a much larger scale than ever before. They allow for an in-depth characterization of protein mutational landscapes by assessing the impact of virtually all possible SAVs. Nevertheless, conducting experimental assays for entire proteomes remains elusive (Atlas of Variant Effects Alliance n.d.).

Leveraging the power of computational models helps to gain insights into the functional consequences of protein variants and prioritize them for further experimental validation. However, the sparse availability of annotated training data challenges the development of such models. While many supervised machine learning (ML) methods have proven accurate (Adzhubei, Jordan, and Sunyaev 2013; Gray et al. 2018; Hecht, Bromberg, and Rost 2015), they are inherently biased towards the limited number of proteins characterized by MAVEs or having annotated disease-associated variants (Atlas of Variant Effects Alliance n.d.; Livesey and Marsh 2023). As a result, different methods tend to correlate highly for the tiny subset of experimental data, while their predictions for, e.g., all possible SAVs in the human proteome correlate very poorly (Hecht, Bromberg, and Rost 2013; Mahlich et al. 2017). Prediction methods are also sensitive to the noise and uncertainty in these data. MAVE annotations, for instance, may vary substantially across experiments, even when measuring the same phenotype for the same protein (Reeb, Wirth, and Rost 2020). These difficulties have stimulated a growing interest in unsupervised or weakly supervised methods predicting variant effects by only exploiting information from protein sequences observed in nature (Ng and Henikoff 2003).

Among the best-performing unsupervised predictors, GEMME explicitly models the evolutionary history of protein sequences (Laine, Karami, and Carbone 2019; Notin et al. 2023). Starting from a multiple sequence alignment (MSA), it determines how protein sites are segregated along the topology of phylogenetic trees to quantify the sensitivity of each site to mutations and the number of changes required to accommodate a substitution. It relies on only a few biologically meaningful parameters and is robust to low variability in the input MSA. GEMME proved instrumental for investigating the

interplay between protein stability and function, and elucidating disease mechanisms (Abildgaard et al. 2023; Cagiada et al. 2023; Gersing et al. 2023; Tiemann et al. 2023; Tsuboyama et al. 2023). Combining GEMME with a fast MSA generation algorithm allows for producing proteome-wide substitution score matrices within a few days (Abakarova et al. 2023).

Other methods rely on protein Language Models (pLMs) pre-trained over large databases of raw sequences (Elnaggar et al. 2021; Lin et al. 2023). Although foundation pLMs are not competitive with state-of-the-art (SOTA) methods out of the box (Meier et al. 2021; Notin et al. 2023), their performances can be boosted through adding evolutionary or structural information (Cheng et al. 2023; Marquet et al. 2022; Meier et al. 2021; Nijkamp et al. 2022; Notin, Dias, et al. 2022; Truong Jr and Bepler 2023). Nevertheless, most pLM-based approaches remain computationally costly due to the necessity for systematically computing probability log-odd ratios for every substitution at every position.

Here, we optimized prediction speed by circumventing this computationally expensive step and directly mapping pLM embeddings to complete SAV mutational landscapes using the evolutionary-informed model GEMME as a teacher (Fig. 1). To this end, we trained a comparatively shallow (660k free parameters) neural network on top of a pre-trained pLM without computing log-odd ratios to learn GEMME predictions. Our strategy overcomes the bottleneck of sparsely annotated experimental training data. Moreover, it avoids the noise and inconsistencies of the experimental assays. We implemented our approach as a lean tool for fast **V**ariant **E**ffect **S**core **P**rediction without **A**lignments enabled by **G**EMME (VespaG). We assessed prediction performance against over 3 million (M) SAVs across diverse protein families. VespaG performed on par with SOTA methods and in some cases, the student even surpassed the teacher GEMME. As we circumvent the need to compute log-odds ratios of substitution probabilities, VespaG enables proteome-wide predictions in less than a half hour on a standard consumer laptop. We also demonstrated VespaG to generalize across organisms and protein families.

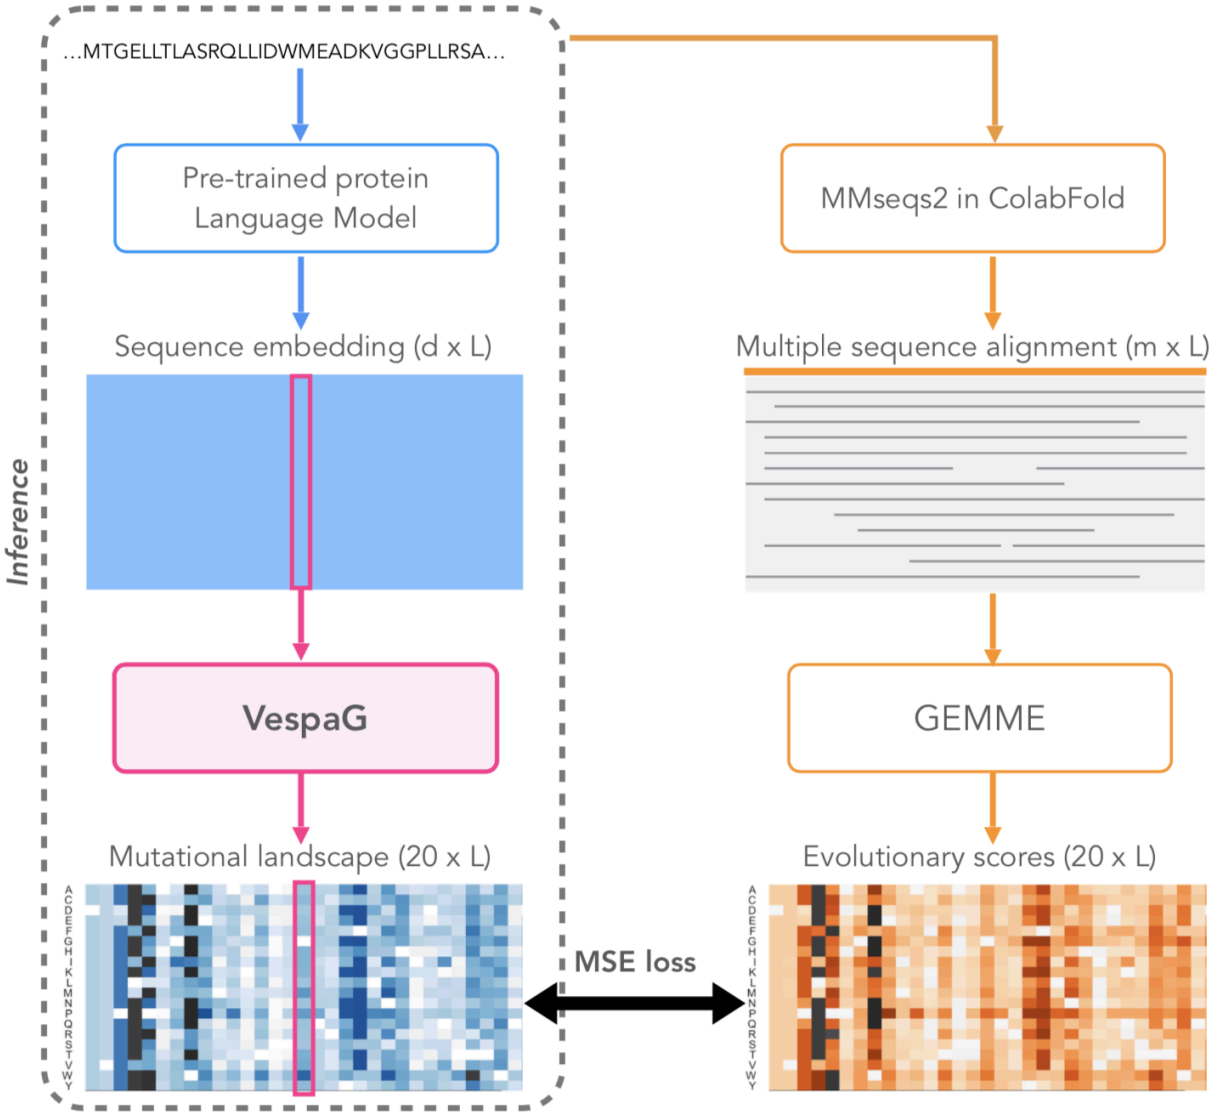

Figure 1: Outline for VespaG's expert-guided approach. VespaG takes as input a d-dimensional vector representation of a wild-type residue in a protein computed by a pre-trained protein language model and outputs a 20-dimensional vector of predicted mutational outcome estimates. The training loss measures the mean squared error between the predicted estimates and the evolutionary scores computed by GEMME (Laine, Karami, and Carbone 2019). We generate millions of training samples through the MMseqs2-based ColabFold protocol for searching and aligning sequences (Abakarova et al. 2023; Mirdita et al. 2022; Steinegger and Söding 2017). We do not use alignments at inference time (dotted rectangle).

Results

The method introduced in this work, VespaG, is a one-layer feed-forward neural network (FNN) inputting sequence embeddings from the protein language model (pLM) ESM-2 (Lin et al. 2023). We trained VespaG on a set of about 5,000 human proteins to learn a mapping between the input pLM embeddings and evolutionary scores computed by GEMME. The latter served as surrogates for mutational phenotypic outcomes. The proteins used for training represented a non-redundant subset of the human proteome (Methods and Supporting Online Material, SOM: Table S1, Fig. S2).

VespaG competitive with state-of-the-art

On a test set of 189 experimental multiplex assays of variant effect (MAVE) of the 161 eukaryotic and prokaryotic proteins in ProteinGym (*ProteinGymOrganismal189*), VespaG reached a average overall Spearman correlation coefficient (ρ) of 0.491±0.012 (±1.96 standard errors, i.e., 95% Confidence Interval, CI; Fig. 2A, Methods). Its average ρ was significantly higher than those of TranceptEVE L ($\Delta\rho$=0.013, one-tailed paired t-test p-value p-val=0.006) and VESPA ($\Delta\rho$=0.027, p-val=1.705e-19). VespaG also slightly outperformed its *teacher* GEMME ($\Delta\rho$=0.001), although the statistical significance of this difference was limited at a p-value higher than 0.1. Despite wide performance variations from one assay or protein to another, the four highlighted methods showed similar trends (SOM Fig.s S6-S7). VespaG's predictive accuracy stood out as less variable from one DMS to another ($\Delta\rho$=-0.011±0.036*)* than its embedding-based predecessor VESPA ($\Delta\rho$=-0.014±0.041). For instance, VESPA performed substantially worse ($\Delta\rho$<-0.1) than all other methods for some assays of the human proteins A4, OTC, PITX2, SOX30, RASH, RASK, RD23A and VKOR1, the mouse protein CSN4, the viral protein NRAM_I33A0, and the yeast protein HIS7 (SOM Fig. S7).

We obtained similar results on a subset of *ProteinGymOrganismal189*, dubbed *ProteinGymOrganismal66*, which contains eukaryotic and prokaryotic proteins of the first iteration of ProteinGym and for which results of the SOTA methods AlphaMissense and PoET were readily available (SOM Fig. S8). On *ProteinGymOrganismal66*, VespaG reached a predictive accuracy comparable with the more sophisticated predictors PoET ($\Delta\rho$=-0.01), and AlphaMissense ($\Delta\rho$=-0.028). Its average performance is slightly better than TranceptEVE L ($\Delta\rho$=0.001), GEMME ($\Delta\rho$=0.004) and the best previous method purely based on pLM embeddings VESPA ($\Delta\rho$=0.012). The statistical significance of these average performance differences is limited at a one-tailed paired t-test p-value of 1e-3 for the comparison with AlphaMissense and higher than 0.01 for all others.

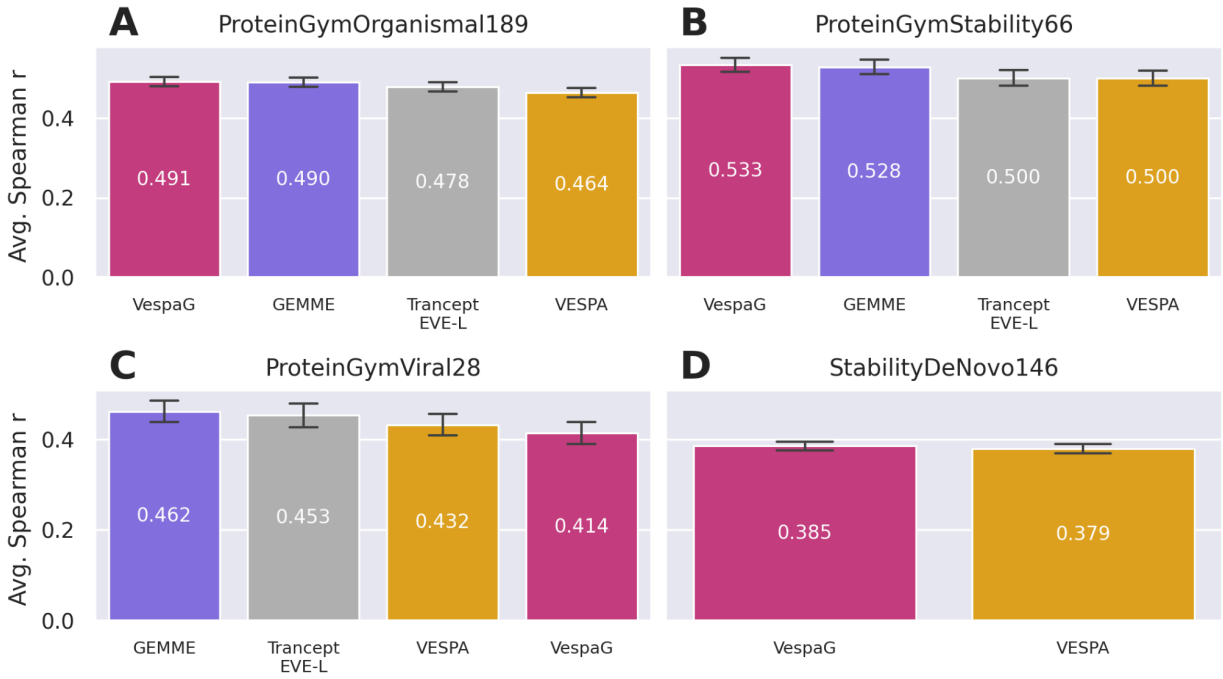

Figure 2: **VespaG accuracy on-par with SOTA.** Each panel corresponds to a different test set (with partially overlapping proteins between the test sets) from the ProteinGym substitution benchmark (Notin et al. 2023) and from a dataset by Tsuboyama et al. 2023 measuring stability under identical conditions for 146 *de novo* designed domains (40–72 residues) which was not included in ProteinGym: (A) *ProteinGymOrganismal189*, containing 189 experimental assays for 161 eukaryotic and prokaryotic proteins; (B) *ProteinGymStability66*, containing 66 assays for 66 eukaryotic, prokaryotic, and viral proteins; (C) *ProteinGymViral28*, containing 28 assays for 26 viral proteins; (D) *StabilityDeNovo146*, containing 146 folding stability assays of 148 *de novo* proteins. For each set, we ordered from best (left) to worst (right). We could not evaluate all methods against all datasets. Specifically, we could not recompute results for TranceptEVE L (Notin, Niekerk, et al. 2022) depicted in gray and directly extracted the predictions from the ProteinGym. In addition, we did not apply GEMME to *de novo* proteins because it relies on MSAs. The error bars show the 95% confidence interval.

The proverbial *student* VespaG bested the *teacher* GEMME by a large margin ($\Delta\rho$ > 0.1) for the human protein LYAM1, the murine MAFG, the bacterial proteins DN7A, F7YBW7, ISDH, NUSA, and SBI, the plant RCD1, and the yeast ubiquitin RL40A (SOM Fig. S7). In particular, VespaG correctly identified the glycines G75 and G76 in the top five ubiquitin residues most sensitive to mutations, whereas GEMME incorrectly predicted them as mildly sensitive (Fig. 3). These two residues play essential roles for E1 activation (Mavor et al. 2016). Reciprocally, VespaG agreed with the experiment on the mild tolerance of K27, whereas GEMME predicted this residue as highly sensitive (Fig. 3). We can interpret these discrepancies in light of previous works showing that ubiquitin stands out from the general trends between evolutionary sequence conservation and the experimentally measured tolerance to substitutions (Mavor et al.

2016, 2018; Roscoe et al. 2013). It challenges the common view that high selection pressure implies high mutational sensitivity. Hence, applying this principle on an input MSA, as GEMME does, leads to a limited accuracy (ρ in the 0.36-0.44 range). VespaG's representation learning-based approach allows overcoming this limitation and capturing key aspects of the peculiar sequence-phenotype ubiquitin relationship (ρ in the 0.48-0.54 range).

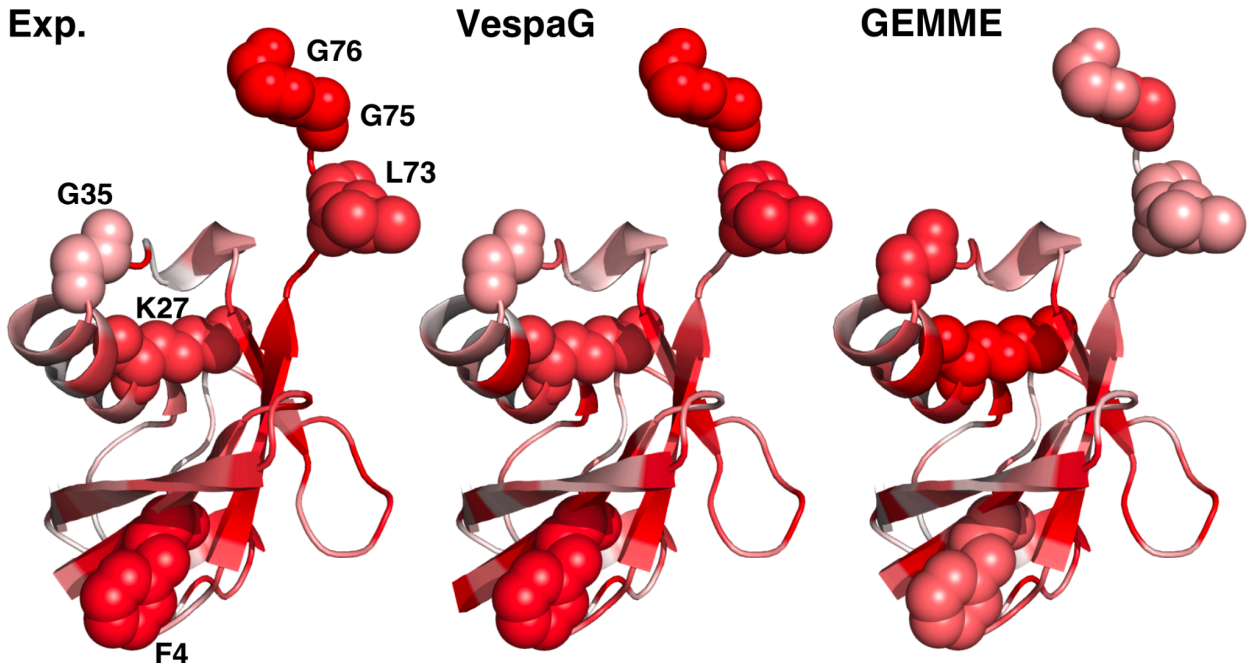

Figure 3: Details of *student* VespaG vs *teacher* GEMME. For the yeast ubiquitin (RL401A_YEAST), we compared experimental measurements (left panel, labeled Exp.; Mavor et al. 2016) with predictions by mapping the per-residue mutational sensitivities onto the 3D structures predicted by AlphaFold2 (AF-P0CH08-F1-model_v4, residues 2 to 76, Jumper et al. 2021). We estimated the extent to which a residue is sensitive to mutations as the rank of its average predicted or measured effect over the 19 possible substitutions. The more reddish the more sensitive. We highlighted six residues (labeled by one-letter amino acid code followed by position in the sequence, e.g., G27: glycine at position 27) for which VespaG agreed with the experiment (rank difference <5) while GEMME strongly disagreed (rank difference >15). The experimental values reflect ubiquitin fitness landscape under normal growth conditions (Mavor et al. 2016).

We observed a highly specific and starkly different trend on a test set of 28 assays characterizing 26 viral proteins, and the subset of 19 proteins *ProteinGymViral19* for which PoET and AlphaMissense predictions were available (Fig. 2C, *ProteinGymViral28*; SOM Fig. S8). For most of these, VespaG reached lower

correlations than the other methods (SOM Fig. S7). Moreover, GEMME and TranceptEVE L systematically achieved similar or higher accuracy than VespaG and VESPA, by a large margin ($\Delta\rho$>0.05) for half of the 28 assays. This result holds true when using other pLM embeddings as input (SOM Fig. S3).

VespaG successful in predicting stability changes for native and *de novo* proteins

We investigated VespaG's ability to assess the impact of mutations on several phenotype subsets from ProteinGym (SOM Table S10). VespaG outperformed GEMME, TranceptEVE L and VESPA, as well as all purely sequence-based methods assessed in ProteinGym (Notin et al. 2023), at predicting mutation-induced folding stability changes (Fig. 2B, *ProteinGymStability66*). It reached an average Spearman correlation ρ_{nat}=0.533±0.02 with experimental values, similar to GEMME ($\Delta\rho$=0.005; p-value=0.28), and better than the other methods we evaluated by $\Delta\rho$=0.033 (one-tailed paired t-test p-value < 0.005).

Out of the 66 assays in *ProteinGymStability66*, 64 originate from a dataset by Tsuboyama et al. 2023, who assessed the thermodynamic folding stability of native and *de novo* designed 40-72 amino acid long protein domains. As the *de novo* designed proteins were not included in ProteinGym, we added the test set *StabilityDeNovo146* to our analysis. The success for native protein sequences partially extended to *de novo* designed proteins: VespaG reached an average Spearman correlation of ρ_{nov}=0.385±0.0.011 on 146 *de novo* proteins (Fig. 2D, *StabilityDeNovo146*). VESPA was also able to handle these protein sequences, leading to a similar performance ($\Delta\rho$=0.006, p-value >0.1). GEMME only produced predictions for four *de novo* proteins due to a lack of sufficient input alignments. Its Spearman correlation on this small subset was nearly zero (0.085), compared to 0.383 for VespaG.

VespaG generalizing across multiple organisms

We further assessed the impact of the training set on VespaG's predictive performance and ability to transfer knowledge across organisms. Specifically, we retrained from scratch the same architecture with the same hyperparameters on non-redundant sets of ~4,000 proteins from the insect *Drosophila melanogaster*, ~2,000 proteins from the bacterium *Escherichia coli*, ~1,500 proteins coming from several viruses, and ~9,000 proteins from a combination of all. Training VespaG on a few thousand diverse proteins from a single organism sufficed to generalize across diverse taxa (Fig. 4). Overall, the performance differences for respective taxa were small across organismal training sets ($\Delta\rho\leq$0.1). However, we consistently observed a lower agreement between VespaG

predictions and the experiments for viral proteins, compared to other taxa, across all training sets. In particular, exclusively learning from ~1,500 viral proteins did not improve performance for viral proteins (Fig. 4). Inputting embeddings from other pLMs did not alter this trend (SOM Fig. S2).

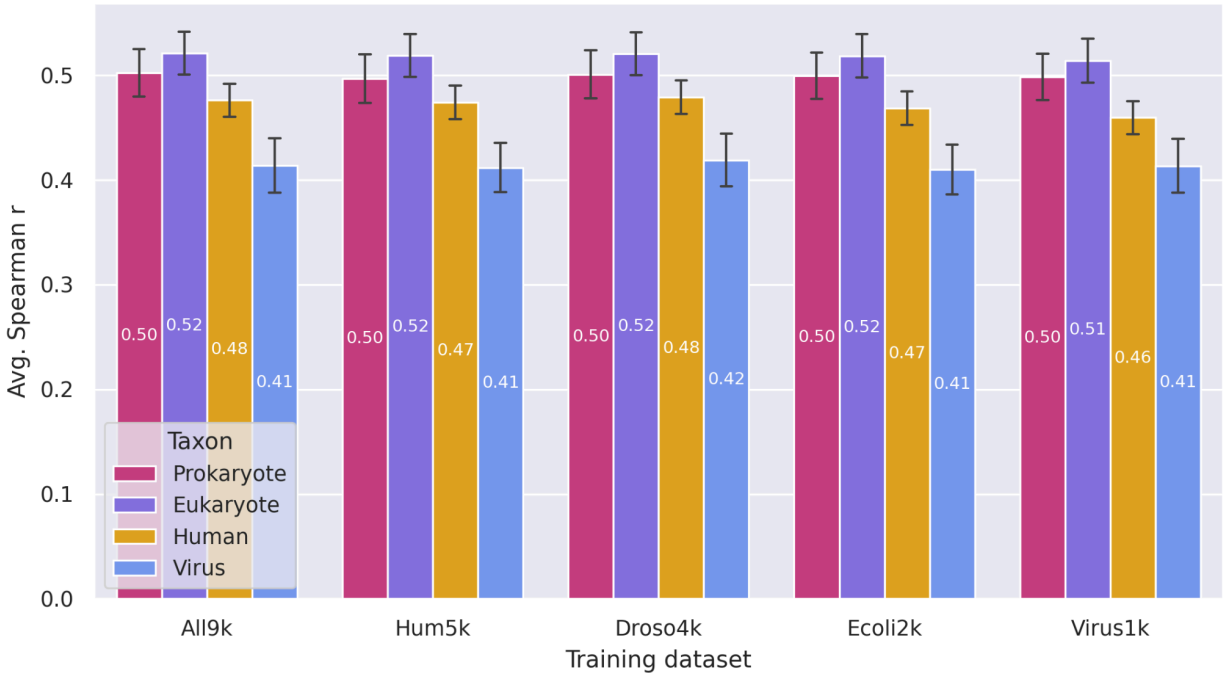

Figure 4: Per-taxon performance of VespaG independent of taxa included in training. For each training set, indicated on the x-axis, we reported the average Spearman correlations computed for each of the five taxa represented in ProteinGym benchmark (Notin et al. 2023). The error bars show the 95% confidence interval. Regardless of training data (all data sets were redundancy reduced, names reflect training and bars the test set; *All9k*: about 9,000 proteins mixing all taxa shown to the right; *Hum5k*: ~9,000 human proteins, *Droso4k*: ~4,000 fruit fly proteins (*drosophila melanogaster*) , *Ecoli2k*: ~2,000 E.coli proteins (*Escherichia coli*), *Virus1k*: mix of ~1,500 viral proteins), VespaG generalized equally well for all taxa assessed. For all training sets, it performed best for prokaryotes and non-human eukaryotes, followed by human. Even when trained explicitly on viral proteins, VespaG performed the worst for viral proteins.

VespaG predictions blazingly fast

Out of the top performing methods evaluated on the ProteinGym benchmark, VespaG was, in our hands, the most scalable for proteome-wide analyses. VespaG needed, on average, 18.0 seconds for the 73 unique proteins of the first iteration of ProteinGym (SOM Fig. S5) on low-end hardware (Intel Xeon Gold 6248 with 32x2.5 GHz, 64GB DDR4 ECC RAM, no GPU; Methods). Thus, VespaG was three orders of magnitude faster than GEMME and VESPA, which in turn substantially outpace other methods: the

authors of PoET observed their method and TranceptEVE L to be three orders of magnitude slower than GEMME (Truong Jr and Bepler 2023). On top, VespaG, unlike other pLM-based methods such as VESPA, does not require a GPU for fast inference, providing the same runtime on only a CPU.

Running VespaG on a consumer-grade laptop (Intel i7-1355U with 12x5 GHz, 16GB DDR4 RAM; Methods) showcases its speed: interpolating its runtime measured on 80 sets of proteins with 7k to 80k total residues (SOM Fig. S8), VespaG needs 24:11 minutes to compute the entire single-site mutational landscape for a Human proteome of 12M amino acids (Sinitcyn et al. 2023) and 6:56:18 hours for the entire 2024-02 release of UniProt/Swiss-Prot (The UniProt Consortium et al. 2023) containing about 206M amino acids. GEMME, taking 771 times longer to run, could predict single-site mutation effects for 741 average-length proteins in the same amount of time and TranceptEVE L or PoET, assuming a runtime three orders of magnitude higher than GEMME, could only generate predictions for a single protein of 268 amino acids in the same time.

Discussion

In this work, we explored the possibility of modeling the sequence-phenotype relationship by learning a simple mapping function from protein Language Model (pLM) representations, or embeddings, to evolutionary scores predicted by an expert method. We demonstrated the validity of this approach on several hundred diverse proteins across different organisms. The performance of the resulting method, dubbed VespaG, reached that of much more sophisticated methods.

Using predicted scores instead of curating a dataset of experimental measures allowed the creation of a larger training set (totaling 39M SAVs) than those used previously. By comparison, the SNAP2 development set contained about 100,000 SAV effect annotations from 10,000 proteins (Hecht, Bromberg, and Rost 2015). In addition, exploiting only the pLM embedding of the wild-type protein of interest, instead of explicitly modeling its mutants through log-odd probability estimations, enabled reaching a much higher efficiency and inference speed than previous pLM-based methods. The feasibility of this strategy also emphasized the usefulness of the information encoded in a wild-type protein query embedding for assessing all its variants.

Our results suggested that even relatively shallow neural networks (660k free parameters) effectively leveraged the knowledge encoded in an unsupervised method such as GEMME (Laine, Karami, and Carbone 2019). The transparency of the GEMME algorithm and the simplicity of VespaG's architecture open ways to interpreting its predictions and getting insights into the features and patterns it has learned. A fundamental difference between student (VespaG) and teacher (GEMME) is its usage of a universal protein representation space. More specifically, VespaG is able to relate

proteins with each other via representations generated by a pLM pre-trained over a huge diversity of natural protein sequences across protein families.

This property allows VespaG to generalize across organisms without considering any specific input generation or training schema. Training VespaG on a few thousand proteins from either *Homo sapiens*, or *Drosophila melanogaster*, or *Escherichia coli* sufficed to produce high quality predictions on a diverse set of proteins. For eukaryotic and prokaryotic proteins, the pLM-based student VespaG performed overall numerically higher and more consistently than the MSA-based teacher GEMME (Fig. 2). For instance, VespaG improved for cases such as ubiquitin which do not follow the general trends between evolutionary conservation and mutational outcomes. In addition, VespaG yielded better accuracy in predicting thermodynamic folding stability over hundreds of single-domain proteins.

VespaG did not perform well for viral proteins, even when explicitly trained on those (Fig. 4). The ProteinGym assessment reported the same trend for other pLM-based methods (Notin et al. 2023). The gap in performance between the student, VespaG, and the teacher, GEMME, for this class of proteins (Fig. 2C), suggests specific properties of the embeddings or the evolutionary scores preventing an accurate mapping between them. Despite retaining high accuracy, GEMME evolutionary scores for viral proteins have a lower resolution than for organismal proteins (SOM Fig. S2). Many mutations are assigned the same score, likely reflecting the comparatively lower variability of the associated input MSAs (Table 1). Nevertheless, the fact that VespaG trained exclusively on viral proteins exhibits a high predictive capability on organismal proteins (Fig. 4) suggests the impact of this resolution loss is limited. It further supports the hypothesis that the embeddings computed by the pre-trained pLMs for viral proteins are intrinsically noisy. Possibly, viral proteins are simply too under-represented in the training of pLMs, due to a comparatively small number and low diversity (Ding and Steinhardt 2024; Elnaggar et al. 2021; Lin et al. 2023; The UniProt Consortium et al. 2023). In addition, the pLMs may struggle to capture the inherent peculiarities of viral protein evolution (Koonin, Dolja, and Krupovic 2022). Structurally and functionally relevant evolutionary constraints are expected to manifest through smaller differences in viral protein sequences compared to other taxa, warranting a special treatment of these sequences for extracting co-variations (Hopf et al. 2017). A future improvement of pLMs could be to develop viral-specific fine-tuning steps. Alternatively, combinations of VespaG with complementary approaches leveraging 3D protein structure representations, *e.g.*, with the efficient expert-guided ML method RaSP for predicting stability changes (Blaabjerg et al. 2023), might improve predictions.

Finally, the fact that VespaG generates protein-specific rather than family-averaged predictions, i.e. is independent from MSAs, enables tackling *de novo* designed proteins.

Most established sequence variant effect prediction methods such as GEMME largely succeed due to evolutionary information derived from MSAs and tend to barely outperform random for single sequences (Hecht, Bromberg, and Rost 2015). Given the absence of reference data for *de novo* proteins, conventional tools tend to provide unreliable estimates of their respective properties, if any at all. Although we observed a drop in performance compared to natural proteins, we can envision using VespaG for fast screens prior to applying future methods adapted to that problem, or as a guide for designing more biocompatible *de novo* proteins.

In conclusion, VespaG closes the gap in performance between the best and the fastest variant effect predictors. For an unprecedentedly small trade-off in performance, it can predict variants several orders of magnitude faster than other state-of-the-art methods.

Methods

Related work

A key component of computational variant effect predictors is the ability to capture arbitrary range dependencies between amino acid residues. A majority of predictors extracts these dependencies from an input multiple sequence alignment (MSA) generated from large protein sequence databases. Some rely on the statistical inference of pairwise couplings (Figliuzzi et al. 2016; Hopf et al. 2017), others on the implicit account for global context with latent variables (Frazer et al. 2021; Riesselman, Ingraham, and Marks 2018). They remain computationally costly due to the large number of inferred parameters and strongly depend on the input alignment's variability.

More recently, using representations generated by protein language models (pLMs) emerged as an alternative to using MSAs as input (Brandes et al. 2023; Cheng et al. 2023; Marquet et al. 2022; Meier et al. 2021; Nijkamp et al. 2022; Notin, Dias, et al. 2022). These high-capacity pLM transformer architectures, borrowed from natural language processing, learn to reconstruct masked or missing amino acids in an input query sequence (Elnaggar et al. 2021; Lin et al. 2023). They model raw protein sequence data over large databases, thereby capturing evolutionary constraints that generalize across protein families. Once trained, they can serve as zero-shot variant effect predictors by estimating the likelihood of each amino acid at each position. The log-odds ratios of the amino acid probabilities computed from the pretext reconstruction task already provide reasonably accurate estimates of variant effects (Meier et al. 2021). Nevertheless, they are not competitive with state-of-the-art (SOTA) methods (Notin et al. 2023).

A limitation of pLMs is that they overlook natural protein sequences' evolutionary history and relationships. Strategies for overcoming this oversimplification aim at encoding evolutionary semantics in the pLMs representation space, *e.g.*, by augmenting the input with a multiple sequence alignment (Rao et al. 2021). The latter informs the model about sequences evolutionary related to the input query and how their amino acids match. The integration of weak labels coming from inter-individual polymorphisms and of a physical prior through supervised learning of the 3D structure, as AlphaMissense does, further enhances the predictive performances (Cheng et al. 2023). Alternatively, the predictor VESPA combines the pLM-derived log-odds substitution scores with evolutionary conservation levels predicted from the learned protein representations (Marquet et al. 2022). The predictors Tranception (Notin, Dias, et al. 2022) and PoET (Truong Jr and Bepler 2023) have also explored retrieval-augmented strategies by

conditioning the predictions of a pre-trained pLM on a set of related raw or aligned protein sequences at inference time.

Comparison to State-of-the-art methods

We chose to compare our methods to the following five state-of-the-art methods on the ProteinGym test set. *GEMME* (Laine, Karami, and Carbone 2019) as it is (1) tied for the best performing method on ProteinGym (Notin et al. 2023), (2) a purely MSA-based method not using machine learning, and (3) was used to annotate VespaG's training data. *TranceptEVE L* (Notin, Niekerk, et al. 2022) as it is the best performing method on ProteinGym next to GEMME, and because it is a hybrid model, making use of both MSAs and pLM embeddings as input, combining the previously developed autoregressive Tranception (Notin, Dias, et al. 2022) with the Bayesian variational autoencoder EVE (Frazer et al. 2021). *PoET* (Truong Jr and Bepler 2023), a recently developed autoregressive generative method modeling protein families as sequences-of-sequences and slightly outperforming other methods against the first version of the ProteinGym set (although not currently included in the benchmark). *AlphaMissense* (Cheng et al. 2023), also recently introduced and building up on the protein structure predictor AlphaFold (Jumper et al. 2021) by incorporating population frequency data. Similar to PoET, AlphaMissense was evaluated on the first ProteinGym iteration but is not included in the updated benchmark. Lastly, we compare to *VESPA* (Marquet et al. 2022), the currently best purely pLM-based method in ProteinGym. VESPA takes pLM embeddings as input to predict per-residue conservation scores and combines them with per-mutation protein-dependent log-odds scores and per-mutation protein-independent substitution scores.

We computed VESPA and GEMME predictions for both test sets, while we downloaded TranceptEVE L predictions from ProteinGym. For PoET, the authors provided Spearman correlation values for each experimental assay in the first iteration of the ProteinGym test set on request. For AlphaMissense, we extracted pre-computed per-assay scores of the first ProteinGym iteration from the supplemental materials.

Due to the different ways of accessing or computing the predictions of GEMME, TranceptEVE, PoET, and AlphaMissense, we did not evaluate the effect of using different MSAs as input. GEMME predictions were generated through the MMseqs2-based ColabFold protocol as described in (Abakarova et al. 2023). Further, we did not benchmark ensemble methods such as PoET+GEMME (Truong Jr and Bepler 2023) on the test sets as ensembling any other method with VespaG would eradicate its speed advantage. To assess thermodynamic folding stability of de novo domains, we only compared VespaG with GEMME and VESPA, since precomputed results were not available for the other methods.

Method development

Datasets

In the following, the datasets applied in this work are described. We utilized ten different datasets, one based on the human proteome (Hum5k) for training and validation, and four re-training datasets (Droso4k, Ecoli2k, Virus1k, All9k). All training data was annotated using evolutionary-informed substitution scores computed by GEMME. Additionally, we compared VespaG against SOTA methods on the two test sets *ProteinGym* (with nine subsets) and *StabilityDeNovo146*.

Test data

ProteinGym Substitution Benchmark. The substitution benchmark ProteinGym (Notin et al. 2023) comprised 217 deep mutational scans (DMS) from 187 unique proteins with diverse lengths (37 - 3,423 residues with a median of 245), protein families (*e.g.*, polymerases, tumor suppressors, kinases, transcription factors), sizes, functions (e.g. drug resistance, ligand binding, viral replication, thermostability), and taxa, totalling about 696k single missense variants and 1.76M multiple missense mutations from 69 of the 217 proteins (SOM Fig. S4a). Protein sequences and experimental scores were accessed from the ProteinGym GitHub repository (Notin 2024). The ProteinGym substitution benchmark was not redundancy reduced, and we divided the ProteinGym into subsets based on organisms, function, and time of availability. The first split by organism and can contain any function: (1) *ProteinGymOrganismal189* with 189 assays on 161 prokaryotic and eukaryotic proteins, and (2) *ProteinGymViral28* with 28 assays on 26 viral proteins. The second split was based on function and can contain prokaryotic, eukaryotic and viral proteins: (3) *ProteinGymStability66* with 66 assays assessing stability, (4) *ProteinGymActivity43* with with 43 assays assessing activity, (5) *ProteinGymBinding13* with 13 assays assessing binding, (6) *ProteinGymExpression18* with 18 assays assessing organismal expression, and (7) *ProteinGymFitness77* with 77 assays assessing fitness. The last splits were added to assess methods for which predictions were available only for the first iteration of ProteinGym. This first iteration contained 87 DMS from 73 unique proteins (72 - 3,423 residues with a median of 379), totalling about 1.5M variants with mostly single and, for 11 proteins, multiple missense mutations (SOM Fig. S4b). The set was divided into (8) *ProteinGymOrganismal66,* containing 66 assays of 54 proteins with 1.4M variants, and (9) *ProteinGymViral21*, containing 21 assays of 19 proteins with 184K variants.

StabilityDeNovo146. The dataset by Tsuboyama et al. 2023 includes DMS measuring stability under identical conditions of all single substitutions, deletions and two insertions in 266 native and 146 *de novo* designed domains (40–72 residues). This is the most comprehensive available dataset assessing how amino acid substitutions affect

thermodynamic folding stability. The ProteinGym substitution benchmark includes 64 native proteins from this dataset (all included in *ProteinGymStability66*). To additionally assess the predictors on *de novo* domains, we created the test set *StabilityDeNovo146* with 123k variants across 146 unique proteins designed using TrRosetta (Yang et al. 2020) with the hallucination protocol described in (Anishchenko et al. 2021; Norn et al. 2021) or the blueprint-based approach described in (Huang et al. 2011; Kim et al. 2022). We selected all mutations with tabulated free energy changes (ΔΔG) annotations, discarded all deletions, insertions, and wild-type sequences, and averaged multiple measurements for the same mutation.

Training data

To generate training data, we constructed a main set based on the *Homo sapiens* proteome and additional sets representing a diverse set of origins. These included proteomes from *Drosophila melanogaster* and *Escherichia coli,* as well as all viral proteins included in Swiss-Prot (Table 1). Each training dataset was curated following the same process of first downloading the UniProt (The UniProt Consortium et al. 2023) reference proteome(s) with one protein sequence per gene for the respective organism or set of organisms and removing any proteins with a sequence length smaller than or equal to 25 residues or longer than or equal to 1024. The dataset sizes are presented in Table 1.

We redundancy reduced the training data in two steps, firstly against the test data to prevent data leakage and secondly against themselves to reduce the number of training samples. To do so, we clustered the protein sequences using UniqueProt (Mika 2003) with an HSSP (homology-derived secondary structure of proteins)-value < 0, corresponding to no pair of proteins in the redundancy reduced data set having over 20% pairwise sequence identity over 250 aligned residues (Rost 1999; Sander and Schneider 1991). This criterion is stricter than the classically used threshold of 30% overall sequence identity. It accounts for the fact that protein structures can show high similarity even at lower sequence similarity levels (Rost 1999). To improve runtime, we modified the original UniqueProt protocol (Olenyi, Tobias et al. n.d.) by replacing BLAST (Altschul et al. 1997) with MMseqs2 (Steinegger and Söding 2017). Additionally, we discarded alignments of fewer than 50 residues for pairs of sequences with more than 180 residues as they provide only a relatively weak support for similarity. Finally, we generated training and validation sets using a random 80/20 split.

To circumvent the need for a large, comprehensive set of experimental variant effect annotations, we employed the established method GEMME (see Evaluation/SOTA Methods) following the protocol introduced in (Abakarova et al. 2023). Specifically, for each protein from the training set, we retrieved and aligned a set of homologous

sequences with the MMseqs2-based multiple sequence alignment (MSA) generation strategy implemented in ColabFold (Mirdita et al. 2022). We then used the generated MSA as input for GEMME. GEMME outputs a complete substitution matrix of dimension L x 20, with L being the length (in residue) of the input query protein sequence. GEMME scores range from -10 to 2. Drawing from our previous findings (Abakarova et al. 2023), we flagged the mutational landscapes derived from fewer than a couple hundred homologous sequences as lowly confident.

Name	**Source organism(s)**	**Number of sequences**				**Number of residues**	
		Reference proteome	**After Length restriction** (Min.25, Max.1024)	**After Redundancy Reduction**	**Confident GEMME predictions**	**Total**	**Min., Max., Median protein length**
Hum5k	*Homo sapiens*	20,357	18,043	5,886	5,305	2,010,031	36 / 1024 / 328
Droso4k	*Drosophila melanogaster*	13,811	12,305	4,809	4,081	1,610,079	40 / 1024 / 346
Ecoli2k	*Escherichia coli*	4,403	4,284	2,544	2,333	652,257	36 / 990 / 243
Virus1k	All viral in Swiss-Prot	17,320	15,954	4,498	1,400	401,204	29 / 1017 / 228
All9k	All of above	55,891	50,586	15,175	9,616	3,291,539	29 / 1024 / 286

Table 1: Size statistics of training datasets used. The dataset used for the development of VespaG, dubbed Hum5k, is a redundancy reduced version of the human proteome. To investigate generalizability across organisms, we further analyzed the performance of VespaG trained on several other datasets. These included the redundancy reduced proteomes of *Drosophila melanogaster* (Droso4k) and *Escherichia coli* (Ecoli2k), a redundancy reduced set of all viral proteins in Swiss-Prot (Virus1k), and a redundancy reduced combination of all (All9k).

Model Specifications

Input features

All developed models rely solely on embeddings computed from pre-trained pLMs as input. Specifically, we used ProtT5-XL-U50 (Elnaggar et al. 2021), an encoder-decoder transformer architecture trained on the Big Fantastic Database (Steinegger and Söding 2017) and fine-tuned on UniRef50, and ESM2-T36-3B-UR50 (Lin et al. 2023), a BERT (Devlin et al. 2019) style 3-billion-parameter encoder-only transformer architecture trained on all clusters from Uniref50, augmented by sampling sequences from the

Uniref90 clusters of the representative chains (excluding artificial sequences). In the following, we refer to these pLMs as ProtT5 and ESM-2, respectively. We used the pre-trained pLMs as is, without fine-tuning their weights.

Configuration

We built the predictors with the following architectures: (1) Linear regression, *i.e.*, a feed-forward neural network (FNN, (LeCun, Bengio, and Hinton 2015)) without any hidden layer, dubbed *LinReg*; (2) FNN with one hidden layer, called *VespaG*; (3) FNN with two hidden layers, called *FNN_2_layer*; (4) Convolutional neural network (CNN, (LeCun, Bengio, and Hinton 2015)) with one 1-dimensional convolution and two hidden dense layers, referred to as *CNN*; and (5) an ensemble of separately optimized FNN and CNN (with the same architecture as the best stand-alone model for each architecture), with the output being the mean of the two networks.

All layers were linked through the LeakyReLu activation function (Agarap 2019; Xu et al. 2015), as well as dropout (Srivastava et al. 2014). No transformation was applied to the output of the final layer. We implemented the models in PyTorch (Paszke et al. 2019) v1.13.1 using Python 3.10.8. They were trained using the AdamW (Kingma and Ba 2017; Loshchilov and Hutter 2019) optimizer with an initial learning rate of 10e-4, decaying with a decay factor of 0.33 and a patience of 8 epochs using mean squared error (MSE) loss between predicted and target GEMME scores. Each batch contained 25K residues. We applied early stopping based on the validation partition loss, with a patience of 10 epochs. The maximal training duration was 200 epochs. Hyperparameters were optimized through exhaustive parameter search on the validation split of the *Hum5k* dataset for each architecture. The same hyperparameters were then used to re-train separate models on each of the additional datasets (Droso4k, Ecoli2k, Virus1k, and All9k), using the same training scheme (SOM Table S1). The training dataset and hyperparameters of the final model were selected based on validation loss.

Evaluation

Performance measures

To assess the performance of the predictors, we relied on the Spearman rank correlation coefficient: $\rho_s = r_{R(X),R(Y)} \frac{cov(R(X),R(Y))}{\sigma_{R(X)}\sigma_{R(Y)}}$, with raw scores X, Y as ranks $R(X)$, $R(Y)$, Pearson correlation coefficient $r_{R(X),R(Y)}$, covariance $cov(R(X), R(Y))$, and standard deviations $\sigma_{R(X)}$, $\sigma_{R(Y)}$. This metric enables quantifying the strength of non-linear relationships between predicted and experimental scores. The Spearman

correlation coefficient was computed for each DMS experiment separately. For proteins covered by multiple DMS experiments, we averaged the correlation over the experiments before computing the mean over all proteins.

Error estimates

To assess the significance of the performance differences between the evaluated predictors, we computed symmetric 95% confidence intervals (CI) $\mu \pm 1.96\frac{\sigma}{\sqrt{n}}$ over n = 1,000 bootstraps using sampling with replacement from the respective datasets. The terms μ and σ are the mean and standard deviation of the bootstrap distribution, respectively. In addition, we performed one-tailed paired t-tests. For any pair of methods, the test statistic is expressed as, $t = \frac{\bar{d}}{\sigma/\sqrt{n}}$, where $\bar{d}$ and σ are the mean and standard deviation of the n performance differences. The associated p-value reflects the probability of observing the test statistic under the null hypothesis (no significant difference, the true mean of differences is zero). It is obtained by comparing t to a t -distribution with n-1 degrees of freedom.

Baseline

To compare model performances on the validation set to an untrained random baseline retaining on the target distribution, we randomly permuted the GEMME target scores for each protein, i.e. randomly shuffled all values of the Lx20 GEMME output matrix.

Multi-mutations

To score combinations of multiple substitutions, we add the scores of their constituent single substitutions, a technique introduced previously (Meier et al. 2021).

Runtime

We measured the wall-clock runtime of VESPA, GEMME, and VespaG for 1.6M SAVs of the 73 unique proteins of the first iteration of the ProteinGym test set using 32 CPU cores of an Intel Xeon Gold 6248 at 2.50 GHz and 64 GB of DDR4 ECC RAM. VESPA, in addition, was allotted an Nvidia Quadro RTX 8000 GPU with 46 GB RAM. Runtimes for other methods could not be obtained due to computational limitations. We did not measure the runtime for the generation of the input MSAs for GEMME, nor the input embeddings for the embedding-based methods since we used pre-computed data.

The relationship between wall-clock runtime of VespaG and number of residues in the input was estimated by creating random subsets of the ProteinGym217 benchmark set, with 10 subsets each (sizes: 25, 50, 75, 100, 125, 150, 175, and 200 sequences) and

measuring the execution time of VespaG for each of the 150 resulting subsets. A linear regression was fitted for the number of residues in each subset and the execution time per subset. To illustrate VespaG's effectiveness, this measurement was done on a consumer laptop with an Intel i7-1355U (12x5 GHz) and 16 GB of DDR4 RAM.

Declarations

Acknowledgements. This work was granted access to the high-performance computing resources of the SACADO MeSU platform at Sorbonne Université. Thanks to Nikita Kugut for invaluable help with administrative aspects of this work, as well as Tim Karl and Tobias Olenyi for technical aspects. Thanks to the group around Pascal Notin for creating the benchmark ProteinGym, and to all developers who make their methods freely available. Thanks to all experimentalists for making their data available - with special thanks to the group around Gabriel J. Rocklin for developing the mega-scale experimental analysis of protein folding stability.

Funding. This work was funded by the Bavarian Ministry of Education through funding to the TUM, by a grant from the Alexander von Humboldt foundation through the German Ministry for Research and Education (BMBF: Bundesministerium für Bildung und Forschung), and by a grant from Deutsche Forschungsgemeinschaft (DFG-GZ: RO1320/4-1). The French National Research Agency provided a salary to M.A. via the grant ANR-20-CE44-0010, Adagio project. This work has also been partially supported by the European Research Council under the European Union's H2020 Framework Programme (2023–2028)/ ERC Grant agreement ID 101087830 awarded to E.L..

Competing interests. The authors declare no conflicts of interest.

Ethics approval. Not applicable

Consent to participate. Not applicable

Consent for publication. Not applicable

Data and Code Availability. The source code and model weights of this work are freely available at https://github.com/JSchlensok/VespaG.

Author Contributions. C.M. supervised the method implementation and evaluation, and co-wrote the manuscript. J.S. implemented and evaluated the method and wrote the initial draft. M.A. computed GEMME scores and assisted in the usage of the tool. B.R. co-conceived and guided the work. E.L. supervised the work and co-wrote the manuscript. J.S., C.M., M.A and E.L. produced and analyzed the results. All authors edited, read, and approved the final manuscript.

Bibliography

Abakarova, Marina, Céline Marquet, Michael Rera, Burkhard Rost, and Elodie Laine. 2023. “Alignment-Based Protein Mutational Landscape Prediction: Doing More with Less.” *Genome Biology and Evolution* 15(11): evad201. doi:10.1093/gbe/evad201.

Abildgaard, Amanda B., Sofie V. Nielsen, Inge Bernstein, Amelie Stein, Kresten Lindorff-Larsen, and Rasmus Hartmann-Petersen. 2023. “Lynch Syndrome, Molecular Mechanisms and Variant Classification.” *British Journal of Cancer* 128(5): 726–34. doi:10.1038/s41416-022-02059-z.

Adzhubei, Ivan, Daniel M. Jordan, and Shamil R. Sunyaev. 2013. “Predicting Functional Effect of Human Missense Mutations Using PolyPhen-2.” *Current protocols in human genetics / editorial board, Jonathan L. Haines ... [et al.]* 0 7: Unit7.20. doi:10.1002/0471142905.hg0720s76.

Agarap, Abien Fred. 2019. “Deep Learning Using Rectified Linear Units (ReLU).” doi:10.48550/arXiv.1803.08375.

Altschul, Stephen F., Thomas L. Madden, Alejandro A. Schäffer, Jinghui Zhang, Zheng Zhang, Webb Miller, and David J. Lipman. 1997. “Gapped BLAST and PSI-BLAST: A New Generation of Protein Database Search Programs.” *Nucleic Acids Research* 25(17): 3389–3402. doi:10.1093/nar/25.17.3389.

Anishchenko, Ivan, Samuel J. Pellock, Tamuka M. Chidyausiku, Theresa A. Ramelot, Sergey Ovchinnikov, Jingzhou Hao, Khushboo Bafna, et al. 2021. “De Novo Protein Design by Deep Network Hallucination.” *Nature* 600(7889): 547–52. doi:10.1038/s41586-021-04184-w.

Armenteros, Jose Juan Almagro, Alexander Rosenberg Johansen, Ole Winther, and Henrik Nielsen. 2020. “Language Modelling for Biological Sequences – Curated Datasets and Baselines.” : 2020.03.09.983585. doi:10.1101/2020.03.09.983585.

“Atlas of Variant Effects Alliance.” *Atlas of Variant Effects Alliance*. https://www.varianteffect.org (October 6, 2023).

Blaabjerg, Lasse M, Maher M Kassem, Lydia L Good, Nicolas Jonsson, Matteo Cagiada, Kristoffer E Johansson, Wouter Boomsma, Amelie Stein, and Kresten Lindorff-Larsen. 2023. “Rapid Protein Stability Prediction Using Deep Learning Representations” eds. José D Faraldo-Gómez, Detlef Weigel, Nir Ben-Tal, and Julian Echave. *eLife* 12: e82593. doi:10.7554/eLife.82593.

Brandes, Nadav, Grant Goldman, Charlotte H. Wang, Chun Jimmie Ye, and Vasilis Ntranos. 2023. “Genome-Wide Prediction of Disease Variant Effects with a Deep Protein Language Model.” *Nature Genetics* 55(9): 1512–22. doi:10.1038/s41588-023-01465-0.

Cagiada, Matteo, Sandro Bottaro, Søren Lindemose, Signe M. Schenstrøm, Amelie Stein, Rasmus Hartmann-Petersen, and Kresten Lindorff-Larsen. 2023. “Discovering Functionally Important Sites in Proteins.” *Nature Communications* 14(1): 4175. doi:10.1038/s41467-023-39909-0.

Cheng, Jun, Guido Novati, Joshua Pan, Clare Bycroft, Akvilė Žemgulytė, Taylor Applebaum, Alexander Pritzel, et al. 2023. “Accurate Proteome-Wide Missense Variant Effect Prediction with AlphaMissense.” *Science* 381(6664): eadg7492.

doi:10.1126/science.adg7492.
Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. “BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding.” *arXiv:1810.04805 [cs]*. http://arxiv.org/abs/1810.04805 (January 31, 2021).
Ding, Frances, and Jacob Steinhardt. 2024. “Protein Language Models Are Biased by Unequal Sequence Sampling across the Tree of Life.” : 2024.03.07.584001. doi:10.1101/2024.03.07.584001.
Elnaggar, Ahmed, Michael Heinzinger, Christian Dallago, Ghalia Rehawi, Yu Wang, Llion Jones, Tom Gibbs, et al. 2021. “ProtTrans: Towards Cracking the Language of Life’s Code Through Self-Supervised Learning.” : 2020.07.12.199554. doi:10.1101/2020.07.12.199554.
Figliuzzi, Matteo, Hervé Jacquier, Alexander Schug, Oliver Tenaillon, and Martin Weigt. 2016. “Coevolutionary Landscape Inference and the Context-Dependence of Mutations in Beta-Lactamase TEM-1.” *Molecular Biology and Evolution* 33(1): 268–80. doi:10.1093/molbev/msv211.
Fowler, Douglas M, and Stanley Fields. 2014. “Deep Mutational Scanning: A New Style of Protein Science.” *Nature Methods* 11(8): 801–7. doi:10.1038/nmeth.3027.
Frazer, Jonathan, Pascal Notin, Mafalda Dias, Aidan Gomez, Joseph K. Min, Kelly Brock, Yarin Gal, and Debora S. Marks. 2021. “Disease Variant Prediction with Deep Generative Models of Evolutionary Data.” *Nature* 599(7883): 91–95. doi:10.1038/s41586-021-04043-8.
Gersing, Sarah, Matteo Cagiada, Marinella Gebbia, Anette P. Gjesing, Atina G. Coté, Gireesh Seesankar, Roujia Li, et al. 2023. “A Comprehensive Map of Human Glucokinase Variant Activity.” *Genome Biology* 24(1): 97. doi:10.1186/s13059-023-02935-8.
Gray, Vanessa E., Ronald J. Hause, Jens Luebeck, Jay Shendure, and Douglas M. Fowler. 2018. “Quantitative Missense Variant Effect Prediction Using Large-Scale Mutagenesis Data.” *Cell systems* 6(1): 116-124.e3. doi:10.1016/j.cels.2017.11.003.
Hecht, Maximilian, Yana Bromberg, and Burkhard Rost. 2013. “News from the Protein Mutability Landscape.” *Journal of Molecular Biology* 425(21): 3937–48. doi:10.1016/j.jmb.2013.07.028.
Hecht, Maximilian, Yana Bromberg, and Burkhard Rost. 2015. “Better Prediction of Functional Effects for Sequence Variants.” *BMC Genomics* 16(S8): S1. doi:10.1186/1471-2164-16-S8-S1.
Hopf, Thomas A., John B. Ingraham, Frank J. Poelwijk, Charlotta P. I. Schärfe, Michael Springer, Chris Sander, and Debora S. Marks. 2017. “Mutation Effects Predicted from Sequence Co-Variation.” *Nature Biotechnology* 35(2): 128–35. doi:10.1038/nbt.3769.
Huang, Po-Ssu, Yih-En Andrew Ban, Florian Richter, Ingemar Andre, Robert Vernon, William R. Schief, and David Baker. 2011. “RosettaRemodel: A Generalized Framework for Flexible Backbone Protein Design.” *PLoS ONE* 6(8): e24109. doi:10.1371/journal.pone.0024109.
Jumper, John, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, et al. 2021. “Highly Accurate Protein Structure Prediction with AlphaFold.” *Nature* 596(7873): 583–89.

doi:10.1038/s41586-021-03819-2.
Kim, Tae-Eun, Kotaro Tsuboyama, Scott Houliston, Cydney M. Martell, Claire M. Phoumyvong, Alexander Lemak, Hugh K. Haddox, Cheryl H. Arrowsmith, and Gabriel J. Rocklin. 2022. "Dissecting the Stability Determinants of a Challenging de Novo Protein Fold Using Massively Parallel Design and Experimentation." *Proceedings of the National Academy of Sciences* 119(41): e2122676119. doi:10.1073/pnas.2122676119.
Kingma, Diederik P., and Jimmy Ba. 2017. "Adam: A Method for Stochastic Optimization." doi:10.48550/arXiv.1412.6980.
Koonin, Eugene V., Valerian V. Dolja, and Mart Krupovic. 2022. "The Logic of Virus Evolution." *Cell Host & Microbe* 30(7): 917–29. doi:10.1016/j.chom.2022.06.008.
Laine, Elodie, Yasaman Karami, and Alessandra Carbone. 2019. "GEMME: A Simple and Fast Global Epistatic Model Predicting Mutational Effects." *Molecular Biology and Evolution* 36(11): 2604–19. doi:10.1093/molbev/msz179.
LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. "Deep Learning." *Nature* 521(7553): 436–44. doi:10.1038/nature14539.
Lin, Zeming, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, et al. 2023. "Evolutionary-Scale Prediction of Atomic-Level Protein Structure with a Language Model." *Science* 379(6637): 1123–30. doi:10.1126/science.ade2574.
Livesey, Benjamin J, and Joseph A Marsh. 2023. "Updated Benchmarking of Variant Effect Predictors Using Deep Mutational Scanning." *Molecular Systems Biology* 19(8): e11474. doi:10.15252/msb.202211474.
Loshchilov, Ilya, and Frank Hutter. 2019. "Decoupled Weight Decay Regularization." doi:10.48550/arXiv.1711.05101.
Mahlich, Yannick, Jonas Reeb, Maximilian Hecht, Maria Schelling, Tjaart Andries Petrus De Beer, Yana Bromberg, and Burkhard Rost. 2017. "Common Sequence Variants Affect Molecular Function More than Rare Variants?" *Scientific Reports* 7(1): 1608. doi:10.1038/s41598-017-01054-2.
Marquet, Céline, Michael Heinzinger, Tobias Olenyi, Christian Dallago, Kyra Erckert, Michael Bernhofer, Dmitrii Nechaev, and Burkhard Rost. 2022. "Embeddings from Protein Language Models Predict Conservation and Variant Effects." *Human Genetics* 141(10): 1629–47. doi:10.1007/s00439-021-02411-y.
Mavor, David, Kyle A. Barlow, Daniel Asarnow, Yuliya Birman, Derek Britain, Weilin Chen, Evan M. Green, et al. 2018. "Extending Chemical Perturbations of the Ubiquitin Fitness Landscape in a Classroom Setting Reveals New Constraints on Sequence Tolerance." *Biology Open* 7(7): bio036103. doi:10.1242/bio.036103.
Mavor, David, Kyle Barlow, Samuel Thompson, Benjamin A Barad, Alain R Bonny, Clinton L Cario, Garrett Gaskins, et al. 2016. "Determination of Ubiquitin Fitness Landscapes under Different Chemical Stresses in a Classroom Setting" ed. Jeffery W Kelly. *eLife* 5: e15802. doi:10.7554/eLife.15802.
Meier, Joshua, Roshan Rao, Robert Verkuil, Jason Liu, Tom Sercu, and Alexander Rives. 2021. "Language Models Enable Zero-Shot Prediction of the Effects of Mutations on Protein Function." : 2021.07.09.450648. doi:10.1101/2021.07.09.450648.
Mika, S. 2003. "UniqueProt: Creating Representative Protein Sequence Sets." *Nucleic*

Acids Research 31(13): 3789–91. doi:10.1093/nar/gkg620.
Mirdita, Milot, Konstantin Schütze, Yoshitaka Moriwaki, Lim Heo, Sergey Ovchinnikov, and Martin Steinegger. 2022. “ColabFold: Making Protein Folding Accessible to All.” *Nature Methods*. doi:10.1038/s41592-022-01488-1.
Murray, J. E., N. Laurieri, and R. Delgoda. 2017. “Chapter 24 - Proteins.” In *Pharmacognosy*, eds. Simone Badal and Rupika Delgoda. Boston: Academic Press, 477–94. doi:10.1016/B978-0-12-802104-0.00024-X.
Ng, Pauline C., and Steven Henikoff. 2003. “SIFT: Predicting Amino Acid Changes That Affect Protein Function.” *Nucleic Acids Research* 31(13): 3812–14. doi:10.1093/nar/gkg509.
Nijkamp, Erik, Jeffrey Ruffolo, Eli N. Weinstein, Nikhil Naik, and Ali Madani. 2022. “ProGen2: Exploring the Boundaries of Protein Language Models.” doi:10.48550/arXiv.2206.13517.
Norn, Christoffer, Basile I. M. Wicky, David Juergens, Sirui Liu, David Kim, Doug Tischer, Brian Koepnick, et al. 2021. “Protein Sequence Design by Conformational Landscape Optimization.” *Proceedings of the National Academy of Sciences of the United States of America* 118(11): e2017228118. doi:10.1073/pnas.2017228118.
Notin, Pascal. 2024. “OATML-Markslab/ProteinGym.” https://github.com/OATML-Markslab/ProteinGym (January 8, 2024).
Notin, Pascal, Mafalda Dias, Jonathan Frazer, Javier Marchena-Hurtado, Aidan Gomez, Debora S. Marks, and Yarin Gal. 2022. “Tranception: Protein Fitness Prediction with Autoregressive Transformers and Inference-Time Retrieval.” doi:10.48550/ARXIV.2205.13760.
Notin, Pascal, Aaron W. Kollasch, Daniel Ritter, Lood van Niekerk, Steffanie Paul, Hansen Spinner, Nathan Rollins, et al. 2023. “ProteinGym: Large-Scale Benchmarks for Protein Design and Fitness Prediction.” : 2023.12.07.570727. doi:10.1101/2023.12.07.570727.
Notin, Pascal, Lood Van Niekerk, Aaron W. Kollasch, Daniel Ritter, Yarin Gal, and Debora S. Marks. 2022. “TranceptEVE: Combining Family-Specific and Family-Agnostic Models of Protein Sequences for Improved Fitness Prediction.” : 2022.12.07.519495. doi:10.1101/2022.12.07.519495.
Olenyi, Tobias, Bernhofer, Michael, Mirdita, Mirlot, Steinegger, Martin, and Rost, Burkhard. “Rostclust Redundancy Reduction.”
Paszke, Adam, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, et al. 2019. “PyTorch: An Imperative Style, High-Performance Deep Learning Library.” doi:10.48550/arXiv.1912.01703.
Rao, Roshan M., Jason Liu, Robert Verkuil, Joshua Meier, John Canny, Pieter Abbeel, Tom Sercu, and Alexander Rives. 2021. “MSA Transformer.” In *Proceedings of the 38th International Conference on Machine Learning*, PMLR, 8844–56. https://proceedings.mlr.press/v139/rao21a.html (June 21, 2022).
Reeb, Jonas, Theresa Wirth, and Burkhard Rost. 2020. “Variant Effect Predictions Capture Some Aspects of Deep Mutational Scanning Experiments.” *BMC Bioinformatics* 21(1): 107. doi:10.1186/s12859-020-3439-4.
Riesselman, Adam J., John B. Ingraham, and Debora S. Marks. 2018. “Deep Generative Models of Genetic Variation Capture the Effects of Mutations.” *Nature*

methods 15(10): 816–22.
Roscoe, Benjamin P., Kelly M. Thayer, Konstantin B. Zeldovich, David Fushman, and Daniel N. A. Bolon. 2013. “Analyses of the Effects of All Ubiquitin Point Mutants on Yeast Growth Rate.” *Journal of molecular biology* 425(8): 1363–77. doi:10.1016/j.jmb.2013.01.032.
Rost, Burkhard. 1999. “Twilight Zone of Protein Sequence Alignments.” *Protein Engineering, Design and Selection* 12(2): 85–94. doi:10.1093/protein/12.2.85.
Sander, C., and R. Schneider. 1991. “Database of Homology-Derived Protein Structures and the Structural Meaning of Sequence Alignment.” *Proteins* 9(1): 56–68. doi:10.1002/prot.340090107.
Sinitcyn, Pavel, Alicia L. Richards, Robert J. Weatheritt, Dain R. Brademan, Harald Marx, Evgenia Shishkova, Jesse G. Meyer, et al. 2023. “Global Detection of Human Variants and Isoforms by Deep Proteome Sequencing.” *Nature Biotechnology* 41(12): 1776–86. doi:10.1038/s41587-023-01714-x.
Srivastava, Nitish, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. “Dropout: A Simple Way to Prevent Neural Networks from Overfitting.” *Journal of Machine Learning Research* 15(56): 1929–58.
Steinegger, Martin, and Johannes Söding. 2017. “MMseqs2 Enables Sensitive Protein Sequence Searching for the Analysis of Massive Data Sets.” *Nature Biotechnology* 35(11): 1026–28. doi:10.1038/nbt.3988.
The UniProt Consortium, Alex Bateman, Maria-Jesus Martin, Sandra Orchard, Michele Magrane, Shadab Ahmad, Emanuele Alpi, et al. 2023. “UniProt: The Universal Protein Knowledgebase in 2023.” *Nucleic Acids Research* 51(D1): D523–31. doi:10.1093/nar/gkac1052.
Tiemann, Johanna Katarina Sofie, Henrike Zschach, Kresten Lindorff-Larsen, and Amelie Stein. 2023. “Interpreting the Molecular Mechanisms of Disease Variants in Human Transmembrane Proteins.” *Biophysical Journal* 122(11): 2176–91. doi:10.1016/j.bpj.2022.12.031.
Truong Jr, Timothy F., and Tristan Bepler. 2023. “PoET: A Generative Model of Protein Families as Sequences-of-Sequences.” http://arxiv.org/abs/2306.06156 (July 20, 2023).
Tsuboyama, Kotaro, Justas Dauparas, Jonathan Chen, Elodie Laine, Yasser Mohseni Behbahani, Jonathan J. Weinstein, Niall M. Mangan, Sergey Ovchinnikov, and Gabriel J. Rocklin. 2023. “Mega-Scale Experimental Analysis of Protein Folding Stability in Biology and Design.” *Nature* 620(7973): 434–44. doi:10.1038/s41586-023-06328-6.
Xu, Bing, Naiyan Wang, Tianqi Chen, and Mu Li. 2015. “Empirical Evaluation of Rectified Activations in Convolutional Network.” doi:10.48550/arXiv.1505.00853.
Yang, Jianyi, Ivan Anishchenko, Hahnbeom Park, Zhenling Peng, Sergey Ovchinnikov, and David Baker. 2020. “Improved Protein Structure Prediction Using Predicted Interresidue Orientations.” *Proceedings of the National Academy of Sciences* 117(3): 1496–1503. doi:10.1073/pnas.1914677117.