Rendering protein mutation movies with MutAmore

Konstantin Weissenow [1,2] & Burkhard Rost [1, 3]

1 TUM (Technical University of Munich) Department of Informatics, Bioinformatics & Computational Biology - i12, Boltzmannstr. 3, 85748 Garching/Munich, Germany
2 TUM Graduate School, Center of Doctoral Studies in Informatics and its Applications (CeDoSIA), Boltzmannstr. 11, 85748 Garching, Germany
3 Institute for Advanced Study (TUM-IAS), Lichtenbergstr. 2a, 85748 Garching/Munich, Germany & TUM School of Life Sciences Weihenstephan (WZW), Alte Akademie 8, Freising, Germany
* Corresponding author: k.weissenow@tum.de

Abstract

The success of *AlphaFold2* in reliable protein three-dimensional (3D) structure prediction, assists the move of structural biology toward studies of protein dynamics and mutational impact on structure and function. This transition needs tools that qualitatively assess alternative 3D conformations. We introduce *MutAmore*, a bioinformatics tool that renders individual images of protein 3D structures for, e.g., sequence mutations into a visually intuitive movie format. *MutAmore* streamlines a pipeline casting single amino-acid variations (SAVs) into a dynamic 3D mutation movie providing a qualitative perspective on the mutational landscape of a protein. By default, the tool first generates all possible variants of the sequence reachable through SAVs (L*19 for proteins with L residues). Next, it predicts the structural conformation for all L*19 variants using state-of-the-art models. Finally, it visualizes the mutation matrix and produces a color-coded 3D animation. Alternatively, users can input other types of variants, e.g., from experimental structures. *MutAmore* samples alternative protein configurations to study the dynamical space accessible from SAVs in the post-AlphaFold2 era of structural biology. As the field shifts towards the exploration of alternative conformations of proteins, *MutAmore* aids in the understanding of the structural impact of mutations by providing a flexible pipeline for the generation of protein mutation movies using current and future structure prediction models.

Key words: protein structure prediction, single amino-acid variant, protein mutation movie

Abbreviations used: **3D**, three-dimensional (coordinates); **3D structure**, three-dimensional coordinates of protein structure; **AI**: Artificial Intelligence; **API**: application programing interface; **embeddings:** fixed-size vectors derived from pre-trained pLMs; **GPU:** graphical processing unit; **PDB:** Protein Data Bank; **PIDE**: percentage pairwise sequence identity; **pLM:** protein Language Model; **PMM**: protein mutation movie

Introduction

AI has changed structural biology and its impact. The remarkable success of *AlphaFold2* [1] in effectively predicting protein three-dimensional (3D) structure from sequence is shifting paradigms in structural biology and beyond. *AlphaFold2* combines advanced Artificial Intelligence (AI) with evolutionary information from multiple sequence alignments (MSAs). Reliable 3D predictions for over 200 million proteins [2] have begun to change molecular biology. Concurrently, protein

Language Models (pLMs) have emerged as a new approach to represent protein sequences [3-5]. Downstream prediction methods based on pLM embeddings (vectors representing the last hidden layers of the pLMs: Fig. 1 [4]) find new solutions, e.g., in structure prediction methods such as *ESMFold* [6]. Embeddings allow predictions at unprecedented speed [7, 8], often outperforming state-of-the-art methods [9] using the combination of evolutionary information introduced three decades ago for secondary structure prediction [10] and eclipsed by *AlphaFold2*.

From available 3D model to dynamics? In the post-*AlphaFold2* era, structural biology moves toward a deeper exploration of alternative protein conformations and mutational landscapes [11]. Many proteins significantly change function upon minor sequence changes [12-14]. Understanding these changes will be critical to unlocking the complexities of protein function, protein evolution, and disease progression at the molecular level.

Despite its immense success, *AlphaFold2* often fails to correctly predict the effect of missense mutations upon 3D structure [7, 8, 15]. Even if it correctly captured such effects, it would still be too resource-intensive to be used to explore the mutational space of an average protein, e.g., by evaluating all possible single amino-acid substitutions (SAVs) even for short proteins [7, 8]. Structure predictions based on pLMs such as ESMFold or EMBER2 [7] offer both the necessary speed and promise higher sensitivity to small changes in input sequences (Fig. 2 [8]).

Many methods have been developed to predict the effects of sequence variants upon protein function, including SIFT [16], PolyPhen [17], SNAP2 [18], GEMME [19], DeepSequence [20], Packpred [21], Tranception [22] and VESPA [23] (for a longer list: [22, 24]). On the other hand, tools such as I-Mutant [25, 26], FoldX [27], PoPMuSiC [28], DUET [29], or INPS-MD [30] aim at predicting the structural impact of mutations by providing estimates for changes in stability, folding or dynamics, which greatly helps our understanding of disease emergence. Despite convincing examples for how to use the numerical output from such methods to rationalize on static images about possible dynamical changes [21, 29, 31], none of these methods directly displays the actual structural change in the 3D conformation.

Imagine we had fast and accurate structure predictions for all L*19 SAV mutants of a protein with L residues. How to visualize those data? Today, no comprehensive and accessible visualization tools for alternative protein conformations are available. To fill this void, we introduced *MutAmore* (**MutA**tion **mo**vie **re**nderer), a tool that provides a pipeline rendering a mutated protein sequence into a dynamic 3D protein mutation movie (PMM), thereby making the analysis of mutational landscapes accessible and tangible. The implementation and results presented highlight the potential of *MutAmore* to fill a growing need within the structural biology field. Its efficient visualizations could aid research of protein dynamics, function, evolution, and disease mechanisms.

Implementation

MutAmore is designed to create an animated visualization of the mutational landscape of proteins. Inputting a protein amino acid sequence (or a set thereof) as a FASTA file [32], *MutAmore* first generates all possible single amino acid variants (SAVs) for this sequence(s) and uses a structure prediction model to predict 3D structure for each SAV (Fig. 1). We provide a ready-made interface to *ESMFold* [6] and *ColabFold/AlphaFold2* [33] along with documentation on how to easily use any other prediction model inputting FASTA files and outputting PDB files. Optionally, a user can

provide experimental structures for some of the mutants to *MutAmore*. These variants are then skipped during the prediction stage.

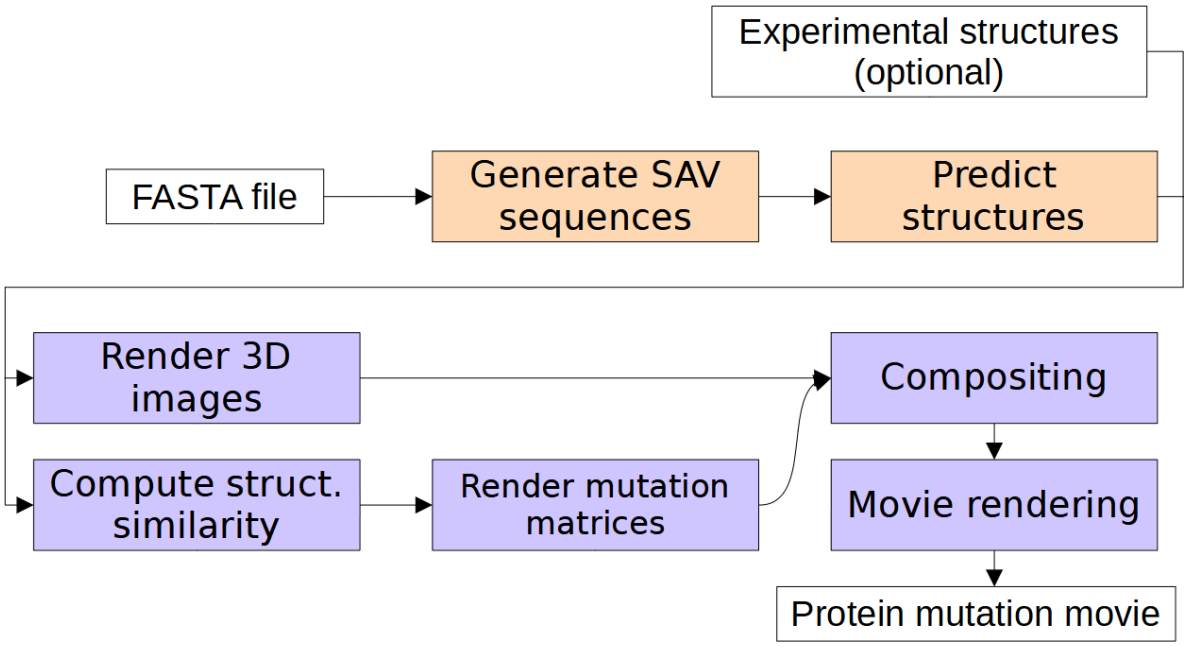

Fig. 1: MutAmore pipeline. The tool generates mutated (all possible SAVs) versions of the input sequences and predicts 3D structure for each, e.g., using ColabFold [33] or ESMFold [6]. Optionally, experimental structures for mutants can be input as an alternative to predictions. *MutAmore* computes the structural difference between each mutant and the predicted wild-type structure and renders mutation profiles along with 3D visualizations. After merging both, the tool renders the final protein mutation movie (PMM). *MutAmore* can be run in a two-step process, e.g., predicting structure (orange) on a server machine and running the rendering steps (blue) on a desktop machine.

MutAmore computes a mutation profile by assessing the structural divergences between each mutated protein and the (predicted) wild-type structure. The *local Distance Difference Test* (lDDT) [34] tallies scores across all residues to derive a single score for each structure pair (typically wild-type/native vs. SAV). Structures similar to the wild-type have scores near one while divergent structures approach zero. These data are then visualized as a mutation matrix (19 x protein length) using Python-Pillow [35].

3D visualizations of all variants are created via the PyMOL API [36] and color-coded to indicate predicted confidence levels (blue for high confidence, yellow-red for low confidence; following the *AlphaFold2* standard). All recent 3D structure prediction methods include predicted lDDT values as confidence scores, including systems which do not build on top of the AlphaFold2 architecture, such as RoseTTAFold [37] and EMBER3D [8]. PyMOL aligns all 19*L SAV structures to the original wild-type prediction prior to rendering to maintain uniformity in viewing angles in the final animated visualization.

MutAmore then assembles the final frames of the animation in residue index order (from first at the N-terminus to the last at the C-terminus), merging the 3D renderings with the mutation profile to create the PMM with ffmpeg [38]. Rendered at a rate of 19 frames per second, each residue remains in focus for one second and shows all potential SAVs for this position in the protein. The

mutated position in each frame is rendered in black in both the 3D visualization and the mutation profile. Details of the SAVs displayed in each frame are indicated in the top-left corner through the standard single-letter amino acid code: XnY meaning that the wild-type amino acid X at residue position n is mutated to amino acid Y (Fig. 2: K76H).

When users provide experimental structures – or otherwise their own labels – in a PDB-formatted file, the visual output highlights the differences between predictions and experimental models by showing the latter at full opacity while applying slight transparency to the former, giving the visual impression of "filling in the gaps" between known structures with predictions.

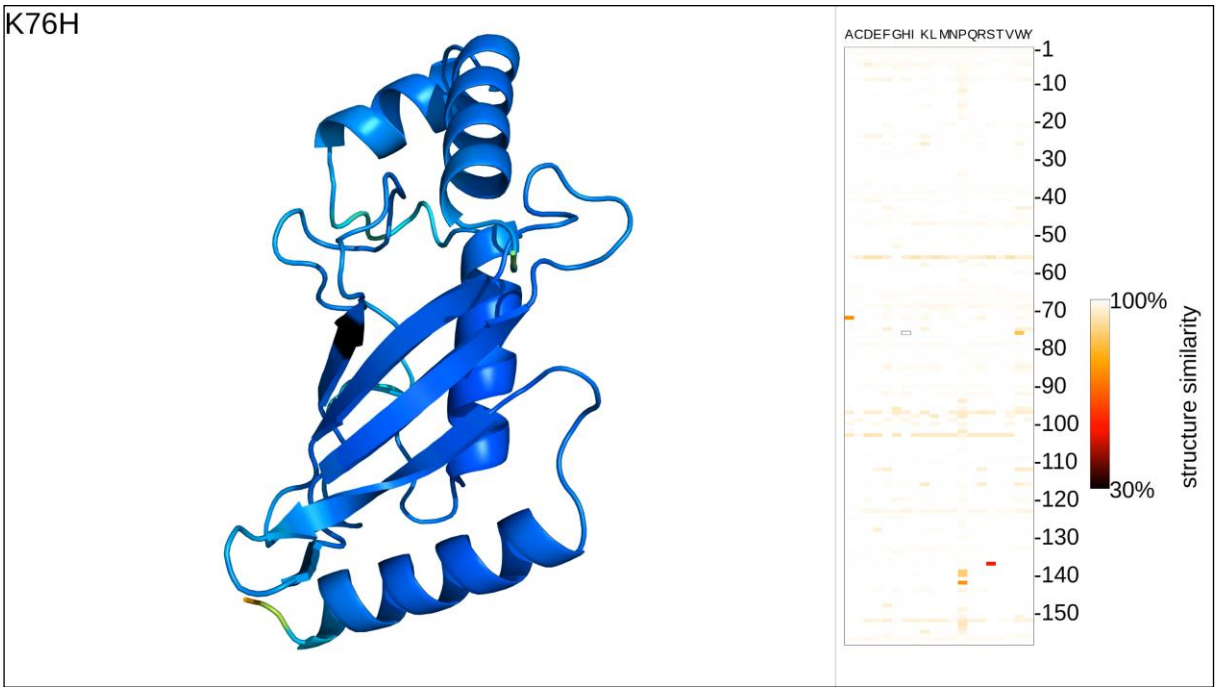

Fig. 2: Single frame of protein mutation movie (PMM). *Left panel:* 3D visualization of the prediction for one single amino-acid variant (SAV). The respective SAV is indicated in the top-left corner (residue position 83: native lysin (K) mutated into isoleucine (I)) and the affected residue is rendered in black in the visualization. All other residues are colored by the *AlphaFold2* predicted confidence in the *AlphaFold2* color scheme (blue: high confidence; yellow-red: low confidence). *Right panel:* the mutation profile shows the structural difference between mutant and wild-type structure with residue indices running along the vertical axis and substitution amino-acids on the horizontal. The currently displayed SAV in each frame is indicated by a black border. The movie frame shown was rendered in a screen resolution of 3840x2160.

Users can adapt the resolution of the final animation to their needs, e.g., choosing high-quality for publications or lower resolution clips for web sharing. The mutation profile automatically scales to the specified vertical resolution for optimal visual interpretability. The 3D visualization rendered by PyMOL automatically chooses a zoom level, which allows enough space to accommodate structural changes caused by mutations. *MutAmore* lets experienced users override the zoom level manually where needed.

MutAmore also lets users render subsets of the most impactful mutations, e.g., top-50: those with the highest effect upon 3D. In this mode, the framerate is slightly reduced (slowed down) for better visual comparison.

Many advanced structure prediction systems require robust and substantial GPU resources. Therefore, *MutAmore* provides an option for a two-step process: computation of all mutation predictions using a server machine, and subsequent processing and rendering of the animated visualization on a desktop computer.

Results

We evaluated *MutAmore* on an Intel Xeon Gold 6248 CPU with a NVidia Quadro RTX 8000 GPU (48GB) using twelve proteins ranging in lengths from 72 to 639 residues, including both globular and membrane proteins. Although *AlphaFold2* [1] or its faster spin-off *ColabFold* [33] outperform *ESMFold* [6], the latter seems slightly better at capturing the effects of SAVs upon structure [7, 8]. We also used ESMFold to predict 3D structures because it is substantially faster which mattered for the 55,879 SAVs in all twelve samples. Then we generated movies at the default resolution of 1280x720 (720p) and 3840x2160 (4K, Table 1).

Table 1: Benchmark of *MutAmore* in 720p and 4K resolution for twelve proteins. *

	Runtime in seconds	
	720p resolution	4K resolution
Prediction (ESMFold)	1,423,814	
3D rendering	26,893	106,829
Structural similarity	4,008	4,224
Mutation profile rendering	4,472	16,871
Compositing final frames	4,730	33,032
Movie rendering	172	1,012
Total (rendering)	40,275	161,968
Total (rendering + prediction)	1,464,089	1,585,782

* Runtime of individual pipeline steps in seconds for screen resolutions 720p and 4K, summed up over twelve proteins with lengths between 72 and 639 residues (total number of SAVs/3D predictions: 55,879; total number of seconds playtime: 2,941). Structure prediction with ESMFold [6] and structural similarity computation using lDDT [34] are naturally unaffected by larger image sizes, but 3D rendering with PyMOL [36], compositing (putting together the individual image pieces) and movie rendering (encoding individual frames to a movie) scale with a factor closer to the increase in screen estate. Tables S1-S12 in the Supporting Online Materials break down the runtime for each of the twelve proteins while Fig. S13 visualizes the scaling of runtime with protein length. Note: the 1.4 million seconds required for the *ESMFold* 3D predictions amounted to about 16 days.

The time required for structure predictions substantially varies with the method used. ColabFold, utilized for the five smallest proteins only, accumulated over a week of GPU time, given the non-linear scale with protein length, this rapidly becomes infeasible for longer proteins. ESMFold obtained predictions for all twelve proteins in about 16 days, with the bulk of the computation time dedicated to the longest samples (11 days for the protein with 639 residues). In contrast, the predictions for the nine smaller proteins were computed by ESMFold in 30 hours. The tremendous increase in runtime by protein length (Fig. S13) is due to GPU memory limitations. Structure

prediction systems such as ESMFold compute multiple samples in parallel, to fill memory as efficiently as possible. For shorter proteins, this allows batches of up to dozens of simultaneous predictions. Longer sequences require more computation time and more memory. This limits the number of samples that can be processed concurrently, leading to an exponential increase in total runtime. Overall, these numbers highlight the need for future structure predictions with both increased speed and memory efficiency to properly explore the mutational landscape of longer proteins.

Given the 3D predictions, creating the animated visualizations is considerably faster. The majority of *MutAmore's* processing time is devoted to the rendering of 3D visualizations, followed by the composition of final frames.

After generating structure predictions on our server hardware, we applied *MutAmore's* rendering pipeline on a consumer grade laptop with an Intel Core i7 6700HQ CPU to compare performance with the server environment. Performance decreased by roughly 30%, showing that *MutAmore* fits to readily available hardware, at least for shorter proteins.

Enhancing the resolution required additional runtime. Rendering at 4K increased the processing time for the 3D visualization, compositing, and rendering by a factor of three to seven over rendering at 720p (Table 1), but generally less than the increase in amount of pixels (4K/720p=9x). The time for computing structural similarity and for generating the profile did not differ much between 720p and 4K (Table 1). We provide a detailed breakdown of the runtime of all pipeline steps on the twelve individual samples in the Supporting Online Materials (SOM: Tables S1-S12). Even at 4K, the total processing time for *MutAmore* remained substantially below that needed to predict structures even with *ESMFold* for all but the shortest protein sequences (Fig. S13).

Limitations. A profound consequence of all attempts to visualize the dynamics of 3D objects lies in the obstruction of internal parts. For instance, buried residues and changes of local regions around these will remain obscured from our 3D movies. For some proteins, such as beta barrels, a cleverly chosen viewing angle might provide a better perspective, but globular proteins do not provide such an alternative. Transparency might address such issues, but transparency tends to cause a lack of depth perception and too much visual clutter to clearly comprehend the visual information being presented. Thus, it remains unclear how to show internal structural changes of proteins in a visually concise manner apart from going back to two-dimensional distance maps, which are only intuitive to a well-trained structural expert.

Another implicit limitation for using 3D predictions for all 19-non-native SAVs is in the substantial demand on computing resources of today's prediction methods. This becomes particularly challenging for long proteins (Fig. S13), *MutAmore* would greatly benefit from future prediction systems that are leaner without having to sacrifice performance [8].

Too few high-resolution experimental structures establish the effect of point mutations upon 3D to evaluate how well prediction methods capture SAV effects. Proxying deep mutational scanning data [39] might suffice to establish correlations between observed and predicted impact upon function without probing how well methods predict the effect of SAVs upon 3D structure and dynamics. Once more experimental data of variant structures will be available, *MutAmore* could be extended to serve as a benchmarking tool for the sensitivity of structure prediction.

A webserver for *MutAmore* might ease the access for users with less experience in computational biology. While we are looking for the resources to realize such a project, we provide a Google Colab Notebook linked on the *MutAmore* website and allows the creation of PMM's without the need of a local installation.

Conclusions

We designed *MutAmore* to bridge a crucial gap in the post-AlphaFold era. By rendering conceivable and visually comprehensible protein mutation movies (PMMs) of single amino-acid substitutions (SAVs), *MutAmore* enhances the exploration and understanding of alternative protein conformations brought on by mutations. This is particularly significant in translating the depth and complexity of the protein mutational landscape. Our benchmark demonstrated the efficiency and versatility of *MutAmore* even for high-definition 4K video settings. The tool effectively balances computational load by allowing multi-step operation across multiple systems, ensuring usability across varied system capabilities. We hope that structural biology increasingly shifts toward routine analysis of alternate protein conformations. *MutAmore* supports such a more dynamic perspective on protein structures and might aid analyzing protein function and studying protein evolution.

Availability

MutAmore is publicly available at https://github.com/kWeissenow/MutAmore.

Acknowledgements

Thanks to Tim Karl (TUM) for invaluable help with hardware and software; to Nikita Kugut (TUM) for support with many other aspects of this work; to Michael Heinzinger (TUM) for helpful discussions and comments on work and manuscript. Furthermore, we thank all those who make experimental and predicted structures and other resources publicly available, in particular, thanks to DeepMind (AlphaFold2) and Meta (ESMFold), as well as to the creators and community of PyMOL and ffmpeg.

References

1. Jumper, J., et al., *Highly accurate protein structure prediction with AlphaFold.* Nature, 2021.
2. Varadi, M., et al., *AlphaFold Protein Structure Database: massively expanding the structural coverage of protein-sequence space with high-accuracy models.* Nucleic Acids Research, 2021. **50**(D1): p. D439-D444.
3. Rao, R., et al., *Transformer protein language models are unsupervised structure learners.* bioRxiv, 2020: p. 2020.12.15.422761.
4. Elnaggar, A., et al., *ProtTrans: Towards Cracking the Language of Lifes Code Through Self-Supervised Deep Learning and High Performance Computing.* IEEE Trans Pattern Anal Mach Intell, 2021.
5. Heinzinger, M., et al., *Modeling aspects of the language of life through transfer-learning protein sequences.* BMC Bioinformatics, 2019. **20**(1): p. 723.

6. Lin, Z., et al., *Language models of protein sequences at the scale of evolution enable accurate structure prediction.* bioRxiv, 2022: p. 2022.07.20.500902.
7. Weissenow, K., M. Heinzinger, and B. Rost, *Protein language-model embeddings for fast, accurate, and alignment-free protein structure prediction.* Structure, 2022. **30**(8): p. 1169-1177.e4.
8. Weissenow, K., et al., *Ultra-fast protein structure prediction to capture effects of sequence variation in mutation movies.* bioRxiv, 2022: p. 2022.11.14.516473.
9. Bordin, N., et al., *Novel machine learning approaches revolutionize protein knowledge.* Trends in Biochemical Sciences, 2023. **48**(4): p. 345-359.
10. Rost, B. and C. Sander, *Prediction of protein secondary structure at better than 70% accuracy.* Journal of Molecular Biology, 1993. **232**: p. 584-599.
11. Sala, D., et al., *Modeling conformational states of proteins with AlphaFold.* Current Opinion in Structural Biology, 2023. **81**: p. 102645.
12. Vedithi, S.C., et al., *Structural Implications of Mutations Conferring Rifampin Resistance in Mycobacterium leprae.* Scientific Reports, 2018. **8**(1): p. 5016.
13. Portelli, S., et al., *Understanding molecular consequences of putative drug resistant mutations in Mycobacterium tuberculosis.* Scientific Reports, 2018. **8**(1): p. 15356.
14. Gerasimavicius, L., B.J. Livesey, and J.A. Marsh, *Loss-of-function, gain-of-function and dominant-negative mutations have profoundly different effects on protein structure.* Nature Communications, 2022. **13**(1): p. 3895.
15. Buel, G.R. and K.J. Walters, *Can AlphaFold2 predict the impact of missense mutations on structure?* Nature Structural & Molecular Biology, 2022. **29**(1): p. 1-2.
16. Ng, P.C. and S. Henikoff, *SIFT: Predicting amino acid changes that affect protein function.* Nucleic Acids Res, 2003. **31**(13): p. 3812-4.
17. Adzhubei, I., D.M. Jordan, and S.R. Sunyaev, *Predicting functional effect of human missense mutations using PolyPhen-2.* Curr Protoc Hum Genet, 2013. **Chapter 7**: p. Unit7.20.
18. Hecht, M., Y. Bromberg, and B. Rost, *Better prediction of functional effects for sequence variants.* BMC Genomics, 2015. **16 Suppl 8**(Suppl 8): p. S1.
19. Laine, E., Y. Karami, and A. Carbone, *GEMME: A Simple and Fast Global Epistatic Model Predicting Mutational Effects.* Molecular Biology and Evolution, 2019. **36**(11): p. 2604-2619.
20. Riesselman, A.J., J.B. Ingraham, and D.S. Marks, *Deep generative models of genetic variation capture the effects of mutations.* Nature Methods, 2018. **15**(10): p. 816-822.
21. Tan, K.P., et al., *Packpred: Predicting the Functional Effect of Missense Mutations.* Front Mol Biosci, 2021. **8**: p. 646288.
22. Notin, P.a.D., Mafalda and Frazer, Jonathan and Hurtado, Javier Marchena and Gomez, Aidan N and Marks, Debora and Gal, Yarin. *Tranception: Protein Fitness Prediction with Autoregressive Transformers and Inference-time Retrieval*. in *Proceedings of the 39th International Conference on Machine Learning*. PMLR.
23. Marquet, C., et al., *Embeddings from protein language models predict conservation and variant effects.* Human Genetics, 2022. **141**(10): p. 1629-1647.
24. Livesey, B.J. and J.A. Marsh, *Updated benchmarking of variant effect predictors using deep mutational scanning.* Mol Syst Biol, 2023. **19**(8): p. e11474.
25. Capriotti, E., P. Fariselli, and R. Casadio, *I-Mutant2.0: predicting stability changes upon mutation from the protein sequence or structure.* Nucleic Acids Res, 2005. **33**(Web Server issue): p. W306-10.

26. Capriotti, E., et al., *A three-state prediction of single point mutations on protein stability changes.* BMC Bioinformatics, 2008. **9**(2): p. S6.
27. Schymkowitz, J., et al., *The FoldX web server: an online force field.* Nucleic Acids Res, 2005. **33**(Web Server issue): p. W382-8.
28. Dehouck, Y., et al., *PoPMuSiC 2.1: a web server for the estimation of protein stability changes upon mutation and sequence optimality.* BMC Bioinformatics, 2011. **12**(1): p. 151.
29. Pandurangan, A.P. and T.L. Blundell, *Prediction of impacts of mutations on protein structure and interactions: SDM, a statistical approach, and mCSM, using machine learning.* Protein Sci, 2020. **29**(1): p. 247-257.
30. Savojardo, C., et al., *INPS-MD: a web server to predict stability of protein variants from sequence and structure.* Bioinformatics, 2016. **32**(16): p. 2542-4.
31. Hecht, M., Y. Bromberg, and B. Rost, *News from the protein mutability landscape.* J Mol Biol, 2013. **425**(21): p. 3937-48.
32. Pearson, W.R. and D.J. Lipman, *Improved Tools for Biological Sequence Comparison.* Proceedings of the National Academy of Sciences of the United States of America, 1988. **85**(8): p. 2444-2448.
33. Mirdita, M., et al., *ColabFold - Making protein folding accessible to all.* bioRxiv, 2021: p. 2021.08.15.456425.
34. Mariani, V., et al., *lDDT: a local superposition-free score for comparing protein structures and models using distance difference tests.* Bioinformatics, 2013. **29**(21): p. 2722-8.
35. Clark, A., *Python-Pillow*. 2010.
36. Schrodinger, LLC, *The PyMOL Molecular Graphics System, Version 1.8*. 2015.
37. Baek, M., et al., *Accurate prediction of protein structures and interactions using a three-track neural network.* Science, 2021. **373**(6557): p. 871-876.
38. Tomar, S., *Converting video formats with FFmpeg.* Linux Journal, 2006. **2006**(146): p. 10.
39. Fowler, D.M. and S. Fields, *Deep mutational scanning: a new style of protein science.* Nature Methods, 2014. **11**(8): p. 801-807.