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

We introduce MutAmore, a bioinformatics tool that renders individual photos of, e.g., protein three-dimensional (3D) structures for protein mutations into an easily understandable movie format. MutAmore provides a streamlined pipeline casting single amino-acid variations (SAVs) into a dynamic 3D protein mutation movie which provides a qualitative perspective on the mutational landscape of a protein. By default, the tool first generates all possible variants of the protein sequence reachable through a 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. MutAmore samples alternative protein configurations to study the dynamical space accessible from SAVs in the post-AlphaFold2 era of structural biology.

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 (Jumper et al., 2021) 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 (Varadi et al., 2021) have begun to change molecular biology. Concurrently, protein Language Model (pLMs) have emerged as a new approach to represent protein sequences (Elnaggar et al., 2021; Heinzinger et al., 2019; Rao et al., 2020). Downstream prediction methods based on pLM embeddings (essentially the vectors

representing the last hidden layers of the pLM: Fig. 1 (Elnaggar et al., 2021)), in particular structure prediction methods such as ESMFold (Lin et al., 2022), allow predictions at unprecedented speed (Weissenow, Heinzinger, & Rost, 2022; Weissenow, Heinzinger, Steinegger, et al., 2022), often outperforming the state-of-the-art methods (Bordin et al., 2023) using the combination of evolutionary information introduced three decades ago for secondary structure prediction (Rost & Sander, 1993) 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 (Sala et al., 2023). Numerous proteins exhibit significant functional diversity through minor alterations in their sequences (Gerasimavicius et al., 2022; Portelli et al., 2018; Vedithi et al., 2018). Understanding these alternate configurations will have critical implications in 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 structure (Buel & Walters, 2022; Weissenow, Heinzinger, & Rost, 2022; Weissenow, Heinzinger, Steinegger, et al., 2022). Even if it correctly captured such effects, it would still be too resource intensive to be used as a tool for the exploration of a protein's mutational space, e.g., by evaluating all possible single amino-acid substitutions (SAVs) even for short proteins (Weissenow, Heinzinger, & Rost, 2022; Weissenow, Heinzinger, Steinegger, et al., 2022). Structure predictions based on pLMs such as ESMFold or EMBER2 (Weissenow, Heinzinger, & Rost, 2022) offer both the necessary speed and promise higher sensitivity to small changes in input sequences.

Imagine we had fast and accurate predictions of all L*19 SAV mutants of a protein with L residues. How could we visualize those predictions? No comprehensive and accessible visualization tools for alternative protein conformations are available at the moment. To fill this void, we introduced *MutAmore* (**MutA**tion **mo**vie **re**nderer), a tool that provides a pipeline that renders 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. The tool's ability to offer efficient visualizations could aid research into 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 (Pearson & Lipman, 1988), *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 (Lin et al., 2022) and ColabFold (AlphaFold2) (Mirdita et al., 2021) along with documentation on how to easily use any other prediction model inputting FASTA files and outputting PDB files.

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)

(Mariani et al., 2013) tallies scores across all residues to derive a single score for each structure pair (typically wild-type/native vs. SAV-variant). Structures similar to the wild-type have scores near 1 while divergent structures have scores closer to 0. These data are then visualized as a mutation matrix (19 x protein length) using Python-Pillow (Clark, 2010).

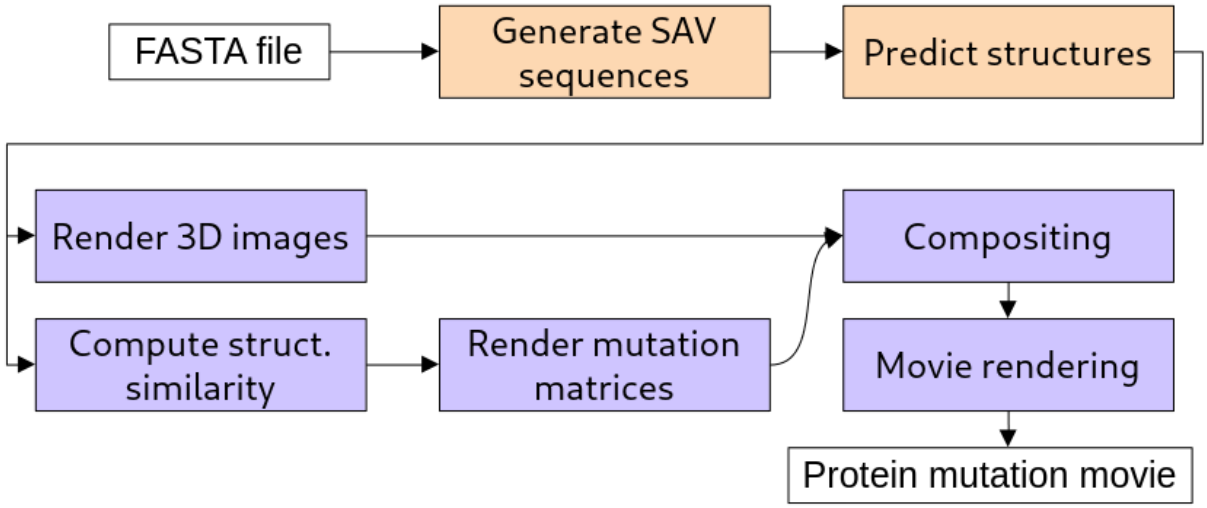

Fig. 1: MutAmore pipeline. The tool generates mutated (all possible SAVs) versions of the input sequences and predicts 3D structure for each, e.g., using ColabFold (Mirdita et al., 2021) or ESMFold (Lin et al., 2022). MutAmore computes the structural difference between each mutant and the predicted wild-type 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.

3D visualizations of all variants are created via the PyMOL API (Schrodinger, 2015) and color-coded to indicate predicted confidence levels (blue for high confidence, yellow-red for low confidence; following the AlphaFold2 standard). 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 (Tomar, 2006). Each frame, rendered at a rate of 19 frames per second, stays for one second and shows all potential SAVs for one residue 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: K83I, i.e., lysine at position 83 mutated to isoleucine).

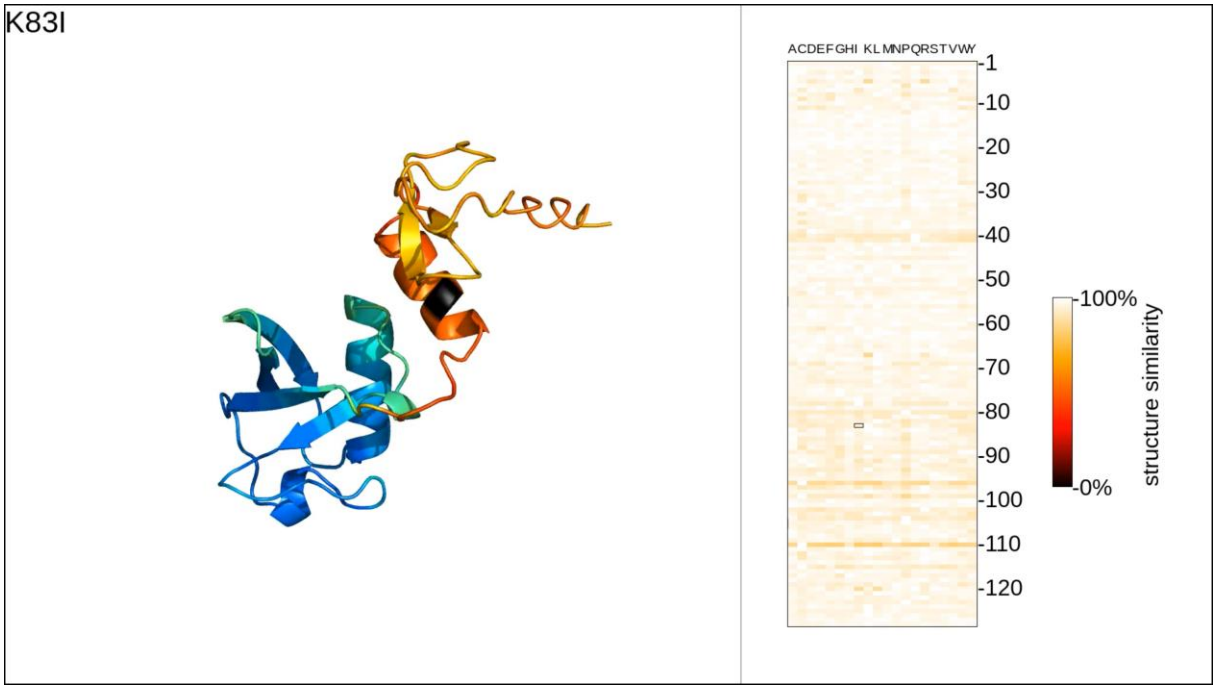

Fig. 2: Single frame of protein mutation movie (PMM). *Left panel:* 3D visualization of a single mutant displaying the prediction for one single amino-acid variant (SAV). The respective SAV is indicated in the top-left corner and the affected residue is rendered in black in the visualization. All other 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 1920x1080.

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.

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 nine proteins ranging in lengths from 72 to 248 residues. Although *AlphaFold2* (Jumper et al., 2021) or its faster spin-off *ColabFold* (Mirdita et al., 2021) outperform *ESMFold* (Lin et al., 2022), the latter seems slightly better at capturing the effects of SAVs upon structure (Weissenow, Heinzinger, & Rost, 2022; Weissenow, Heinzinger, Steinegger, et al., 2022). We also used ESMFold to predict 3D structures because it is substantially faster which mattered for the 29,108 SAVs in all nine 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 nine proteins. *

	Runtime (in sec)	
	720p resolution	4K resolution
Prediction (ESMFold)	110,334	
3D rendering	13,962	61,652
Structural similarity	1,506	1,494
Mutation profile rendering	2,219	2,535
Compositing final frames	2,738	17,057
Movie rendering	107	654
Total (rendering)	20,532	83,392
Total (all)	130,866	193,726

* Runtime of individual pipeline steps in seconds for screen resolutions 720p and 4K, summed up over nine proteins with lengths between 72 and 248 residues (total number of SAVs/3D predictions: 29,108; total number of seconds playtime: 1,532). Structure prediction with ESMFold (Lin et al., 2022) and structural similarity computation using lDDT (Mariani et al., 2013) are naturally unaffected by larger image sizes, but 3D rendering with PyMOL (Schrodinger, 2015), 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-S9 break down the runtime for each of the nine proteins while Fig. S10 visualizes the scaling of runtime with protein length.

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. ESMFold obtained all predictions for the nine proteins in about 30 hours. 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 also ran MutAmore's rendering pipeline on a consumer grade laptop with an Intel Core i7 6700HQ CPU to compare performance with the server environment and observed roughly the same runtime, presumably due to PyMOL mostly scaling with single-core performance.

Enhancing the resolution required additional runtime. Rendering at 4K increased the processing time for the 3D visualization, compositing, and rendering by a factor of six to eight over rendering at 720p (Table 1), i.e., slightly less than the increase in amount of pixels (4K/720=9). The time for computing structural similarity and for generating the profile did not differ much between 720p and 4K (Table 1). Even at 4K, the total processing time for *MutAmore* remained substantially below that needed to predict structures even with ESMFold. We provide a detailed breakdown of the runtime of all pipeline steps on the nine individual samples in Tables S1-S9.

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 optimal usability across varied system capabilities. We hope that structural biology increasingly shifts toward routine analysis of alternate protein conformations. *MutAmore* is ready to support such a more dynamic perspective on structures. It is already ready to have widespread application in analyzing protein function, studying protein evolution and dynamics.

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 for helpful discussions and comments on the 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

Bordin, N., Dallago, C., Heinzinger, M., Kim, S., Littmann, M., Rauer, C., Steinegger, M., Rost, B., & Orengo, C. (2023). Novel machine learning approaches revolutionize protein knowledge. *Trends in Biochemical Sciences*, *48*(4), 345-359. https://doi.org/10.1016/j.tibs.2022.11.001

Buel, G. R., & Walters, K. J. (2022). Can AlphaFold2 predict the impact of missense mutations on structure? *Nature Structural & Molecular Biology*, *29*(1), 1-2. https://doi.org/10.1038/s41594-021-00714-2

Clark, A. (2010). *Python-Pillow*. In (Version 9.4.0) python-pillow.org

Elnaggar, A., Heinzinger, M., Dallago, C., Rehawi, G., Yu, W., Jones, L., Gibbs, T., Feher, T., Angerer, C., Steinegger, M., Bhowmik, D., & Rost, B. (2021). ProtTrans: Towards Cracking the Language of Lifes Code Through Self-Supervised Deep Learning and High Performance Computing. *IEEE Trans Pattern Anal Mach Intell*. https://doi.org/10.1109/TPAMI.2021.3095381

Gerasimavicius, L., Livesey, B. J., & Marsh, J. A. (2022). Loss-of-function, gain-of-function and dominant-negative mutations have profoundly different effects on protein structure. *Nature Communications*, *13*(1), 3895. https://doi.org/10.1038/s41467-022-31686-6

Heinzinger, M., Elnaggar, A., Wang, Y., Dallago, C., Nechaev, D., Matthes, F., & Rost, B. (2019). Modeling aspects of the language of life through transfer-learning protein sequences. *BMC Bioinformatics*, *20*(1), 723. https://doi.org/10.1186/s12859-019-3220-8

Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., . . . Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*. https://doi.org/10.1038/s41586-021-03819-2

Lin, Z., Akin, H., Rao, R., Hie, B., Zhu, Z., Lu, W., Santos Costa, A. d., Fazel-Zarandi, M., Sercu, T., Candido, S., & Rives, A. (2022). Language models of protein sequences at the scale of evolution enable accurate structure prediction. *bioRxiv*, 2022.2007.2020.500902. https://doi.org/10.1101/2022.07.20.500902

Mariani, V., Biasini, M., Barbato, A., & Schwede, T. (2013). lDDT: a local superposition-free score for comparing protein structures and models using distance difference tests. *Bioinformatics*, *29*(21), 2722-2728. https://doi.org/10.1093/bioinformatics/btt473

Mirdita, M., Schütze, K., Moriwaki, Y., Heo, L., Ovchinnikov, S., & Steinegger, M. (2021). ColabFold - Making protein folding accessible to all. *bioRxiv*, 2021.2008.2015.456425. https://doi.org/10.1101/2021.08.15.456425

Pearson, W. R., & Lipman, D. J. (1988). Improved Tools for Biological Sequence Comparison. *Proceedings of the National Academy of Sciences of the United States of America*, *85*(8), 2444-2448. http://www.jstor.org/stable/31450

Portelli, S., Phelan, J. E., Ascher, D. B., Clark, T. G., & Furnham, N. (2018). Understanding molecular consequences of putative drug resistant mutations in Mycobacterium tuberculosis. *Scientific Reports*, *8*(1), 15356. https://doi.org/10.1038/s41598-018-33370-6

Rao, R., Meier, J., Sercu, T., Ovchinnikov, S., & Rives, A. (2020). Transformer protein language models are unsupervised structure learners. *bioRxiv*, 2020.2012.2015.422761. https://doi.org/10.1101/2020.12.15.422761

Rost, B., & Sander, C. (1993). Prediction of protein secondary structure at better than 70% accuracy. *Journal of Molecular Biology*, *232*, 584-599. http://www.rostlab.org/papers/1993_jmb/

Sala, D., Engelberger, F., McHaourab, H. S., & Meiler, J. (2023). Modeling conformational states of proteins with AlphaFold. *Current Opinion in Structural Biology*, *81*, 102645. https://doi.org/https://doi.org/10.1016/j.sbi.2023.102645

Schrodinger, LLC. (2015). *The PyMOL Molecular Graphics System, Version 1.8*.

Tomar, S. (2006). Converting video formats with FFmpeg. *Linux Journal*, *2006*(146), 10.

Varadi, M., Anyango, S., Deshpande, M., Nair, S., Natassia, C., Yordanova, G., Yuan, D., Stroe, O., Wood, G., Laydon, A., Žídek, A., Green, T., Tunyasuvunakool, K., Petersen, S., Jumper, J., Clancy, E., Green, R., Vora, A., Lutfi, M., . . . Velankar, S. (2021). AlphaFold Protein Structure Database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic Acids Research*, *50*(D1), D439-D444. https://doi.org/10.1093/nar/gkab1061

Vedithi, S. C., Malhotra, S., Das, M., Daniel, S., Kishore, N., George, A., Arumugam, S., Rajan, L., Ebenezer, M., Ascher, D. B., Arnold, E., & Blundell, T. L. (2018). Structural Implications of Mutations Conferring Rifampin Resistance in Mycobacterium leprae. *Scientific Reports*, *8*(1), 5016. https://doi.org/10.1038/s41598-018-23423-1

Weissenow, K., Heinzinger, M., & Rost, B. (2022). Protein language-model embeddings for fast, accurate, and alignment-free protein structure prediction. *Structure*, *30*(8), 1169-1177.e1164. https://doi.org/10.1016/j.str.2022.05.001

Weissenow, K., Heinzinger, M., Steinegger, M., & Rost, B. (2022). Ultra-fast protein structure prediction to capture effects of sequence variation in mutation movies. *bioRxiv*, 2022.2011.2014.516473. https://doi.org/10.1101/2022.11.14.516473

Figure captions

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 (Mirdita et al., 2021) or ESMFold (Lin et al., 2022). 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.

Fig. 2: Single frame of protein mutation movie (PMM). *Left panel:* 3D visualization of a single mutant displaying the prediction for one single amino-acid variant (SAV). The respective SAV is indicated in the top-left corner 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 1920x1080.

Tables

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

	Runtime (in sec)	
	720p resolution	4K resolution
Prediction (ESMFold)	110,334	
3D rendering	13,962	61,652
Structural similarity	1,506	1,494
Mutation profile rendering	2,219	2,535
Compositing final frames	2,738	17,057
Movie rendering	107	654
Total (rendering)	20,532	83,392
Total (all)	130,866	193,726

* Runtime of individual pipeline steps in seconds for screen resolutions 720p and 4K, summed up over nine proteins with lengths between 72 and 248 residues (total number of SAVs/3D predictions: 29,108; total number of seconds playtime: 1,532). Structure prediction with ESMFold (Lin et al., 2022) and structural similarity computation using lDDT (Mariani et al., 2013) are naturally unaffected by larger image sizes, but 3D rendering with PyMOL (Schrodinger, 2015), 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-S9 break down the runtime for each of the nine proteins while Fig. S10 visualizes the scaling of runtime with protein length.