

## AI-Predicted Protein Deformation Encodes Energy Landscape Perturbation

John M. McBride<sup>1,\*</sup> and Tsvi Tlusty<sup>1,2,†</sup><sup>1</sup>Center for Algorithmic and Robotized Synthesis, Institute for Basic Science, Ulsan 44919, South Korea<sup>2</sup>Departments of Physics and Chemistry, Ulsan National Institute of Science and Technology, Ulsan 44919, South Korea

(Received 27 November 2023; revised 27 February 2024; accepted 24 July 2024; published 26 August 2024)

AI algorithms have proven to be excellent predictors of protein structure, but whether and how much these algorithms can capture the underlying physics remains an open question. Here, we aim to test this question using the AlphaFold2 (AF) algorithm: We use AF to predict the subtle structural deformation induced by single mutations, quantified by strain, and compare with experimental datasets of corresponding perturbations in folding free energy  $\Delta\Delta G$ . Unexpectedly, we find that physical strain alone—without any additional data or computation—correlates almost as well with  $\Delta\Delta G$  as state-of-the-art energy-based and machine-learning predictors. This indicates that the AF-predicted structures alone encode fine details about the energy landscape. In particular, the structures encode significant information on stability, enough to estimate (de-)stabilizing effects of mutations, thus paving the way for the development of novel, structure-based stability predictors for protein design and evolution.

DOI: 10.1103/PhysRevLett.133.098401

AI has ushered in a revolution in structural biology, yet we are still in uncharted waters [1,2]. In particular, it is not clear whether AI algorithms that predict protein structure from sequence, such as AlphaFold (AF) [3] or RoseTTAFold [4], owe their unprecedented accuracy to highly sophisticated pattern recognition or these algorithms can capture some of the many-body physics underlying protein folding. Recent studies provide evidence suggesting that AF has learned an *effective* energy functional that is searched in order to accurately predict the native structure [5], even if it includes uncommon structural motifs [6].

A stringent test for whether an AI algorithm has learned the actual *physical* energy landscape would be the capacity to probe, from the predicted structure alone, changes in the thermodynamic free energy ( $\Delta\Delta G$ ) due to single mutations. This would indicate that the predicted structure encodes fine details about the physical energy landscape. Besides this fundamental interest, such capacity may be impactful in applications: Of particular importance for protein design and sequence generation [7–10], and for protein evolution [11–13] is the ability to predict whether a given protein sequence will lead to a stably folded structure [3,4]. AF has been reported to predict folded structures for proteins that are not stable [14–16]. Despite this, some analyses suggest that AF and RoseTTAFold can be used for predicting stability changes upon mutation [14,17–20]. All this motivates us to systematically investigate here the question of whether AI-predicted structures can be used to infer changes to free energy landscapes.

To this end, we study a curated subset of 2499 measurements of stability change ( $\Delta\Delta G$ ) due to single mutations, taken from ThermoMutDB (TMDB) [21]. Strain is a simple general measure of deformation in proteins [22–27]. We find that the deformation upon mutation—as measured by the effective strain [14]—correlates well with  $\Delta\Delta G$ . We show it is essential to average over an ensemble of multiple AF-predicted structures to get precise estimates of strain due to mutation, and that most of the relevant information is gleaned from the residues within 15 Å of the mutated residue. Our initial motivation was to examine whether the AF structures encode any information about stability. Surprisingly, we found that correlations between strain and  $\Delta\Delta G$  compare well against those obtained using state-of-the-art  $\Delta\Delta G$  predictors, suggesting that AF predictions are highly informative of stability changes. We propose that new energy-based force fields can be developed that may provide a mechanistic understanding of the effects of mutations on stability. Since stability measurements are easier than structure determination, such a development could settle the question of whether AI algorithms have truly learned the physics of protein folding by testing on massive stability datasets [5,6,28].

*Strain correlates with  $\Delta\Delta G$  in Thermonuclease*—We first examine thermonuclease (NUC, Uniprot ID, P00644; *staphylococcus aureus*)—the protein that has the highest number of mutants in TMDB (491 after applying controls, Appendix A). NUC consists of a folded region (starting around K88) and an extended disordered region near the *N* terminus [Fig. 1(a)], indicated by the low pLDDT (AF-predicted confidence score) values [Figs. 1(a)–1(b)]. The NUC mutants are all sampled from the folded domain [Fig. 1(c)]. Note that we define  $\Delta\Delta G$  such that an increase in  $\Delta G$  upon mutation is *destabilizing*.

\*Contact author: jmmcbride@protonmail.com

†Contact author: tsvitlusty@gmail.com

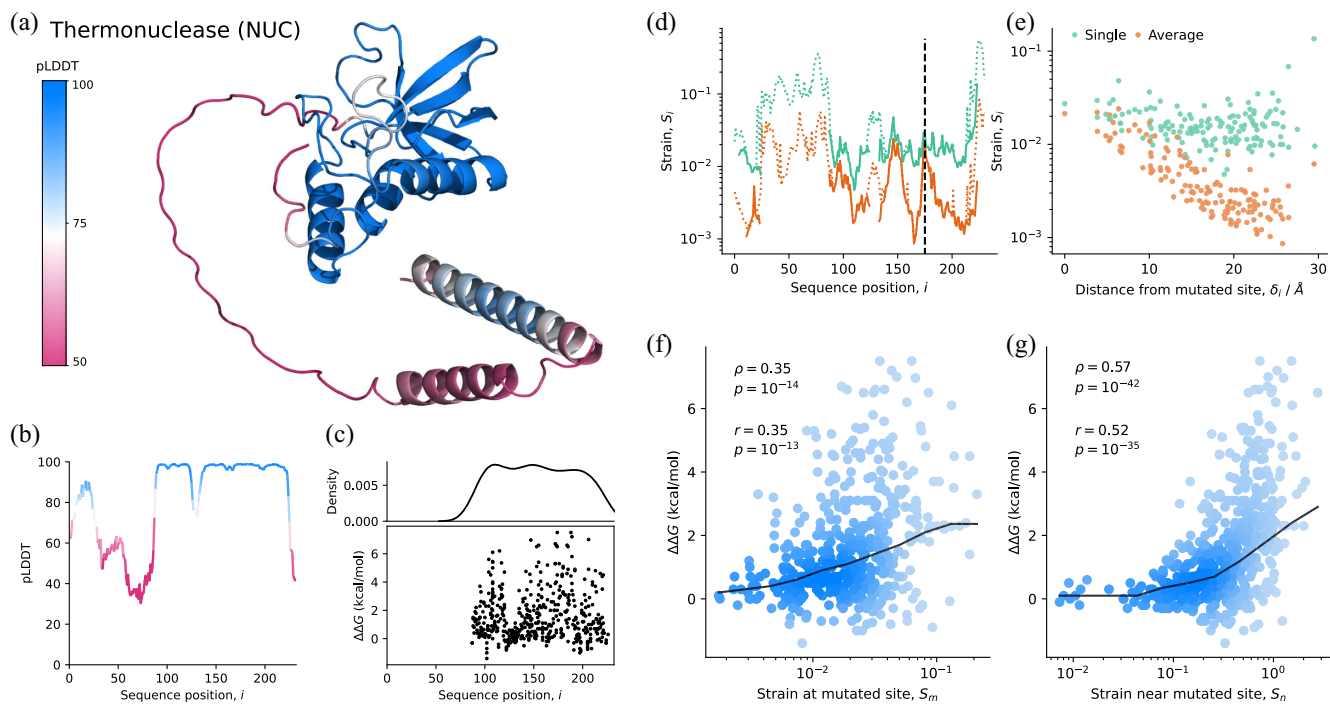


FIG. 1. Strain calculated using AF-predicted structures correlates with  $\Delta\Delta G$ . (a) AF-predicted structure of thermonuclease from *staphylococcus aureus* (NUC)—the most common protein in TMDB. Residues are colored according to pLDDT (AF-predicted confidence score). (b) pLDDT per residue. (c) Distribution of 491 mutation sites along the sequence, and corresponding changes in stability  $\Delta\Delta G$ . (d) Strain upon mutation (A176G,  $\Delta\Delta G = 2.4$  kcal/mol) per residue; the mutated site is indicated by the black dashed line. The strain calculation either includes (dotted line) or excludes (solid line) residues with pLDDT  $< 70$ , and uses either a single pair of structures (green) or pairs of averaged structures (orange). (e) Strain as a function of distance from the mutated site,  $\delta_i$ , for single pairs and ensemble-averaged pairs of structures. (f)–(g) Empirical  $\Delta\Delta G$  vs strain at the mutated site,  $S_m$  (f), and the sum of strain over all residues within 15 Å of the mutated site,  $S_n$  (g). Solid line shows the median; Spearman’s  $\rho$ , Pearson’s  $r$ , and corresponding  $p$  values are shown; circles are shaded by density.

To measure deformation, we calculate effective strain (ES; Appendix B) per residue,  $S_i$ , between wild-type (WT) and mutant structures predicted by AF [Fig. 1(d)]. Disordered residues always show high ES [Fig. 1(d), dotted lines] due to prediction noise, regardless of mutations [14]. We therefore exclude residues whose pLDDT  $< 70$  from the ES calculations [Fig. 1(d), solid lines]. Even without disordered residues, if we only compare two static structures [Figs. 1(d) and 1(e), teal], we still see residual ES in regions far from the mutated site. This occurs since regions in proteins with high flexibility tend to have high variability across repeat AF predictions (Supplemental Material Fig. 1) [29]. To achieve a more accurate estimate of deformation *due to mutation*, we calculate deformation using “averaged” AF structures (Appendix B) [14]. This drastically reduces ES in most regions (which originates chiefly from noise and fluctuations), except for regions near the mutated site [Figs. 1(d) and 1(e), orange].

Using this more precise measure of deformation due to mutation, we find a significant correlation (Spearman’s  $\rho = 0.35$ ) between strain at the mutated site  $S_m$  and change in stability  $\Delta\Delta G$  [Fig. 1(f)]. This correlation is even higher (Spearman’s  $\rho = 0.57$ ) when calculating the sum of strain

over all residues within a spherical neighborhood of radius  $\gamma = 15$  Å around the mutated site,  $S_n$  [Fig. 1(g)]. See Supplemental Material Fig. 2 for similar analyses without excluding low pLDDT residues and without using average structures [29]. Our rationale for choosing  $\gamma = 15$  Å will become clear in a following section. For now, we highlight that, in this particular example (NUC) it appears that AF-predicted deformation correlates quite well with empirical measurements of changes in stability.

*Strain correlates with  $\Delta\Delta G$  within protein families*—We expand our analysis of NUC to more protein families, again focusing on the families that have the highest coverage in TMDB. For the second-to-fifth most common proteins, we find correlations between stability change and local deformation ranging from  $0.39 \leq \rho \leq 0.61$  [Fig. 2(a)]. Extending this analysis to the 40 most common proteins reveals that there are insufficient samples to show significant strain-stability correlations in most cases. Nevertheless, in the 16 correlations that are statistically significant, the average Spearman’s  $\rho$  is 0.54 [Fig. 2(b)] with an overall range 0.29–0.78. We see that for all cases with sufficient data, there appears to be a consistent correlation between strain and changes in stability.

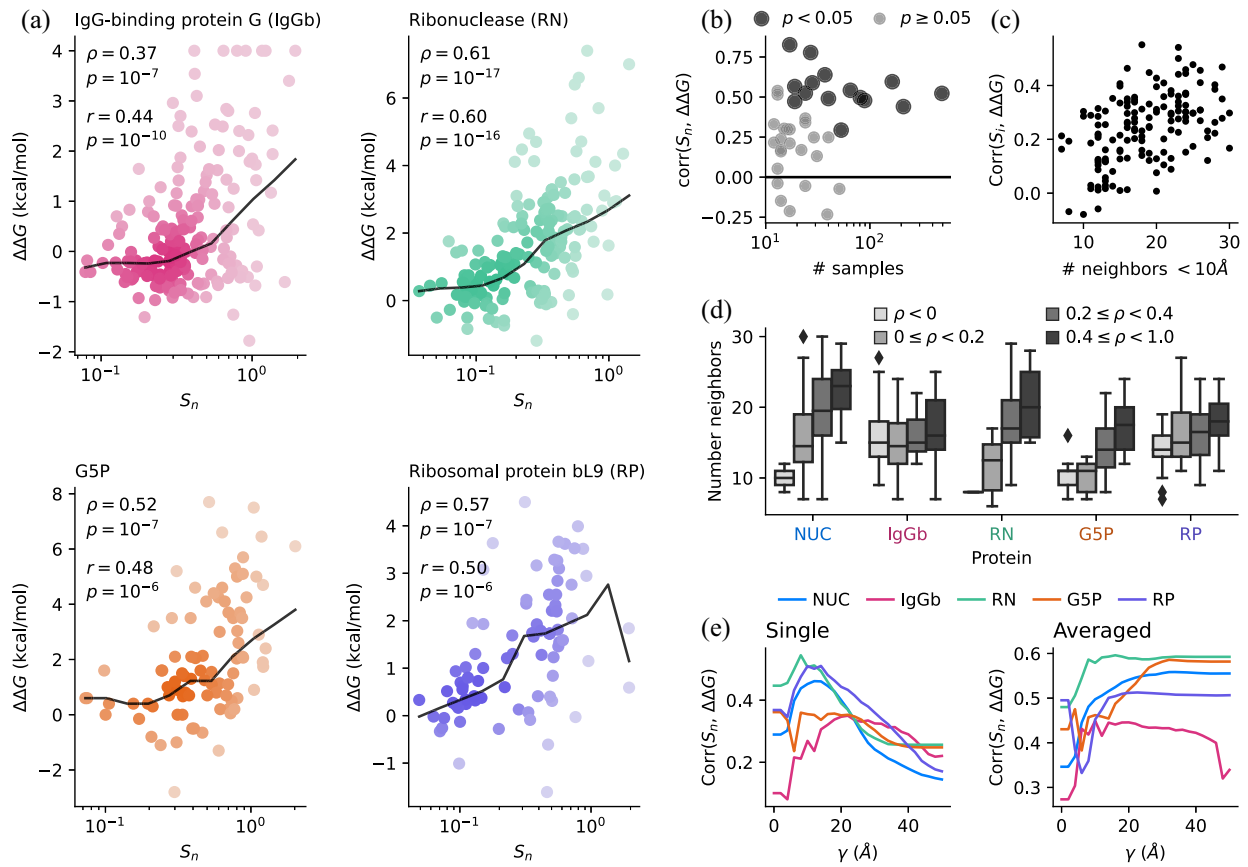


FIG. 2. When and why does  $S_n$  correlate with  $\Delta\Delta G$ ? (a) Change in stability  $\Delta\Delta G$  against strain near mutated residue  $S_n$  for four proteins (the second to fifth most common in TMDB; the most common one, NUC, is shown in Fig. 1); Spearman’s  $\rho$ , Pearson’s  $r$ , and corresponding  $p$  values are shown; black line is the median; circles are shaded by density. (b)  $S_n - \Delta\Delta G$  correlation for each of the 40 most common proteins in TMDB as a function of the number of samples; results with  $p < 0.05$  are shown by dark circles. (c)  $S_i - \Delta\Delta G$  (Pearson’s) correlation for each residue  $i$  as a function of the number of neighbors (for NUC only). (d) Distributions of numbers of neighbors, grouped by  $S_i - \Delta\Delta G$  Spearman’s correlation, for the five most common proteins. (e)  $S_n - \Delta\Delta G$  (Pearson’s) correlation for each residue  $i$  as a function of the neighborhood threshold value  $\gamma$  for calculating  $S_n$ , for both pairs of single structures and averaged structures, for five proteins.

*Determinants of strain-stability correlation*—To better understand why strain is correlated with changes in stability, we examine the correlations between strain at individual residues  $S_i$  (not necessarily the mutated ones) and  $\Delta\Delta G$ , and compare this with the number of neighbors within  $10\text{\AA}$  of each residue  $i$ . For NUC, we find that mutations in buried regions (those with many neighbors) tend to have an outside impact on stability [Fig. 2(c)], as expected, given the standard paradigm of buried residues having low mutation rates [36]. In general, one expects mutations of buried residues to affect more bonds and therefore inflict larger stability changes. Indeed, across the five most common proteins in the TMDB, we see a clear trend whereby the residues with the highest strain-stability correlations are amongst the most buried within that protein [Fig. 2(d)].

Evidently, when only a few points are sampled, the resulting correlation is not particularly informative. Also, in a large protein, more points are needed to achieve sufficient

sampling. This is demonstrated most clearly in IgG-binding protein *G* (IgGb, P06654): this protein has a length  $L = 448$ , so even though we have 225 samples, many buried residues do not correlate with stability changes [Fig. 2(d)]. This is because many regions have no mutations sampled from them, so deformation remains low no matter how buried the residues are. Another complication arises from the abundance of disordered regions (Supplemental Material, Fig. 3 [29]) in which AF appears to have little capacity to predict mutation effects, while well-folded regions are small. As a result, for IgGB the link between number of neighbors and effect on stability is weak. This case highlights the need for a nuanced approach to understanding the relationship between strain and changes in stability.

*Range of informative residues*—We find that the  $S_n - \Delta\Delta G$  correlation increases with  $\gamma$ , the radius of neighborhood used to calculate  $S_n$ , up to about  $10 \leq \gamma \leq 20\text{\AA}$ , depending on the protein family [Fig. 2(e)]. This optimal

length scale of  $\sim 15$  Å, could be the outcome of two possible effects:

First, we note that the residues that are most informative of stability changes are the most buried ones [Fig. 2(c)], and not necessarily the mutated residues. Hence, the optimal  $\gamma$  needs to be large enough to include some of these informative, buried residues in the  $S_n$  calculation. In TMDB, for example (Supplemental Material Fig. 4 [29]), the average distance between buried and mutated residues is 10 Å. This length scale can be rationalized by a simple geometric argument: the average distance to the center of a sphere (the most buried part) of diameter  $d$  is  $r_0 = (3/8)d$ ; the average volume of an amino acid is  $\sim 144$  Å<sup>3</sup> [37], so in a folded, approximately spherical domain of  $\sim 100$  amino acids,  $r_0 \sim 11.3$  Å.

Second, when we compare pairs of single AF-predicted structures, which exhibit significant noise unrelated to mutations (Supplemental Material Fig. 1 [29]), strain-stability correlations reach a maximum as a function of  $\gamma$  and then decrease; this decrease is not observed for averaged structures [Fig. 2(e)]. This is consistent with reports that AF strain predictions are usually indistinguishable from noise after about 15 Å, while this range could be extended to 25 Å (and noise reduced by a factor of 4) by averaging over structures [14,22].

In summary, we propose that the reason for the maximum information from strain around a range of 15 Å is due to both the need to increase  $\gamma$  to include buried residues and also because long-range mutation effects may be masked by prediction noise. In this context, we note that conformational dynamics of proteins were reported to exhibit correlations with similar and even longer ranges, as demonstrated by normal mode analysis [22,38,39]. However, one needs to remember that displacement and strain—the tensorial derivative of the displacement—may show different correlation lengths [22].

*Strain correlates with  $\Delta\Delta G$  across protein families*—We find a moderate correlation ( $\rho = 0.36$ ) between  $S_n$  and  $\Delta\Delta G$  when comparing all the measurements together [Fig. 3(a)]. This correlation is expected to be limited due to the inability of strain to differentiate between stabilizing and destabilizing mutations. Since Eq. (B1) measures the absolute relative change in distances,  $S \geq 0$  by definition, and thus the strain is invariant to reversing the reference and target structures. Indeed, we find a somewhat higher correlation ( $\rho = 0.41$ ) between  $S_n$  and the magnitude of stability changes,  $|\Delta\Delta G|$  [Fig. 3(b)]. This indicates that more generally, large structural changes lead to large stability changes, independent of the sign of the change.

We do not expect a simple mapping between  $S_n$  and  $\Delta\Delta G$ , given the complexity of protein structures and intramolecular interactions. We obviate protein size effects to an extent by only looking at mutation effects within  $\gamma = 15$  Å, but there are other protein-specific factors—such as the degree of disorder, protein shape, flexibility, and




FIG. 3.  $S_n - \Delta\Delta G$  correlations across protein's. Strain near mutated site  $S_n$  vs  $\Delta\Delta G$  (a) and magnitude of stability change  $|\Delta\Delta G|$  (b) for all mutations in our reduced TMDB sample of 2499 unique mutants. Pearson's  $r$ , and corresponding  $p$  values are shown.

amino acid packing—that may alter the relationship between strain and  $\Delta\Delta G$  for different proteins. Nonetheless, the strain-stability correlations shown here indicate that strain due to mutation contains considerable information about stability changes that may be leveraged in subsequent development of predictors of  $\Delta\Delta G$ .

*Strain correlates with  $\Delta\Delta G$  almost as well as tailored  $\Delta\Delta G$  predictors*—To put the strain-stability correlations in context, we compare them with two state-of-the-art  $\Delta\Delta G$  predictors, DDMut and FoldX (Appendix C) [40,41]. FoldX predicts  $\Delta G$  from structure using empirical energy-based potentials; it is used to calculate  $\Delta\Delta G$  by first a generating structure for the mutant based on a reference WT structure, which enables calculation of  $\Delta G$  for both WT and mutant structures. DDMut uses a neural network to predict  $\Delta\Delta G$ , using a reference structure and a mutation as input.

We find that  $S_n - |\Delta\Delta G|$  correlations are almost as high as correlations obtained using FoldX on average, but lags behind the more recent algorithm, DDMut (Fig. 4). We were genuinely surprised by this performance since strain is a *simple* general measure of deformation and not designed specifically for stability, in contrast to FoldX and DDMut. However, there is clearly room for improvement: simply by counting the number of neighbors at the mutation site we see a correlation of  $\rho = 0.30$  [Fig. 4(d)]; the mean absolute error obtained from a linear fit is 1.2 kcal/mol for both  $S_n$  and DDMut (Supplemental Material Fig. 5 [29]). We note that our aim here is not to use strain to predict  $\Delta\Delta G$ , but rather to see whether the strain predicted by AF is informative about stability changes. But given the surprisingly high correlation, we speculate that AF-predicted structures can be leveraged to produce even better  $\Delta\Delta G$  predictors.

We expected that since FoldX calculates  $\Delta G$  for a structure, it would give us a more accurate estimate of  $\Delta\Delta G$  than strain if we use it on AF-predicted structures. Surprisingly, we find that this method (AF-FoldX)

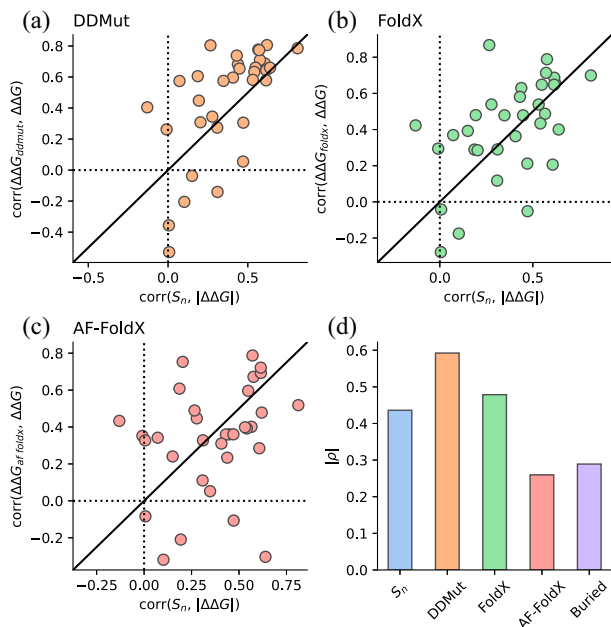


FIG. 4. Comparison with state-of-the-art algorithms. (a)–(c) Correlation (Spearman’s  $\rho$ ; linear correlations are shown in Supplemental Material, Fig. 8 [29]) between strain  $S_n$  and magnitude of stability change  $|\Delta\Delta G|$ , vs correlation between  $\Delta\Delta G$  and predictions of DDMut (a), FoldX (b), and FoldX using AF-predicted structures (c); separate points are shown for each of the 40 most-common proteins. (d) Correlations (as above) for the full sample of 2499 unique mutants. Correlation between number of neighbors within 10 Å of the mutated site is shown as “buried.”.

performs worse than strain in this case ( $\rho = 0.26$ ). Likewise, we find that calculating strain using the structures generated by FoldX results in a lower correlation with  $\Delta\Delta G$  ( $\rho = 0.30$ , Supplemental Material Fig. 6 [29]). We considered also the possibility that AF-predicted structures are not as accurate as FoldX-generated structures. However, we find that strain calculated using FoldX structures is not correlated with distance from the mutated site (Supplemental Material Fig. 7 [29]), indicating that FoldX is less accurate than AF in predicting the effect of mutations on structure. These results suggest that there is a promising path for generation of new energy-based methods for  $\Delta\Delta G$  prediction using AF-predicted structures.

*Can AF be used to predict stability changes?*—We emphasize that our aim here is not to outright develop a  $\Delta\Delta G$  predictor but rather to investigate whether AF predictions are informative of stability changes. We have found that a general measure of deformation, strain, correlates quite well with  $\Delta\Delta G$ . Although it was not designed to be a  $\Delta\Delta G$ -predictor,  $S_n$  appears to be almost as good at predicting the magnitude of stability changes as state-of-the-art  $\Delta\Delta G$  predictors. Of course, this needs to be tested on a larger set of measurements, and more structures are needed for experimental validation of the relationship

between strain and stability (Supplemental Material, Sec. 1C [29]). We also note that higher correlations have been observed for these predictors on different datasets, so this analysis should be repeated on larger sets of  $\Delta\Delta G$  measurements [40,42]. Yet, within these limitations, it seems clear that there is sufficient information in AF-predicted structures to make  $\Delta\Delta G$  predictions. It stands to reason that new physics-based models can be developed to achieve even better predictions, off the back of AF-predicted structures. Alternatively, these structures could be used to reparametrize existing force fields such as FoldX. The recent explosion of high-throughput measurements [28,43] will certainly lead to more machine-learning and sequence-based approaches. Nevertheless, we feel that physics-based methods are essential to offer a detailed view into the mechanistic effects of mutations on stability. The breakthrough by AF in structure prediction might offer the key to this future.

*Strain-energy relation*—Our predictions allow us to deduce an effective strain-energy relation by modeling the protein as an elastic spring network (as detailed in Supplemental Material, Sec. 2 [29]) [23,25]. Within this framework, mutation can be seen as introducing a point defect that perturbs the network [26], thereby inducing energy change  $\Delta\Delta G$ , which includes contributions from entropy, changing topology (breaking and forming bonds), and energy release in prestressed frustrated bonds [44]. The energy change is quadratic in the strain, with average *effective* spring constants  $K$  ranging between 1–10 kcal/mol/Å<sup>2</sup> (Supplemental Material, Fig. 9 [29]). These typical values of  $K$  [27] estimate the average curvature of the high-dimensional energy landscape [45].  $K$  varies widely in and between protein families, reflecting the anisotropic geometry and heterogeneity of the landscapes.

*Advice for using AF to predict mutation effects*—A previous study examined the same ThermoMutDB dataset, yet they concluded that AF cannot be used to predict stability. This is due to using changes in pLDDT to measure mutation effects, which does not appear to be reliable for this purpose (Supplemental Material, Fig. 10 [29]) [14]. We recommend using strain as a more robust measure of the effect of mutations on structure, particularly when using averaged structures. We recommend using about 10 to 20 (i.e., 5 models, 2–4 repeats) structures to get averages, as there are diminishing returns on performance gains (Supplemental Material, Fig. 11 [29]). Finally, we tested the strain predicted by the recently released AlphaFold3 [46] in three TMDB mutants by comparing with AF2 predictions (Supplemental Material, Fig. 12–13 [29]), finding no major differences, except in loops and disordered regions.

We have studied the correspondence between AlphaFold (AF) predictions of mutation effects on structure (measured using strain) and changes in stability. We find that strain correlates well with  $\Delta\Delta G$ , almost as well as state-of-the-art

$\Delta\Delta G$  predictors. Altogether, our findings suggest that new algorithms can be developed to extract more information from AF structures to produce accurate physics-based models of stability change upon mutation. Furthermore, quantitative mapping of the energy landscape may pave the way for more realistic modeling of functional conformational dynamics of protein machines, beyond coarse-grained elastic-network models [22,47].

*Acknowledgments*—This work was supported by the Institute for Basic Science, Project Code IBS-R020-D1.

*Data availability*—Code used to filter the ThermoMutDB can be found at [48]. We used the PDAnalysis PYTHON package to calculate strain. All AF2-predicted PDB structures are available at [49], compressed using FoldComp [50].

- 
- [1] J. M. Thornton, R. A. Laskowski, and N. Borkakoti, AlphaFold heralds a data-driven revolution in biology and medicine, *Nat. Med.* **27**, 1666 (2021).
- [2] S. Subramaniam and G. J. Kleywegt, A paradigm shift in structural biology, *Nat. Methods* **19**, 20 (2022).
- [3] J. Jumper *et al.*, Highly accurate protein structure prediction with AlphaFold, *Nature (London)* **596**, 583 (2021).
- [4] M. Baek *et al.*, Accurate prediction of protein structures and interactions using a three-track neural network, *Science* **373**, 871 (2021).
- [5] J. P. Roney and S. Ovchinnikov, State-of-the-art estimation of protein model accuracy using AlphaFold, *Phys. Rev. Lett.* **129**, 238101 (2022).
- [6] O. Herzberg and J. Moult, More than just pattern recognition: Prediction of uncommon protein structure features by ai methods, *Proc. Natl. Acad. Sci. U.S.A.* **120**, e2221745120 (2023).
- [7] A. Goldenzweig and S. J. Fleishman, Principles of protein stability and their application in computational design, *Annu. Rev. Biochem.* **87**, 105 (2018).
- [8] B. Kuhlman and P. Bradley, Advances in protein structure prediction and design, *Nat. Rev. Mol. Cell Biol.* **20**, 681 (2019).
- [9] D. Sgarbossa, U. Lupo, and A.-F. Bitbol, Generative power of a protein Language model trained on multiple sequence alignments, *eLife* **12**, e79854 (2023).
- [10] C. Malbranke, D. Bikard, S. Cocco, R. Monasson, and J. Tubiana, Machine learning for evolutionary-based and physics-inspired protein design: Current and future synergies, *Curr. Opin. Struct. Biol.* **80**, 102571 (2023).
- [11] N. Tokuriki and D. S. Tawfik, Stability effects of mutations and protein evolvability, *Curr. Opin. Struct. Biol.* **19**, 596 (2009).
- [12] N. Ota, R. Kurahashi, S. Sano, and K. Takano, The direction of protein evolution is destined by the stability, *Biochimie* **150**, 100 (2018).
- [13] E. Mauri, S. Cocco, and R. Monasson, Mutational paths with sequence-based models of proteins: From sampling to mean-field characterization, *Phys. Rev. Lett.* **130**, 158402 (2023).
- [14] J. M. McBride, K. Poley, A. Abdirasulov, V. Reinharz, B. A. Grzybowski, and T. Tlusty, AlphaFold2 can predict single-mutation effects, *Phys. Rev. Lett.* **131**, 218401 (2023).
- [15] G. R. Buel and K. J. Walters, Can AlphaFold2 predict the impact of missense mutations on structure?, *Nat. Struct. Mol. Biol.* **29**, 1 (2022).
- [16] M. A. Pak, K. A. Markhieva, M. S. Novikova, D. S. Petrov, I. S. Vorobyev, E. S. Maksimova, F. A. Kondrashov, and D. N. Ivankov, Using AlphaFold to predict the impact of single mutations on protein stability and function, *PLoS One* **18**, 1 (2023).
- [17] M. Akdel *et al.*, A structural biology community assessment of AlphaFold2 applications, *Nat. Struct. Mol. Biol.* **29**, 1056 (2022).
- [18] S. Mansoor, M. Baek, D. Juergens, J. L. Watson, and D. Baker, Accurate mutation effect prediction using rosetta-fold, *bioRxiv*, 10.1101/2022.11.04.515218 (2022).
- [19] S. Iqbal, F. Ge, F. Li, T. Akutsu, Y. Zheng, R. B. Gasser, D.-J. Yu, G. I. Webb, and J. Song, Prost: AlphaFold2-aware sequence-based predictor to estimate protein stability changes upon missense mutations, *J. Chem. Inf. Model.* **62**, 4270 (2022).
- [20] F. Peccati, S. Alunno-Rufini, and G. Jiménez-Osés, Accurate prediction of enzyme thermostabilization with rosetta using AlphaFold ensembles, *J. Chem. Inf. Model.* **63**, 898 (2023).
- [21] J. S. Xavier, T.-B. Nguyen, M. Karmarkar, S. Portelli, P. M. Rezende, J. P. L. Velloso, D. B. Ascher, and D. E. V. Pires, ThermoMutDB: A thermodynamic database for missense mutations, *Nucleic Acids Res.* **49**, D475 (2020).
- [22] J. M. McBride and T. Tlusty, The physical logic of protein machines, *J. Stat. Mech.* (2024) 024001.
- [23] J. P. Eckmann, J. Rougemont, and T. Tlusty, Colloquium: Proteins: The physics of amorphous evolving matter, *Rev. Mod. Phys.* **91**, 031001 (2019).
- [24] M. R. Mitchell, T. Tlusty, and S. Leibler, Strain analysis of protein structures and low dimensionality of mechanical allosteric couplings, *Phys. Natl. Acad. Sci. USA* **113**, E5847 (2016).
- [25] T. Tlusty, A. Libchaber, and J.-P. Eckmann, Physical model of the genotype-to-phenotype map of proteins, *Phys. Rev. X* **7**, 021037 (2017).
- [26] S. Dutta, J.-P. Eckmann, A. Libchaber, and T. Tlusty, Green function of correlated genes in a minimal mechanical model of protein evolution, *Phys. Natl. Acad. Sci. U.S.A.* **115**, E4559 (2018).
- [27] J. M. McBride, J.-P. Eckmann, and T. Tlusty, General theory of specific binding: Insights from a genetic-mechanochemical protein model, *Mol. Biol. Evol.* **39**, msac217 (2022).
- [28] K. Tsuboyama, J. Dauparas, J. Chen, E. Laine, Y. Mohseni Behbahani, J. J. Weinstein, N. M. Mangan, S. Ovchinnikov, and G. J. Rocklin, Mega-scale experimental analysis of protein folding stability in biology and design, *Nature (London)* **620**, 434 (2023).
- [29] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.133.098401> for detailed information about methods used in this Letter, which includes Refs. [30–35].

- [30] L. D. Landau, E. M. Lifšic, E. M. Lifshitz, A. M. Kosevich, and L. P. Pitaevskii, *Theory of Elasticity: Volume 7* (Elsevier, New York, 1986).
- [31] J. S. Merkel, J. M. Sturtevant, and L. Regan, Sidechain interactions in parallel beta sheets: The energetics of cross-strand pairings, *Structure* **7**, 1333 (1999).
- [32] D. G. Isom, B. R. Cannon, C. A. Castañeda, A. Robinson, and E. Bertrand García-Moreno, High tolerance for ionizable residues in the hydrophobic interior of proteins, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 17784 (2008).
- [33] D. G. Isom, C. A. Castañeda, B. R. Cannon, P. D. Velu, and E. Bertrand García-Moreno, Charges in the hydrophobic interior of proteins, *Proc. Natl. Acad. Sci. U.S.A.* **107**, 16096 (2010).
- [34] A. C. Robinson, J. L. Schlessman, and B. García-Moreno E, Dielectric properties of a protein probed by reversal of a buried ion pair, *J. Phys. Chem. B* **122**, 2516 (2018).
- [35] D. Shortle, W. E. Stites, and A. K. Meeker, Contributions of the large hydrophobic amino acids to the stability of staphylococcal nuclease, *Biochemistry* **29**, 8033 (1990).
- [36] J. Echave, S. J. Spielman, and C. O. Wilke, Causes of evolutionary rate variation among protein sites, *Nat. Rev. Genet.* **17**, 109 (2016).
- [37] S. J. Perkins, Protein volumes and hydration effects, *Eur. J. Biochem.* **157**, 169 (1986).
- [38] Q.-Y. Tang and K. Kaneko, Long-range correlation in protein dynamics: Confirmation by structural data and normal mode analysis, *PLoS Comput. Biol.* **16**, 1 (2020).
- [39] Q.-Y. Tang and K. Kaneko, Dynamics-evolution correspondence in protein structures, *Phys. Rev. Lett.* **127**, 098103 (2021).
- [40] Y. Zhou, Q. Pan, D. E. V. Pires, C. M. Rodrigues, and D. Ascher, DDMut: Predicting effects of mutations on protein stability using deep learning, *Nucleic Acids Res.* **51**, W122 (2023).
- [41] J. Schymkowitz, J. Borg, F. Stricher, R. Nys, F. Rousseau, and L. Serrano, The FoldX web server: An online force field, *Nucleic Acids Res.* **33**, W382 (2005).
- [42] A. Broom, K. Trainor, Z. Jacobi, and E. M. Meiering, Computational modeling of protein stability: Quantitative analysis reveals solutions to pervasive problems, *Structure* **28**, 717 (2020).
- [43] A. S. Dunham and P. Beltrao, Exploring amino acid functions in a deep mutational landscape, *Mol. Syst. Biol.* **17**, e10305 (2021).
- [44] D. U. Ferreira, E. A. Komives, and P. G. Wolynes, Frustration in biomolecules, *Q. Rev. Biophys.* **47**, 285 (2014).
- [45] J.-P. Eckmann and T. Tlusty, Dimensional reduction in complex living systems: Where, why, and how, *BioEssays* **43**, 2100062 (2021).
- [46] J. Abramson *et al.*, Accurate structure prediction of biomolecular interactions with AlphaFold 3, *Nature (London)* **630**, 493 (2024).
- [47] H. Flechsig and A. S. Mikhailov, Simple mechanics of protein machines, *J. R. Soc. Interface* **16**, 20190244 (2019).
- [48] J. M. McBride, [github.com/jomimc/AF2\\_Stability\\_PRL\\_2024](https://github.com/jomimc/AF2_Stability_PRL_2024) (2024).
- [49] J. M. McBride, Protein structure data for AI-predicted protein deformation encodes energy landscape perturbation (2024), .
- [50] H. Kim, M. Mirdita, and M. Steinegger, Foldcomp: A library and format for compressing and indexing large protein structure sets, *Bioinformatics* **39**, btad153 (2023).
- [51] H. M. Berman, J. Westbrook, Z. Feng, G. Gilliland, T. N. Bhat, H. Weissig, I. N. Shindyalov, and P. E. Bourne, The protein data bank, *Nucleic Acids Res.* **28**, 235 (2000).
- [52] T. U. Consortium, Uniprot: A worldwide hub of protein knowledge, *Nucleic Acids Res.* **47**, D506 (2018).
- [53] M. Mirdita, K. Schütze, Y. Moriwaki, L. Heo, S. Ovchinnikov, and M. Steinegger, Colabfold: Making protein folding accessible to all, *Nat. Methods* **19**, 679 (2022).
- [54] J. Delgado, L. G. Radusky, D. Cianferoni, and L. Serrano, FoldX 5.0: Working with RNA, small molecules and a new graphical interface, *Bioinformatics* **35**, 4168 (2019).
- [55] J. H. Carra and P. L. Privalov, Energetics of denaturation and  $m$  values of staphylococcal nuclease mutants, *Biochemistry* **34**, 2034 (1995).

## End Matter

*Appendix A: Stability data*—We select all proteins from ThermoMutDB (TMDB) [21] with sequence length,  $50 \leq L \leq 500$ , and measurements for single mutants made within  $293 \leq T \leq 313$  K and  $5 \leq \text{pH} \leq 8$ ; this amounts to 5078 out of 13 337 measurements. One problem we found with TMDB is that the indices in mutation codes can be associated with either Protein Data Bank (PDB) [51], Uniprot [52], or “unsigned,” yet we wanted to match mutations to Uniprot sequences. Hence, if indices were matched to PDB indices that differed from Uniprot indices, we used a custom script to convert the mutation codes to the correct Uniprot indices. This script only works when the mutation code index refers to the “\_atom\_site.label\_seq\_id” entry of the mmCIF file; we

excluded many cases where the TMDB mutation indices refer to idiosyncratic indexing (i.e., not starting from one) in the “\_atom\_site.auth\_seq\_id” entries in the mmCIF files. Out of caution, if no correct matches were found for a protein we excluded all TMDB entries that were related to this protein (i.e., if they have the same value in the “PDB\_wild” column). We excluded “unsigned,” since these were ambiguous, and we discovered that some of these were labeled with incorrect Uniprot accession ids. After this procedure we are left with 3236 measurements. When multiple  $\Delta\Delta G$  measurements are available we report the average  $\Delta\Delta G$  across all measurements. We leave out a mutant if the standard deviation of the  $\Delta\Delta G$  measurements is higher than 1 kT; this occurred in only

about 2% of cases. We examined why some of these cases had such high variance, and found occasional errors where the  $\Delta\Delta G$  sign differed from the reference it was taken from. A full list of errors and corrections can be found in Supplemental Material, Sec. 1D [29]. The final set of  $\Delta\Delta G$  measurements includes 2499 unique mutants. We study correlations within individual wild-type (WT) proteins (and their single mutants), and across the full set of measurements.

*Appendix B: Structure analysis*—We use the ColabFold implementation of AF to predict protein structures [53]. We run all 5 models, without using templates, and using 6 recycles, and minimization using the Amber forcefield. We run 10 repeat predictions for each sequence, for each model, and create averaged structures following [14]. We calculate effective strain, ES, which measures the deformation between mutants and WT proteins, for both averaged and nonaveraged structures. ES can be described simply as the average relative change of  $C_\alpha - C_\alpha$  distances between neighboring residues; we define neighbors as residues whose  $C_\alpha$  positions are within 13 Å. We calculate ES as [14],

$$S_i = \left\langle \frac{|\Delta \mathbf{r}_{ij}|}{|\mathbf{r}_{ij}|} \right\rangle = \frac{1}{n_i} \sum_{j \in N_i} \frac{|\mathbf{r}_{ij} - \mathbf{r}'_{ij}|}{|\mathbf{r}_{ij}|}, \quad (\text{B1})$$

where  $\mathbf{r}_{ij}$  is the distance vector between  $C_\alpha$  positions in residue  $i$  and neighbor  $j$  in a reference structure,  $\mathbf{r}'_{ij}$  is the corresponding distance vector in a target structure which has been aligned to  $\mathbf{r}_{ij}$ ,  $N_i$  is the set of neighbors, and  $n_i = |N_i|$  is the number of neighbors. We refer to ES measured at the mutated site  $m$  as  $S_m$ . We also calculate the neighborhood sum of strain  $S_n$  over residues whose distance to the mutation site  $\delta_i$ , is shorter than some threshold  $\gamma$ ,

$$S_n = \sum_{\delta_i \leq \gamma} S_i. \quad (\text{B2})$$

We mainly use  $\gamma = 15$  Å. We typically only include AF-predicted residues in strain calculations if pLDDT > 70, and treat the rest as disordered, except where otherwise noted.

*Appendix C:  $\Delta\Delta G$  predictors*—We use two state-of-the-art methods to predict  $\Delta\Delta G$ : FoldX 5 [54], and DDMut [40]. We use the API for the DDMut web server to predict  $\Delta\Delta G$ . For FoldX, we use the BuildModel command to generate structures of the WT and mutant sequences and  $\Delta\Delta G$  predictions. We use five runs as recommended, and report average values of  $\Delta\Delta G$ . For each algorithm, we provide the top-ranked (by pLDDT) AF-predicted WT structure as an input, along with a list of all mutants for that protein.