

# Proteome-wide prediction of the functional impact of missense variants with ProteoCast

Received: 14 September 2025

Accepted: 3 April 2026

Published online: 27 April 2026

 Check for updates

Marina Abakarova <sup>1,2</sup>, Maria Inés Freiberger<sup>1</sup>, Arnaud Liehrmann<sup>1</sup>,  
Michael Rera <sup>2</sup>  & Elodie Laine <sup>1,3</sup> 

Dissecting the functional impact of genetic mutations is essential to advancing our understanding of genotype-phenotype relationships and identifying therapeutic targets. Despite progress in sequencing and genome editing technologies, proteome-wide mutation effect prediction remains challenging. Here we show that evolutionary information alone enables accurate prediction of mutation effects across entire proteomes. ProteoCast is a scalable and interpretable computational method that leverages protein sequence conservation to classify genetic variants and identify functionally important protein sites. We apply ProteoCast to the complete *Drosophila melanogaster* proteome (22,000 isoforms, 300 million mutations) and validate it against nearly 400,000 natural and experimental variants. It correctly classifies 85% of known lethal mutations as functionally impactful versus 13-18% of population variants. ProteoCast-guided genome editing experiments confirm these predictions. Moreover, ProteoCast successfully identifies functionally important protein modification sites and binding motifs. ProteoCast provides a publicly available resource and deployable pipeline for studying gene function and mutations in any organism.

More than 50,000 single-nucleotide mutations are associated with genetic disorders in humans<sup>1</sup>, and correcting them at the genomic level could provide curative treatments for various diseases<sup>2</sup>. The advent of CRISPR (Clustered Regularly Interspaced Short Palindromic Repeats) technology has revolutionised genetics by enabling precise and efficient DNA editing as well as establishing targeted mutational screens as a powerful forward genetic approach to identify novel druggable targets<sup>3</sup>. When combined with high-throughput sequencing, which enables the analysis of inter-individual and inter-tissue genomic variability, these single-point genome editing tools allow unprecedented exploration of the functional significance of genetic variants. They provide effective means for assessing their contributions to disease susceptibility, drug responses and, more generally, to complex traits.

However, the vast number of possible substitutions makes it challenging to select the optimal amino acid target(s) for genome

editing. To prioritise single-point genetic modifications for CRISPR experiments, researchers often use conservation scores derived from multiple sequence alignments (MSAs), whole-exome sequencing to identify patient-specific mutations, or statistical tools like Genome-Wide Association Studies (GWAS) to link variants to specific diseases or phenotypes. While invaluable, these methods provide only a limited view of impactful mutations across the entire proteome. Computational predictive models can address these limitations, offering a more comprehensive approach to understanding the functional impact of mutations.

The last decades have witnessed a profusion of protein variant effect predictors. Modern methods generate a mutational landscape - namely, the predicted effects of all possible substitutions at every position within a query protein by leveraging protein sequence data across species. They vary widely in complexity and number of parameters, encompassing explicit evolutionary models<sup>4-7</sup>, family-specific

<sup>1</sup>Sorbonne Université, CNRS, IBPS, Department of Computational, Quantitative and Synthetic Biology (CQSB UMR7238), Paris, France. <sup>2</sup>Université Paris Cité, CNRS, Functional and Adaptive Biology (UMR8251), Paris, France. <sup>3</sup>Institut Universitaire de France (IUF), Paris, France. ✉ e-mail: [michael.rera@cnrs.fr](mailto:michael.rera@cnrs.fr); [elodie.laine@sorbonne-universite.fr](mailto:elodie.laine@sorbonne-universite.fr)

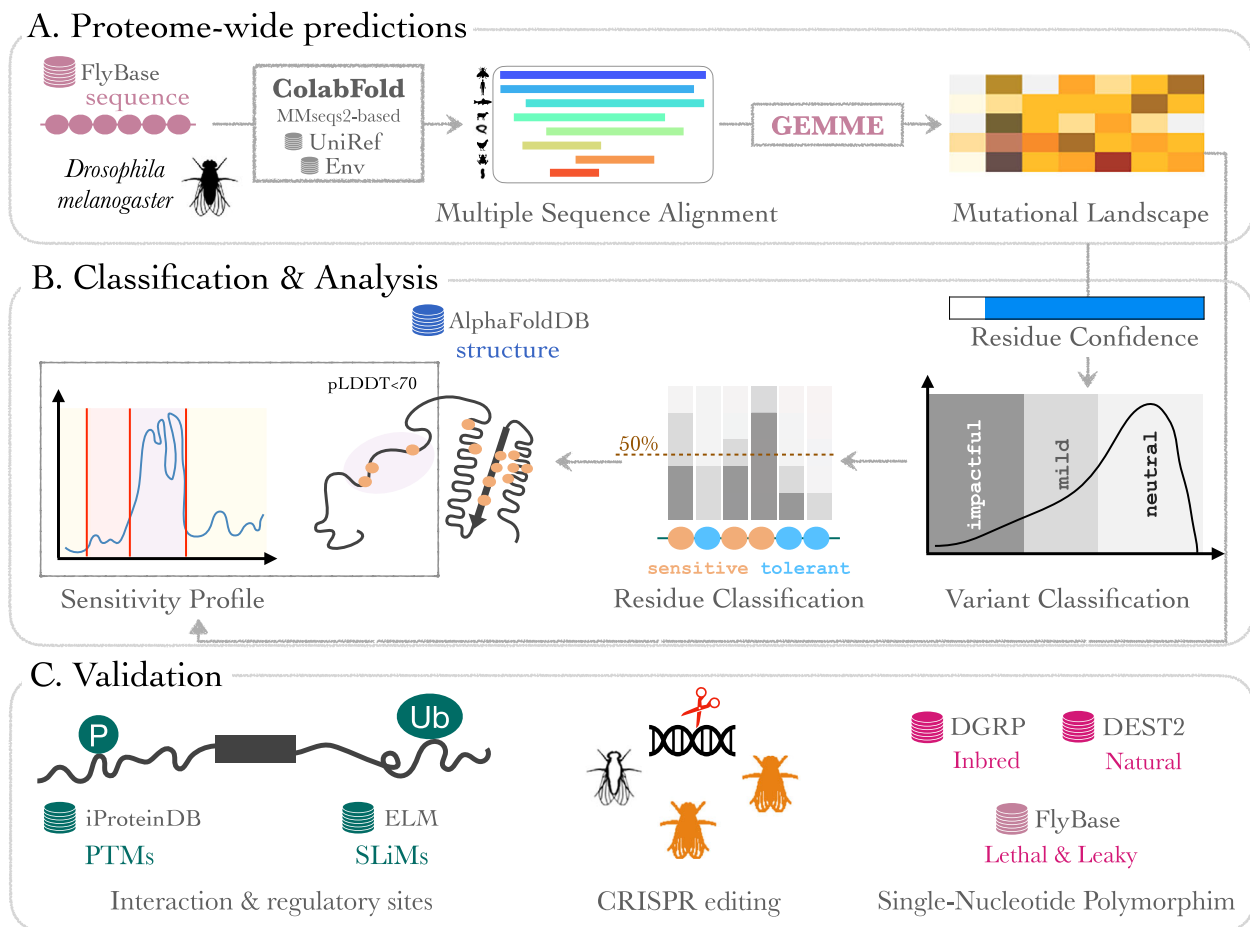
deep learning models<sup>8,9</sup>, and universal protein language models (pLMs)<sup>10–15</sup>. A couple of methods focusing on human applications are equipped with variant classification and uncertainty estimation to facilitate clinical interpretation. Specifically, the EVE method interprets family-specific variational autoencoder score distributions as a mixture of two Gaussians, one representing pathogenic variants and the other benign variants. This approach assumes a direct correspondence between the underlying data distribution and biological variant classes.

Additionally, some predictors are trained on variants with known phenotypes or polymorphisms observed in the population<sup>12,16</sup>. These population-based predictors, however, do not typically perform better than population-free methods and tend to suffer from circularity issues and pervasive ancestry biases<sup>17</sup>. Furthermore, a new generation of predictors<sup>12,18–21</sup> systematically integrates high-quality 3D structural models<sup>22</sup>. However, they have shown limited improvement over purely sequence-based methods<sup>23</sup>, and the structural contribution is sometimes unclear.

Systematic benchmarking studies<sup>23–25</sup> have demonstrated that predictive performance against deep mutational scans is a good indicator of a variant effect predictor's ability to identify pathogenic

mutations in humans, with modern methods outperforming traditional approaches such as SIFT<sup>4</sup> and PROVEAN<sup>26</sup> at this task. The results highlighted GEMME, a fully interpretable, unsupervised, MSA-based method, as one of the top-performing models<sup>7</sup>. By focusing on how protein residues segregate along the topology of evolutionary trees, GEMME has proven instrumental in studying protein stability, function, and disease mechanisms<sup>27–30</sup>. Notably, GEMME is 100 to 10,000 times faster than deep learning predictors such as ESM-2<sup>11</sup>, EVE<sup>8</sup>, and PoET<sup>14</sup>, and it does not require high-end GPU resources<sup>7,12</sup>. When paired with the ColabFold protocol, which provides cost-effective, small, and diverse MSAs, GEMME enables computing mutational landscapes for tens of thousands of proteins within just a few days<sup>31</sup>, making it particularly attractive for large-scale studies.

In this study, we expand on GEMME by developing ProteoCast, a scalable fully automated computational framework for facilitating proteome-wide single-point genetic modification prioritisation, and demonstrate its application in *Drosophila melanogaster*. ProteoCast seamlessly estimates the impact of missense mutations on protein function, computes MSA-based confidence metrics, and classifies variants by accounting for the regimes underlying the predicted score distribution (Fig. 1). It determines which residues are sensitive to



**Fig. 1 | Schematic representation of ProteoCast workflow.** **A** The workflow starts by generating a Multiple Sequence Alignment (MSA) for each *Drosophila melanogaster* proteoform sequence using the ColabFold MMseqs2-based pipeline against UniRef and environmental sequence databases. GEMME then produces an Lx20 matrix representing the proteoform complete single-site Mutational Landscape, with L the number of amino acid residues. **B** ProteoCast filters out unreliable predictions (Residue Confidence, based on input MSA assessment) and then classifies the remaining variants as impactful, mild or neutral (Variant Classification) by fitting a Gaussian mixture model to the raw score distribution. It then determines

whether each residue is sensitive or tolerant to mutations depending on its proportion of impactful variants (Residue Classification). ProteoCast further maps per-residue mutational sensitivity (average raw score) to 3D models from AlphaFold DB. It segments the mutational sensitivity profile to identify potential binding and regulatory sites in unstructured regions. **C** ProteoCast validation through CRISPR editing experiments and benchmarking against known post-translational modification (PTM) sites and short linear motifs (SLiMs), as well as single-nucleotide polymorphism, either observed in inbred or wild populations, or associated with deleterious outcome such as developmental lethality or partial loss of function.

mutations and maps this information on protein 3D structures. Additionally, it identifies putative binding and regulatory sites in unstructured regions by segmenting the predicted mutational landscape. We validated ProteoCast on a comprehensive benchmark dataset we compiled of over 386,000 single-point mutations from *D. melanogaster*, encompassing genetic variation under natural or artificial selection and known hypomorph and lethal mutations.

Although much genetic research focuses on humans, we selected *D. melanogaster* as our model organism due to its rich and well-documented genetic and phenotypic data accumulated over a century of research. This extensive body of knowledge offers unique insights into natural genetic variability and the functional consequences of mutations. We used three key public resources: the Drosophila Genetic Reference Panel (DGRP)<sup>32</sup>, a collection of 205 inbred lines, the community-generated population genomic resource Drosophila Evolution over Space and Time (DEST v2.0)<sup>33</sup>, a collection of 530 high-quality pooled libraries from >32,000 flies collected across six continents over more than a decade, and FlyBase<sup>34</sup>, a comprehensive database on Drosophila genetics that has aggregated the phenotypic effects of point mutations for over 30 years.

With approximately 75% of human genetic diseases being associated with genes conserved in fruit flies, the relevance of *Drosophila* as a model organism for human disease is clear<sup>35</sup>. Our findings, therefore, are advancing our understanding and ability to intervene in functional genomics, with a broad applicability to human health and diseases. ProteoCast and all generated and analysed data are freely available to the community at <https://proteocast.ijm.fr/drosophiladb/>. ProteoCast can be run on any protein sequence from any organism using the Docker image at <https://hub.docker.com/r/marinaabakarova/proteocast/> or through the webserver at <https://proteocast.ijm.fr/>.

## Results

We used ProteoCast to generate full single-site mutational landscapes over the entire *D. melanogaster* proteome, including all referenced protein isoforms (proteoforms) resulting from alternative splicing (Fig. 1A, see also Methods). We estimated the functional impact of every possible substitution at every position for 22,169 unique proteoforms, covering 99% of the Drosophila proteome (Table S1). We provide predictions of variant effects for about 293 million missense mutations enriched with detailed information to help guide users in interpreting predictions more effectively. These metadata include confidence scores, variant functional classification as neutral, mild, or impactful, as well as each residue's sensitivity to mutations. Our online database <https://proteocast.ijm.fr/drosophiladb/> allows for querying any Drosophila proteoform and interactively exploring the associated ProteoCast predictions. Below, we provide a graphical summary of the prediction pipeline (Fig. 1) and functional validation of the predicted functional impacts.

### Sufficient evolutionary information for most of the variants

To help users assess the reliability of predictions based on input data quality, ProteoCast provides confidence metrics that flag cases, proteins or residues, where insufficient evolutionary information may compromise prediction accuracy (Fig. 1B, *Residue Confidence*). At the protein level, we consider that mutational landscapes predicted from less than a couple hundred sequences may be of low quality. This cutoff is based on our previous findings<sup>31</sup> and is calibrated for the MMseqs2-based MSA generation protocol implemented in ColabFold<sup>36,37</sup> that maximises diversity while minimising the number of sequences.

At the residue level, ProteoCast uses the predicted scores' dispersion over the 19 possible substitutions as a proxy for the input alignment information content. A small dispersion indicates that ProteoCast is not able to distinguish substitutions from one another. On

the one hand, this situation typically occurs upon strong selection pressure, when a position is highly conserved in evolution and thus does not tolerate any substitution. A high conservation level implies a limited subset of amino acids observed in the input MSA and a low gap content. On the other hand, a narrow distribution of predicted scores may reflect lack of evidence from the input MSA, for instance in the case of highly gapped or highly divergent positions. In this situation, the conservation level will be low (Fig. S1). Hence, ProteoCast considers that positions with a small score dispersion, a low conservation level, and a low coverage or diversity in the input MSA lack supporting evidence and flags them as unreliable (see Methods for a more formal definition and Fig. S1, insert).

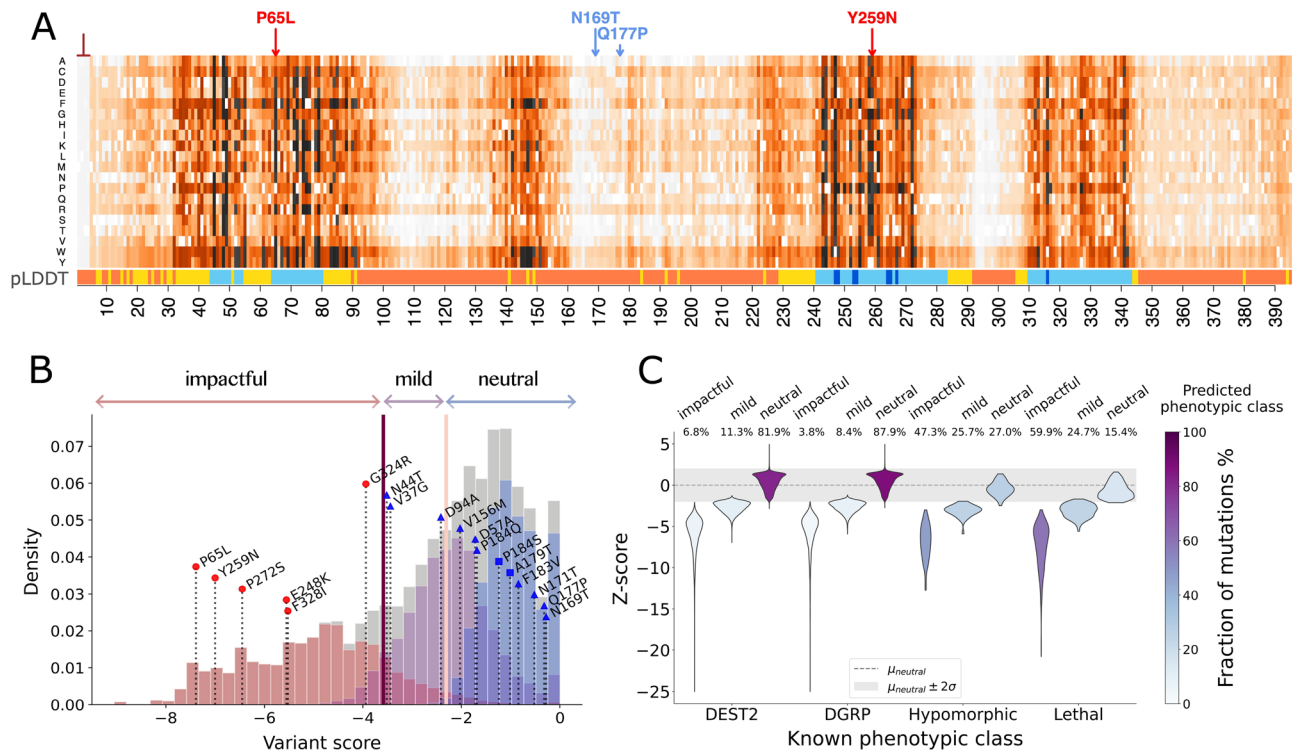
Over the whole *D. melanogaster* proteome, ProteoCast filtered out about 12% of all unique proteoforms according to the global (per-protein) criterion (Fig. S1). It further detected unreliable predictions for -460,000 residues (-3%) in the remaining proteoforms via the local (per-residue) confidence metric (Fig. S1). Overall, the predictions labelled as unreliable based on the global or local confidence metrics concerned less than 10% of all residues. Moreover, for most of the proteoforms, more than 70% of all possible amino acid substitutions (19 multiplying the protein length) are in fact observed in the input alignment, ensuring high-resolution predictions. These results underscore the quality and abundance of the evolutionary information available for the fly proteome.

### High correlation between predicted effects and organismal fitness

ProteoCast predicts a protein's complete single mutational landscape based solely on evolutionary information specific to that protein, without considering its biological function nor interactions with other proteins. To evaluate the relevance of these protein-centered predictions for organismal fitness, we created a custom benchmark of missense mutations resulting from single-nucleotide polymorphisms (Fig. 1C, *Single-Nucleotide Polymorphism*). This benchmark comprises all inbred missense DGRP polymorphisms, all natural missense DEST2 polymorphism as well as all ethyl methanesulfonate (EMS) point mutations referenced in FlyBase. We used FlyBase's experimentally supported annotations labelled as hypomorph - indicating reduced gene function—or as developmentally lethal in flies (Table S2 and Fig. S2). The benchmark is highly imbalanced, comprising approximately -137,000 DGRP and -270,000 DEST2 missense mutations spread over more than 10,000 genes, -1000 lethal mutations over almost 500 genes, and 148 hypomorph mutations (Fig. S2B).

We illustrate ProteoCast analysis with the Yorkie protein (Fig. 2A, B), a key component of the Hippo pathway involved in apoptosis and cell proliferation, among other developmental processes<sup>38</sup>. We highlight four mutations on the landscape predicted by ProteoCast in Fig. 2A: the lethal mutations P65L and Y259N rank among the top 5% most negative predictions (dark cells), clearly distinguishable from the naturally occurring DEST2 variants N169T and Q177P (light cells). We map all annotations on the predicted score distribution in Fig. 2B, where mutations inducing lethality (in red) are shifted toward lower (more negative) scores compared to population variants (in blue).

To systematically and automatically segregate variants, ProteoCast applies a thresholding strategy tailored to each protein by fitting a mixture of Gaussians to the predicted raw score distributions (Fig. 1B, *Variant Classification*, Fig. 2B). We set the number of components to three as it represented an optimal choice for maximising likelihood while avoiding overfitting and overparametrisation across 20 trials of 100 randomly selected proteoforms (Fig. S3A, B, and see Methods). We take the mean of the central component as a threshold to distinguish between *neutral* mutations (scores towards 0) and mutations altering the protein function (lower, more negative scores, Fig. 2B). The latter span a substantially wider range of scores than the former (Fig. S3C, D), suggesting varying degrees of evolutionary constraints against them.



**Fig. 2 | Variant classification using ProteoCast. A** Full single-mutational landscape computed for Yorkie protein (*yki* gene, FlyBase proteoform id: FBpp0288697), a key downstream effector in the Hippo/SWH (Sav/Wts/Hpo) signalling pathway. Each cell of the matrix contains a numerical estimate of the functional impact of an amino acid substitution (y-axis) at a given residue position (x-axis) on the protein function. Darker colours indicate stronger effects, and white cells correspond to the wild-type amino acids. The inverted T-shaped symbol at the N-terminus marks low-confidence predictions (Residue Confidence). The high-lighted mutations provoke developmental lethality—as referenced in FlyBase (P65L and Y259N, in red) or occur in DEST2 (N169T, and Q177P, in blue). The pLDDT (predicted local distance difference test) score, shown as a colour bar below the matrix, provides a confidence measure of the AlphaFold2-predicted 3D structure,

orange: very low, yellow: low, light blue: high, dark blue: very high. Residues 241–274 and 310–343 correspond to the well-characterised and well-structured WW domains. **B** ProteoCast classification of Yorkie mutations, where the pale pink and deep purple vertical lines indicate the neutral/mild and mild/impactful frontiers, respectively. All 18 known mutations in our benchmark are indicated with red dots for lethal, blue triangles for DEST2, and blue squares for DGRP. **C** Distributions of Z-scores for DGRP and DEST2 polymorphisms, as well as hypomorph and lethal mutations annotated in FlyBase, depending on whether they are classified as impactful, mild or neutral. We computed the Z-scores with respect to the distribution of estimated effects for the neutral class within each protein. Three mutations with scores below -25 were excluded from the plot for graphical representation clarity. Source data are provided as a Source Data file.

To reflect this richness in the evolutionary data, we define two subgroups of variants, namely *mild* and *impactful*, using the intersection between the left-most and central Gaussians (Fig. 2B). The three variant classes exhibit overlapping predicted score distributions with increasing means, from impactful through mild to neutral (Fig. S4).

This classification procedure exhibited high discriminative capability over our benchmark. On the one hand, ProteoCast detected 85% of the lethal mutations and 73% of the hypomorph ones as impactful or mild. On the other hand, it identified 88% of the DGRP variants and 82% of the DEST2 ones as neutral (Fig. 2C and Table S3). The difference between the two datasets may reflect the stronger selection over flies inbred in the laboratory compared to natural variants in living flies. Results obtained for individual proteins are in line with these overall statistics: ProteoCast correctly classified over 80% of the lethal, DGRP and DEST2 mutations in 78%, 82%, and 66% of the cases, respectively (Fig. S5). To further examine the internal structure of our predicted classes, we computed Z-scores measuring how each variant's raw score deviates from the neutral class distribution (Fig. 2C). Lethal and hypomorph mutations consistently exhibited more negative Z-scores than population-based variants, even within the same predicted class, revealing a continuous gradation of predicted severity that aligns with experimental annotations. Additionally, population polymorphisms with lower Minor Allele Frequency (MAF) tend to be assigned more negative raw scores and Z-scores (Fig. S6A).

By adapting to the shape of each protein's score distribution, ProteoCast variant classification strikes a better balance between detecting lethal mutations as impactful or mild (recall) and assigning neutral impact on population variants (specificity) than universal thresholds (Fig. S6B, C, Table S3, and Fig. S7). This advantage is especially visible for proteins with near-complete evolutionary sampling of substitutions (> 95% actually observed in the input MSA), where the predicted scores are shifted toward higher values (Fig. S8A, B). For instance, applying Yorkie's thresholds to the serine/threonine phosphatase Flapwing (Flw), essential for larval development, would make its two lethal mutations (Y264F and D222N, scores of -1.38 and -0.97) appear as neutral. Thanks to its adaptive strategy, ProteoCast identified these mutations as mild and correctly distinguished them from the DGRP mutation E314D, classified as neutral (Fig. S8C, D).

Taken together, these results show that variant effects inferred from the evolutionary history of individual proteins correlate well with experimentally characterised organismal fitness. The vast majority of mutations observed in viable strains, such as the DGRP lines and DEST2 pooled libraries, are not strongly selected against in evolution. Conversely, higher severity tends to be associated with higher evolutionary pressure detectable by our approach. The consistency between overall performance metrics and individual protein results further demonstrates that our approach generalises well across the entire proteome.

**Table 1 | Residues selected for CRISPR introduction of single amino acid missense mutations**

Naprt Variant	Variant class	Raw score	Rank overall <sup>a</sup>	Rank at residue	Residue class	Residue rank <sup>*</sup>	SIFT scores
G497P	Impactful	-8.75	1 ↑	1 ↑	Sensitive	1 ↑	0.00
K532G	Impactful	-8.73	2 ↑	1 ↑	Sensitive	11 ↑	0.00
E147V	Impactful	-8.33	52 ↑	1 ↑	Sensitive	21 ↑	0.00
I618P	Neutral	-1.54	3,322 ↓	1 ↑	Tolerant	5 ↓	0.21
A262F	Neutral	-0.84	1,670 ↓	3 ↑	Tolerant	2 ↓	0.05

For each selected variant, we report its class (among impactful, mild and neutral), its raw predicted score and associated overall and relative ranks. In addition, we report the mutational sensitivity class (either tolerant or sensitive) and rank of the corresponding wild-type residue. For comparison, we indicate SIFT predicted normalised probabilities for the selected variants.

<sup>\*</sup>Residues with pLDDT < 70 are excluded from the ranking. This leaves us with 533 residues and 10,127 possible variants. The arrows indicate how the values are ordered to compute the ranks. For instance, a variant ranked 1 with an up arrow has the lowest (most negative) raw score. When some values are equal, they are assigned the same (minimum) rank. Sequencing validations of the five mutants are presented in Annex 1.

The distinction between mild and impactful categories was not systematically evaluated, as available experimental annotations are predominantly qualitative and heterogeneous rather than providing standardised severity gradations. Moreover, the distinction between lethal and hypomorph mutations in our benchmark may be confounded by experimental conditions, as some hypomorph mutations have only been characterised in heterozygous states or under permissive conditions, and could actually be lethal. This complexity explains why the correspondence between experimental categories (lethal/hypomorph) and our predicted classes (impactful/mild) is not straightforward. Nevertheless, our graduated classification captures meaningful severity gradations, as illustrated by insulin receptor mutations where the mild variants R1467C and G1599R cause moderate phenotype (reduced head size) while the impactful variant G1539E produces severe effects (strongly reduced head size, putative null allele)<sup>39</sup>.

### ProteoCast effective in guiding genome editing strategies

We used ProteoCast predictions to prioritise variants for experimental genome editing, introducing single-point missense mutations in *D. melanogaster* (Fig. 1C, CRISPR editing). First, we selected nine candidate genes with high-quality input alignments, consisting of 10,000 to 30,000 sequences each. These genes were chosen based on the absence of lethal missense mutations reported in FlyBase and a broad range of predicted raw scores, including very negative values (below -8), indicating strong effects (see Methods). We first tested the developmental lethality of RNAi-induced knockdowns for these genes using transgenics under control of the ubiquitous *daughterless::GeneSwitch* inducible driver. The RNAi-mediated knockdown of Naprt (CG3714), which encodes nicotinate phosphoribosyltransferase, an enzyme involved in NAD biosynthesis, resulted in 100% developmental lethality, identifying it as an essential gene during development (Fig. S9).

These results led us to prioritise Naprt for genome editing. We selected five variants based on ProteoCast predictions computed for Naprt-PH (662 residues), the longest proteoform cross-referenced in UniProt (Table 1). We first chose the two variants identified by ProteoCast as the most impactful, namely G497P and K532G. Both are located near the putative active sites, as inferred by homology transfer (Fig. 3 and Table S4). We chose the third target residue, E147, among the 25 most sensitive to mutations, as estimated by averaging scores over the 19 possible substitutions (Table 1). The 3D structure of Naprt's human homologue suggests it could be involved in the protein's homo-dimeric interface (Fig. 3D and Fig. S10). We substituted it with the amino acid predicted to have the strongest effect, namely Valine (V).

As negative controls, we looked for residues predicted to be tolerant to mutations, with more than 10 out of the 19 possible substitutions predicted as neutral (Fig. 1B, Residue Classification). To ensure a rigorous test of ProteoCast's predictions, we excluded the N-terminal extremity (residues 1–21) and an inter-domain linker

(residues 301–407). These regions are poorly conserved and unstructured, and they comprise the vast majority (17 out of 20) of the naturally occurring missense mutations referenced in DGRP and DEST2. We focused on the five most tolerant residues in the structured protein fraction. The top-4 residues form a cluster in 3D space (Fig. 3B), from which we randomly selected a representative residue, A262.


Additionally, we retained I618, the fifth-ranked residue, located in a solvent-exposed helix. To challenge the robustness of ProteoCast's neutral predictions, we substituted A262 and I618 with amino acids that exhibit starkly contrasting physico-chemical properties (Table 1). Specifically, we mutated the small A262 into Phenylalanine (F), a large aromatic amino acid, and I618 into Proline (P), which has the tendency to break secondary structures. This design allowed us to test whether ProteoCast could correctly predict neutral effects despite substantial changes in amino acid properties.

The single-point mutation genome editing was conducted in collaboration with the company WellGenetics following the designs presented in Annex 1, in a *w<sup>1118</sup>* genetic background. The obtained results were perfectly in line with our predictions. Specifically, the introduction of either of the mutations G497P, K532G or E147V did not allow for homozygotization, while introducing I618P or A262F mutations led to viable homozygotes.

To further demonstrate the added value of ProteoCast, we compared it with the popular tool SIFT<sup>40</sup> (Fig. S11). Predictions from the two methods showed overall good agreement (Fig. S11A, C). However, SIFT assigned a score of 0.05 to the variant A262F (Table 1). While this value falls at the boundary of SIFT's deleterious threshold, both ProteoCast and the CRISPR experiment clearly indicate A262F is neutral. Moreover, ProteoCast offers much finer resolution and a higher dynamic range of the mutational effects (Fig. S11B, D). For instance, SIFT assigns the minimal score of 0 to 2382 variants, representing about a quarter of all possible variants. Among these, ProteoCast classifies 188 variants as neutral, demonstrating its ability to discriminate within what SIFT considers uniformly deleterious. This ability of ProteoCast to discriminate between substitutions across protein sites and also at each protein site constitutes a significant advantage for experimental design and clinical prioritisation.

### Discovery of interaction and regulatory sites in unstructured regions

ProteoCast maps predicted variant scores to proteoform 3D models from the AlphaFold database<sup>22,41</sup>, leveraging its extensive protein structural data. We observed that when GEMME faces difficulty making accurate predictions due to limited evolutionary information, AlphaFold2 similarly struggles to produce reliable 3D models. Specifically, 96% of the residues flagged by ProteoCast as having unreliable predictions are located within protein regions with low- or very low-confidence AlphaFold2-predicted 3D coordinates (predicted Local Distance Difference Test scores below 70, Fig. S12A). Conversely however, ProteoCast estimated that 94% of the residues with pLDDT scores



**Fig. 3 | 3D mapping of Naprt mutational sensitivity and selected variants.** The AlphaFold 3D model of the Naprt-PH proteoform (FlyBase id: FBpp0306840, 662 residues) is coloured according to ProteoCast-predicted mutational sensitivity, from blue (low sensitivity) through white to red (high sensitivity). The unstructured regions 1–21 and 301–407, with pLDDT < 70, are omitted for clarity. The five variants selected for CRISPR validation are highlighted in spheres and labelled. **A, B** Complex obtained by superimposing the biologically relevant homo-dimeric arrangement of the human Naprt (PDB id: 4YUB, see also Table S4 and Fig. S10). The

two yellow compounds, nicotine mononucleotides, indicate the putative location of the active sites. They were positioned by superimposing the structure of TaNAPRTase from *Thermoplasma acidophilum* (PDB id: 1YTK). **A** Front view. **B** Back view. **C, D** Transparent surfaces indicate regions potentially involved in binding the substrate or in forming the homo-dimer. **C** Residues within 10 Å from the nicotine mononucleotides. **D** Residues within 5 Å from the superimposed human Naprt's second monomer.

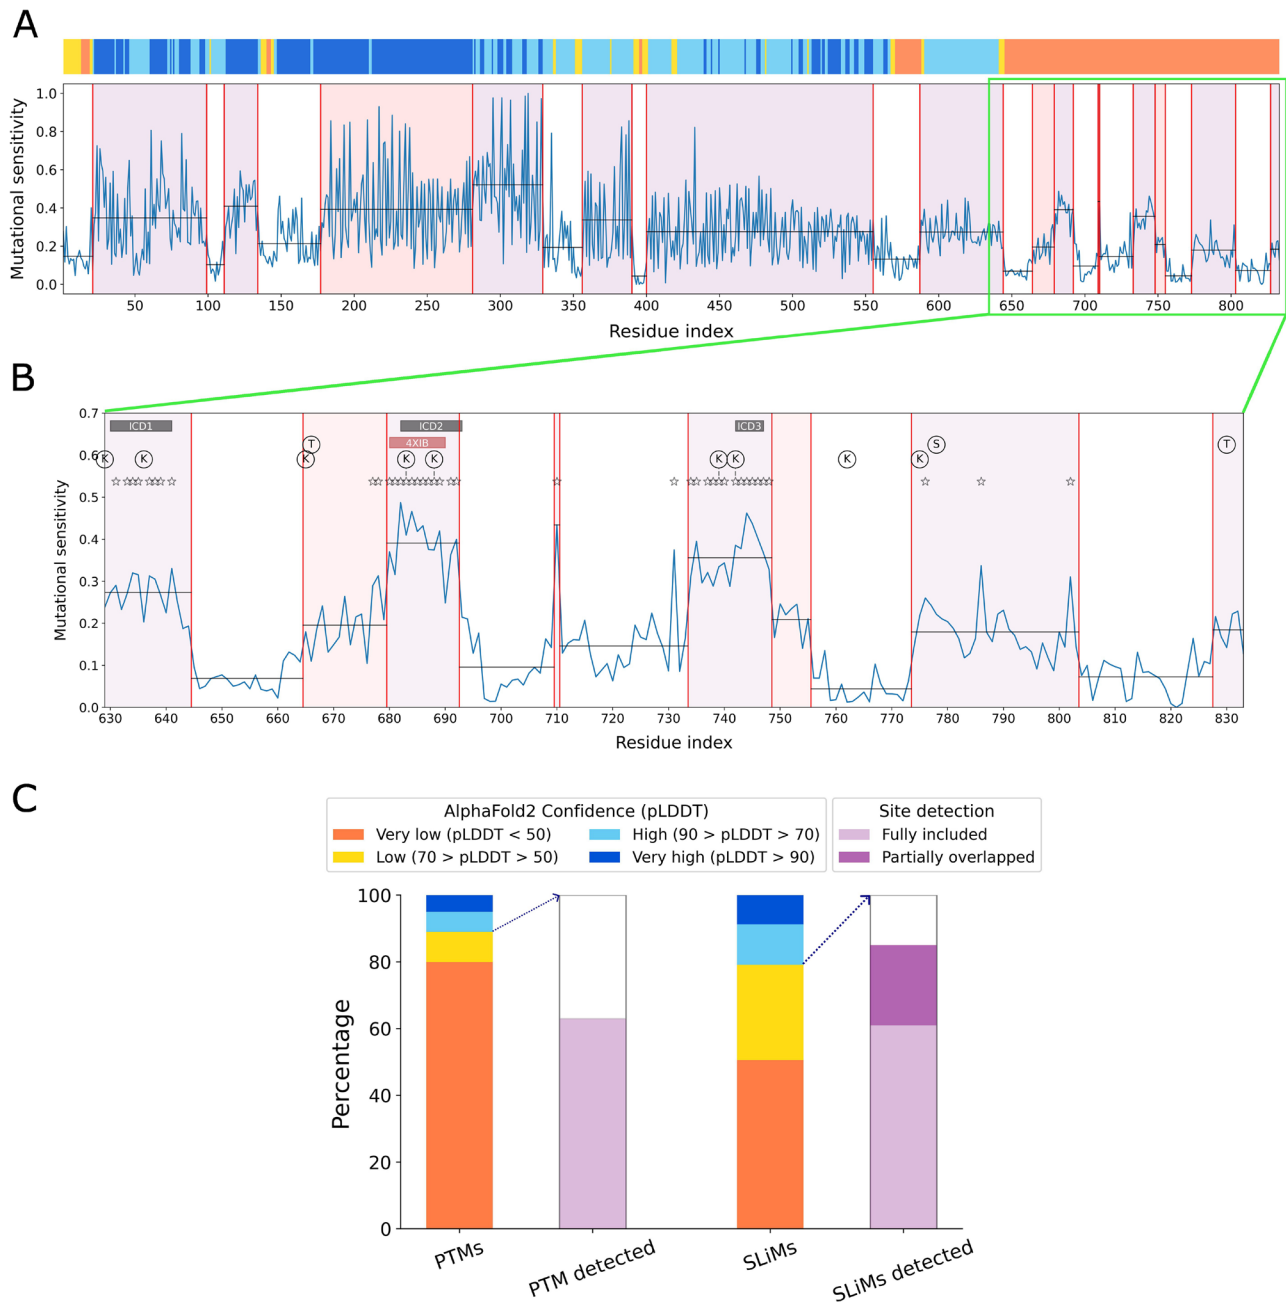
below 70 had confident variant predictions (Fig. S12B). Hence, AlphaFold2's inability to model 3D structures does not imply that the evolutionary signal is too weak for ProteoCast to predict variant effects.

Yorkie's mutational landscape exhibits several local patterns of strong predicted variant effects within low-pLDDT regions (Fig. 2A). While UniProt does not provide any information about these regions, beyond intrinsic disorder detected by MobiDB-lite<sup>42</sup>, the high mutational sensitivity of segments 32–98 and 135–160 suggests that they are important for the protein function. To validate this hypothesis, we analysed post-translational modification (PTM) site annotations from iProteinDB<sup>43</sup> and experimental structural data from the Protein Data Bank (PDB)<sup>44</sup>. Both segments exhibit a high concentration of known phosphorylation sites—10 out of 20 annotated (Fig. S13A), and segment 32–98 corresponds to the TEAD binding domain of Yorkie's human homologue, YAPI<sup>45</sup>. This domain is natively unfolded and conditionally adopts a stable conformation upon binding to its partner that is closely mirrored by Yorkie's AlphaFold2-predicted model (Fig. S13A). Hence, the evolutionary signals detected without any supervision by ProteoCast point at unstructured segments directly involved in protein regulation and interactions.

To systematically and automatically detect such sites, ProteoCast models mutational sensitivity as a piecewise constant continuous parameter and detects both the number of changepoints and their location using the FPOP algorithm<sup>46,47</sup> (Fig. 1B, *Sensitivity Profile*, and see Methods). The pLDDT profile is used to tune granularity and avoid over-segmenting well-structured regions. For ease of interpretability, the detected segments are colour-coded (Fig. 4A, B and Fig. S13B) depending on whether they have elevated sensitivity compared to both neighbouring segments (purple), only one (red) or none (white).

We quantitatively assessed this segmentation over the entire fly proteome, focusing on regions poorly modelled by AlphaFold2 (Fig. 1C, *Interaction & regulatory sites*). Our evaluation benchmark comprises 53,737 PTM sites annotated in iProteinDB database<sup>43</sup> and 72 short linear motifs (SLiMs) from the ELM resource<sup>48</sup>. ProteoCast successfully identified 63% of the PTM sites in segments showing elevated mutational sensitivity (Table S5 and Fig. 4C). This estimate is in line with previous studies reporting between 35% and 65% of phosphosites under evolutionary constraints across eukaryotic proteomes<sup>49,50</sup>. Regarding SLiMs, 85% partially overlapped with the segments detected by ProteoCast and 61% are fully included in them (Table S5 and Fig. 4C). Overall, these results confirm that SLiMs, and to a lesser extent PTMs, are under selective pressure. Moreover, the strategy of defining segments whose mutational sensitivity is homogeneous and elevated compared to their surrounding context yields a higher detection rate than relying on the absolute mutational sensitivities computed for individual residues (Table S5 and Fig. S14).

We further used ProteoCast profiling to decrypt the sequence of the unstructured C-terminal part of the Delta protein (833 residues), a ligand of the Notch signalling pathway (Fig. 4A, B). The residues predicted as sensitive to mutations (> 50% substitutions classified as mild or impactful, indicated as stars) are mostly clustered in three segments identified by ProteoCast as having higher sensitivity than their background context (in purple). While UniProt does not provide any functional information about these segments, we found evidence in the literature for their direct implication in regulating Delta's ubiquitination and activation by Mindbomb1 and Neuralized (Fig. 4B, black and pink boxes)<sup>51</sup>. By contrast, the ubiquitination sites themselves (marked with "K" in Fig. 4B) show varying levels of mutational



**Fig. 4 | Structural confidence and mutational sensitivity of the Delta protein.**

**A** Top. The pLDDT confidence score from the AlphaFold2-predicted 3D structure is shown as a colour bar. Orange and yellow indicate very low to low confidence, while light and dark blue represent high to very high-confidence regions. Bottom. Mutational sensitivity as estimated by the average of the scores predicted for the 19 possible substitutions. The curve is segmented using the FPOP algorithm (see Methods). The identified segments are coloured red if their mean value (horizontal line) is higher than one neighbouring segment, and purple if it is higher than both neighbouring segments. **B** Close-up view of ProteoCast sensitivity profile for the C-terminal region from residues 629 to 833, where the stars highlight residues sensitive to mutations, according to ProteoCast classification. The black boxes

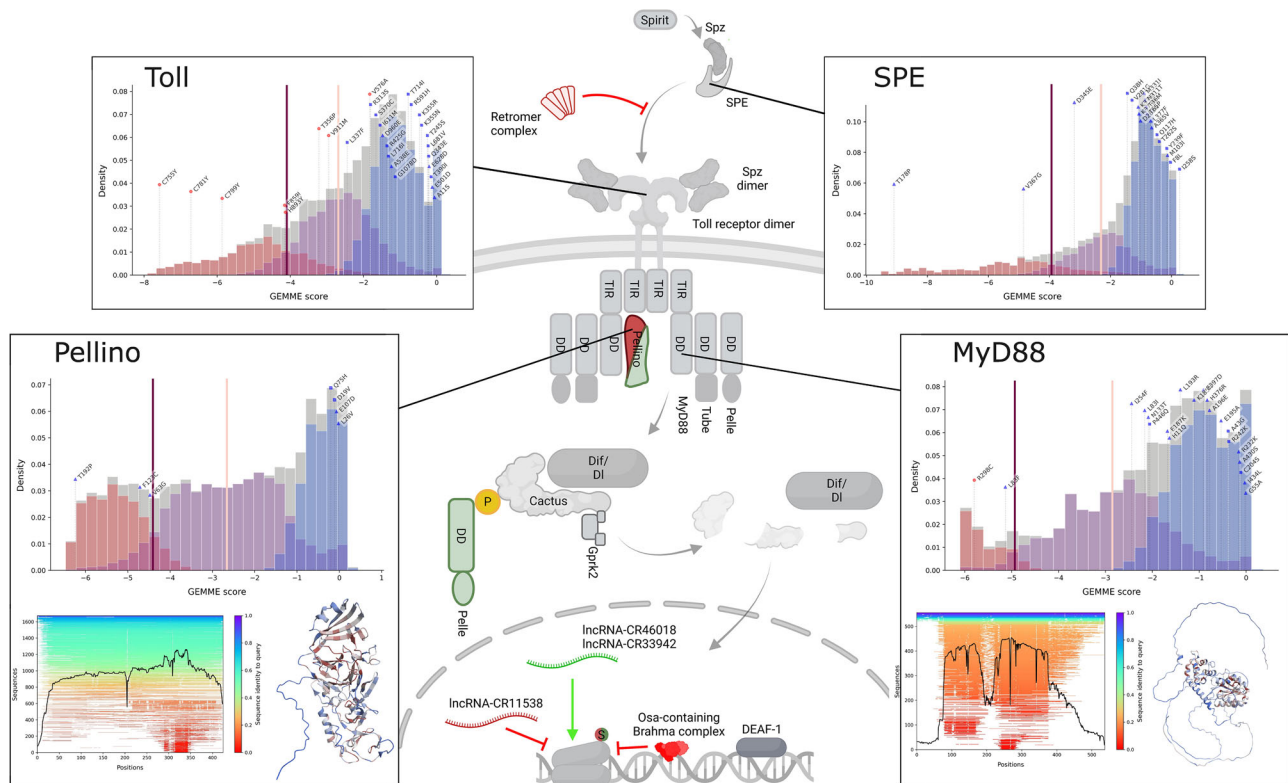
mark experimentally validated binding motifs for ubiquitylation by Mindbomb1 and Neuralized<sup>51</sup>. The pink box indicates the region whose structure was experimentally resolved in complex with Mind bomb 1 (MIB1) (PDB code: 4XIB, residues 680 to 690). Ubiquitination sites are indicated by circled "K", and phosphorylations by circled "S" and "T". **C** ProteoCast detection of PTM sites and SLIMs over the *D. melanogaster* proteome. The left bars indicate the proportion of known PTM sites and SLIMs in the different AlphaFold pLDDT bins. The right bars give the proportions of PTM sites and SLIMs partially overlapping with or fully included in segments detected by ProteoCast as having elevated sensitivity. Source data are provided as a Source Data file.

sensitivity, consistent with experimental evidence suggesting their dispensability - ubiquitin does not necessarily need to be added to these lysines for proper regulation<sup>51</sup>. In addition, ProteoCast analysis highlighted four cysteines, namely C710, C731, C786, and C802, in the mutational sensitivity profile (Fig. 4B). These predictions, along with experimental studies supporting the essential regulatory role of cysteines within intrinsically disordered regions in response to

Reactive Oxygen Species (ROS)<sup>52,53</sup>, suggest this region may function in a redox-dependent manner, with conditional disorder<sup>54</sup>.

### Case study on Toll pathway proteins

We finally conducted a focused analysis to examine the relationship between natural polymorphism, available cross-species sequence data, and ProteoCast predictions, using a multi-protein case study from the



**Fig. 5 | ProteoCast predicted mutational landscape for proteins from the Toll pathway.** ProteoCast score distributions and variant classification are represented for four proteins of the Toll pathway, the membrane receptor Toll (FBpp0084431, 1097 residues), the serine protease SPE (FBpp0083832, 400 residues), the transducer MyD88 (FBpp0087679, 537 residues) and its negative regulator Pellino (FBpp0083913, 424 residues). The three fitted Gaussians are highlighted in colours and the vertical bars indicate the variant classification thresholds, as in Fig. 2B.

Sampled natural polymorphisms (inbred or wild) and reference lethal mutations are depicted as blue and red points, respectively. For Pellino and MyD88, we also show the input alignment plots and the AlphaFold-predicted 3D models coloured according to ProteoCast-predicted mutational sensitivity, from blue through white to red. The schematic representation of the Toll pathway is adapted from Valanne et al.<sup>83</sup> with permission from Oxford University Press.

Toll pathway (Fig. 5 and Table S6). This pathway is essential to the developmental patterning of the dorso-ventral axis as well as immunity. Consistent with our proteome-wide analysis, ProteoCast misclassified only a small fraction of population polymorphisms (48 out of 344) in these proteins. The Toll membrane receptor, for example, exhibits 54 mutations in the DEST2 and DGRP populations, all but 4 deemed neutral by ProteoCast. The vast majority of the misclassified mutations (38 out of 48) are exclusively observed in the DEST2 pool (Table S6). None of the mutations exclusively observed in the DGRP are predicted as impactful. This result highlights the purifying selection operating on flies inbred in the laboratory for more than 20 generations - highly homozygote - compared to wild individuals - mostly heterozygote. Moreover, mutations observed exclusively in the DEST2 pool are about four times more numerous than those exclusively observed in the DGRP, in line with proteome-wide estimates (DEST2/DRGP ratio of 3.96 for all 22,392 representative proteoforms). Larger proteins tend to harbour more polymorphisms, due to more positions being mutated while the per-residue polymorphism rate remains constant (mean = 1.04, s.d. = 0.06), a tendency also confirmed at the level of the entire proteome (adjusted  $R^2 = 0.33$  for 18,238 proteoforms spanning 40–1000 residues).

The fraction of mutations observed in the input alignment varies widely across proteins, from less than 10% for Mtk - suggesting its specificity to a few organisms - to over 90% for the extracellular serine proteases Snake and SPE, both cleaving the Toll ligand Spatzle. This high fraction reflects the ubiquity of this family of enzymes in eukaryotes and prokaryotes. Intermediate fractions (50–65%) were observed for both the intra-cellular proteins MyD88, a transducer of

the Toll receptor, and its negative regulator Pellino. However, the mutational landscapes of these two proteins contrast strikingly. MyD88 contains unstructured regions that are poorly covered in the input alignment and prone to accumulating natural polymorphisms (19 out of 21 classified as neutral), and thus less likely to alter the protein function. Conversely, Pellino is largely structured, and the alignment is dominated by homologues sharing high sequence similarity (> 40%). Its low number of natural variants in DEST2 (5 mutations) likely stems from strong structural constraints.

Overall, the systematic application of ProteoCast can assist experimentalists and theorists in exploring the role of natural polymorphism in the function of specific proteins or entire pathways, offering complementary information for interpreting Genome Wide Association Studies (GWAS).

### Comparison with other state-of-the-art methods

Having demonstrated ProteoCast's accuracy on *Drosophila*-specific benchmarks and case studies, we further assessed its broader applicability through independent evaluation on standard benchmarks (Table 2). We evaluated ProteoCast without any parameter adjustment against established methods originally developed using these or similar datasets.

**Distinguishing pathogenic from benign variants in ClinVar.** ProteoCast provided confident predictions for 99% of the 63,000 human variants annotated in ClinVar, correctly identifying 77.1% of the pathogenic ones as mild or impactful with a specificity of 86.9% (Table S7). Moreover, it prioritised some pathogenic variants among

**Table 2 | Independent evaluation of ProteoCast on standard benchmarks compared to established methods**

ProteoCast functionality	Evaluation benchmark	Comparison methods	Supervised (Y/N)	Used the benchmark (Y/N)	Results summary
Variant classification	ClinVar <sup>82</sup> (63k)	EVE <sup>8</sup>	N	Y	higher coverage for ProteoCast, similar AUC-ROC for both methods, slightly higher recall for EVE on a variant subset (Table S7, Fig. S15)
Functional site identification in unstructured regions	curated set from the yeast proteome <sup>55</sup> (526) CAID3 Binding-IDR <sup>56</sup> (64)	phylo-HMM <sup>55</sup> top-10 CAID3 predictors	N Y	Y N <sup>a</sup>	substantially higher detection rate for ProteoCast (Table S8, Fig. S16) ProteoCast on par with the other predictors (Tables S9 and S10, Fig. S18)

For each dataset, we indicate the ProteoCast functionality evaluated, the number of data points, the comparison methods used, whether these methods are supervised, and whether the evaluation benchmark was used in the original method publications (indicating potential use during development). Main results are summarised with references to corresponding tables and figures.

<sup>a</sup>The methods were trained on other CAID versions or other binding datasets.

the top 10 strongest predicted effects or among the top 5 most sensitive residues for 632 out of 2525 proteins. For instance, ProteoCast identified the disease mutation R578H, associated with Pitt-Hopkins syndrome, as the most impactful variant in transcription factor 4 (TCF4) among 12,749 possible substitutions. These results demonstrate that ProteoCast effectively captures human genetic and disease patterns. For comparison with the state of the art, we chose EVE, which, like ProteoCast, performs unsupervised variant classification based on predicted evolutionary raw scores<sup>8</sup>. We made the comparison on the subset of 1736 human proteins from ClinVar that are covered by EVE (data downloaded from <https://evemodel.org/download/bulk>). EVE provided predictions for only 88% of the 47,022 variants annotated in these proteins based on its input MSA quality assessment. This assessment is conceptually similar to our confidence score but results in a much lower coverage of the mutational landscape for EVE (median of 80%, minimum of 5%) compared to ProteoCast (median of 99% and minimum of 79%, Fig. S15B). EVE further excluded 25% of the remaining variants because of high uncertainty in its predictions (S15C, D). This protocol led to lower recall (69.0%) and specificity (72.5%) compared to ProteoCast (78.3% and 85.2%) on the full dataset (Table S7). Both EVE and ProteoCast achieved similar specificity (~88) on the subset of 33,274 variants confidently classified as pathogenic or benign by EVE, with a slightly higher recall for EVE (86.6 versus 81.9 for ProteoCast, Table S7). Overall, ProteoCast provides a good balance between sensitivity and specificity across virtually all mutations, while EVE, despite achieving higher recall on the subset it classifies, leaves a substantial fraction (20%) of annotated variants uncertain.

**Identifying functional sites in the disordered fraction of the yeast proteome.** We focused on the curated set of 500 phosphorylation sites, localization signals, degradation signals, SUMO sites, and interaction motifs defined in Nguyen Ba et al. (2012) (see Supplementary Data S1). Half of the sites are shorter than 4 residues, with a maximum length of 65 residues (Fig. S17A). Among the 437 low-pLDDT sites, ProteoCast fully included 186 sites within segments showing locally elevated mutational sensitivity—more than twice the number detected by phylo-HMM (99 sites, Table S8). When considering the 347 sites labelled as disordered by ref. 55, ProteoCast achieved detection rates of 38% compared to 30% for phylo-HMM (Table S8). These results demonstrate that modelling mutational sensitivity as a continuous signal from large sequence datasets allows recovering evolutionary constraints that cannot be detected by phylogenetic HMMs over a dozen closely related species. For instance, in SCD5, phylo-HMM detects only one out of three copies of the motif “LKPTATG”. Despite the relatively low conservation of the two others, ProteoCast contrasts them with high precision against their background (Fig. S16).

**Identifying disordered binding sites from CAID3.** The 62 annotated binding sites in CAID3<sup>56</sup> dataset (<https://caid.idpcentral.org/assets/sections/challenge/static/references/3>) are substantially longer than PTM sites or SLIMs, with a median length of 19 residues and extending up to 989 (Fig. S17C). 40% of them were fully included in ProteoCast segments with elevated sensitivity (Table S9). Using sensitivity differences between the identified segments as a predictive metric, we achieved residue-wise performance comparable with the top 10 performing predictors in the CAID3 challenge – our approach is ranked 6th based on AUC (Fig. S18 and Table S10). Relying solely on per-residue mutational sensitivity yielded lower performance. These results demonstrate that our unsupervised MSA-based approach identifies disordered binding sites as effectively as supervised machine learning methods trained specifically on this type of data, including those leveraging embeddings computed by pre-trained protein language models.

## Discussion

### Biological insights from *Drosophila* proteome analysis

Proteocast has three primary purposes for the *Drosophila* community and beyond. Firstly, it facilitates the identification of genome-editing targets by enabling rapid prioritisation of the hundreds of millions of possible missense mutations over the fly proteome. Secondly, it sheds light on the evolutionary conservation and 3D structural context of hundreds of thousands of missense variants found in inbred and wild fly populations and lacking annotations. Thirdly, it can aid in elucidating the role of unstructured regions of the fly proteome by pointing at potential interaction and regulatory sites within these regions. It offers freely accessible, user-friendly, and interactive online web services for a detailed examination of any single-point variant of interest at <https://proteocast.ijm.fr/drosophiladb/>.

We have provided functional validations of our predictions using the rich datasets available for *Drosophila* as well as CRISPR experiments. These results demonstrated ProteoCast's ability to accurately predict organism-level mutational impact, specifically in terms of binary phenotypic outcomes such as developmental viability. Moreover, they allowed for discovering three recessive developmentally lethal mutations in a protein involved in the biosynthesis of a coenzyme central to metabolism. Ongoing and future work will aim at full biochemical characterisation of the generated mutants.

Furthermore, our analysis has revealed that a substantial fraction of variants found in inbred and wild fly populations are counter-selected in evolution, suggesting that they may alter protein or organismal fitness. Some of them occur in proteins with known lethal missense mutations, thus likely essential for the organism. For instance, ProteoCast identified impactful DEST2 mutations clustered near the putative DNA-binding interface of the nucleosome remodeler Mi-2, and within the kinase domain of the Doa somatic sex determination protein, where the natural variant V364G scored more severely than the nearby lethal mutation C275Y (Fig. S19). These findings call for biochemical characterisation of the putative impact of these naturally occurring mutations on protein activity and interactions. More broadly, these observations raise questions about how such potentially deleterious variants are maintained differently across population types. ProteoCast predictions could facilitate multiple comparative approaches: examining the differences between the polymorphism fixed in DGRP, inbred in the laboratory for more than 20 generations ( $F = 0.986$ ), with that of natural populations in DEST2, or studying seasonality, geography and natural polymorphism by focusing on impactful DEST2 variants rather than genome-wide scans. We explored the latter approach for two xenobiotics detoxifying genes, *Cyp6a9* and *Cyp317a1*, with geographically influenced incidence (Fig. S20). Reciprocally, while ProteoCast helps interpret natural variants, these population data may in turn help refine annotations for the fly proteome. We identified 32 mutations referenced as developmentally lethal in FlyBase, yet observed in inbred or wild fly populations. ProteoCast provided high-confidence predictions for 21 mutations, classifying all but one as neutral, highlighting the need for better curation of lethality annotations. For example, the R10K mutation in hopscotch kinase is annotated as lethal but likely owes its phenotype to an accompanying premature stop codon that eliminates the entire kinase domain<sup>57</sup>.

Overall, the evolutionary data available for the fly allowed obtaining high-resolution for most of the proteins. A subset of 725 unique proteoforms from 593 genes even have their mutational landscapes nearly entirely covered by the input alignment, with more than 95% of all possible substitutions actually observed in natural sequences (Fig. S21). This natural mutational saturation highlights the ubiquity of the corresponding protein families across life kingdoms. These proteoforms have many identifiable homologues (at least 4000 sequences in the alignment) that retain detectable similarity while exhibiting extensive variation at individual sites. They are more likely to harbour

mutations impairing organismal fitness compared to the rest of the proteome—1.8-fold enrichments in lethal mutations. They tend to be more connected in the fly interactome, with a 2.8-fold higher likelihood of having > 20 protein partners. Furthermore, their genomic loci are significantly enriched for regions associated with transcription factor (TF) binding sites recognised by early development genes such as *CTCF*, *BEAF-32* or *Dref* that are also involved in chromatin domain organisation and early development (Fig. S22). Collectively, these findings suggest the functional importance of these proteins to the organism through their involvement in key developmental processes, and that such importance can, to some extent, be inferred directly from its evolutionary history, without needing additional information about population, genetics or interactomics data.

Turning to structural applications, we systematically mapped ProteoCast predictions on AlphaFold-predicted 3D models over the fly proteome. While both methods exploit multiple sequence alignments, ProteoCast extracts evolutionary signals unsupervised whereas AlphaFold2 maps alignments to 3D coordinates. Regions with low pLDDT (indicating mapping difficulties) can still harbour sites under selection detectable by ProteoCast, thus remaining important for protein interactions and functions. Our results demonstrate ProteoCast's potential for transparent screening of phosphoproteomics datasets to identify functionally relevant PTM sites without expert intervention. Its adaptive detection of evolutionary signals in unstructured regions through an automated segmentation may prove advantageous over methods using absolute thresholds<sup>58</sup>. Beyond filtering known PTM sites and linear motifs, ProteoCast can discover uncharacterised functional sites in disordered regions.

### Methodological innovations and comparative advantages

This study builds upon established approaches for accurate, interpretable, and scalable prediction of variant effects that exploit evolutionary information from protein sequences<sup>7,31</sup>. ProteoCast introduces several methodological innovations compared to GEMME by (i) integrating evolutionary scores with protein structural features, (ii) identifying regulatory sites within unstructured protein regions, and (iii) providing calibrated confidence metrics and systematic variant classifications to guide experimental prioritisation. While individual methods in the field address some of these aspects separately, ProteoCast uniquely integrates all three within a single evolutionary-based framework.

ProteoCast provides a graduated classification of variants as neutral, mild, and impactful, aimed at helping users interpret the predicted mutational landscape and prioritise variants for experimental validation. This classification encompasses all possible variants of a given protein, in marked contrast with approaches like EVE that use GMM probabilities to estimate uncertainty and exclude uncertain variants from evaluation. Our approach is pragmatic and transparent, avoiding potential model or dataset biases that may lead to inconsistencies between statistical probabilities and actual biological uncertainties. ProteoCast further leverages its continuous mutational sensitivity scores through a segmentation strategy designed to identify functional motifs including regulatory and binding sites in unstructured protein regions. This strategy represents a substantial conceptual advance over previous foundational works for addressing this problem<sup>55,59</sup>.

First, we exploit mutational sensitivity rather than conservation, as we observed that the former provides higher resolution and better agreement with experimental measurements of protein stability<sup>30</sup>. Second, our segmentation approach models mutational sensitivity as a piecewise constant continuous parameter, providing a quantitative framework for capturing and disentangling different levels of evolutionary constraints across protein segments. As Figure 4 clearly illustrates, mutational sensitivity contains substantially richer information than binary signals, and our model is considerably more expressive

than conservation-based models relying on binary state distinctions<sup>55,59</sup>. This conceptual framework is supported by previous studies demonstrating the superior accuracy of segmentation approaches over simpler “two-state” models on diverse biological datasets, including RNA-seq<sup>60</sup> and ChIP-seq data<sup>61</sup>, reinforcing their suitability for capturing the complexity and heterogeneity of biological signals.

### Limitations and perspectives

Despite the robustness of our approach, alignment errors might result in variant or residue misclassification. Protein extremities are typically highly gapped in the input alignment and ambiguities in gap opening/closing positions may arise at the borders of poorly covered regions. Nevertheless, ProteoCast confidence metrics provide a means for flagging potentially problematic positions. For instance, half of the 22 lethal or hypomorph mutations affecting the first methionine are misclassified as neutral, but most of these incorrect predictions (7/11) are flagged as unreliable due to poor alignment quality. This demonstrates that our confidence metrics effectively identify problematic predictions. More broadly, while alignment quality parameters and thresholds might theoretically be optimized differently for distinct applications, our approach prioritises practical usability with minimal parameter tuning. Our validation across diverse tasks and the entire *Drosophila* proteome demonstrates the robustness of using standardized alignment generation protocols for evolutionary constraint-based analyses. Future versions of ProteoCast will explore integrating complementary approaches based on pLMs to further improve prediction accuracy in challenging regions. This is a promising avenue as we previously showed that the pLM-based predictor VespaG helps to regularise GEMME scores thanks to its knowledge of a universal protein representation space<sup>12</sup>.

Future research will aim at a tighter integration of evolutionary and structural data. For instance, applying ProteoCast segmentation strategy to AlphaFold2 pLDDT profiles revealed that around 34% of the known PTM sites and 35% of the known SLiMs fall within regions modelled with higher confidence than their surroundings. This detection likely reflects AlphaFold2's ability to recapitulate and generalise experimental knowledge about conditional folding. It only partially overlaps with ProteoCast detection from the predicted mutational sensitivity profiles, suggesting avenues for combining the two types of information. Another direction will consist in extending the resource to other common model organisms, such as *X. laevis*, *C. elegans*, *S. cerevisiae*, and *E. coli*, as well as *H. sapiens*. The wealth of mutational data and associated phenotype will help further assess the quality of our predictions. The generated data will facilitate the evaluation of large-scale predictions against phenotypic knowledge at both organismal and molecular levels as well as significantly ease the identification of residue-specific genome editing. Our findings will serve as a starting point toward identifying the molecular determinants of continuous phenotypic traits controlled by multiple genes.

## Methods

### Drosophila experiments

*Lines used. daGS driver*<sup>62</sup> (provided by the Monnier lab, Université de Paris Cité).

*Vienna Drosophila Resource Center stock (gene name):* 42571 (*JIL-1*, FBgn0020412), 31712 (*brm*, FBgn0000212), 31053 (*HDAC6*, FBgn0026428), 61875 (*Naprt*, FBgn0031589), 35704 (*smid*, FBgn0016983), 67758 (*Gpo3*, FBgn0028848), 31395 (*Rm62*, FBgn0003261), 28038 (*nAChRβ2*, FBgn0004118), and 27503 (*CadN*, FBgn0015609).

*RNAi-induced developmental lethality.* All the flies are kept in closed vials in incubators at controlled temperature, humidity (60%) and 12 hours light cycle. Experiments are carried at 26 °C. Five individual replicate vials were made per condition. Each RNAi was induced

using two distinct concentrations of RU486 0 μg/mL (negative control) and 20 μg/mL, the maximum inducer concentration usable during development as previously described<sup>63</sup>. Each contains 6 *daGS* virgin females and 5 RNAi males left to cross and lay eggs for 48 hours on the following food composition: 8.33% (w/v) yeast, 8.33% (w/v) corn, 1.13 (w/v) agar and Methyl 4-hydroxybenzoate (Moldex) at a final concentration of 5.3 g/L to prevent fungal contamination. Eclosing flies were then counted daily for 10 days.

### Datasets

**FlyBase.** We retrieved the *Drosophila melanogaster* proteome version 6.44 as a FASTA file (*dmel-all-translation-r6.44.fasta*), available at FlyBase<sup>34</sup> (<https://ftp.flybase.net/genomes/>). This transcriptome contains a total of 30,738 sequences, of which 22,392 are unique. We focused our attention on single-point mutations generated by ethyl methylsulfonate (EMS) and made a list of those associated with developmental lethality - using the Phenotype QuickSearch tool provided by Flybase - as well as a second list of hypomorph mutations. We filtered out the 250 mutations overlapping between the two.

**Drosophila genetic reference panel.** We obtained all polymorphisms from the 205 inbred lines of the *Drosophila* Genetic Reference Panel (DGRP)<sup>32</sup> in a text file (*dgrp.fb557.annot.txt*). This dataset contains approximately 4.5 million genomic variations, including deletions, insertions, single-nucleotide polymorphisms (SNPs), and multi-nucleotide polymorphisms (MNPs). SNPs are the most prevalent type of variations (89% of all detected variants), 4.5% of which result in missense mutations leading to amino acid changes in proteins. Mapping these missense mutations to the fly proteome sequences resulted in 178,264 unique variants affecting 20,137 (89.9%) proteoforms from 12,682 genes (90.8%).

**Drosophila evolution over space and time.** The *Drosophila* Evolution over Space and Time (DEST)<sup>33</sup> resource is a large-scale dataset that compiles genomic polymorphism data from natural *Drosophila melanogaster* populations sampled across diverse geographic locations and time periods. We obtained all annotated SNPs from the August 2024 version of the dataset in VCF format (*dest.all.PoolSNP.001.50.3-May2024.ann.vcf.gz*). It includes approximately 4.8 million SNPs classified into 27 different annotation types. For this study, we focused specifically on missense SNPs, identifying 332,848 variants that affect 94% of *Drosophila* genes and 92.4% of proteoforms. After resolving mismatches in gene and transcript identifiers, we retained 285,771 unique missense SNPs for further analysis.

**Post-translational modifications.** We retrieved all PTM sites annotated for the *Drosophila melanogaster* proteome in the iProteinDB<sup>43</sup> database at <https://www.flyrnai.org/tools/iproteindb/> through its API. We identified PTM sites in 8530 representative proteoforms (38%) from 4276 genes (31%). The vast majority of these sites are targeted by phosphorylation (-92%), the remaining sites being oxidized or acetylated.

**Short linear motifs.** We obtained 113 short linear motifs (SLiMs) ranging from 3 to 20 amino acids, from the Eukaryotic Linear Motif (ELM) resource (<http://elm.eu.org/>)<sup>48</sup>. These motifs were retrieved from a downloadable file, *elm\_instances.tsv*, using the following search path: ELM DB → ELM instances → Filter by organism (*Drosophila melanogaster*).

### ProteoCast workflow

The ProteoCast workflow unfolds in five main steps.

1. *Multiple sequence alignment generation* (Fig. 1A). Given an input query protein sequence, it retrieves related sequences using the MMseqs2-based protocol implemented in ColabFold<sup>37</sup>.

2. **Full single-mutational landscape prediction** (Fig. 1A). It uses the variant effect predictor GEMME<sup>7</sup> to estimate the impact of all possible amino acid substitutions at every position in the query.
3. **Prediction reliability assessment** (Fig. 1B). It flags the query as having unreliable predictions due to insufficient evolutionary information if the input alignment contains less than 200 sequences. Otherwise, it assigns binary confidence scores that reflect alignment quality and diversity to each residue in the query.
4. **Variant and residue classification** (Fig. 1B). It classifies variants as impactful, mild or neutral by fitting a Gaussian mixture model<sup>64</sup> to the predicted raw scores. It further classifies residues as sensitive to mutations if at least 10 out of the 19 possible substitutions are mild or impactful. Otherwise, residues are considered tolerant to mutations.
5. **3D structural mapping and segmentation** (Fig. 1B). It maps residue mutational sensitivity scores onto the 3D model predicted by AlphaFold2 for the query protein. It decomposes the mutational sensitivity profile into segments using the FPOP algorithm<sup>46,47</sup>. The granularity of the segmentation accounts for AlphaFold pLDDT profile.

In the following, we give details about the different steps, illustrating their application over the entire fly proteome.

**Step 1. Multiple sequence alignment generation.** For every unique proteoform from the fly proteome (FlyBase version 6.44), we generated a multiple sequence alignment (MSA) using the ColabFold (CF) protocol<sup>37</sup>. We previously showed that this protocol compares favourably with workflows relying on profile hidden Markov models for the prediction of mutational effects<sup>31</sup>. It is much faster, thanks to its usage of the highly efficient sequence MMseqs2 search algorithm<sup>36</sup>, and it yields similar or better predictive accuracy. CF protocol produced two alignments for each input sequence by searching against UniRef30, a clustered version at 30% sequence identity of UniRef100<sup>65</sup>, and a curated database comprising various environmental sequence sets<sup>37</sup> (ColabFold environmental database, ColabFold\_Env). We used UniRef30 v2103 and ColabFold\_Env v202108. We concatenated the alignments and gave the resulting one as input for GEMME. The CF protocol did not manage to find any homologue for 9 proteoforms.

**Step 2. Full single-mutational landscape prediction.** We utilised the *reformat.pl* tool from the HH-suite<sup>66</sup> to convert A3M alignment files into FASTA format. We handled sequences containing undefined residues (represented by symbols “X” or “U”) by employing the following strategy: when the undefined residue was located at the beginning of the sequence, the rest of the corresponding column in the alignment was always filled with gaps, and thus we removed that column. For other instances of undefined residues, we replaced them with the most frequently occurring amino acid at the respective position(s) within the alignment. We ran GEMME with default parameters except that (1) we increased the memory (RAM) to handle very long proteins (*Xmx9126m*), (2) we increased the number of JET<sup>67</sup> iterations to 7 to ensure robustness. The JET method subsamples the initial input set of sequences, generates a tree for each subset and determines each position’s evolutionary trace, reflecting the level in the tree where an amino acid appeared at that position and remained fixed thereafter. It estimate a conservation value  $T_{JET}(s^*, i)$  for each residue  $i$  in the query  $s^*$  by averaging evolutionary traces over the generated trees. Subsampling is performed using a Gibbs sampling strategy and hence, multiple runs may produce different conservation estimates. We found that the Pearson correlation between the conservation values computed from two independent runs could be as low as 0.78 over a subset of 200 protein sequences. Setting the number of iterations to 7 and retaining the maximum conservation value for each

residue improved reproducibility, with correlation coefficients systematically above 0.95 for any two runs. Another parameter in JET controls the number of sequences from the input alignment that are used to estimate residue evolutionary conservation. By default, it is set to 20,000 sequences, a threshold that effectively allows us to retain all sequences for the vast majority (97%) of proteoforms. In total, GEMME produced some predictions for 22,169 unique proteoforms (Table S1).

**Step 3. Prediction reliability assessment.** A popular measure for estimating the diversity of an alignment is the number of effective sequences<sup>9,25,68</sup>. It is computed by summing up sequence-specific weights that reflect their overall similarity. However, we previously showed that this metric is not a good indicator of the quality of GEMME predictions<sup>31</sup>. Instead, we found that ColabFold’s MSA generation protocol, which maximises diversity and minimises the number of sequences, allows to set a reliability threshold based on the number of sequences in the MSA. Building on these previous results, ProteoCast flags proteins with fewer than 200 homologous sequences in the input alignment as likely to have unreliable predictions. For the remaining proteins, it estimates a binary-valued confidence score for each residue. Instead of simply computing the percentage of gaps in the alignment, as is often done, ProteoCast metric integrates this measure with the dispersion of predicted scores, the estimated conservation, and the fraction of substitutions actually observed in the alignment. This multifactorial approach better captures the complexity of alignment quality than simple gap counting or sequence diversity metrics alone. The confidence score  $c(s^*, i)$  for residue  $i$  in the query sequence  $s^*$  is expressed as,

$$c(s^*, i) = \begin{cases} 0, & \text{if } \sigma_{\text{pred}}(s^*, i) < 0.15 \text{ and } T_{JET}(s^*, i) < 0.2 \text{ and } cov(s^*, i) \cdot F_{\text{obs}}(s^*, i) < 0.1, \\ 1, & \text{otherwise,} \end{cases} \quad (1)$$

where  $\sigma_{\text{pred}}(s^*, i)$  is the standard deviation of the 19-long vector of predicted variant effects,  $T_{JET}(s^*, i)$  is the conservation value,  $cov(s^*, i)$  is proportion of sequences that have an amino acid at position  $i$  (as opposed to a gap), and  $F_{\text{obs}}(s^*, i)$  the fraction of observed substitutions at position  $i$  expressed as,

$$F_{\text{obs}}(s^*, i) = \frac{1}{19} \sum_{a \in \mathcal{A} \setminus \{s_i^*\}} \mathbb{1}_{f_i(a) > 0}, \quad (2)$$

where  $f_i(a)$  is the frequency of occurrence of the amino acid  $a$  at position  $i$  in the MSA. The resolution of the predicted mutational landscape directly depends on  $F_{\text{obs}}(s^*, i)$ . Indeed, GEMME computes an evolutionary distance between the query sequence and the closest sequence in the input alignment displaying the mutation of interest. Mutations not observed in the alignment are assigned the maximum evolutionary distance.

**Threshold tuning.** We assessed the influence of changing the thresholds for  $\sigma_{\text{pred}}(s^*, i)$ ,  $T_{JET}(s^*, i)$ , and  $score = cov(s^*, i) \cdot F_{\text{obs}}$  on the number of segments  $N_{\text{seg}}$  detected as unreliable (low-confidence score) and the proportion of segments of length one (isolated unreliable residues,  $f_{\text{seg}1}$ ). We tested several combinations of threshold values in the intervals:  $\sigma_{\text{pred}} \in \{0.1, 0.11, 0.12, 0.13, 0.14, 0.15\}$ ;  $T_{JET} \in \{0.1, 0.15, 0.2, 0.25, 0.3\}$ ;  $score \in \{0.05, 0.1, 0.15, 0.2\}$  (Fig. S23). We obtained very similar results for the tested combinations, indicating that they are robust within these value ranges. We observed a slight tendency for fewer segments and relatively more single-residue detections upon increasing the thresholds (Fig. S23).

**Smoothing.** We further smooth the per-residue confidence scores over the sequence to mitigate the discontinuities that may arise from using deterministic cutoffs. The final confidence score for the  $i^{\text{th}}$

residue of the query sequence  $s^*$  is thus computed as:

$$\text{confidence}(s^*, i) = \begin{cases} 1, & \text{if } \frac{\sum_{k=-2}^2 c(s^*, i+k) \cdot w[k]}{\sum_{k=-2}^2 w[k]} \geq 0.5 \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where the weight vector  $\mathbf{w} = [2, 3, 4, 3, 2]$  (Fig. S24).

**Step 4. Variant and residue classification.** The range of variant predicted raw scores may vary substantially from one protein to another, making direct comparisons over the entire proteome difficult. This motivated us to devise an adaptive variant classification strategy tailored to each protein-specific raw score distribution. Namely, ProteoCast uses a Gaussian Mixture Model (GMM), a probabilistic approach that assumes the data are generated from a mixture of Gaussian distributions<sup>64</sup>. The GMM parameters are estimated through the Expectation-Maximization (EM) algorithm<sup>69</sup>. Similar strategies with two or three Gaussians were previously proposed to classify experimental measurements<sup>28</sup> or scores predicted by a variational auto-encoder<sup>9</sup>.

We used the mixture module from python *scikit-learn* library<sup>70</sup> with component-specific general covariance matrices ("full" type), 30 initialisations, a convergence threshold of  $10^{-5}$ , and a maximum number of EM iterations of 1000. We quantified the effect of the number of Gaussian components on model fit across 20 trials of 100 randomly selected proteoforms. We excluded scores with low MSA-based local confidence to reduce risk of bias they might introduce. For each number of components  $k = 1 \dots 5$ , we computed the mean Akaike Information Criterion (AIC) and the log-likelihood (LL) per trial, normalised by the number of data points per protein. AIC balances model fit (*i.e.*, the log-likelihood) and model complexity (number of parameters) and thus penalises overfitting. Both metrics indicate improved model quality with increasing components (Fig. S3). Log-likelihood increases by 7% and AIC decreases correspondingly from 1 to 2 components, and by 2% from 2 to 3 components, on average. All proteins benefit from the increased number of components up to three. Beyond three Gaussians, the average gains become negligible (< 1%) and the goodness-of-fit even deteriorates for a substantial fraction of proteins, and hence, we fixed the number of components to three.

Given the three identified components, G1, G2 and G3, with mean values  $\mu_{G1}$ ,  $\mu_{G2}$ , and  $\mu_{G3}$ , respectively, and  $\mu_{G1} < \mu_{G2} < \mu_{G3}$ , ProteoCast defines three variant classes as follows,

$$\begin{aligned} & \text{if } x_{a_i \rightarrow b_i} > \mu_{G2}, \text{ then } a_i \rightarrow b_i \text{ is neutral,} \\ & \text{if } P_{G1}(x_{a_i \rightarrow b_i}) > P_{G2}(x_{a_i \rightarrow b_i}), \text{ then } a_i \rightarrow b_i \text{ is impactful,} \\ & \text{otherwise, } a_i \rightarrow b_i \text{ is mild.} \end{aligned}$$

with  $x_{a_i \rightarrow b_i}$  the raw score predicted for substituting the amino acid  $a$  by  $b$  at position  $i$ . Moreover, for each variant, ProteoCast computes a Z-score relative to the distribution of neutral mutations, Z-score ( $a_i \rightarrow b_i$ ) =  $(x_{a_i \rightarrow b_i} - \mu_{\text{neutral}}) / \sigma_{\text{neutral}}$ , where  $\mu$  and  $\sigma$  stand for the mean and standard deviation, respectively.

**Step 5. 3D structural mapping and segmentation.** For every unique proteoform from the fly proteome (Flybase version 6.44), we retrieved its 3D model either by batch downloading (<https://www.uniprot.org/proteomes/UPO00000803>) from the AlphaFold Database (13,475 structures), or by using Python *requests* library<sup>71</sup> (5119 structures). For the missing structures, we generated all those with fewer than 4000 amino acids (3345 structures) by running ColabFold with default parameters. ProteoCast then mapped the computed per-residue mutational sensitivity values (average of raw scores) and classes onto each protein's 3D model. It further used AlphaFold pLDDT scores to segment the mutational sensitivity profile, as described in the following.

The problem of segmenting a one-dimensional continuous measure along a protein sequence can be formulated as follows. We consider  $D$  changepoints  $\tau_1 < \dots < \tau_D$  within the range 1 and  $n - 1$ . These changepoints correspond to unknown positions along a protein of length  $n$ , where a shift in the mean of the measure is observed. We adopt the convention that  $\tau_0 = 0$  and  $\tau_{|T|} = n$ . These changepoints define  $|T| = D + 1$  distinct segments. The  $j$ -th segment includes the data  $\{\tau_{j-1} + 1, \dots, \tau_j\}$ . Each segment is premised on the assumption that the  $Y_i$  therein are independent and follow the same Gaussian distribution, with a mean  $\mu_j$  specific to that segment and a common variance  $\sigma^2$ . Expressed mathematically, we have:

$$\forall i \in \{\tau_{j-1} + 1, \dots, \tau_j\}, Y_i \sim \mathcal{N}(\mu_j, \sigma^2)$$

The parameters of the model, including  $\tau_1 < \dots < \tau_D$ , can be estimated using penalized maximum likelihood inference. To achieve this, we used the FPOP algorithm<sup>46,47</sup> which solves the inference problem exactly with dynamic programming. Informally, the FPOP algorithm works by considering the penalized maximum likelihood of the data from observation 1 to  $t$  as a function of the parameter (the mean) of the last segment. This concept is referred to as 'functional pruning'. In the Gaussian case, the resulting function is piecewise quadratic. When a new observation at time  $t + 1$  is made, this function can be efficiently updated (*i.e.*, compute the penalized likelihood penalized maximum likelihood function from observation 1 to  $t$ ) using a formula similar to that of the Viterbi algorithm. This formula is applied piece by piece or interval by interval. At each step, the algorithm searches for the optimal value of the parameter for the last segment that maximizes the penalized likelihood.

The number of changepoints  $D$  estimated by FPOP is a decreasing function of the penalty  $\alpha \sigma^2 \log(n)$ . The constant  $\alpha$  is a hyperparameter set to 1.4, determined after assessing the fraction of detected PTMs in the resulting segmentations while ensuring the visual goodness-of-fit (Fig. S25). The variance  $\sigma^2$  is estimated on the data using the Hall estimator<sup>72</sup>. We observed that the variations in mutational sensitivity were higher in protein regions modelled with high confidence by AlphaFold2 (pLDDT > 70) compared to other regions. To avoid over-segmenting these areas (by introducing artificial changepoints), we applied lower weights to the observations in these regions, with  $w$  set to 0.1, while keeping it at 1 elsewhere. After estimating the segments, we assign discrete scores to each of them. A score of 0 indicates that the means of both adjacent segments are greater than that of the current segment; a score of 1 that one adjacent segment has a lower mean and the other a higher mean than the current segment; and a score of 2 that the means of both adjacent segments are lower than that of the current segment.

### Implementation and computational details

To generate the input we run ColabFold protocol, multiple sequence alignment generation part, through the following command: `colabfold_search -threads=8 -db-load-mode=2 $fasta_file ${COLABFOLD_DB} $out_folder`. For a set of 3000 sequences comprising a total of 1,185,789 amino acids it took 1.5 hours to generate MSAs on the supercomputer "MeSU" of Sorbonne University (16 CPUs from AMD EPYC Milan processors, 400GB shared RAM memory).

We implemented ProteoCast in Python and R. ProteoCast relies on *scikit-learn* library<sup>70</sup> to fit a Gaussian mixture model, the *gemmi* library<sup>73</sup> to parse PDB files and manipulate the B-factors associated to the structure, the *re* module to parse annotation data files, and *pandas* for data manipulation. For the same set of the protein sequences the generation of mutational landscapes, classification and analysis (Fig. 1A *Mutational Landscape* and Fig. 1B) took 4 hours on 40 CPUs.

The online database (<https://proteocast.ijm.fr/drosophiladb/>) is a Django-based project running on a virtual machine hosted by Institut

Jacques Monod (Dell PowerEdge R740 server). The backend was developed in Python using *Django*, while the frontend was built with HTML, CSS, and JavaScript to provide a user-friendly interface. The application was deployed on an Apache2 web server with Django's WSGI integration. It incorporates *plotly* for interactive data visualization and *Mol\** (*Molstar*)<sup>74</sup> for 3D structure visualisation. Data handling and processing were facilitated using *pandas*.

## Evaluation

We estimated the performance with the balanced accuracy metric, suited to unbalanced datasets and defined as  $\text{BalancedAccuracy} = \frac{\text{Sensitivity} + \text{Specificity}}{2}$  where  $\text{Specificity} = \frac{TN}{TN+FP}$  is the fraction of DGRP or DEST2 polymorphisms predicted as neutral, and  $\text{Sensitivity}(\text{Recall}) = \frac{TP}{TP+FN}$  is the fraction of lethality-inducing mutations predicted as impactful. If several proteoforms are affected by a SNP, we take the one with the highest  $F_{obs}$ .

## Comparison with other methods

**ClinVar dataset: variant classification and comparison with EVE.** We retrieved the ClinVar annotations from the ProteinGym benchmark<sup>25</sup>. They encompass approximately 63,000 missense mutations across 2525 proteins, each labelled as either pathogenic or benign. For this analysis, we directly took the MSAs and raw GEMME scores we computed in a previous study<sup>31</sup> (<https://doi.org/10.5061/dryad.vdncjsz1s>) and applied the subsequent ProteoCast analyses to them. For 316 ClinVar sequences that did not match our existing dataset, we computed predictions from scratch with ProteoCast Docker image. As for EVE, we downloaded their bulk predictions from: <https://evemodel.org/download/bulk>.

ProteinGym provides RefSeq IDs (e.g., NP\_001750.1) and sequences for the ClinVar proteins. We performed ID mapping to UniProt accessions for ProteoCast (e.g., P30279) and to UniProt entry names for EVE (e.g., CCND2\_HUMAN), using the UniProt ID-mapping tool. EVE predictions were unavailable for 789 proteins due to ID-mapping issues (7), missing predictions for the mapped IDs (506), or sequence mismatches (276), leaving a subset of 1736 proteins covering 47,022 mutations. Additionally, EVE does not provide predictions for certain residues lacking its internal confidence score, reducing the set to 41,329 mutations (88% of the mapped mutations). If we finally remove the non-confident predictions from ProteoCast, we obtain a final common evaluation set of 41,282 mutations (99.9%).

**Functional sites in the disordered fraction of *Saccharomyces cerevisiae*.** We compared ProteoCast with phylo-HMM on a dataset of 526 yeast protein sites curated from the literature and described in Supplementary Data S1. These sites were classified into five functional categories: phosphorylation sites (the majority group, representing 82% of all sites), localization signals, degradation signals, SUMO sites, and interaction motifs. We retrieved the 174 corresponding full-length protein sequences thanks to their ORF identifiers (e.g., YKLI12W) from the *Saccharomyces* Genome Database (SGD) archive at [http://sgd-archive.yeastgenome.org/sequence/S288C\\_reference/orf\\_protein/orf\\_trans\\_all.fasta](http://sgd-archive.yeastgenome.org/sequence/S288C_reference/orf_protein/orf_trans_all.fasta). We slightly adjusted the coordinates provided by Nguyen Ba et al. (2012) for five sites, verifying them against the cited literature sources, to obtain exact matches between Supplementary Data S1 and the SGD FASTA file for all sites. We then ran ProteoCast on all proteins, generating MSAs with the ColabFold protocol (uniref30\_2302, colabfold\_envdb\_202108) and using 3D models from AlphaFold Database. ProteoCast did not provide predictions for two proteins due to highly gapped MSAs. We subsequently focused on three subsets: (1) 437 sites from 146 proteins displaying average pLDDT < 70, (2) 347 sites from 113 proteins labelled as “disordered” in Nguyen Ba et al., and (3) 333 sites from 112 proteins intersecting sets (1) and (2) (see Table S8).

**Disordered binding sites from CAID3.** We retrieved the CAID3 Binding-IDR dataset (15/07/2025 update) from [https://caid.idpcentral.org/assets/sections/challenge/static/references/3/binding\\_idr.fasta](https://caid.idpcentral.org/assets/sections/challenge/static/references/3/binding_idr.fasta). It comprises 64 annotated binding sites within 52 protein sequences. We excluded two protein sequences, one because it was too short (< 20 residues) and the other one because the MSA was too gapped. We generated 3D models for all sequences using AlphaFold2 within ColabFold.

CAID3 annotations label each residue as structured (“-”), unstructured (“0”), or unstructured and involved in binding (“1-”). Following the official CAID3 protocol, residues annotated as “1-” were treated as positives (binding), residues annotated as “0” as negatives (non-binding), while structured residues (“-”) were excluded from the evaluation for a more balanced benchmark and a fairer test of predictor performance.

Binding site identification was then assessed at two complementary levels: segment detection and per-residue prediction. For the first one, segments annotated as states 1 or 2 were considered detected sites (Table S9). For per-residue prediction, designed to enable comparison with other methods providing continuous scores, we used three residue-level metrics: raw mutational sensitivity (ProteoCast-mu\_residue); segment sensitivity contrast, defined as the maximum difference in mutational sensitivity between a given segment (states 1 or 2) and its two neighbours, set to 0 otherwise (ProteoCast-delta-mu\_segment); and segment labels mapped to per-residue positives (Table S10 and Fig. S18). Continuous predictions were evaluated using ROC curves, with area under the curve (AUC), average precision (APS), maximum F1 score (F1 max), and balanced accuracy (bac). Metrics were computed with the vectorized\_cls\_metrics library (available at [https://github.com/marnec/vectorized\\_cls\\_metrics.git](https://github.com/marnec/vectorized_cls_metrics.git)), as recommended by CAID3.

## Analysis of case studies

We used iCisTarget to search for enrichment in transcription factor (TF) binding sites<sup>75,76</sup>. In brief, iCisTarget searches existing databases of known TF binding sites as well as experimental chromatin immunoprecipitation experiments. We considered two categories of genes defined based on the fraction of observed mutations, either  $F_{obs} \geq 95$  or  $85 < F_{obs} < 90$ . We extracted a subset of 635 genes for the  $85 < F_{obs} < 90$  category to have roughly the same cardinality as the  $F_{obs} \geq 95$  set (593). We considered the enrichment to be significant with a number of enriched features above 7 for a threshold set at Normalised Enrichment Score (NES) threshold of 5.

We visualised 3D structures with Pymol<sup>77</sup> and we used the *align* command to superimpose structures resolved for homologues. For the Naprt example, after careful examination of the publication accompanying the PDB entry 4YUB<sup>78</sup>, we realised that the dimeric arrangement annotated as *Biological Assembly* did not match the one described as biologically relevant in the publication. Hence, we generate symmetry mates with Pymol to identify and visualise the correct arrangement (Fig. S10).

## Database description

For each proteoform, we provide its sequence and the corresponding multiple sequence alignment (downloadable in FASTA format), along with a visual representation, inspired by ColabFold<sup>37</sup>. The mutational landscape can be interactively browsed. It includes raw predicted scores, variant predicted classification, known annotations, and local confidence scores (downloadable in CSV format). Additionally, we provide a visualization of the raw score distribution with variant classification thresholds, and the lethal, inbred, and natural missense mutations mapped on it (downloadable as a CSV file). The segmented mutational sensitivity profile can be interactively zoomed in and out and is accompanied with a colour bar for pLDDT values (if a structure is available). Finally, we provide access to the mapping of AF2 pLDDT

scores, ProteoCast sensitivity scores and residue classification on the 3D structure (if available, all data downloadable as PDB files with values on the B-factor column).

### Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

### Data availability

The ProteoCast predictions and processed datasets generated in this study have been deposited in the Zenodo repository under accession code 10.5281/zenodo.18488840<sup>79</sup> [<https://zenodo.org/records/18488840>]. Interactive visualization and query of the predictions are provided through the ProteoCast web server at <https://proteocast.ijm.fr/drosophiladb/>. Source data are provided with this paper. The public datasets used in this study are available from the following repositories: pre-computed gene annotations from FlyBase ([https://ftp.flybase.net/releases/FB2022\\_01/precomputed\\_files/genes/](https://ftp.flybase.net/releases/FB2022_01/precomputed_files/genes/)), variant annotations from DGRP2 based on FB5.57 (<http://dgrp2.gnets.ncsu.edu/data/website/dgrp.fb557.annot.txt>), variant annotations from DEST2.0 (<https://dest.bio/data-files/SNP-tables/dest.all.PoolSNP.001.50.24Aug2024.ann.vcf.gz>), and the *Drosophila melanogaster* proteome from AlphaFold DB ([https://ftp.ebi.ac.uk/pub/databases/alphafold/latest/UPO00000803\\_7227\\_DROME\\_v4.tar](https://ftp.ebi.ac.uk/pub/databases/alphafold/latest/UPO00000803_7227_DROME_v4.tar)). AlphaFold 3D models missing from the proteome archive were retrieved manually from the AlphaFold DB (<https://alphafold.ebi.ac.uk/>). Source data are provided with this paper.

### Code availability

ProteoCast is an open-source software released under the GNU General Public License. The software can be run locally using the Docker image available at <https://hub.docker.com/r/marinaabakarova/proteocast/> or accessed through the ProteoCast web server at <https://proteocast.ijm.fr/><sup>80</sup>. All scripts used for data processing, analysis, and figure generation are publicly available on GitHub at <https://github.com/abakarovaMarina/ProteoCast><sup>81</sup>.

### References

- Abbasi, J. DNA base editing could reverse most disease-causing point mutations. *JAMA* **318**, 2173 (2017).
- Vertex and CRISPR therapeutics announce authorization of the first CRISPR/Cas9 gene-edited therapy, CASGEVY™ (exagamglogene autotemcel), by the United Kingdom MHRA for the treatment of sickle cell disease and transfusion-dependent beta thalassemia — Vertex Pharmaceuticals. <https://investors.vrtx.com/news-releases/news-release-details/vertex-and-crispr-therapeutics-announce-authorization-first>.
- Ruetz, T. J. et al. CRISPR-cas9 screens reveal regulators of ageing in neural stem cells. *Nature* **634**, 1150–1159 (2024).
- Ng, P. C. & Henikoff, S. SIFT: predicting amino acid changes that affect protein function. *Nucleic Acids Res.* **31**, 3812–3814 (2003).
- Figliuzzi, M., Jacquier, H., Schug, A., Tenaillon, O. & Weigt, M. Coevolutionary landscape inference and the context-dependence of mutations in beta-lactamase TEM-1. *Mol. Biol. Evol.* **33**, 268–280 (2016).
- Hopf, T. A. et al. Mutation effects predicted from sequence co-variation. *Nat. Biotechnol.* **35**, 128–135 (2017).
- Laine, E., Karami, Y. & Carbone, A. GEMME: a simple and fast global epistatic model predicting mutational effects. *Mol. Biol. Evol.* **36**, 2604–2619 (2019).
- Frazer, J. et al. Disease variant prediction with deep generative models of evolutionary data. *Nature* **599**, 91–95 (2021).
- Riesselman, A. J., Ingraham, J. B. & Marks, D. S. Deep generative models of genetic variation capture the effects of mutations. *Nat. Methods* **15**, 816–822 (2018).
- Marquet, C. et al. Embeddings from protein language models predict conservation and variant effects. *Hum. Genet* **141**, 1629–1647 (2022).
- Lin, Z. et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science* **379**, 1123–1130 (2023).
- Marquet, C., Schlensok, J., Abakarova, M., Rost, B. & Laine, E. Expert-guided protein language models enable accurate and blazingly fast fitness prediction. *Bioinformatics* **40**, btae621 (2024).
- Cheng, J. et al. Accurate proteome-wide missense variant effect prediction with AlphaMissense. *Science* **381**, eadg7492 (2023).
- Truong Jr, T. & Bepler, T. PoET: A generative model of protein families as sequences-of-sequences. : *Adv. Neural Inf. Process. Syst.* **36**, 77379–77415 (2023).
- Su, J. et al. SaProt: protein language modeling with structure-aware vocabulary. <https://doi.org/10.1101/2023.10.01.560349v5> (2024).
- Kircher, M. et al. A general framework for estimating the relative pathogenicity of human genetic variants. *Nat. Genet* **46**, 310–315 (2014).
- Pathak, A. K. et al. Pervasive ancestry bias in variant effect predictors. <https://www.biorxiv.org/content/10.1101/2024.05.20.594987v3> (2025).
- Tekpinar, M., David, L., Henry, T. & Carbone, A. PRESCOTT: a population aware, epistatic, and structural model accurately predicts missense effects. *Genome Biol.* **26**, 113 (2025).
- Jagota, M. et al. Cross-protein transfer learning substantially improves disease variant prediction. *Genome Biol.* **24**, 182 (2023).
- Tsishyn, M., Hermans, P., Pucci, F. & Rooman, M. Residue conservation and solvent accessibility are (almost) all you need for predicting mutational effects in proteins. <https://doi.org/10.1101/2025.02.03.636212v1> (2025).
- Tan, Y., Wang, R., Wu, B., Hong, L. & Zhou, B. Retrieval-enhanced mutation mastery: Augmenting zero-shot prediction of protein language model. <http://arxiv.org/abs/2410.21127> (2024).
- Varadi, M. et al. AlphaFold protein structure database in 2024: providing structure coverage for over 214 million protein sequences. *Nucleic Acids Res.* **52**, D368–D375 (2024).
- Livesey, B. J. & Marsh, J. A. Variant effect predictor correlation with functional assays is reflective of clinical classification performance. <https://doi.org/10.1101/2024.05.12.593741v1> (2024).
- Livesey, B. J. & Marsh, J. A. Updated benchmarking of variant effect predictors using deep mutational scanning. *Mol. Syst. Biol.* **19**, e11474 (2023).
- Notin, P. et al. ProteinGym: large-scale benchmarks for protein fitness prediction and design. *Adv. Neural Inf. Process. Syst.* **36**, 64331–64379 (2023).
- Choi, Y. & Chan, A. P. PROVEAN web server: a tool to predict the functional effect of amino acid substitutions and indels. *Bioinformatics* **31**, 2745–2747 (2015).
- Abildgaard, A. B. et al. Lynch syndrome, molecular mechanisms and variant classification. *Br. J. Cancer* **128**, 726–734 (2023).
- Cagiada, M. et al. Discovering functionally important sites in proteins. *Nat. Commun.* **14**, 4175 (2023).
- Gersing, S. et al. A comprehensive map of human glucokinase variant activity. *Genome Biol.* **24**, 97 (2023).
- Tsuboyama, K. et al. Mega-scale experimental analysis of protein folding stability in biology and design. *Nature* **620**, 434–444 (2023).
- Abakarova, M., Marquet, C., Rera, M., Rost, B. & Laine, E. Alignment-based protein mutational landscape prediction: doing more with less. *Genome Biol. Evol.* **15**, evad201 (2023).
- Mackay, T. F. C. et al. The *Drosophila melanogaster* genetic reference panel. *Nature* **482**, 173–178 (2012).
- Nunez, J. C. B. et al. Footprints of worldwide adaptation in structured populations of *Drosophila melanogaster* through the expanded DEST 2.0 genomic resource. *Mol. Biol. Evol.* **42**, msaf132 (2025).

34. Öztürk Çolak, A. et al. FlyBase: updates to the *Drosophila* genes and genomes database. *Genetics* **227**, iyad211 (2024).
35. Rubin, G. M. et al. Comparative genomics of the eukaryotes. *Science* **287**, 2204–2215 (2000).
36. Steinegger, M. & Söding, J. MMseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nat. Biotechnol.* **35**, 1026–1028 (2017).
37. Mirdita, M. et al. ColabFold: making protein folding accessible to all. *Nat. Methods* **19**, 679–682 (2022).
38. Oh, H. & Irvine, K. D. Yorkie: the final destination of hippo signaling. *Trends Cell Biol.* **20**, 410–417 (2010).
39. Brogiolo, W. et al. An evolutionarily conserved function of the *Drosophila* insulin receptor and insulin-like peptides in growth control. *Curr. Biol.* **11**, 213–221 (2001).
40. Sim, N.-L. et al. SIFT web server: predicting effects of amino acid substitutions on proteins. *Nucleic Acids Res.* **40**, W452–457 (2012).
41. Jumper, J. et al. Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583–589 (2021).
42. Piovesan, D. et al. MobiDB: intrinsically disordered proteins in 2021. *Nucleic Acids Res.* **49**, D361–D367 (2021).
43. Hu, Y. et al. iProteinDB: an integrative database of *Drosophila* post-translational modifications. *G3* **9**, 1–11 (2019).
44. Berman, H. M. et al. The protein data bank. *Nucleic Acids Res.* **28**, 235–242 (2000).
45. Mesrouze, Y. et al. Dissection of the interaction between the intrinsically disordered YAP protein and the transcription factor TEAD. *Elife* **6**, e25068 (2017).
46. Maidstone, R., Hocking, T., Rigaiil, G. & Fearnhead, P. On optimal multiple changepoint algorithms for large data. *Stat. Comput* **27**, 519–533 (2017).
47. Rigaiil, G. *fpopw: Weighted Segmentation using Functional Pruning and Optimal Partitioning*. <https://CRAN.R-project.org/package=fpopw> (2022).
48. Kumar, M. et al. ELM-the eukaryotic linear motif resource-2024 update. *Nucleic Acids Res.* **52**, D442–D455 (2024).
49. Gray, V. E. & Kumar, S. Rampant purifying selection conserves positions with posttranslational modifications in human proteins. *Mol. Biol. Evol.* **28**, 1565–1568 (2011).
50. Landry, C. R., Levy, E. D. & Michnick, S. W. Weak functional constraints on phosphoproteomes. *Trends Genet.* **25**, 193–197 (2009).
51. Kalodimou, K. et al. Separable roles for neur and ubiquitin in delta signalling in the *Drosophila* CNS lineages. *Cells* **12**, 2833 (2023).
52. Montero, L. et al. Conserved cysteine-switches for redox sensing operate in the cyclin-dependent kinase inhibitor p21(CIP/KIP) protein family. *Free Radic. Biol. Med.* **224**, 494–505 (2024).
53. Carroll, B. et al. Oxidation of SQSTM1/p62 mediates the link between redox state and protein homeostasis. *Nat. Commun.* **9**, 256 (2018).
54. Erdős, G., Mészáros, B., Reichmann, D. & Dosztányi, Z. Large-scale analysis of redox-sensitive conditionally disordered protein regions reveals their widespread nature and key roles in high-level eukaryotic processes. *Proteomics* **19**, e1800070 (2019).
55. Nguyen Ba, A. N. et al. Proteome-wide discovery of evolutionary conserved sequences in disordered regions. *Sci. Signal.* **5**, rs1 (2012).
56. Mehdiabadi, M. et al. Critical assessment of protein intrinsic disorder (caid) round 3 – predicting disorder in the era of protein language models. *Proteins*. <https://doi.org/10.1002/prot.70045> (2025).
57. Luo, H. et al. The *Drosophila* Jak kinase hopscotch is required for multiple developmental processes in the eye. *Dev. Biol.* **213**, 432–441 (1999).
58. Ochoa, D. et al. The functional landscape of the human phosphoproteome. *Nat. Biotechnol.* **38**, 365–373 (2020).
59. Davey, N. E. et al. SLiMPrints: conservation-based discovery of functional motif fingerprints in intrinsically disordered protein regions. *Nucleic Acids Res.* **40**, 10628–10641 (2012).
60. Liehrmann, A. et al. DiffSegR: an RNA-seq data driven method for differential expression analysis using changepoint detection. *NAR Genom. Bioinform.* **5**, lqad098 (2023).
61. Liehrmann, A., Rigaiil, G. & Hocking, T. D. Increased peak detection accuracy in over-dispersed ChIP-seq data with supervised segmentation models. *BMC Bioinforma.* **22**, 323 (2021).
62. Tricoire, H. et al. The steroid hormone receptor EcR finely modulates *Drosophila* lifespan during adulthood in a sex-specific manner. *Mech. Ageing Dev.* **130**, 547–552 (2009).
63. Rera, M., Monnier, V. & Tricoire, H. Mitochondrial electron transport chain dysfunction during development does not extend lifespan in *Drosophila melanogaster*. *Mech. Ageing Dev.* **131**, 156–164 (2010).
64. Hand, D. J. Review of mixture models: Inference and applications to clustering. *J. R. Stat. Soc. C* **38**, 384–385 (1989).
65. Suzek, B. E. et al. UniRef clusters: a comprehensive and scalable alternative for improving sequence similarity searches. *Bioinformatics* **31**, 926–932 (2015).
66. Steinegger, M. et al. HH-suite3 for fast remote homology detection and deep protein annotation. *BMC Bioinform.* **20**, 473 (2019).
67. Engelen, S., Trojan, L. A., Sacquin-Mora, S., Lavery, R. & Carbone, A. Joint evolutionary trees: a large-scale method to predict protein interfaces based on sequence sampling. *PLoS Comput. Biol.* **5**, e1000267 (2009).
68. Notin, P. et al. TranceptEVE: combining family-specific and family-agnostic models of protein sequences for improved fitness prediction. <https://doi.org/10.1101/2022.12.07.519495v2> (2022).
69. Dempster, A. P., Laird, N. M. & Rubin, D. B. Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. B (Methodological)* **39**, 1–38 (1977).
70. Pedregosa, F. et al. Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
71. requests: Python HTTP for humans. <http://python-requests.org>.
72. HALL, P., KAY, J. W. & TITTERINGTON, D. M. Asymptotically optimal difference-based estimation of variance in nonparametric regression. *Biometrika* **77**, 521–528 (1990).
73. Wojdyr, M. GEMMI: a library for structural biology. *J. Open Source Softw.* **7**, 4200 (2022).
74. Sehnał, D. et al. Mol\* viewer: modern web app for 3d visualization and analysis of large biomolecular structures. *Nucleic Acids Res.* <https://academic.oup.com/nar/article/49/W1/W431/6270780?login=false> (2021).
75. Herrmann, C., Van de Sande, B., Potier, D. & Aerts, S. i-cisTarget: an integrative genomics method for the prediction of regulatory features and cis-regulatory modules. *Nucleic Acids Res.* **40**, e114 (2012).
76. Imrichová, H., Hulselmans, G., Kalender Atak, Z., Potier, D. & Aerts, S. i-cisTarget 2015 update: generalized cis-regulatory enrichment analysis in human, mouse and fly. *Nucleic Acids Res.* **43**, W57–W64 (2015).
77. DeLano, W. L. The PyMOL molecular graphics system. <https://pymol.org> (2002).
78. Marletta, A. S. et al. Crystal structure of human nicotinic acid phosphoribosyltransferase. *FEBS Open Bio* **5**, 419–428 (2015).
79. Abakarova, M., Freiberger, M. I., Liehrmann, A., Rera, M. & Laine, E. Proteome-wide prediction of the functional impact of missense variants with proteocast. <https://doi.org/10.5281/zenodo.18488840> (2026).
80. Abakarova, M., Freiberger, M. I., Liehrmann, A., Rera, M. & Laine, E. ProteoCast: Proteome-wide variant effect prediction web server. <https://proteocast.ijm.fr/> (2026).
81. Abakarova, M., Freiberger, M. I., Liehrmann, A., Rera, M. & Laine, E. Proteome-wide prediction of the functional impact of missense

- variants with ProteoCast. <https://github.com/abakarovaMarina/ProteoCast>. Available at <https://github.com/abakarovaMarina/ProteoCast> (2026).
82. Landrum, M. J. et al. ClinVar: public archive of interpretations of clinically relevant variants. *Nucleic Acids Res.* **44**, D862–D868 (2016).
83. Valanne, S., Vesala, L., Maasdorp, M. K., Salminen, T. S. & Rämetsä, M. The *Drosophila* Toll pathway in innate immunity: from the core pathway toward effector functions. *J. Immunol.* **209**, 1817–1825 (2022).

## Acknowledgements

This work was granted access to the HPC resources of the SACADO MeSU platform at Sorbonne Université. Further, thanks to all contributing to open-source programming libraries. Last, not least, thanks to all who deposit experimental data in public databases, to those who maintain these databases, and those who make methods available enriching experimental data. This work has been funded by ADAGIO - Ageing and Natural DeATH Genetic cOntrollers (ANR-20-CE44-0010) and co-funded by the European Union (ERC, PROMISE, 101087830). Views and opinions expressed are however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. We thank Joël Marchand for providing support with the online database and web platform.

## Author contributions

E.L. and M.R. conceptualised the project. M.A., E.L., and M.R. designed the computational framework. M.A. implemented the ProteoCast framework, performed all analyses and data curation. A.L. implemented the segmentation algorithm. M.A., E.L., M.R., and A.L. analysed the results. M.A. and M.I.F. developed the web platform. M.A., E.L., and M.R. wrote the manuscript. All authors reviewed and approved the final manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41467-026-72140-1>.

**Correspondence** and requests for materials should be addressed to Michael Rera or Elodie Laine.

**Peer review information** *Nature Communications* thanks the anonymous reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026