AlphaRING: Missense Variant Classification via Protein Modelling and Residue Interaction Network Analysis

Aaron Logsdon[1], Roghaiyeh Safari[1,2], Anna Barnard[3], Goedele Maertens[1*], Evangelos Bellos[1,4*]
[1]Department of Infectious Disease, Imperial College London, London, United Kingdom
[2]RS & RS Scientific Ltd, Sandwich, United Kingdom
[3]Department of Chemistry, Imperial College London, London, United Kingdom
[4]School of Human Development and Health, University of Southampton, Southampton, United Kingdom
[*]evangelos.bellos09@imperial.ac.uk; g.maertens@imperial.ac.uk

Abstract

Recent advancements in protein modelling, particularly with AlphaFold, have opened new avenues for assessing the structural impacts of missense variants on protein function. Tools like AlphaMissense leverage AlphaFold2 to predict missense variant pathogenicity with state-of-the-art performance. However, AlphaMissense is deep learning-based, and its process to predict pathogenicity is opaque and does not explicitly predict the impact of missense variants based on additional biophysical properties. Additionally, AlphaMissense is not publicly available, limiting accessibility. Here, we present AlphaRING, a publicly available tool that leverages AlphaFold2 and captures real physico-chemical properties of variant residues to predict the pathogenicity of both characterised and uncharacterised missense variants. AlphaRING models wild-type and variant proteins using AlphaFold2 and identifies atomic-level non-covalent interactions via residue interaction network analysis using RING4. Unlike deep learning models, our scoring captures real physico-chemical properties of mutated residues, providing interpretable insights into mutation effects. We developed novel formulas for each non-covalent bond type (hydrogen, ionic, π-cation, π-π stacking, and π-hydrogen), incorporating bond-specific energy and geometry values to calculate unique weighting scores per residue. The absolute $\log_2$ fold change between wild-type and variant residue scores yields a non-negative AlphaRING score—the higher the score, the greater the predicted pathogenicity. Evaluating AlphaRING on over 1,300 human gold-standard missense variants curated by ClinVar expert panels, our results show effective distinction between benign and pathogenic variants, with a receiver operating characteristic area under the curve of 0.79. This strong performance suggests that AlphaRING is a reliable, open-source tool suitable for high-throughput structural predictions and large-scale variant characterisation, accessible to non-experts. By enabling detailed exploration of the relationship between genetic variation and protein structure, AlphaRING addresses key limitations of previous models and paves the way for broader applications in genetic research and precision medicine.

Keywords

Missense variants, Variant classification, Pathogenicity prediction, AlphaFold2, RING4, AlphaRING, Protein modelling, Residue interaction network analysis, Non-covalent bonds, ClinVar, Computational biology.

Background

Missense variants are among the most prevalent genetic alterations impacting human health, contributing to a wide variety of diseases, including cancer, cardiovascular disorders, and neurodegenerative conditions. Predicting the pathogenicity of these variants is essential for analysing and prioritising them in both research and clinical settings. The functional consequences of missense variants often hinge on their location within key protein regions, such as active sites, binding domains, or structural motifs, which are essential for proper protein function. As a result, these variants are frequently enriched in functionally critical areas of proteins (Sun et al., 2024), making it crucial to consider their impact on protein structure in pathogenicity prediction. Recent advancements in protein modelling have created new opportunities for assessing how these structural disruptions can contribute to disease.

Among these advancements is AlphaMissense (Frazer et al., 2023), which is built upon the revolutionary protein modelling tool, AlphaFold2 (Jumper et al., 2021), fine-tuning the deep learning-based model on human proteins with an additional variant pathogenicity classification objective. On annotated missense variants from ClinVar (Landrum et al., 2018), AlphaMissense classified 32% as likely pathogenic and 57% as likely benign using a cutoff yielding 90% precision (Frazer et al., 2023). The state-of-the-art performance by AlphaMissense delineates the power of using AlphaFold2 technology to predict the pathogenicity of missense variants. However, AlphaMissense does not predict differences in structural coordinates for proteins harbouring missense substitutions. Rather, wild-type structural predictions are used to guide the AlphaMissense model. Additionally, being deep learning-based, AlphaMissense's process to predict pathogenicity is opaque and does not predict the impact of missense variants explicitly on more biophysical properties such as protein stability. Finally, to date, the AlphaMissense model has not been publicly released, impeding researchers and clinicians from leveraging AlphaFold2 to evaluate the pathogenicity of new missense variants.

To address these issues, we developed AlphaRING, a publicly available tool that leverages the AlphaFold2 model and captures real physico-chemical properties of mutated residues to predict the pathogenicity of characterised and uncharacterised missense variants alike. For any given missense variant, AlphaRING uses AlphaFold2 to predict the structure of the variant and wild-type protein. Next, in lieu of the AlphaMissense model, AlphaRING rapidly captures the changes in protein structure by using RING4 (Del Conte et al., 2024), a residue interaction network (RIN) analysis (del Sol & Carbonell, 2007) tool that uses distance, geometric constraints, and other physico-chemical properties to calculate the bonds formed between residues at the atomic level. Finally, for each non-covalent bond, novel formulas incorporating bond-specific values for energy and geometry are applied to calculate a weighting score per residue. From here, the absolute log2 of the fold change between the weighting scores of the wild-type and variant residues can be readily calculated, generating the final non-negative AlphaRING score. The higher the score, the greater the degree of predicted pathogenicity.

To evaluate AlphaRING's ability to predict missense variant pathogenicity, we benchmarked it against over 1,300 human gold-standard missense variants curated by expert panels in ClinVar

(Landrum et al., 2018). Our results demonstrate that AlphaRING effectively distinguishes between benign and pathogenic variants, validating its utility as a classifier. By providing a publicly available tool that leverages protein modelling and RIN analysis, AlphaRING empowers researchers and clinicians to analyse and prioritise missense variants. This advancement sets the stage for large-scale characterisation of the relationship between genetic variation and protein structure.

AlphaRING

A significant challenge in missense variant classification is the limited availability of cutting-edge tools. Although AlphaMissense (Frazer et al., 2023) has demonstrated exceptional performance in predicting variant pathogenicity using protein structural modelling through AlphaFold2 (Jumper et al., 2021), it does not provide structural coordinate predictions for proteins with missense mutations. Instead, it leverages wild-type structural predictions to inform its model. Due to its deep learning foundation, AlphaMissense's pathogenicity predictions are somewhat opaque, lacking explicit assessments of how missense variants may affect biophysical properties such as protein stability. Additionally, the AlphaMissense model remains unreleased, preventing researchers and clinicians from using AlphaFold2 to assess the pathogenic potential of new missense variants. AlphaRING addresses this challenge by offering an accessible, publicly available pipeline (Fig. 1) for missense variant classification, leveraging the structural modelling capabilities of AlphaFold2 and capturing real physico-chemical properties of variant residues. By integrating multiple layers of structural and interaction-based analysis, AlphaRING ensures a comprehensive assessment of each variant's impact.

To classify the pathogenicity of a missense variant, AlphaRING begins by inputting both the wild-type and variant protein sequences in FASTA format into AlphaFold2. AlphaFold2 predicts the three-dimensional structures of both proteins, considering evolutionary structural constraints to ensure accuracy. Once both models are generated, AlphaRING applies RIN analysis using RING4 (Del Conte et al., 2024) to capture non-covalent interactions, which are essential for maintaining protein stability and functionality as they mediate critical physical properties of the protein. By analysing how a missense variant alters these interactions, AlphaRING can assess its potential contribution to pathogenicity. Even small changes in bond geometry and energy can lead to significant changes in protein behaviour, emphasising the importance of this detailed interaction analysis.

Following the identification of non-covalent bonds, AlphaRING assigns each residue a weighting score using novel bond-specific (hydrogen, ionic, π-cation, π-π stacking, and π-hydrogen) formulas that consider bond type, energy, and geometry. The weightings of all bonds associated with each residue are summed to provide a per-residue weighting. The degree of disruption between the wild-type and variant residues is then quantified by calculating the fold change (FC) between their respective weightings. Pathogenic variants are expected to exhibit fold changes significantly deviating from 1, reflecting their ability to destabilise key structural interactions, while benign variants should have FC values closer to 1, indicating minimal disruption.

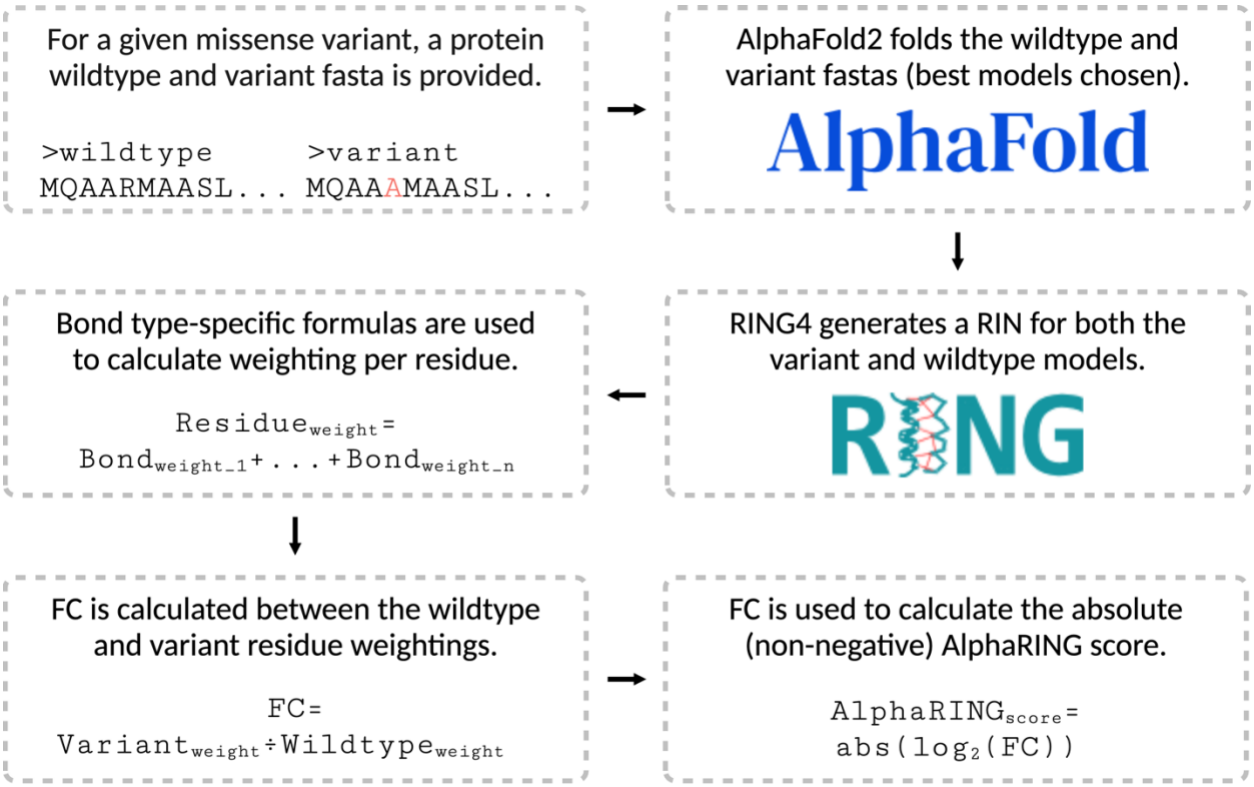

Fig. 1. Flowchart of the AlphaRING pipeline for missense variant classification. For any given missense variant, the input consists of the wild-type and variant protein sequences, both in FASTA format. AlphaFold2 (Jumper et al., 2021) models the structures of both proteins, followed by residue interaction network (RIN) analysis using RING4 (Del Conte et al., 2024). Next, non-covalent bonds (hydrogen, ionic, π-cation, π-π, and π-hydrogen) are assigned non-negative weights using bond-specific formulas, which are summed per residue to calculate residue-level weighting. Finally, the fold change (FC) between the variant and wild-type residue weightings is calculated, with the absolute $\log_2$ taken to yield a non-negative AlphaRING score. The higher the AlphaRING score, the greater the predicted pathogenicity.

Finally, a non-negative AlphaRING score is computed by taking the absolute $\log_2$ of the fold change. Variants causing minimal structural disruptions will have AlphaRING scores close to 0, while those causing more substantial changes will result in higher scores. This scoring approach provides a clear and interpretable measure, with the intention that the more pathogenic a variant is, the greater the AlphaRING score. This score offers a quantitative and scalable method to assess missense variant pathogenicity, effectively filling the gap left by the inaccessibility of AlphaMissense and providing researchers and clinicians with a practical, high-throughput solution for missense variant classification. The availability of such a tool has the potential to advance our understanding of genetic variation and its implications for human health.

Benchmarking

To determine AlphaRING's performance as a classifier of missense variants, we aimed to rigorously validate it against a gold-standard benchmark dataset. To achieve this, we devised a strategy to curate our own dataset of missense variants (Fig. 2). We began by collecting a comprehensive set of human missense variants from ClinVar (Landrum et al., 2018), focusing

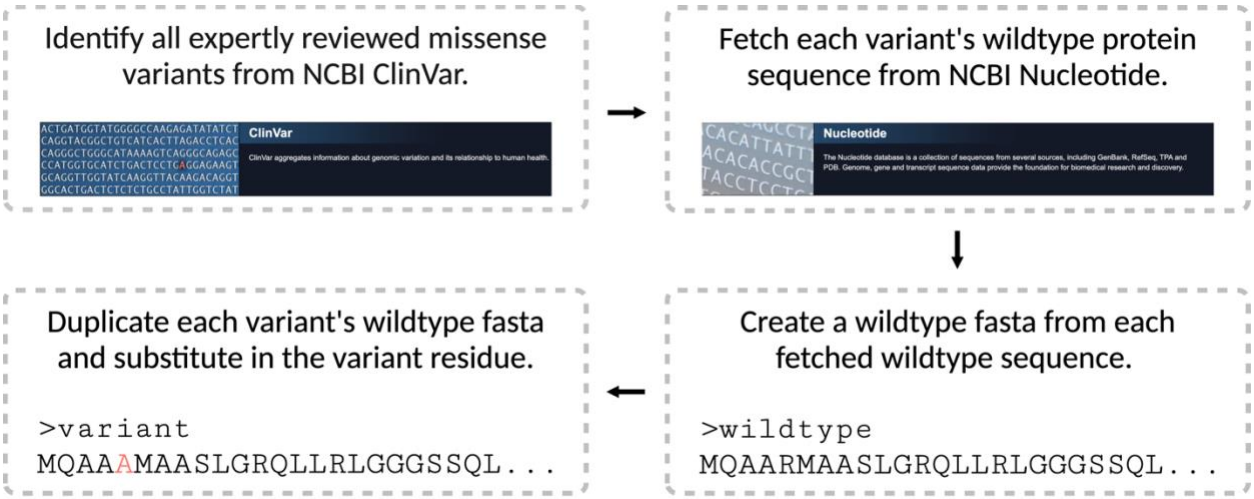

Fig. 2. Flowchart of the curation process for the AlphaRING benchmark dataset. Expertly reviewed benign and pathogenic missense variants were identified through NCBI ClinVar (Landrum et al., 2018), and their wild-type protein sequences were retrieved from NCBI Nucleotide (Sayers et al., 2024). Each wild-type sequence was then used to create a wild-type FASTA. For each variant, the wild-type FASTA was duplicated, with the variant residue substituted in, completing the dataset.

specifically on those reviewed by expert panels to ensure high-quality and accurate classifications, narrowing down to those definitively labelled as either benign or pathogenic. Next, we used Entrez Direct (Kans, 2013) to retrieve each variant's wild-type protein sequence from NCBI Nucleotide (Sayers et al., 2024), which was converted to FASTA format. Finally, each benign and pathogenic variant's wild-type FASTA was duplicated to substitute in the variant residue, creating over 1,300 wild-type-variant pairs to be processed by AlphaRING. Additionally, alongside AlphaRING itself, we have made this dataset and its AlphaRING results available to the research community, along with code that automates the generation of the dataset, allowing researchers to easily replicate or expand on our benchmarking strategy.

After processing the benchmark dataset through AlphaRING and obtaining AlphaRING scores for over 1,300 variants, we performed an exploratory analysis of the score distributions for benign and pathogenic variants. Given the complex nature of missense variant impacts, the aim was to determine whether AlphaRING's scoring mechanism could identify statistically significant separation between the two variant classes. The resulting distribution revealed a strong distinction, with pathogenic variants typically displaying higher AlphaRING scores than benign ones. A Mann-Whitney U test confirmed this difference with a highly significant result ($p < 2.2\times10^{-16}$; Fig. 3), reinforcing AlphaRING's potential as a classifier. This significant separation not only reflects the model's sensitivity to structural disruptions but also provided us with the statistical evidence to proceed with more formal classification assessments. Establishing this clear, statistically significant distinction between classes marks an important milestone in verifying AlphaRING's predictive reliability.

To further quantify AlphaRING's classification ability, we conducted a receiver operating characteristic (ROC) analysis. The ROC curve provided a comprehensive view of AlphaRING's performance across a range of score thresholds, enabling us to evaluate its effectiveness in

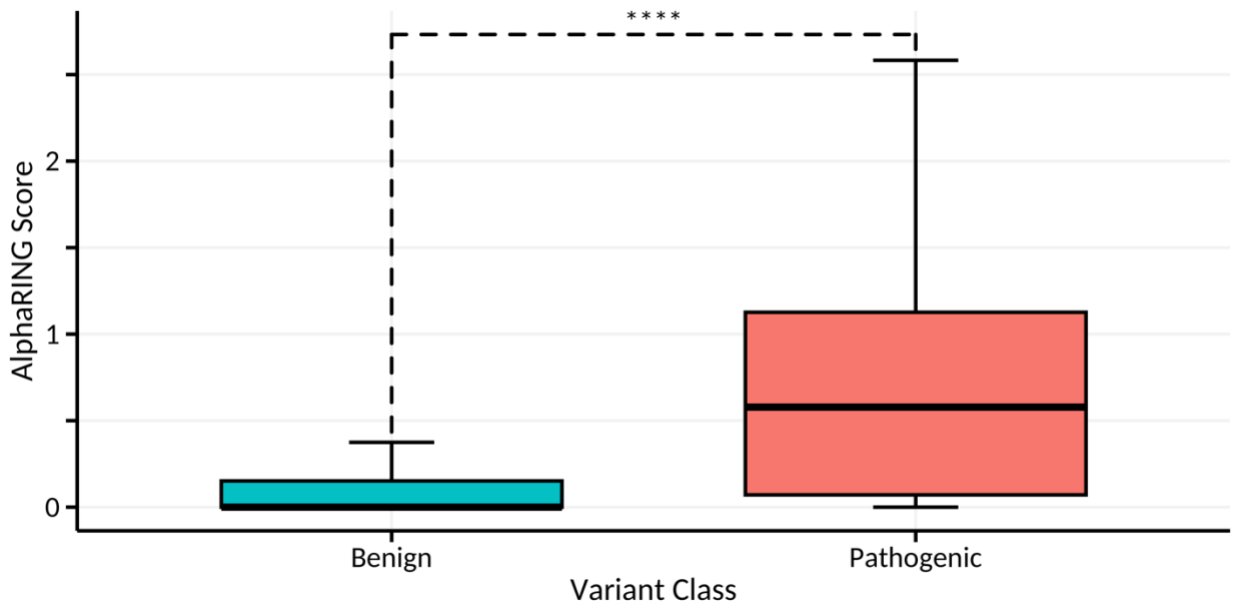

Fig. 3. Box plot representation of the distribution difference in AlphaRING scores for human benign and pathogenic missense variants. Missense variants—benign (n=451) and pathogenic (n=871)—were identified using ClinVar (Landrum et al., 2018) and filtered to retain only those reviewed by an expert panel. For each variant, AlphaRING was used to calculate an AlphaRING score. A Mann-Whitney U test was performed to assess whether the distribution difference in AlphaRING scores between benign and pathogenic variants was statistically significant. Asterisks indicate statistically significant differences: (****) $p < 2.2\times10^{-16}$.

distinguishing pathogenic from benign variants. With an area under the curve (AUC) of 0.79 (Fig. 4), AlphaRING demonstrated robust classification capabilities, offering a high level of discriminatory power. The analysis also allowed us to identify an optimal cut-off AlphaRING score of 0.014, derived using Youden's Index, which maximised both sensitivity and specificity. This optimal threshold provides a clear, actionable score that can be consistently applied for variant classification across diverse datasets. The AUC score, coupled with the identified cut-off point, validates AlphaRING's practical applicability as a classification tool, reinforcing its role as an accessible, open-source resource in missense variant interpretation. By quantifying AlphaRING's classifier strength, we affirm its potential utility in both research and clinical settings, supporting informed variant prioritisation and fostering advancements in precision medicine.

Conclusions

In conclusion, AlphaRING offers a powerful, publicly accessible tool for the classification of missense variants through cutting-edge protein modelling and RIN analysis. By leveraging AlphaFold2 (Jumper et al., 2021) and RING4 (Del Conte et al., 2024), AlphaRING enables a detailed examination of the structural impacts of missense variants on protein function. Our benchmarking against over 1,300 expertly curated missense variants from ClinVar (Landrum et al., 2018) demonstrates the tool's efficacy, with a ROC analysis yielding an AUC score of 0.79. Importantly, Youden's Index identified an optimal cut-off point of 0.014, providing a clear threshold for distinguishing between pathogenic and benign variants.

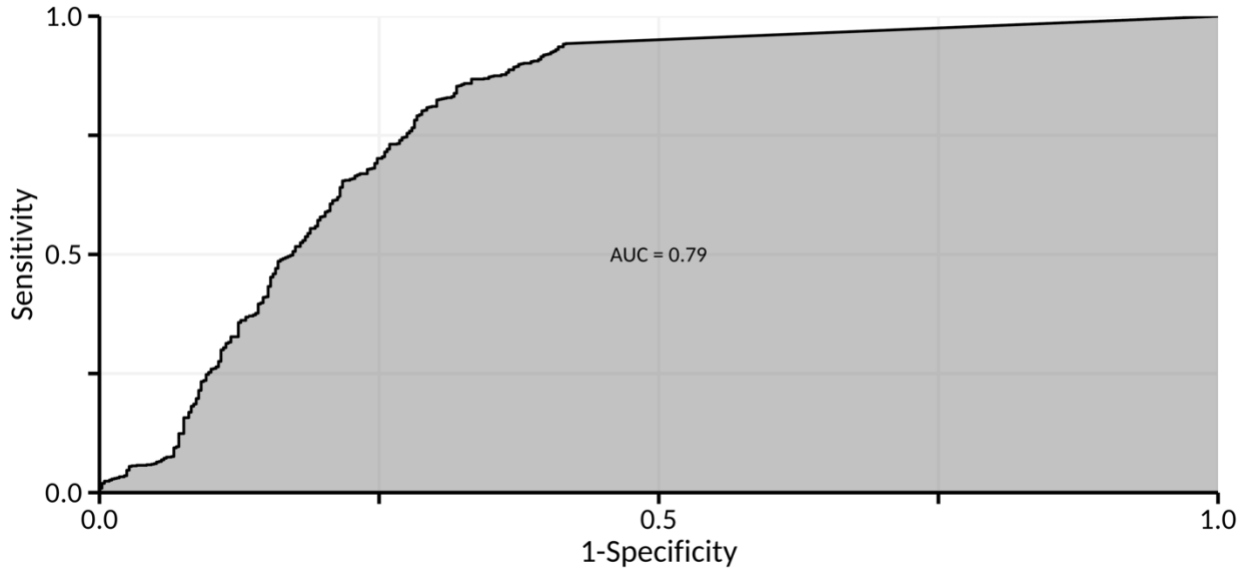

Fig. 4. Receiver operating characteristic (ROC) curve for the classification of benign and pathogenic missense variants by AlphaRING. Sensitivity (true positive rate) is plotted against 1-Specificity (false positive rate) across varying AlphaRING score thresholds. The shaded area represents the area under the curve (AUC), with the AUC score indicated within this region.

A key contribution of this work is the development and release of a gold-standard benchmark dataset, made available as a community resource. This dataset, alongside the AlphaRING platform, equips researchers to undertake large-scale missense variant classification with ease and accuracy. The open availability of these resources is intended to accelerate research in genetic variant interpretation and foster collaboration in the field.

Looking ahead, AlphaRING holds immense potential for diverse applications. In clinical genomics, it can be employed to prioritise variants for further functional validation or guide therapeutic decisions in precision medicine. In research, AlphaRING can facilitate the discovery of novel disease mechanisms by exploring the structural consequences of genetic variation across different proteins. Although currently focused on assessing the effects of variants in isolation, the framework is easily generalisable to account for protein-protein and protein-ligand interactions, providing a richer context for understanding the complex interplay of biomolecular relationships. As the field of computational biology advances, AlphaRING could be further integrated with machine learning frameworks and multi-omics data, enhancing its predictive power and broadening its scope in personalised healthcare and beyond.

By providing a scalable, open-source solution for high-throughput variant characterisation, AlphaRING is poised to play a pivotal role in advancing our understanding of genetic variation and its implications for human health.

Data Availability

All the code and data curated or generated during the course of this study are available at https://github.com/loggy01/alpharing.

Funding

A. L. was supported by Imperial College London (President's PhD Scholarship). G. M. and R. S. were supported by the Medical Research Council (MR/W00206X/1 to G. M.). A. B. was supported by a Sir Henry Dale Fellowship jointly funded by the Wellcome Trust and the Royal Society (213435/Z/18/Z).

Conflict of interest statement: None declared.

References

Cheng, J., Novati, G., Pan, J., Bycroft, C., Žemgulytė, A., Applebaum, T., Pritzel, A., Wong, L. H., Zielinski, M., Sargeant, T., Schneider, R. G., Senior, A. W., Jumper, J., Hassabis, D., Kohli, P. & Avsec, Ž. (2023) Accurate proteome-wide missense variant effect prediction with AlphaMissense. *Science*. 381 (6664), eadg7492. https://doi.org/10.1126/science.adg7492.

Del Conte, A., Camagni, G. F., Clementel, D., Minervini, G., Monzon, A. M., Ferrari, C., Piovesan, D. & Tosatto, S. C. E. (2024) RING 4.0: faster residue interaction networks with novel interaction types across over 35,000 different chemical structures. *Nucleic Acids Research*. 52, W306–W312. https://doi.org/10.1093/nar/gkae337.

del Sol, A. & Carbonell, P. (2007) The Modular Organization of Domain Structures: Insights into Protein–Protein Binding. *PLOS Computational Biology*. 3 (12), e239. https://doi.org/10.1371/journal.pcbi.0030239.

Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., Back, T., Petersen, S., Reiman, D., Clancy, E., Zielinski, M., Steinegger, M., Pacholska, M., Berghammer, T., Bodenstein, S., Silver, D., Vinyals, O., Senior, A. W., Kavukcuoglu, K., Kohli, P. & Hassabis, D. (2021) Highly accurate protein structure prediction with AlphaFold. *Nature*. 596 (7873), 583–589. https://doi.org/10.1038/s41586-021-03819-2.

Kans J. Entrez Direct: E-utilities on the Unix Command Line. 2013 Apr 23 [Updated 2024 Oct 21]. In: Entrez Programming Utilities Help [Internet]. Bethesda (MD): National Center for Biotechnology Information (US); 2010-. Available from: https://www.ncbi.nlm.nih.gov/books/NBK179288/

Landrum, M. J., Lee, J. M., Benson, M., Brown, G. R., Chao, C., Chitipiralla, S., Gu, B., Hart, J., Hoffman, D., Jang, W., Karapetyan, K., Katz, K., Liu, C., Maddipatla, Z., Malheiro, A., McDaniel, K., Ovetsky, M., Riley, G., Zhou, G., Holmes, J. B., Kattman, B. L. & Maglott, D. R. (2018) ClinVar: improving access to variant interpretations and supporting evidence. *Nucleic Acids Research*. 46, D1062–D1067. https://doi.org/10.1093/nar/gkx1153.

Sayers, E. W., Beck, J., Bolton, E. E., Brister, J. R., Chan, J., Comeau, D. C., Connor, R., DiCuccio, M., Farrell, C. M., Feldgarden, M., Fine, A. M., Funk, K., Hatcher, E., Hoeppner, M., Kane, M., Kannan, S., Katz, K. S., Kelly, C., Klimke, W., Kim, S., Kimchi, A., Landrum, M., Lathrop, S., Lu, Z., Malheiro, A., Marchler-Bauer, A., Murphy, T. D., Phan, L., Prasad, A. B., Pujar, S., Sawyer, A., Schmieder, E., Schneider, V. A., Schoch, C. L., Sharma, S., Thibaud-Nissen, F., Trawick, B. W., Venkatapathi, T., Wang, J., Pruitt, K. D. & Sherry, S. T. (2024) Database resources of the National Center for Biotechnology Information. *Nucleic Acids Research*. 52 (D1), D33–D43. https://doi.org/10.1093/nar/gkad1044.

Sun, K. Y., Bai, X., Chen, S., Bao, S., Zhang, C., Kapoor, M., Backman, J., Joseph, T., Maxwell, E., Mitra, G., Gorovits, A., Mansfield, A., Boutkov, B., Gokhale, S., Habegger, L., Marcketta, A., Locke, A. E., Ganel, L., Hawes, A., Kessler, M. D., Sharma, D., Staples, J., Bovijn, J., Gelfman, S., Di Gioia, A., Rajagopal, V. M., Lopez, A., Varela, J. R., Alegre-Díaz, J., Berumen, J., Tapia-Conyer, R., Kuri-Morales, P., Torres, J., Emberson, J., Collins, R., Regeneron Genetics Center, RGC-ME Cohort Partners, Cantor, M., Thornton, T., Kang, H. M., Overton, J. D., Shuldiner, A. R., Cremona, M. L., Nafde, M., Baras, A., Abecasis, G., Marchini, J., Reid, J. G., Salerno, W., & Balasubramanian, S. (2024). A deep catalogue of protein-coding variation in 983,578 individuals. *Nature*. 631 (8021), 583–592. https://doi.org/10.1038/s41586-024-07556-0.