

SHINE: protein language model-based pathogenicity prediction for short inframe insertion and deletion variants

Xiao Fan, Hongbing Pan, Alan Tian, Wendy K. Chung and Yufeng Shen

Corresponding authors: Xiao Fan, E-mail: xf2193@cumc.columbia.edu; Wendy K. Chung, E-mail: wkc15@cumc.columbia.edu and Yufeng Shen, E-mail: ys2411@cumc.columbia.edu

Abstract

Accurate variant pathogenicity predictions are important in genetic studies of human diseases. Inframe insertion and deletion variants (indels) alter protein sequence and length, but not as deleterious as frameshift indels. Inframe indel interpretation is challenging due to limitations in the available number of known pathogenic variants for training. Existing prediction methods largely use manually encoded features including conservation, protein structure and function, and allele frequency to infer variant pathogenicity. Recent advances in deep learning modeling of protein sequences and structures provide an opportunity to improve the representation of salient features based on large numbers of protein sequences. We developed a new pathogenicity predictor for SHort Inframe iNsertion and dEletion (SHINE). SHINE uses pretrained protein language models to construct a latent representation of an indel and its protein context from protein sequences and multiple protein sequence alignments, and feeds the latent representation into supervised machine learning models for pathogenicity prediction. We curated training data from ClinVar and gnomAD, and created two test datasets from different sources. SHINE achieved better prediction performance than existing methods for both deletion and insertion variants in these two test datasets. Our work suggests that unsupervised protein language models can provide valuable information about proteins, and new methods based on these models can improve variant interpretation in genetic analyses.

Keywords: inframe indel, variant pathogenicity, transformer, protein language model, transfer learning

Introduction

Inframe insertion and deletion variants (indels) are abundant but are under studied in genetic analyses. In the ~500 000 UK biobank whole exome sequencing dataset, median numbers of the inframe and frameshift variants per individual were 115 and 90, respectively [1]. A recent deep mutational scan study of DDX3X showed 49% of 625 inframe deletions caused cell depletion [2], potentially causal for a rare disease—DDX3X syndrome characterized by motor and language delays and autism spectrum disorder. Although many inframe indels contribute to human diseases [3–5], the impact of rare individual inframe indels is usually uncertain. In ClinVar [6], 64% and 18% of ~12 000 inframe indels are reported as variants of uncertain significance (VUS) and pathogenic/likely pathogenic (P/LP), compared with 12 and 86% of 53 000 frameshift indels are VUS and P/LP [7]. This issue is due to the difficulty in determining pathogenicity of inframe indels: there are fewer clinical cases with pathogenic inframe indels. Accurate pathogenicity predictions for inframe indels can facilitate return of results in genetic testing, refine variant

prioritization for downstream analyses in genetic research study and strengthen the computational evidence used in the American College of Medical Genetics and Genomics and Association for Molecular Pathology guidelines.

Several computational methods were developed to predict pathogenicity for inframe indels [8–16], but the accuracy is limited by the availability of training data. There are only 2208 P/LP and 1599 benign/likely benign (B/LB) inframe indels with two stars (multiple submitters with assertion criteria and no conflicting interpretation) in ClinVar [6]. In addition, nearly all existing computational methods use machine learning with manually encoded predictive features, such as DNA/protein conservation, protein structure and function, occurrence in a repeat region, size of the indels, local amino acid composition, distance to the nearest splice site and allele frequency. Although these features are correlated with variant pathogenicity, they are based on our limited understanding of how variants exert their function. Recent advances in deep learning, particularly protein language models [17, 18, 25], enable us to study latent

Xiao Fan, PhD is an Associate Research Scientist in the Department of Pediatrics and the Department of Systems Biology. She studies computational and statistical methods to understand genetic architecture of rare human diseases and variant effects.

Hongbing Pan, BS is a Master student in the Department of Biomedical Informatics. He studies computational genomics and works on applying machine learning techniques in the analysis of genomic data.

Alan Tian is a high school student at Lynbrook High School. He was a summer visiting student at Dr. Shen's lab in 2021 and 2022.

Wendy Chung, MD, PhD is the Kennedy Family Professor of Pediatrics in Medicine at Columbia University and Chief of the Division of Clinical Genetics in the Department of Pediatrics. Her research relates to the genetic basis of a variety of human diseases including obesity, type 2 diabetes, congenital heart disease, cardiomyopathies, arrhythmias, Long QT syndrome, pulmonary hypertension, endocrinopathies, congenital diaphragmatic hernias, cleft lip/cleft palate, seizures, intellectual disabilities, autism, inherited metabolic conditions, rare disorders, and breast and pancreatic cancer susceptibility.

Yufeng Shen, PhD is an Associate Professor in the Departments of Systems Biology and the Department of Biomedical Informatics and Associate Director of Columbia Genome Center. He studies human biology and disease using genomic and computational approaches. His research group is developing new methods to identify genetic causes of human diseases and to understand the dynamics of the adaptive immune system.

Received: July 18, 2022. **Revised:** November 4, 2022. **Accepted:** November 29, 2022

© The Author(s) 2022. Published by Oxford University Press. All rights reserved. For Permissions, please email: journals.permissions@oup.com

representations for protein sequences that can potentially capture salient information not encoded by existing methods. This opens up an opportunity to improve the input salient features used in protein-related bioinformatic applications.

Here we present SHINE (SHort Inframe iNsertion and dEletion), a transfer learning model to leverage protein language models and limited available pathogenicity data for inframe indels. The protein language models take protein sequences or multiple protein sequence alignments (MSA), and generate latent representations capturing protein features. Previous studies have shown the linear projections of the representations generated from the protein language models encode information about protein secondary and tertiary structure [17]. The same group used zero-shot and few-shot transfer of the protein language models to predict variant effects [19]. SHINE uses the representations as features to separate pathogenic from benign inframe indels. It is trained on curated inframe indels from ClinVar, and short and common indels in gnomAD; and compared with the other computational methods on two independent test datasets. SHINE achieved better prediction performance than existing methods for both deletion and insertion variants in these two test datasets. Finally, we interpreted the salient features from SHINE by comparing them with known predictive protein features, such as protein secondary structure, intrinsically disordered protein regions and relative solvent accessibility. Our work suggests that unsupervised protein language models can provide valuable information about proteins, and new methods based on these models can improve variant interpretation in genetic analyses.

Materials and methods

Data sets

We obtained P/LP and B/LB inframe indels from ClinVar [6] and gnomAD, and excluded inframe indels that have a length greater than three amino acids or allele frequency $< 0.1\%$ in gnomAD. Seventy-five inframe indels are overlapped between ClinVar and gnomAD, and only one is P/LP, suggesting it is reasonable to use gnomAD variants as B/LB. We divided this dataset into training and validation datasets based on the number of deleted/inserted amino acids. Indels with one-amino-acid deletion/insertion are used for training and the rest are for validation. The training dataset includes 1040 pathogenic and 1111 benign deletions, and 142 pathogenic and 537 benign insertions. The validation dataset includes 640 pathogenic and 896 benign deletions, and 272 pathogenic and 662 benign insertions. These datasets are used to select the optimal number of features, prediction models and parameters for the models.

To evaluate the performance, we created two independent datasets for neurodevelopmental disorder (NDD) genes and cancer driver genes. Overlapping indels with the training-validation datasets were removed from the test datasets. The two test datasets are from different sources than ClinVar, which avoids inflated performance when testing on data from the same source. Besides, the two datasets were not used to train any of the previous prediction models for indels, so are good benchmarks for method comparison. Six hundred and eighty-six NDD risk genes were collected based on previous studies [20, 21]. *De novo* inframe indels in the 686 NDD genes in two NDD cohorts were used as proxies of pathogenic variants, and short inframe indels with at most three-amino-acid deletion/insertion in the UK biobank (200 000 exomes) were considered benign. One NDD cohort includes 16 877 trios in Simons Powering Autism Research (SPARK; [20, 22, 23]) and the other includes 13 058 trios




Figure 1. Architecture of the SHINE.

with developmental disorders [21]. This test dataset includes 146 pathogenic and 2808 benign deletions, and 35 pathogenic and 1504 benign insertions. The second test dataset is composed of inframe indels discovered in >1000 cancer mutational hotspots [24]. We identified 307 deletions and 119 insertions in 36 genes as proxies of pathogenic variants. We extracted short inframe indels including 132 deletions and 54 insertions from the UK biobank in the same 36 genes as benign. To be noted, the variants from both cases (NDD cohorts and cancer mutational hotspots) and controls (UK biobank) are mixtures of pathogenic and benign variants, but the *de novo*/somatic variants in NDD risk genes/cancer mutational hotspots identified in cases are much more likely to be pathogenic compared with germline variants in controls. We used the case/control status as a proxy of pathogenic/benign labels in the test.

Model architecture

SHINE uses a transfer learning architecture (Figure 1), leveraging pretrained protein language models and limited available pathogenicity labels for inframe indels. We used two protein language models: ESM-1b [17] and MSA transformers [25]. The two transformers aim to learn biological properties from protein sequences or MSAs with unsupervised learning or without prior knowledge. The ESM-1b transformer was trained on 250 million protein sequences [17]. The MSA transformer was trained on 26 million MSAs [25]. Both generated latent representations containing information about biological properties of input proteins. We elaborate on how we selected proper transformers, ways to handle multiple-amino-acid indels, input dimensions (number of features), prediction models and their parameters. As the latent representations from the transformers are high dimensional and correlated, we first performed feature reduction using principal components analysis (PCA) on the 1024 and 768 latent representations from the last layer of the ESM-1b and the MSA transformer. The optimal number of remaining principal components (nPCs) was selected using linear regression as a base predictor. The transformed PCs were fed as salient features to a supervised machine learning model. We further tested different

supervised machine learning models including random forest, support vector machine, gradient boosting and elastic net. The preselected optimal nPCs were used. Their optimal parameters were selected based on 5-fold cross validation on the training dataset using regression coefficients as optimized scores. Finally, we repeated the selection of the nPCs for each model given the tuned parameters. The above process was done on our training-validation datasets using sklearn [26] for deletions and insertions separately. For multiple-amino-acid indels, we calculated the predictive scores for each amino acid and then tested the maximum, mean and sum of the scores as the final predictive score.

Input features

ESM-1b and MSA transformers take protein primary sequences and MSAs as inputs. Limited by the pretrained transformers, we divided long protein sequences into equal-size pieces fewer than 1023 amino acids. We downloaded MSA data from Ensembl Compara using REST API (<https://rest.ensembl.org/documentation/info/genetree>). The median and mean MSA depth is 211 and 320.2, respectively. As many proteins included in the MSA are not similar enough to the human proteins of interest, we trimmed the phylogenetic trees to remove the less similar proteins and included a maximum of 300 proteins per MSA. This also speeds up the pretraining process to generate latent representations. [Supplementary Figure S1](#) shows the distribution of MSA depth before and after the trim. The median and mean MSA depth after the trim is 199 and 184.3, respectively. For deletions, we fed wild-type protein sequences or MSAs to the pretrained transformers, and extracted latent representations of amino acids that were deleted. For insertions, wild-type MSAs were used for the MSA transformer. Extracted were latent representations of the position (amino acid or gap) followed by the amino acid where the insertion occurs. Mutated protein sequences with inserted amino acid(s) were inputted for ESM-1b transformer and latent representations of the inserted amino acid(s) were used as features. The dimensions of the latent representations from the last layer of the two transformers are fixed as 1024 and 768 for ESM-1b and MSA transformers, respectively.

Measures of predictive quality

Quantitative predictive scores were evaluated using the receiver operating characteristic (ROC) curve and the difference of two score distributions from cases and controls. R package ‘pROC’ was used to estimate area under the ROC curve (AUC) and the level of significant difference between two ROC curves. AUC is used as the quality measure to select the optimal number of features, models and the way to handle multiple-amino-acid indels. The difference between the two distributions from cases and controls was measured using Wilcoxon rank sum test in R. Smaller *P*-values indicate better separation of benign and pathogenic variants. The binary predictions were assessed using balanced accuracy, sensitivity/recall and specificity.

Existing pathogenicity predictors for inframe indels

We selected all available methods that predict pathogenicity for inframe indels and provide a web server for convenient access to the end users. They included VEST-indel [12], CAPICE [16] and CADD [15]. VEST-indel applied random forest included 24 features. CADD used L2 regularized logistic regression with more than 60 features. CAPICE used gradient boosting on decision trees with the same set of features as CADD. The default thresholds for SHINE,

CAPICE, VEST-indel and CADD are 0, 0.02, 0.8 and 20, respectively. Inframe indels with predictive scores larger than the threshold are considered as pathogenic.

Results

SHINE model

Using transfer learning, SHINE takes advantage of the pretrained protein language models to maximize the use of limited available pathogenicity labels for inframe indels. We first tested the performance of the two pretrained models: ESM-1b and MSA transformers. PCA was applied to their latent representations separately and then to combined representations to generate different sizes of salient features (nPCs). Linear regression projecting the salient features to pathogenicity label, was trained on the training dataset and tested on the validation dataset. We picked the highest (worst) predictive score for multiple-amino-acid insertions/deletions as the final predictive score. [Supplementary Figure S2](#) shows the AUC values for the MSA transformer are higher than those for the ESM-1b transformer. Combining the representations from both models does not improve predictions compared with the MSA transformer only. The two transformers capture different properties about proteins: representations from the ESM-1b transformer consider amino acid sequence patterns and representations from the MSA transformer consider cross-species conservation. Therefore, we used both transformers in SHINE.

The following analyses were performed using representations from both ESM-1b and MSA transformers. We evaluated different ways (maximum, mean and sum) to handle multiple-amino-acid indels. [Supplementary Figure S3](#) shows maximum score is the best. Although intuitively, long indels are more likely to be pathogenic than short indels, sum of the individual scores for each amino acid does not work. Thus, SHINE model uses maximum score for multiple-amino-acid indels.

PCA transformed the combined latent representations from the two transformers into low dimensional features (PCs). [Figure 2](#) shows the scatterplot of the first two PCs for pathogenic and benign variants on the training dataset. PC1 and PC2 correlated with pathogenicity for deletions with correlation coefficients of -0.506 and -0.436 . The correlation coefficients for insertions were 0.365 and -0.455 for PC1 and PC2, respectively. The first 10 PCs explained 38.7 and 42.6% of the variance for deletion and insertion representations, respectively. Scree plots are in [Supplementary Figure S4](#). Fifty and 10 PCs were chosen for deletions and insertions as they give the highest AUC values based on the linear regression model ([Supplementary Figure S3](#)).

We tested different supervised machine learning models and tuned their parameters given the same nPCs optimized in the last step as input. At the end, each of the tuned models was used to select their own optimal nPCs based on the best AUC values. Elastic Nets were the best models as they provided consistent good performance, were not sensitive to the number of input PCs and were not likely to overfit on the training dataset ([Figure 3](#)). Parameters alpha and l1_ratio of the Elastic Nets were 1.0 and 0.1 for both deletions and insertions. The Elastic Net for the deletions included 100 PCs, and 32 of them had non-zero coefficients. The Elastic Net for the insertions included 30 PCs, and 14 of them had non-zero coefficients.

Evaluation on the NDD test dataset

Instead of splitting data for training and test from the same source, we used independent test datasets to avoid the inflated performance on data from the same source. We compared SHINE

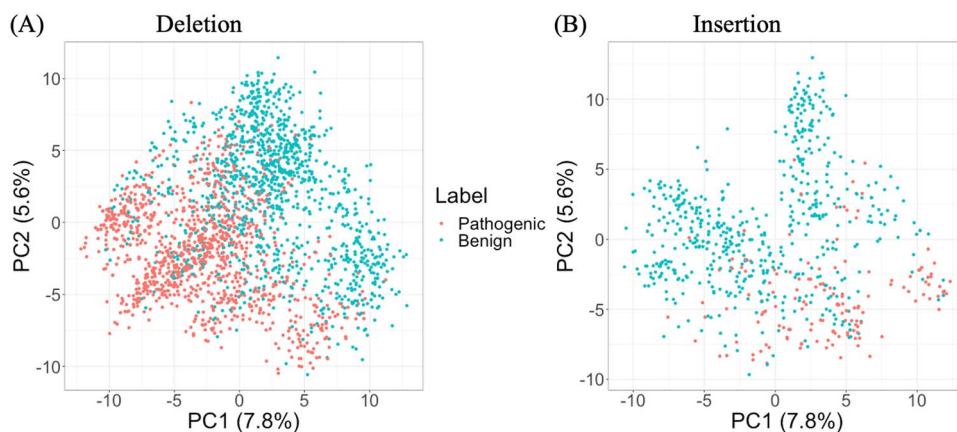


Figure 2. Scatterplot of the first two principal components for pathogenic and benign inframe indels.

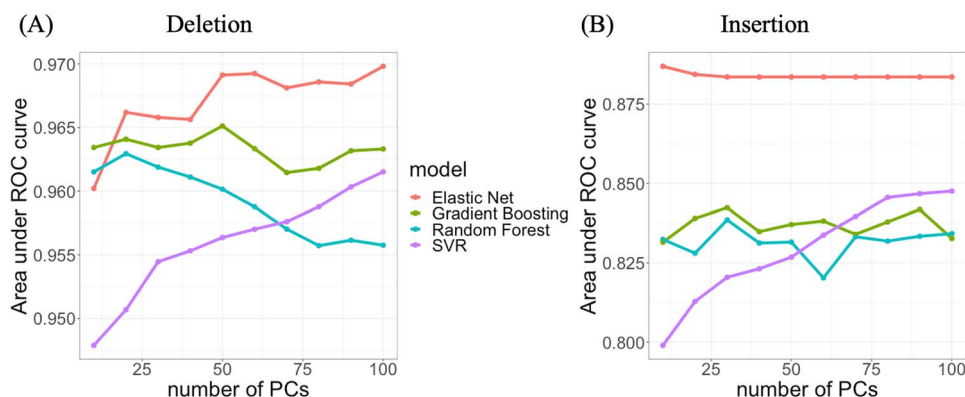


Figure 3. Area under ROC curve for a set of machine learning models with varying numbers of principal components as input features.

with three methods including VEST-indel [12], CAPICE [16] and CADD [15] by classifying indels in the 686 NDD risk genes from NDD cases and UK biobank controls. We note the indels from both cases and controls are mixtures of pathogenic and benign variants, but the *de novo* variants in NDD risk genes identified in cases are more likely pathogenic compared with germline variants in the general population. Thus, we could use the case/control status as a proxy of pathogenic/benign labels in the test. We first examined the predictive power of the autism (SPARK) and NDD cohorts separately. [Supplementary Figure S5](#) shows SHINE scores have distinct modes for cases versus controls for both cohorts, so we combine autism and NDD cases for the following analysis. Compared with the other computational methods, SHINE has the lowest *P*-value ([Supplementary Figure S6](#)), i.e. separates pathogenic from benign the best.

From the ROC curves, [Figure 4](#) shows that SHINE has the highest AUC values of 0.846 and 0.819, improves over the second-best VEST-indel by 0.04 (relative improvement of 4.4%) and 0.05 (relative improvement of 6.6%) for deletions and insertions, respectively. The improvement over all three methods is significant (P -value < 0.05) for deletion. SHINE and VEST-indel are insignificant for insertion predictions ([Supplementary Table S1](#)). The sensitivity rises quickly at a low false positive rate for SHINE. It implies SHINE scores correlate with likelihood of pathogenicity. We evaluated their binary predictions using their default thresholds. [Supplementary Table S2](#) shows the evaluation for the binary prediction including balanced accuracy, sensitivity and specificity. SHINE provides balanced accuracy scores of 0.777 and 0.699 for deletions and insertions. All the methods over-predict. CAPICE predicts the most, so has the highest sensitivity. SHINE, VEST-indel

and CADD offer the best specificity. We note SHINE provides much higher specificity compared with sensitivity for insertion. This is probably due to more benign variants than pathogenic variants in insertion training dataset. Overall, SHINE and VEST-indel give good balanced accuracy balancing sensitivity and specificity. Our analysis suggests SHINE can distinguish pathogenic from benign indels well. For users interested in high-precision predictions, high SHINE scores provide a good solution (high sensitivity given a low false positive rate in [Figure 4](#)).

None of the computational methods used allele frequency in their predictive models, so we tested their sensitivity to allele frequency. We repeated the analysis using a subset of benign variants with gnomAD allele frequency > 10^{-4} (146 deletions and 35 insertions in cases, and 601 deletions and 339 insertions in controls). All the computational methods perform better on this dataset with more restricted benign indels. SHINE still has the highest AUC value ([Supplementary Table S1](#)).

One-amino-acid inframe indels are most abundant among indels in the population and are also the hardest to predict. We extracted indels from our NDD test dataset with only one deleted or inserted amino acid. This subset includes 118 deletions and 18 insertions in cases, and 2207 deletions and 747 insertions in controls. We reported the AUC values on this subset in [Supplementary Table S1](#). All methods provide similar AUC values compared with that from the full NDD dataset.

Evaluation on the cancer mutational hotspot test dataset

We compared SHINE with the other three methods on indels in the other independent test dataset: cancer mutational hotspots

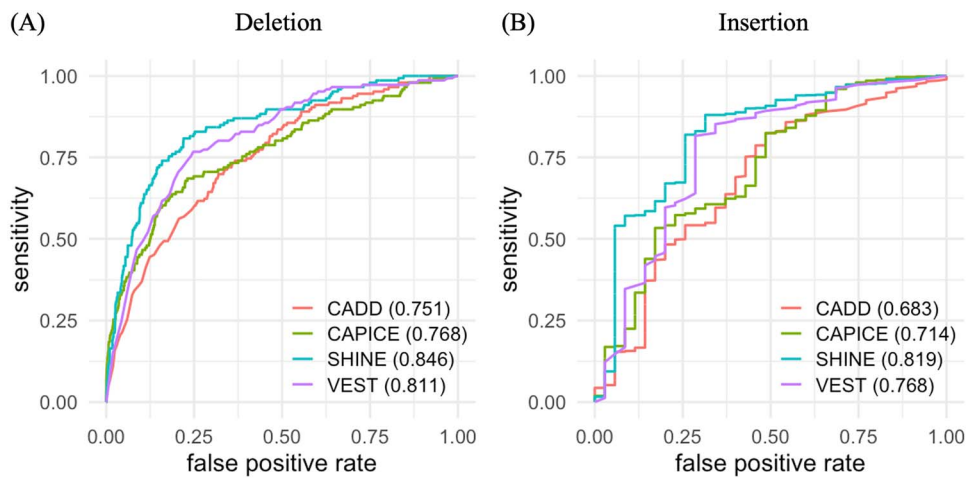


Figure 4. Evaluation of predictive scores using receiver operating characteristic curves for four prediction methods in NDD cases and UK biobank controls.

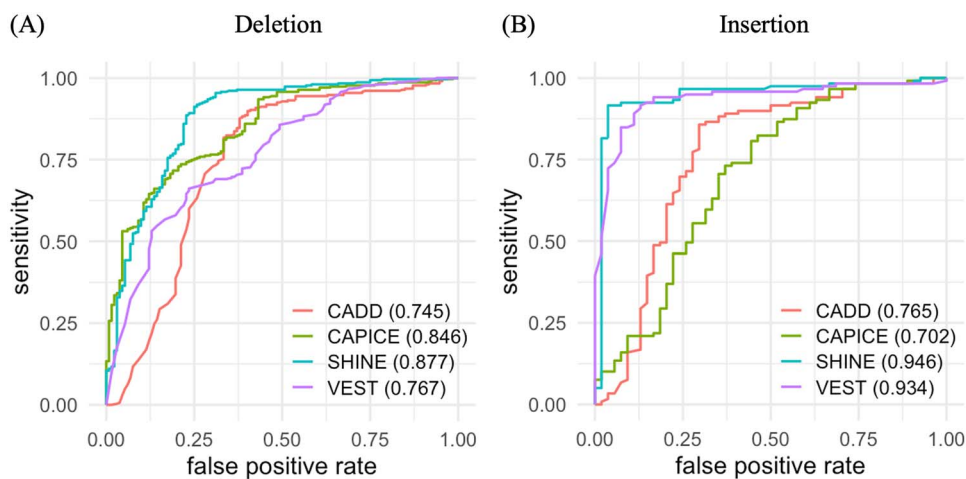


Figure 5. Evaluation of predictive scores using receiver operating characteristic curves for four prediction methods in cancer mutational hotspot cases and UK biobank controls.

and UK biobank. [Supplementary Figure S7](#) shows SHINE has the lowest P -value to separate distributions of predictive scores for cases and controls. VEST-indel provides the similar performance for insertions, but its performance on deletions is inferior. AUC results are consistent ([Figure 5](#)). SHINE with an AUC value of 0.877, is significantly better than VEST-indel and CADD for deletions (P -value < 0.05). It improves over CAPICE by 0.03 (relatively improvement of 3.7%), but the difference is not significant. On the insertion dataset, SHINE with the highest AUC values of 0.946 is insignificantly better than VEST-indel with AUC value of 0.934. Both outperform the other two methods significantly by over 0.17 (22.2%) in AUC values. All in all, our method is consistently the best on the two independent test datasets, despite different disease types.

Interpretation of the input features

Previous studies identified a few top performing features for indel pathogenicity prediction, such as intrinsically disordered protein region, protein secondary structure and solvent accessible surface area [9]. We investigated the correlation between the input features of SHINE—principal components and the three protein-structure features. [Supplementary Figure S8](#) shows the PC1 is significantly correlated with protein secondary structure (coil), intrinsically disordered residues and relative solvent accessibility.

Pathogenic variants are more likely to be seen in structured protein regions and embedded in the core of proteins. Our observations are consistent with the previous studies [9]. Even though we did not encode protein features as previous methods such as protein secondary structure, these features are likely captured in the latent representations generated from the pretrained transformers. Moreover, the latent representations have potential to carry information that we have not discovered. More efforts are still needed to interpret the protein structure and function to better understand how genetic variations affect the proteins.

Discussion

Accurate pathogenicity predictions for inframe indels are important. Available computational approaches to address this question are insufficient, primarily due to the limitations of available pathogenicity labels for training. We take advantage of the protein language models trained using millions of protein sequences to learn protein statistics/features in an unsupervised fashion. We for the first time transfer the knowledge about proteins to pathogenicity prediction for inframe indels. Evaluating on two independent test datasets, our method is consistently better than the current computational methods. SHINE scores provide a pathogenicity likelihood, which can facilitate selection

of top pathogenic inframe indels in research and clinical interpretation of variant pathogenicity.

Most of the current computational studies use pathogenic indels from ClinVar and Human Gene Mutation Database, and common and low-frequency indels from population datasets as benign. The issue is that pathogenic indels have been identified only in a few hundred known disease genes, whereas benign indels are largely extracted from unconstrained genes. The two sets of genes do not overlap much. Our training-validation datasets include pathogenic indels in 917 genes and benign indels in 1566 genes, and only 206 genes contributing both pathogenic and benign indels. Previous methods built on similar datasets tended to select gene-level features, which separate disease genes from the rest. These methods lack the power to distinguish pathogenic from benign indels when testing the same gene. Our test datasets are designed to eliminate the effect from the ascertainment bias. We included variants in the same set of genes from different sources. Neither the NDD nor the cancer mutational hotspot dataset was used to train any of the previous methods. So, our test datasets are good benchmarks to test the predictive power for variant interpretation and compare across different approaches.

We note the labels in our test datasets are not true pathogenic or benign classifications, instead case/control status for the origin of the variants. Not all variants found in cases are pathogenic; not all variants found in controls are benign. In our sensitivity analysis removing ultra-rare variants from controls, all methods perform better. It suggests it is easier to separate pathogenic variants from common benign variants compared with rare benign variants. We also acknowledge a small portion of rare variants in controls could be pathogenic, which affects the predictive performance of the computational methods.

Protein language models generate protein latent representations from the protein sequences or multiple sequence alignments. Based on recent publications about protein language models and their applications in protein-structure prediction, these latent representations should capture information related to manually calibrated features such as sequence conservation, amino acid properties, and protein-structure context. In addition, the language model may capture non-linear interaction of the features and potentially new information not well represented in those features, as it has much better performance for a range of prediction tasks than conventional methods based on manually calibrated features. We acknowledge interpretation of deep learning models is a challenging problem in general. However, we found the PCs (combinations of the protein representations) correlate with the previously discovered predictive features and provide better predictive power. Moreover, these representations are protein signatures learnt from unsupervised approaches and not biased to any specific task, so can be used for many other bioinformatic applications.

Conclusions

SHINE is the first protein language model-based method to predict pathogenicity of inframe indels. The protein language models generate unbiased protein statistics in an unsupervised fashion. Future research should consider expanding the variant types for pathogenicity prediction using similar approaches. Benchmark datasets from functional data will be highly appreciated as deep mutational scan data are becoming available for inframe indels.

Key Points

- We present SHINE, a transfer learning-based method to take advantage of pretrained representation about protein sequence and homologous alignment. The pre-training process is based on unsupervised deep learning protein language models, trained on millions of protein sequences. This approach allows us to overcome the limitation about the small number of known pathogenic inframe indels, a bottleneck in conventional supervised machine learning methods.
- We created two benchmark test datasets, which are independent from commonly used ClinVar variants. The performance evaluation on these two datasets is robust and fair.
- SHINE achieves better predictive performance than existing methods on the two test datasets. SHINE scores provide a pathogenicity likelihood, which can facilitate selection of top pathogenic inframe indels in research and clinical interpretation of variant pathogenicity.

Acknowledgment

We thank Dr Chen Wang for helpful discussions. We are grateful to all of the families in SPARK study, the SPARK clinical sites and SPARK staff. We appreciate obtaining access to the genomic data on SFARI Base. Approved researchers can obtain the SPARK population dataset described in this study by applying at <https://base.sfari.org>. We appreciate obtaining access to recruit participants through SPARK research match on SFARI Base. The SPARK initiative is funded by the Simons Foundation as part of SFARI.

Supplementary data

Supplementary data are available online at <https://academic.oup.com/bib>.

Data availability

The training, validation and test datasets, and codes for processing the MSA data and running the prediction models are available on GitHub: <https://github.com/xf-omics/SHINE>.

Authors' contributions

XF and YS conceptualized and designed the study; XF and HP were involved in data curation; XF and AT took the responsibility of formal analysis and visualization; XF, YS and WKC led the drafting of the original manuscript; YS and WKC were involved in supervision. All authors reviewed the manuscript and approved the submission of this manuscript.

Funding

National Institutes of Health (grant numbers K99HG011490 to X.F. and R01GM120609 to Y.S.); and Columbia University Precision Medicine Joint Pilot Grants Program to Y.S.

References

1. Backman JD, Li AH, Marcketta A, et al. Exome sequencing and analysis of 454,787 UK biobank participants. *Nature* 2021;**599**: 628–34.

2. Radford EJ, Tan HK, Andersson MHL, et al. Saturation genome editing of DDX3X clarifies pathogenicity of germline and somatic variation. medRxiv. 2022. <https://doi.org/10.1101/2022.06.10.2227617917>.
3. Mills RE, Luttig CT, Larkins CE, et al. An initial map of insertion and deletion (INDEL) variation in the human genome. *Genome Res* 2006;**16**:1182–90.
4. Berning D, Adams H, Luc H, et al. In-frame indel mutations in the genome of the blind Mexican Cavefish. *Astyanax mexicanus*. *Genome Biol Evol* 2019;**11**:2563–73.
5. Sergouniotis PI, Barton SJ, Waller S, et al. The role of small in-frame insertions/deletions in inherited eye disorders and how structural modelling can help estimate their pathogenicity. *Orphanet J Rare Dis* 2016;**11**:125.
6. ClinVar. 2021 <https://doi.org/10.1186/s13023-016-0505-0>.
7. Landrum MJ, Lee JM, Benson M, et al. ClinVar: improving access to variant interpretations and supporting evidence. *Nucleic Acids Res* 2018;**46**:D1062–7.
8. Hu J, Ng PC. SIFT Indel: predictions for the functional effects of amino acid insertions/deletions in proteins. *PLoS One* 2013;**8**: e77940.
9. Zhao H, Yang Y, Lin H, et al. DDIG-in: discriminating between disease-associated and neutral non-frameshifting micro-indels. *Genome Biol* 2013;**14**:R23.
10. Bermejo-Das-Neves C, Nguyen H-N, Poch O, et al. A comprehensive study of small non-frameshift insertions/deletions in proteins and prediction of their phenotypic effects by a machine learning method (KD4i). *BMC Bioinf* 2014;**15**:111.
11. Zhang N, Huang T, Cai YD. Discriminating between deleterious and neutral non-frameshifting indels based on protein interaction networks and hybrid properties. *Mol Genet Genomics* 2015;**290**:343–52.
12. Douville C, Masica DL, Stenson PD, et al. Assessing the pathogenicity of insertion and deletion variants with the variant effect scoring tool (VEST-Indel). *Hum Mutat* 2016;**37**:28–35.
13. van der Velde KJ, de Boer EN, van Diemen CC, et al. GAVIN: gene-aware variant Interpretation for medical sequencing. *Genome Biol* 2017;**18**:6.
14. Pagel KA, Antaki D, Lian A, et al. Pathogenicity and functional impact of non-frameshifting insertion/deletion variation in the human genome. *PLoS Comput Biol* 2019;**15**: e1007112.
15. Rentzsch P, Witten D, Cooper GM, et al. CADD: predicting the deleteriousness of variants throughout the human genome. *Nucleic Acids Res* 2019;**47**:D886–94.
16. Li S, van der Velde KJ, de Ridder D, et al. CAPICE: a computational method for consequence-agnostic pathogenicity interpretation of clinical exome variations. *Genome Med* 2020;**12**:75.
17. Rives A, Meier J, Sercu T, et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proc Natl Acad Sci* 2021;**118**: e2016239118.
18. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. *Nature* 2021;**596**:583–9.
19. Meier J, Rao R, Verkuil R, et al. Language models enable zero-shot prediction of the effects of mutations on protein function. *NeurIPS*. 2021; 2021.2007.2009.450648.
20. Zhou X, Feliciano P, Wang T, et al. Integrating de novo and inherited variants in 42,607 autism cases identifies mutations in new moderate-risk genes. *Nat Genet* 2022;**54**:1305–19. <https://doi.org/10.1038/s41588-022-01148-2>.
21. Kaplanis J, Samocha KE, Wiel L, et al. Evidence for 28 genetic disorders discovered by combining healthcare and research data. *Nature* 2020;**586**:757–62.
22. pfeliciano@simonsfoundation.org SCEa, Consortium S. SPARK: a US cohort of 50,000 families to accelerate autism research. *Neuron* 2018;**97**:488–93.
23. Feliciano P, Zhou X, Astrovskaya I, et al. Exome sequencing of 457 autism families recruited online provides evidence for autism risk genes. *NPJ Genom Med* 2019;**4**:19.
24. Chang MT, Bhattarai TS, Schram AM, et al. Accelerating discovery of functional mutant alleles in cancer. *Cancer Discov* 2018;**8**: 174–83.
25. Rao RM, Liu J, Verkuil R, et al. MSA Transformer. In: Marina M, Tong Z (eds). *Proceedings of the 38th International Conference on Machine Learning*. Proceedings of Machine Learning Research: PMLR, 2021, 8844–56.
26. Pedregosa F, Varoquaux G, Gramfort A, et al. Scikit-learn: machine learning in Python. *J Mach Learn Res* 2011;**12**:2825–30.