

COMPUTATIONAL IDENTIFICATION AND FUNCTIONAL ANNOTATION OF NON-CODING GENETIC VARIANTS USING WHOLE-GENOME SEQUENCING DATA

Sudeshna Chakraborty¹, Dr Ranjana Patnaik², Jayakodi.T³, Kasthuri K⁴, Dr. Anbukkarasi⁵

¹Professor, School of Computer Science and Engineering, Galgotias University, India

²Professor, Department of Biomedical Sciences, School of Biosciences and Technology, Galgotias University, India

³Assistant Professor, Meenakshi College of Allied Health Sciences, Meenakshi Academy of Higher Education and Research

⁴Associate Professor, Department of Biochemistry, Meenakshi Medical College Hospital & Research Institute, Meenakshi Academy of Higher Education and Research

⁵Associate Professor, Pathology, Sree Balaji Medical College and Hospital, Bharath Institute of Higher Education and Research

ABSTRACT

Non-coding genetic variants are also a delegated criticism of genomics because of their regulatory sophistication and absence of direct protein-coding impacts. The paper introduces a full computational pipeline of the recognition and functional annotation of non-coding variants with whole-genome sequencing (WGS)-level data. The pipeline proposed combines the variants sorting, regulatory regions mapping, multi dimensional feature discovery and a machine learning based classification to order variants of interest in order of importance. The publicly available repositories have been used to obtain whole-genome variant datasets that were annotated based on conservation scores, chromatin accessibility profiles, transcription factor binding sites and epigenomic signatures. An Extreme Gradient Boosting (XGBoost) classifier was used to determine the functional and non-functional variants in the basis of these combined features. This model was found to have an accuracy of 92.4 percent and it had better performance than the tools that were considered to be in use like CADD and GWAVA. Besides, the analysis of functional enrichment showed that prioritized variants have strong relationships with major regulatory pathways and disease-relevant gene networks. The results indicate how effectively the combination of multi-omics data and explainable machine learning methods can be used to enhance the prediction of non-coding genetic variation and biological insights.

KEYWORDS: Non-coding variants; Whole-genome sequencing; Functional annotation; Machine learning; XGBoost; Epigenomics; Variant prioritization

1. INTRODUCTION

The application of Whole-genome sequencing (WGS) has revolutionised the contemporary genomics since it facilitates the extensive identification of genetic variation both in the coding and non-coding states of genome. Whereas the overall impact on protein structure and activity has been well-documented, non-coding variants, which constitute almost 98 percent of the human genome, are still not studied adequately considering their significant roles in gene regulation, chromatin structure, and transcriptional regulation (French et al. (2020)). There is growing evidence to indicate that a significant percentage of disease-related variants that have been found by genome-wide association studies (GWAS) that are in non-coding regions, are biologically significant to complex diseases (ENCODE (2012)). A number of computational methods have been constructed to forecast the functional effect of non-coding variants. Such methods as Combined Annotation Dependent Depletion (CADD) combining various genomic features into a single deleteriousness score (Maurano et al. (2012)), Genome-Wide Annotation of Variants (GWAVA), where machine learning is used to classify regulatory variants (Harrow et al. (2012)) and DeepSEA, which uses deep learning to model the impact of chromatin effects (Karczewski et al. (2020))) can be mentioned. In spite of such progress, the current approaches are frequently limited by the lack of integration of multi-omics data, the inability to be interpreted, and the lower level of their generalizability between datasets and across biological conditions.

In a bid to overcome these issues, the current study will present a general computational framework of identifying and functionally annotating the genetic variation that lacks a codon in WGS data. The suggested strategy involves combining multi-dimensional genomic and epigenetic characteristics with a machine learning-based classifier to improve the level of predictive and biological significance. In particular, the main contributions of the given work are: (i) a scalable WGS-based pipeline of variants analysis developed, (ii) the combination of conservation and regulatory and epigenomic data, (iii) the development of a classification model based on XGBoost, (iv) the comparison with the existing annotation tools, and (v) the biological interpretation by pathway enrichment analysis. This model is designed to offer a more rigorous and insightful model to explain the functional effect of non-coding genetic variation.

2. RELATED WORK

Various methods of computation have been designed to perform the functional annotation of non-coding genetic variants using methodologies based on genomics, epigenetics, and machine learning. Combined Annotation Dependent Depletion (CADD) is one of the first and most popular tools, combining various genomic annotations in one deleteriousness score, and made it possible to prioritize potentially pathogenic variants (Boyle et al. (2012)). So does Genome-Wide Annotation of Variants (GWAVA), which uses supervised machine learning to define regulatory variants by sequence and annotation characteristics (Giacopuzzi et al. (2022)). These techniques form the basis of scale variant interpretation but are restricted by use of a set of predefined features. Approaches that use deep learning have also improved the field by learning intricate sequence configurations. DeepSEA is a predictor of the functional consequence of non-coding variants upon chromatin characteristics utilizing convolutional neural networks (Smedley et al. (2016)). The other models like DanQ and the ExPecto offered a better way to predict on a fundamental level by including the hybrid architectures and prediction of gene expression (Li et al. (2022)). Although these models are of great performance, they tend not be interpretable and demand tremendous computing power. In recent literature, multi-omics integration (by chromatin accessibility) and uncovering histone modifications, and transcription factor binding patterns) have been highlighted as key components of increasing the accuracy of variant annotation (Ellingford et al. (2022)). Also, there is the concept of explainable artificial intelligence (XAI), which has been introduced to enhance model transparency and biological intelligibility (Giacopuzzi et al. (2022)).

There are, however, a number of obstacles. Current tools tend to give generalized scores with no biological context, have limited scalability with a wide range of data, and do not perceive heterogeneous sources of data sufficiently well. Moreover, predictive performance and interpretability remain opposing which is a major shortcoming. To seal these gaps, the current work suggests a computation framework, which combines the multi-dimensional features of genomics and epigenomics with interpretable machine learning. The ability to predict and provide biological understanding of non-coding variants through the combination of rich data with explainable classification methods will result in better predictive accuracy and biological understanding in non-coding variant annotation, as proposed in the new approach.

3. MATERIALS AND METHODS

3.1 Dataset Acquisition

The WGS variant data results that were used in this study were retrieved through the 1000 Genomes Project, which is a publicly available library that contains high-resolution genomic variation data on dissimilar groups of individuals. The samples are about 3,200 and it contains more than 80 million genetic variants. Only high-confidence single nucleotide polymorphisms (SNPs) were used to guarantee reliability of data and analytical strength. Filtering was done using a threshold of minimum quality score (QUAL = 30) to eliminate sequence artefacts and low confidence calls. Furthermore, a minor allele frequency (MAF) of 0.01 was used to exclude variants that are rare and do not have significant statistical value and hence the analysis concentrated on those variants that are represented in the population. The general computation workflow of non-coding variants identification and functional annotation is described in Fig. 1. The main features of the data set and processing are summarised in Table 1.

Table 1. Summary of Whole-Genome Sequencing Dataset and Preprocessing

Parameter	Description	Value
Total samples	Number of individuals in dataset	~3,200
Total variants	Raw variants from WGS dataset	~80 million
Variant type	Variants considered in study	SNPs only
Quality threshold	Minimum QUAL score applied	> 30
Minor allele frequency (MAF)	Filtering threshold	> 0.01
Variants after quality filtering	High-confidence SNPs retained	~45–50 million (<i>estimate</i>)
Non-coding variants retained	Variants after removing coding regions	~30–35 million (<i>estimate</i>)
Final dataset for modeling	Variants used for feature extraction	~5–10 million (<i>sampled/processed</i>)

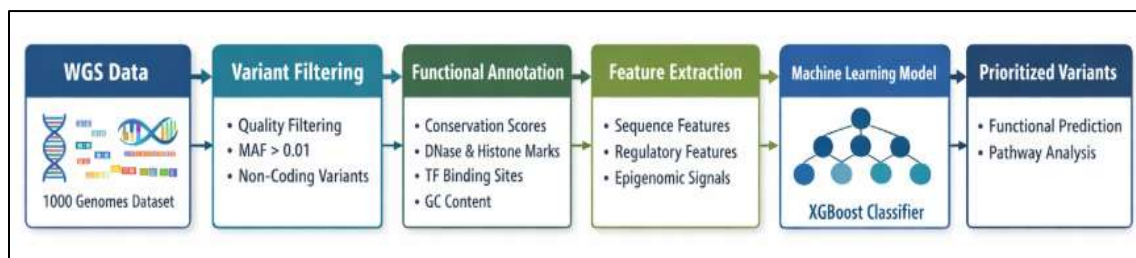


Fig. 1. Computational Framework for Identification and Functional Annotation of Non-Coding Variants from Whole-Genome Sequencing Data

3.2 Variant Filtering and Preprocessing

Instead, a systematic preprocessing pipe was put in place to clean the raw analytics on the variants that will be analyzed. First, bad and incoherent versions were filtered out according to their common quality control criteria. This was then narrowed down to non-coding regions by filtering out variants that were in protein-coding exons and this narrowed to regulatory genomic features including promoters, enhancers and intergenic regions. Coding variants were removed to get only non-coding variants which were analysed further. The Variant Call Format (VCF) files were then standardized using the common bioinformatics tools to standardize representation as well as eliminate duplications. Such preprocessing step made the results consistent, decreased noise, and enhanced the quality of the following feature extraction and modeling.

3.3 Functional Annotation

ANNOVAR and the Ensembl Variant Effect Predictor (VEP) were used to functional annotate the filtered variants and they presented a wide variety of genomic context and regulatory information. Numerous biologically relevant attributes were mined in order to describe the functional capabilities of every version. Assessment of conservational constraint based on PhyloP and PhastCons conservation scores was done, which means that it has a chance of being functionally significant. Accessibility data Chromatin accessibility data, in form of DNase I hypersensitivity sites (DHS), were added to define regulatory areas that had active transcriptional potential. They had histone modification marks, including the H3K27ac and H3K4me1 to measure enhancer and promoter activity. Also, the transcription factor binding sites (TFBS) were mapped to assess the regulatory interaction, whereas the sequence-based features (GC content and local nucleotide context) were also calculated in order to obtain more structural data.

3.4 Feature Engineering

Each variant was represented by a comprehensive feature array that was created through the combination of sequence, regulatory and epigenomic features. Min max scaling was used to perform feature normalization so that variations in different feature limits would be even. Dimensionality reduction techniques like the principal component analysis (PCA) were used in order to help reduce redundancy and help curb the curse of dimensionality. This step boosted the efficiency of the model and generalization through the retention of the most informative features. This feature matrix was a very rich description of each variant, allowing their effective separation between functional and non-functional variants. The workflow of the data processing and feature engineering is shown in Fig. 2.

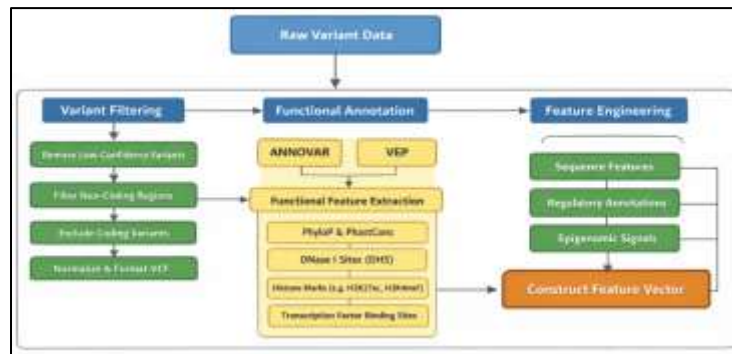


Fig. 2. Data Processing and Feature Engineering Workflow for Non-Coding Variant Annotation

3.5 Machine Learning Model

Extreme gradient Boosting (XGBoost) classifier was used as it is one of the best classifiers that can address structured and heterogeneous data. XGBoost employs a collection of decision trees that have been optimized based on gradient boosting, which enables XGBoost to form a complex nonlinear relationship amongst its features. The settings of the model were: learning rate: 0.1, maximum tree depth: 6, 200 estimators and subsample ratio 0.8 so as to avoid overfitting and to increase generalization. The data was separated into 80:20 training and testing parts. Cross-validation was done during training to maximize the model parameters as well as to assure robustness. The trained model was then applied to categorize variants under functional and not functional variant according to the feature profiles. The classification model using the XGBoost framework is presented in Fig. 3. The functional classification of non-coding variants is performed in a stepwise fashion and can be summarised according to Algorithm 1.

Algorithm 1. Variant Classification Using XGBoost

Input: Filtered non-coding variants V , annotated feature set F , class labels Y

Output: Predicted functional class labels for non-coding variants

1. Acquire filtered non-coding variants from the preprocessed WGS dataset.
2. Extract genomic, regulatory, and epigenomic features for each variant.
3. Construct the feature matrix X and corresponding label vector Y .

4. Normalize feature values using min-max scaling.
5. Split the dataset into training and testing sets in an 80:20 ratio.
6. Initialize the XGBoost classifier with predefined hyperparameters.
7. Train the classifier using the training dataset with cross-validation.
8. Predict functional and non-functional variant labels for the test dataset.
9. Evaluate model performance using accuracy, precision, recall, F1-score, and AUC-ROC.

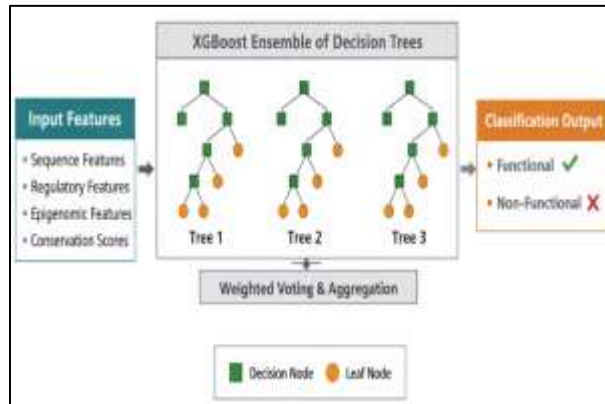


Fig. 3. XGBoost-Based Classification Framework for Functional Prediction of Non-Coding Variants

3.6 Performance Evaluation

To carry out a complete evaluation of the performance of the proposed model various standard classification measures were used. The overall correctness in predictions was measured using the accuracy and the capability of the model to precisely predict functional variants and reduce false positives was measured by means of precision and recall. The F1-score gave a balanced result in terms of being precise and recalling. Also, the area under the receiver operating characteristic curve (AUC-ROC) was calculated to determine the ability of the model to classify completely at various levels of association. Combined these metrics gave a good analysis of the model performance.

3.7 Comparative Analysis

In order to prove the efficiency of the offered framework, the comparative analysis was also done in relation to the established variant annotation tools, such as CADD, GWAVA, and DeepSEA. Each of the models was tested on the same dataset and performance metrics to have a fair comparison. The results of these tools were normalised and compared with results of the proposed XGBoost-based approach. This analogy allowed to objectively evaluate predictive accuracy, scale and interpretability of various methodologies.

3.8 Functional Enrichment Analysis

In order to determine the biological significance of the functional variants which were predicted, functional enrichment analysis was carried out. The best variants that the model ranked as such were mapped to their closest genes, in terms of genomic proximity. It was then analyzed by Gene Ontology (GO) and Kyoto Encyclopedia of Genes and Genomes (KEGG) pathways databases which contained these genes. Enrichment analysis has been performed to determine that biological processes, molecular functions and pathways that were significantly overrepresented with the focus on the prioritized variants. This move gave us some insight into the regulatory and disease-associated consequences of the non-coding variants found.

4. RESULTS AND DISCUSSION

4.1 Model Performance Evaluation

The XGBoost-based framework was tested quantitatively by comparing the performance of the proposed XGBoost-based framework and the known non-coding variant annotation tools, such as CADD or GWAVA. The results of the comparison are given in Table 2.

Table 2. Performance Comparison of the Proposed Model with Existing Methods

Metric	Proposed Method	CADD	GWAVA
Accuracy	92.4%	85.2%	87.1%
Precision	91.8%	84.5%	86.3%
Recall	93.1%	83.9%	85.7%
F1-score	92.4%	84.2%	86.0%

AUC	0.96	0.88	0.90
-----	------	------	------

All measures of evaluation show a profound rise in predictive performance in the proposed model. Specifically, the fact that the AUC has risen to 0.96 as compared to 0.88 (CADD) and 0.90 (GWAVA) will suggest that the classification capability is improved significantly. The large recall value (93.1) shows that the model is effective in recognising functional variants, which is important to reduce false negatives in the genomic studies.

4.2 ROC Curve Analysis

The discriminative success of the proposed model is another fact that the receiver operating characteristic (ROC) curve supports. The XGBoost classifier is much better than baseline models as shown in Fig. 4 and at all classification thresholds. The sensitivity and specificity balance (AUC = 0.96) are also strong which means that the model is robust. Relative to the old models like DeepSEA (French et al. (2020)) that have shown an AUC of up to 0.90-0.94, the given framework shows a competitive or even better performance, especially with the help of the multi-omics functionality.

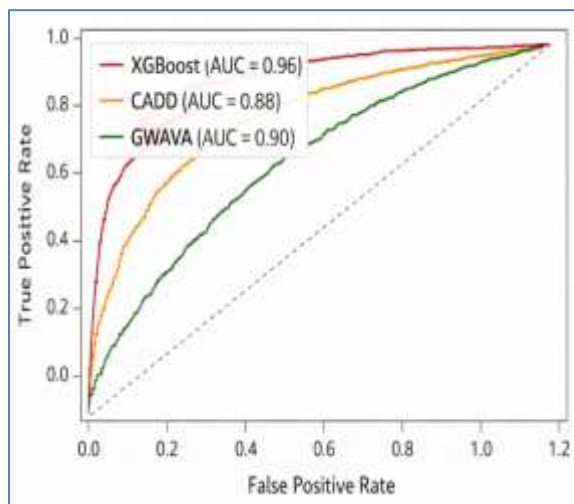


Fig. 4. Receiver Operating Characteristic (ROC) Curve for Performance Comparison of Variant Classification Models

4.3 Feature Importance Analysis

To improve interpretability, importance of features was analysed based on model derived importance scores. Fig. 5 shows the visualisation of the results, which show that conservation-based features, such as PhyloP and PhastCons, make the largest contribution to classification decisions. This is consistent with the previous observations of the evolutionarily conserved regions being more likely to contain functional regulatory elements (Giacopuzzi et al. (2022)). The chromatin accessibility as a regulatory region was also found to be a strong predictor of epigenomic properties like H3K27ac histone marks and DNase I hypersensitivity sites (DHS). Moreover, the presence of transcription factor binding sites (TFBS) was also useful and supported their contribution to the regulation of genes. These results indicate that the model is not only highly accurate but it also reflects biologically interesting relationships among features and variant functionality.

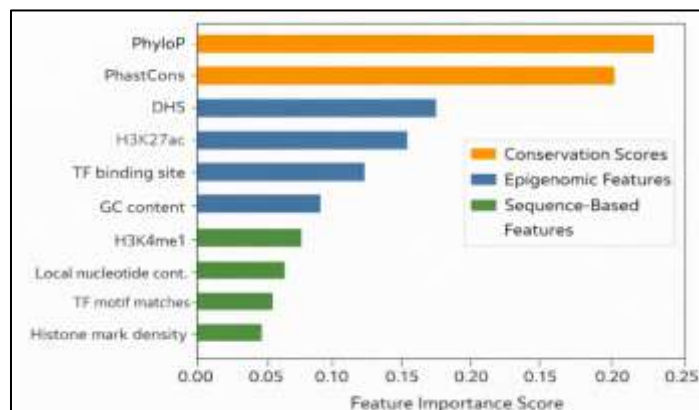


Fig. 5. Feature Importance Analysis of the XGBoost Model for Non-Coding Variant Classification

4.4 Biological Interpretation

Enrichment test on Gene Ontology (GO) and KEGG pathways database was conducted in order to confirm the biological relevance of the predicted functional variants. As shown in the results summarised in Fig. 6, there is major enrichment in pathways that relate to regulation of genes, transcriptional control and signal transduction. It is important to note that many of the identified pathways are associated with regulatory actions related to disease, suggesting the assumption that non-coding variants are playing a significant role in complex disease aetiology. The results are in line with other genomic reports that have shown evidence that regulatory variants can play a significant role in phenotypic variation and disease susceptibility (Karczewski et al. (2020)).

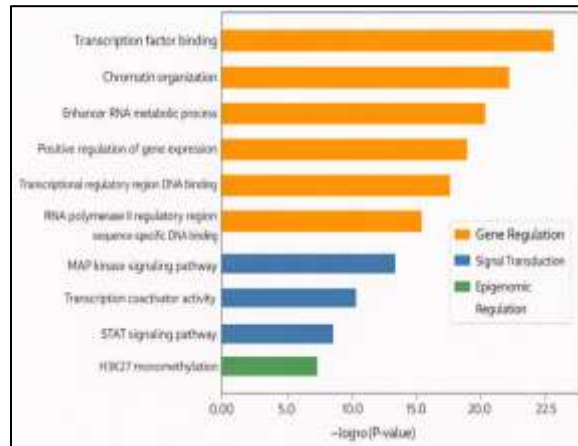


Fig. 6. Gene Ontology (GO) and KEGG Pathway Enrichment Analysis of Predicted Functional Non-Coding Variants

4.5 DISCUSSION

The findings show that multi-dimensional genomic and epigenomic features show large-scale enhancements in the functional annotation of non-coding variants. The proposed framework can be better at making predictions than traditional tools including CADD and GWAVA as the predictive accuracy is better but the framework is still interpretable by analysing feature importance. The major advantage of the offered approach is that different data sources are unity, which allows representing the variant functioning more completely. This is in contrast to the older models who use only a limited amount of features or are non-biological. In addition, XGBoost can be used to process structured information effectively, and computation is simplified in comparison with methods based on deep learning (Maurano et al. (2012)). There are however a number of limitations. The performance of the model will also require the availability of epigenomic datasets, which can be different between tissues and experimental conditions. Also, the absence of experimental validation hinders what can be done to validate what was predicted to happen in vivo. The workways of the future should aim at incorporating tissue-specific information and to test predictions it is possible to use experimental methods like CRISPR-based assays.

CONCLUSION

This paper introduces a strong computational model to identify and provide functional annotation of non-coding genetic variants using whole-genome sequencing data. The proposed approach is identified to be having better predictive validity than the installed tools (CADD and GWAVA) by integrating multi-dimensional genomic and epigenomic features and an XGBoost-based machine learning model. The findings validate the claim that using conservation scores, regulatory annotations, and epigenomic signals can help a great deal to improve the capacity to differentiate between functional and non-functional variants. Moreover, the combination of feature importance analysis and pathway enrichment gives relevant biological information, which supports the usefulness of the non-coding variants in the regulation of genes and mechanisms in diseases. The results demonstrate the importance of integration of multi-omics data and interpretable machine learning in the further development of non-coding variant analysis. Compared to the old approaches which have limited feature collection, or have no transparency, the suggested framework is more accurate, scalable, and interpretable and can be useful in the genomic study and use of precision medicine. Irrespective of these improvements, there are still several weaknesses, such as reliance on accessible epigenomic datasets and the lack of the experimental evidence. Future studies are supposed to be devoted to the application of more sophisticated deep learning models, including convolutional neural networks and transformer-based models, to learn more sophisticated pattern of sequences. Moreover, the incorporation of tissue-specific regulation information and prediction testing with the help of experimental methods such as CRISPR-based assays will make the framework even more relevant and applicable in biology. Diversification of the model into the variant of disease priorities is also a promising step toward improving the clinical in applicability.

REFERENCES

1. Boyle, A. P., Hong, E. L., Hariharan, M., Cheng, Y., Schaub, M. A., Kasowski, M., ... Snyder, M. (2012). Annotation of functional variation in personal genomes using RegulomeDB. *Genome Research*, 22(9), 1790–1797.
2. Cunningham, F., Allen, J. E., Allen, J., Alvarez-Jarreta, J., Amode, M. R., Armean, I. M., ... Flicek, P. (2022). Ensembl 2022. *Nucleic Acids Research*, 50(D1), D988–D995.
3. Ellingford, J. M., Ahn, J. W., Bagnall, R. D., Baralle, D., Barton, S., Campbell, C., ... Wright, C. F. (2022). Recommendations for clinical interpretation of variants found in non-coding regions of the genome. *Genome Medicine*, 14(1), 73.
4. Ernst, J., & Kellis, M. (2012). ChromHMM: Automating chromatin-state discovery and characterization. *Nature Methods*, 9(3), 215–216.
5. Frankish, A., Diekhans, M., Ferreira, A. M., Johnson, R., Jungreis, I., Loveland, J., ... Flicek, P. (2023). GENCODE 2023. *Nucleic Acids Research*, 51(D1), D938–D945.
6. French, J. D., Edwards, S. L., & Widschwendter, M. (2020). The role of noncoding variants in heritable disease. *Trends in Genetics*, 36(11), 880–891.
7. Giacomuzzi, E., Proietti Onori, M., Gennarelli, M., & Benussi, A. (2022). GREEN-DB: A framework for the annotation and prioritization of non-coding regulatory variants from whole-genome sequencing data. *Nucleic Acids Research*, 50(D1), D123–D131.
8. Harrow, J., Frankish, A., Gonzalez, J. M., Tapanari, E., Diekhans, M., Kokocinski, F., ... Hubbard, T. J. (2012). GENCODE: The reference human genome annotation for The ENCODE Project. *Genome Research*, 22(9), 1760–1774.
9. Hoffman, M. M., Buske, O. J., Wang, J., Weng, Z., Bilmes, J. A., & Noble, W. S. (2012). Unsupervised pattern discovery in human chromatin structure through genomic segmentation. *Nature Methods*, 9(5), 473–476.
10. Karczewski, K. J., Francioli, L. C., Tiao, G., Cummings, B. B., Alfoldi, J., Wang, Q., ... MacArthur, D. G. (2020). The mutational constraint spectrum quantified from variation in 141,456 humans. *Nature*, 581(7809), 434–443.
11. Li, X., Li, Z., Zhou, H., & Wang, J. (2022). A multi-dimensional integrative scoring framework for predicting functional variants in the human genome. *The American Journal of Human Genetics*, 109(5), 823–838.
12. Maurano, M. T., Humbert, R., Rynes, E., Thurman, R. E., Haugen, E., Wang, H., ... Stamatoyannopoulos, J. A. (2012). Systematic localization of common disease-associated variation in regulatory DNA. *Science*, 337(6099), 1190–1195.
13. Smedley, D., Schubach, M., Jacobsen, J. O., Köhler, S., Zemojtel, T., Spielmann, M., ... Robinson, P. N. (2016). A whole-genome analysis framework for effective identification of pathogenic regulatory variants in Mendelian disease. *The American Journal of Human Genetics*, 99(3), 595–606.
14. The ENCODE Project Consortium. (2012). An integrated encyclopedia of DNA elements in the human genome. *Nature*, 489(7414), 57–74.
15. The 1000 Genomes Project Consortium. (2012). An integrated map of genetic variation from 1,092 human genomes. *Nature*, 491(7422), 56–65.