

IMMUNOGENICITY PREDICTION WITH DUAL ATTENTION ENABLES VACCINE TARGET SELECTION

Song Li^{1,3}* Yang Tan²* Song Ke³ Liang Hong^{1,3} Bingxin Zhou¹ †

¹ School of Physics and Astronomy, Shanghai Jiao Tong University.

² School of Information Science and Engineering, East China University of Science and Technology.

³ Shanghai Matwings Technology Co., Ltd.

ABSTRACT

Immunogenicity prediction is a central topic in reverse vaccinology for finding candidate vaccines that can trigger protective immune responses. Existing approaches typically rely on highly compressed features and simple model architectures, leading to limited prediction accuracy and poor generalizability. To address these challenges, we introduce VENUSVACCINE, a novel deep learning solution with a dual attention mechanism that integrates pre-trained latent vector representations of protein sequences and structures. We also compile the most comprehensive immunogenicity dataset to date, encompassing over 7000 antigen sequences, structures, and immunogenicity labels from bacteria, viruses, and tumors. Extensive experiments demonstrate that VENUSVACCINE outperforms existing methods across a wide range of evaluation metrics. Furthermore, we establish a post-hoc validation protocol to assess the practical significance of deep learning models in tackling vaccine design challenges. Our work provides an effective tool for vaccine design and sets valuable benchmarks for future research. The implementation is at <https://github.com/songlee/VenusVaccine>.

1 INTRODUCTION

Immunogenicity prediction is a pivotal step in reverse vaccinology. It aims at identifying specific proteins or peptide segments (protective antigens) that can produce humoral and/or cell-mediated immune responses and lead to memory cells production in the host organism (Adu-Bobie et al., 2003; Doneva et al., 2021). Protective antigens of pathogenic origin are potential vaccine candidates (PVCs) (Arnon, 2011). Rapid and accurate prediction of protective antigens can reduce research and development costs, minimize vaccine development risks, and provide safe and effective prevention strategies against emerging and reemerging infectious diseases that pose ongoing threats (Pizza & M., 2000; Bagnoli et al., 2011; Carvalho et al., 2021).

Current machine learning approaches commonly exploit the physicochemical properties of amino acid residues, as characterized by E-descriptors (Venkatarajan & Braun, 2001) or Z-descriptors (Hellberg et al., 1987), and transform these descriptors into a uniform vector representation using auto- and cross-covariance (ACC) methods. This vectorized format serves as input for machine learning models, which are then trained to predict whether antigens are protective or non-protective. (Doytchinova et al., 2007; Rawal et al., 2022). Due to the limited quantity of labeled data and the high complexity of the prediction problem, these methods excessively compress protein information in order to fit a relatively simple model, making it difficult to capture the complex relationship between antigens and their immunogenicity. As a result, these models often exhibit **limited prediction accuracy** and **inadequate generalizability** (e.g., in cross-species immunogenicity predictions).

This study proposes VENUSVACCINE, a deep learning method that interprets immunogenicity based on the multimodal encoding of antigens, including their sequences, structures, and physicochemical properties. To supplement the inherent amino acid (AA) sequence information, two levels of structural tokens are employed to comprehensively capture antigen construction patterns, and the physicochemical properties are hand-crafted based on Z-descriptor and E-descriptor. When allocating protein representations from different modalities, we establish a *dual-attention mechanism*

*Equal contribution first authors. † Corresponding author (bingxin.zhou@sjtu.edu.cn).

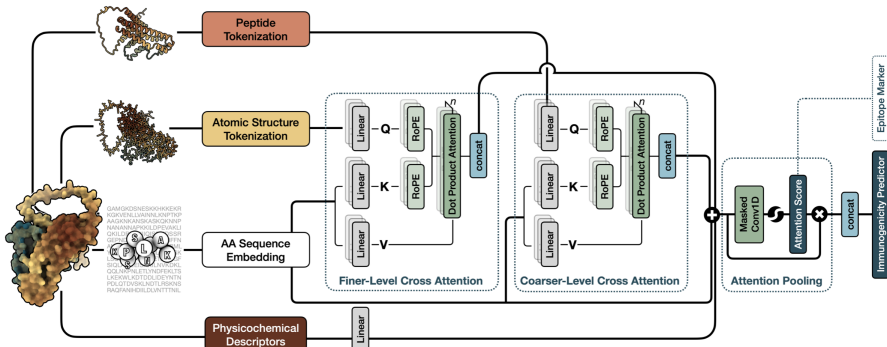


Figure 1: Illustrative framework of VENUSVACCINE. The model encodes sequence and structural representations using a dual-attention mechanism, followed by aggregation layers to incorporate global physicochemical attributes and perform binary classification of immunogenicity prediction.

(Figure 1) to facilitate communication between neighboring AAs at different scales. The use of tokenized structures at the atomic level (Van Kempen et al., 2023) and peptide level (Hayes et al., 2024) aligns with two biological intuitions: protein binding affinity primarily depends on atomic-level interactions, and local protein structures determine their functionality. Subsequently, we apply attention pooling to AA-level hidden representations to integrate them into protein-level vector representations, which incorporate the hand-crafted features that are commonly used to characterize the immunogenicity of proteins to further enhance the model’s capacity for immunogenicity prediction.

The reliability of deep learning models requires thorough empirical validation, such as statistically significant tests on high-quality datasets and comparisons with similar baseline methods. While wet-lab experiments are considered the most reliable validation due to the inherent noise of high-throughput labels, the significant time and financial costs make it impractical for most research works to implement. For this purpose, we establish two types of evaluation protocols. First, we compile an immunogenicity database (**ImmunoDB**) with 7,216 labeled antigens from bacteria, viruses, and tumors. This benchmark enables large-scale validation of prediction methods and provides valuable training data for developing new models. Secondly, we propose two post-hoc validation protocols involving significant subjects (*Helicobacter pylori* and SARS-CoV-2), comprehensive evaluation datasets, and biologically meaningful evaluation criteria.

We verify the effectiveness of VENUSVACCINE using the two established types of validation approaches and provide extensive analysis. Based on **ImmunoDB**, we designed corresponding variations of training and test datasets, and performed comprehensive comparisons with existing methods across various evaluation metrics, focusing on prediction accuracy and generalizability of the trained model. The results demonstrate significant superiority of the proposed VENUSVACCINE from different perspectives. In post-hoc validations, we emphasize the practical significance of VENUSVACCINE by demonstrating that its assigned probability to protective antigens effectively helps users identify potential vaccine candidates for vaccine design.

In summary, this study addresses the key challenges of immunogenicity prediction in vaccine development by (i) introducing VENUSVACCINE, a deep learning supervised learning scheme with a dual attention mechanism; (ii) presenting the most comprehensive dataset of labeled antigens from different sources for training and testing immunogenicity prediction models; and (iii) constructing assessment protocols for evaluating prediction performance using benchmark datasets and post-hoc analyses. Empirically, we provide extensive evidence to validate that VENUSVACCINE meets the challenges of vaccine target selection by providing accurate and robust predictions. Moreover, the key AAs identified by VENUSVACCINE correlate with antigenic epitopes, suggesting more potential of our model to assist in vaccine development. Our work not only offers a reliable tool for advancing vaccine design but also provides crucial resources and insights for further research.

2 IMMUNOGENICITY PREDICTION AND PROTEIN LANGUAGE MODELS

Problem Formulation Reverse vaccinology is an extensively utilized methodology for identifying prospective vaccine candidates by computationally screening the proteome of a pathogen (Dalsass




Figure 2: Data collection, redundancy processing, and dataset construction steps of **ImmunoDB**.

et al., 2019). This research is concerned with a binary classification task aimed at predicting the immunogenicity of proteins from three pathogenic sources: bacteria, viruses, and tumors. The primary goal is to develop a model proficient in discerning whether a protein or peptide segment is a protective antigen or a non-protective antigen. Similar to most protein property prediction tasks, immunogenicity prediction also suffers from the limitation of scarce labeled data, making it crucial to develop a powerful and generalizable model. To achieve this goal, a feasible solution is to adopt a two-step training approach, where pre-trained models are first used to extract protein embeddings, followed by training the output module specifically for the prediction task.

Pre-trained Protein Language Model for AA-level Encoding Protein language models (PLMs) have become the dominant method for learning representations of protein sequences (Zhou et al., 2024a). The two primary pre-training strategies are masked language model (Meier et al., 2021; Rao et al., 2021) and autoregressive generation (Ferruz et al., 2022; Madani et al., 2023). The former involves predicting masked tokens in the input sequence while capturing co-evolutionary relationships between AAs within the same sequence. In contrast, the latter learns to iteratively generate the next token based on previous tokens, which is particularly useful for tasks like sequence design. As we focus on property prediction tasks, we adopt the first strategy. During training, random tokens in the input sequence are masked, and the objective is to optimize the model’s parameters, denoted as θ , to minimize the error between the predicted and actual AAs at the masked positions, *i.e.*,

$$\arg \min_{\theta} \mathbb{E}_{x \sim \mathcal{X}} \mathbb{E}_M - \sum_{i \in M} \log P(x_i | x_{/M}; \theta). \quad (1)$$

The conditional probability $P(x_i | x_{/M})$ for the i th token x_i in the sequence is determined by the unmasked tokens $x_{/M}$. The model learns to capture the interactions between AAs within the protein sequence. Notably, such a sequential representation extraction method is not only applicable to AA sequence tokenization but can also be extended to capture protein structural features. A straightforward approach to achieve this is by expanding the output tokens to include not only AA types but also discretized structural tokens (Su et al., 2023; Heinzinger et al., 2023; Li et al., 2024).

3 IMMUNODB: NEW BENCHMARK FOR TRAINING AND EVALUATION

We curate three datasets with immunogenic (positive) and non-immunogenic (negative) labels from bacterial, viral, and human tumor sources. All positive samples and a part of negative samples are compiled through an exhaustive literature search with manual validation. We check publicly available sources, including PubMed, UniProt, NCBI, Protegen, IEDB, and other publications up to October 2023. The search is strictly limited to proteins proven to be immunogenic in humans. For all three datasets, the preparation steps include collecting positive samples, collecting negative samples, and filtering redundant and tail-region samples. An illustrative workflow is visualized in Figure 2. In total, we have 913/1562 positive/negative instances in **Immuno-Bacteria**, 2078/1886 in **Immuno-Virus**, and 300/477 in **Immuno-Tumor**.

Positive Samples All positive samples are from the Protegen database and previous publications. For **Immuno-Bacteria**, data are from Dimitrov et al. (2020); Edison et al. (2020); Rawal et al. (2022). For **Immuno-Virus**, positive samples are gathered from Rawal et al. (2022); Doneva & Dimitrov (2024). For **Immuno-Tumor**, data are sourced from Sotirov & Dimitrov (2024). Notably, for **Immuno-Tumor**, both positive and negative samples are collected in this step. Protein or peptide sequences are from NCBI (Wheeler et al., 2007) and UniProtKB (Consortium, 2019).

Negative Samples Due to the scarcity of data on non-immunogenic proteins in human studies, a part of the negative sets are collected from previous studies (Rawal et al., 2022; Sotirov & Dimitrov, 2024; Doneva & Dimitrov, 2024; Dimitrov et al., 2020; Li et al., 2023a; Edison et al., 2020), and IEDB which the protein were validated by more than two B cell assays and identified as negative antigen in Homo sapiens, and others are curated following established methods reported in prior research (Dimitrov et al., 2020). For instance, when processing **Immuno-Bacteria**, we first download the proteomes of all of the bacteria involved from the UniprotKB, then conduct a BLAST search to identify sequences with less than 30% sequence identity to the protective antigens (positive samples). To further improve the possibility that the selected sequences are true-negative samples, we scored these sequences with VAXIJEN and kept the sequence where it is predicted to be `Probable Non-Antigen`. Finally, we exclude sequences shorter than 25 AAs or longer than 1024 AAs, and the remaining sequences are the negative samples for **Immuno-Bacteria**. A similar preprocessing approach is applied to the preparation of the negative dataset for **Immuno-Virus**. For **Immuno-Tumor**, we only perform the redundancy removal step, as the negative instances have already been provided by the literature and source database. Note that BLAST and VAXIJEN are used here to ensure that these selected sequences are as true-negative as possible, but this approach cannot be directly extrapolated to predict the immunogenicity of antigens, because it will predict the vast majority of sequences to be positive, which is a clear deviation from the empirical reality.

In summary, the new benchmark dataset for immunogenicity prediction is distinguished by its rigorous quality control, comprehensive data sourcing, and cross-species diversity. **ImmunoDB** undergoes a series of rigorous quality control measures, which include data cleaning and preprocessing steps. It helps in minimizing missing information, erroneous entries, and inconsistencies in data structure, which are critical for reducing noise and enhancing reliability. Meanwhile, **ImmunoDB** provides the largest and most comprehensive benchmark datasets in terms of size and species. This richness in data enhances the benchmark’s applicability for developing and testing immunogenicity prediction models across various contexts. Moreover, including data from bacteria, viruses, and humans positions our dataset as a unique resource for cross-species analysis, thus supporting potential contributions to the robustness and generalizability of future models trained on them.

4 VENUSVACCINE: MODEL TRAINING AND INFERENCE

4.1 INPUT FEATURES

Sequence Embedding and Structure Tokenization We use a pre-trained PLM to extract embeddings for protein sequences and structures. The *AA sequence embedding* can be extracted from, an encoder-only model (e.g., ESM2; Lin et al. (2023)) or an encoder-decoder model (e.g., ANKH; El-naggar et al. (2023)). For instance, given a protein sequence of length L , ESM2-650M delivers an $L \times 1280$ representation, with each AA being represented as a 1280-dimensional vector. For protein structures, we use two tokenizers to encode the backbone predicted by ESMFOLD (Lin et al., 2023). The first *atomic structure tokenization* uses FOLDSEEK (Van Kempen et al., 2023) to analyze the spatial relationships between atoms of neighboring AAs at a finer scale and summarize them into 20-dimensional tokens. The second *peptide tokenization* uses the ESM3 (Hayes et al., 2024) structure tokenizer to map coarser-scale substructures to 4,096-dimensional structure tokens.

Physicochemical Descriptors We incorporate five hand-crafted E-descriptors (Venkatarajan & Braun, 2001) and three Z-descriptors (Hellberg et al., 1987) to quantitatively characterize protein sequences. These descriptors provide an AA-level summary of physicochemical properties, such as hydrophobicity and secondary structure. They have been validated to assist in immunogenicity prediction (Doytchinova et al., 2007; Rawal et al., 2022). Details are in Appendix A.

4.2 DUAL ATTENTION

A protein’s properties are determined by its sequence and structure. For immunogenicity prediction, focusing on local topological structures and molecular interactions may help identify exposed epitopes, thereby enhancing the extraction of critical information from the antigen. We combine hand-crafted features with intrinsic sequence and structural information derived from pre-trained models, as described above. These different categories of protein features are integrated using a dual attention mechanism, which guides the model to focus on important local regions of the proteins. The

features are then passed through attention pooling and propagated to the prediction head to obtain the predicted labels. The overall architecture of the model is shown in Figure 1. In the following sections, we detail the propagation rules of the dual attention and attention pooling modules.

Finer-lever and Coarser-level Attention Module We construct the dual attention mechanism with two components, including *finer-level attention* and *coarser-level attention*. Both components construct a multi-head cross-attention framework to integrate *keys* and *values* from the sequence and *queries* from the structure. Following Lin et al. (2023), we replace the relative positional encoding of queries and keys with *Rotary Position Embeddings* (RoPE) (Su et al., 2024). For a protein sequence of L AAs, *i.e.*, $\mathbf{x} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_L\}$, we employ a pre-trained language model (*e.g.*, ESM2) to extract the sequence encoding $\mathbf{E}_{\text{seq}} \in \mathbb{R}^{L \times d}$. Structural encodings are extracted by structure tokenizers of different scales. We define the finer-level structural encoding from FOLDSEEK as $\mathbf{E}_{\text{fine}} \in \mathbb{R}^{L \times d}$ and the coarser-level structural encoding from ESM3 as $\mathbf{E}_{\text{coarse}} \in \mathbb{R}^{L \times d}$. Taking the finer-level attention as an example, the module’s input is $\{\mathbf{E}_{\text{seq}}, \mathbf{E}_{\text{fine}}\}$, and the output is \mathbf{H}_{fine} . The queries \mathbf{Q} , keys \mathbf{K} , and values \mathbf{V} for the attention layers are defined as follows:

$$\mathbf{Q} = \mathbf{E}_{\text{fine}} \mathbf{W}_Q, \quad \mathbf{K} = \mathbf{E}_{\text{seq}} \mathbf{W}_K, \quad \mathbf{V} = \mathbf{E}_{\text{seq}} \mathbf{W}_V, \quad (2)$$

where $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V \in \mathbb{R}^{d \times d_k}$ are learnable projection matrices, and d_k is the hidden dimension of queries and keys. Next, the positional information is attached to \mathbf{Q} and \mathbf{K} with RoPE, which captures relative positional relationships by applying a rotation to the embeddings based on their positions. The modified queries and keys are:

$$\tilde{\mathbf{Q}}_i = \text{RoPE}(\mathbf{Q}_i, i), \quad \tilde{\mathbf{K}}_j = \text{RoPE}(\mathbf{K}_j, j), \quad (3)$$

where i and j denote the positions in the respective sequences, and $\text{RoPE}(\cdot)$ is the RoPE function. Finally, the cross-attention $\text{Attn}(\cdot)$ is formulated as:

$$\text{Attn}(\mathbf{E}_{\text{seq}}, \mathbf{E}_{\text{fine}}) = \text{softmax} \left(\frac{\tilde{\mathbf{Q}} \tilde{\mathbf{K}}^\top}{\sqrt{d_k}} \right) \mathbf{V}, \quad (4)$$

where $\tilde{\mathbf{Q}} \in \mathbb{R}^{L \times d_k}$ and $\tilde{\mathbf{K}} \in \mathbb{R}^{L \times d_k}$. The coarser-level attention applies the same construction. The only difference is that the input queries are replaced with $\mathbf{E}_{\text{coarser}}$ from ESM3.

Attention Pooling The processed AA-level features $\mathbf{H} = \text{concat}(\mathbf{E}_{\text{seq}}, \mathbf{H}_{\text{fine}}, \mathbf{H}_{\text{coarser}}, \mathbf{E}_{\text{ez}})$ are compressed using attention pooling and fed into the classifier for label prediction. Here, \mathbf{H} for a protein consists of four components: the original sequence embedding, the joint sequence and structure representations processed by coarser-level and finer-level attention, and the handcrafted E and Z descriptors. Consider a processed feature sequence $\mathbf{H} \in \mathbb{R}^{L \times D}$, L is the sequence length, and D is the dimension of concatenated features. The attention pooling on \mathbf{H} is defined as:

$$\text{AttnPool}(\mathbf{H}) = \sum_{i=1}^L \alpha_i \cdot \mathbf{H}_i, \quad \text{where } \alpha = \text{softmax}(\text{MaskedConv1D}(\mathbf{H})). \quad (5)$$

To obtain the final representation for the sequence, the model first computes $L \times 1$ attention scores using a one-dimensional masked convolution with a kernel size of 1 (Yang et al., 2023), then applies an input mask and passes it through a softmax function to obtain the attention weights α . These attention weights not only guide the pooling process but also indicate the importance of each position for the final prediction. In subsequent analysis, we extract these weights as epitope markers. Relevant analysis can be found in the post-hoc analysis section of the Experiments.

Prediction Head The output of the attention pooling $\text{AttnPool}(\mathbf{H})$ is sent directly to the final prediction head to produce the classification logits. This process involves projecting the pooled representation through a linear transformation, followed by a dropout layer, a ReLU activation layer, and a final linear layer: $\text{Logits}(\mathbf{H}) = \text{Linear}(\text{ReLU}(\text{DropOut}(\text{Linear}(\text{AttnPool}(\mathbf{H}))))))$.

5 EXPERIMENTS

5.1 EXPERIMENTAL PROTOCOL

Setup We built and trained the model according to the framework introduced in Section 4. For the sequence embedding, we consider three pre-trained PLMs, including ESM2 (Lin et al., 2023),

Table 1: Average performance of baseline methods in three immunogenicity prediction tasks from 10 random splits. The complete results with standard deviation can be found in Appendices.

Model	Bacteria					Virus					Tumor				
	ACC	T30	MCC	F1	KS	ACC	T30	MCC	F1	KS	ACC	T30	MCC	F1	KS
Random Forest	81.1	77.7	57.1	69.7	0.60	88.3	98.7	76.7	88.5	0.80	66.8	63.5	27.0	46.4	0.42
Gradient Boosting	80.7	75.9	56.8	71.2	0.62	86.0	97.3	72.3	86.5	0.74	65.4	60.0	26.0	53.1	0.38
XGBoost	80.8	76.1	57.0	71.1	0.61	89.1	98.8	78.2	89.2	0.80	69.2	62.2	33.7	56.1	0.40
SGD	78.6	72.4	51.0	65.3	0.56	77.3	88.0	56.0	74.0	0.66	52.6	50.9	29.6	62.2	0.31
Logistic Regression	65.2	67.3	1.8	0.5	0.47	61.4	82.9	35.6	71.9	0.61	60.4	45.2	0.0	0.0	0.25
MLP	78.1	71.6	51.5	68.1	0.57	85.0	90.6	70.5	85.7	0.71	60.4	39.1	0.0	0.0	0.26
SVM	56.7	57.3	18.9	53.0	0.33	61.6	88.7	25.1	67.3	0.53	51.4	61.3	15.2	57.3	0.33
KNN	80.4	73.8	55.3	68.6	0.52	87.4	86.8	74.8	87.3	0.75	57.4	49.6	10.8	45.3	0.12
Vaxi-DL	68.1	55.5	41.7	66.9	0.42	65.3	62.2	33.4	72.9	0.29	54.9	41.3	8.5	47.7	0.10
VaxiJen2.0	75.7	62.2	54.6	72.1	0.57	82.0	74.8	66.6	84.2	0.64	-	-	-	-	-
VaxiJen3.0	83.3	58.7	63.2	75.1	0.63	68.1	62.5	42.3	75.4	0.35	39.0	35.7	0.0	55.9	0.00
VirusImmu	-	-	-	-	-	58.8	59.5	18.1	63.6	0.17	-	-	-	-	-
ProVaccine-Ankh	82.3	78.5	60.3	73.3	0.64	92.2	99.7	84.3	92.3	0.85	76.9	73.5	55.0	73.5	0.61
ProVaccine-ESM2	80.6	77.0	57.0	71.7	0.66	90.3	99.5	80.6	90.6	0.82	74.0	61.7	46.1	68.2	0.58
ProVaccine-ProtBert	84.5	84.5	65.9	77.0	0.66	91.4	97.4	82.8	91.2	0.84	71.5	68.7	44.7	67.6	0.54

† The top three are highlighted by **First, Second, Third**.

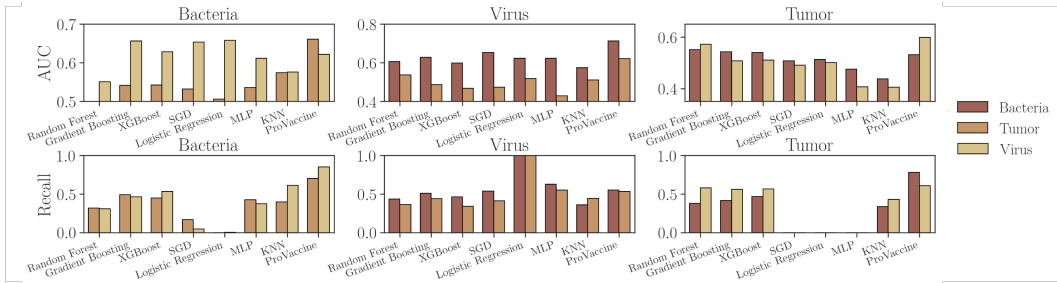


Figure 3: Generalizability of VENUSVACCINE by the AUC and Recall of cross-test evaluations. For instance, the top left figure reports the test performance on **Immuno-Bacteria** by different models trained on **Immuno-Virus** (yellow) and **Immuno-Tumor** (orange).

Ankh-large (Elnaggar et al., 2023), and ProtBert (Elnaggar et al., 2021). The source are listed in Appendix B. The model is optimized using ADAMW (Loshchilov et al., 2017) with a learning rate of 0.0005 and a weight decay of 0.01. In the attention pooling layer, dropout was set to 0.1. To ensure stable GPU memory usage during training, we limit the maximum number of tokens per batch to 4000, with a gradient accumulation of 1. The maximum training epoch was set to 50, with early stopping based on validation accuracy and patience of 5. All implementations were done using PyTorch (version 1.7.0), and experiments were run on an NVIDIA[®] RTX 3090 GPU with 24GB VRAM, mounted on a Linux server. The training process was tracked using WanDB. All source data and the implementation to reproduce the results will be publicly available upon acceptance.

Baseline Methods The performance of the proposed VENUSVACCINE is compared with existing machine learning methods and web server prediction tools. We select commonly used machine learning methods from the literature, including Random Forest, Gradient Boosting, XGBoost, SGD, Logistic Regression, MLP, SVM, and KNN, and follow the usual feature engineering approach by using ACC-transformed datasets to classify proteins as protective antigens (positive class) or non-protective antigens (negative class). We also include four popular web server methods specialized in antigen and vaccine candidate prediction: VAXIJEN2.0 (Doytchinova et al., 2007), VAXIJEN3.0 (Dimitrov et al., 2020), VAXI-DL (Rawal et al., 2022), and VIRUSIMMU (Li et al., 2023a). VAXI-JEN2.0 employs an ACC transformation and PLS algorithm for antigen prediction. VAXIJEN3.0 extends it to include machine learning for broader biological entities. VAXI-DL utilizes deep learning models tailored for different disease-causing agents, and VIRUSIMMU applies an ensemble machine learning method with soft voting for viral immunogenicity prediction.

5.2 RESULTS ANALYSIS ON IMMUNODB

We evaluate the overall performance of the models on VENUSVACCINE. In light of the problems with existing methods, we focus on the accuracy and generalization of each method’s predictive performance. The performance of the models is assessed using a range of metrics for a comprehensive evaluation, including AUC-ROC, accuracy, precision, recall, F1 score, MCC, Top K accuracy, cross-entropy, and KS statistics. Details are provided in Appendix C.

5.2.1 ACCURACY

We first investigate the overall performance of each model on the three datasets in **ImmunoDB**. For each dataset, we randomly split the data in a 7 : 1 : 2 ratio and select the model with the highest accuracy on the 10% validation set for evaluation. We repeat the assessment ten times for each test set, randomly selecting 50% of the test data to calculate the scores for each metric, and report the average scores in Table 1. See Tables 4-6 of the Appendix for additional details.

All machine learning and PLM methods we retrained ensured non-overlapping among the training, validation, and test sets. However, the four web servers do not have open-source code available for retraining on each task, which may introduce some unfair advantages to them. Additionally, during data preparation, we follow the convention in the field by using VAXIJEN3.0 to check negative data labels, giving this method a natural advantage in predicting true negative data. Furthermore, since VAXI-DL is not available for predicting antigens from tumors, we did not test its performance on the **Immuno**-tumor dataset. Similarly, VIRUSIMMU is intended only for testing viruses, so we did not report its performance on Bacteria and Tumors for fairness consideration.

It is clearly demonstrated by the performance comparison of baseline methods across different metrics and datasets that VENUSVACCINE has a significant advantage in prediction accuracy (indicated by top ACC). It achieves high accuracy for both positive and negative samples (indicated by high MCC and F1 scores) and effectively distinguishes the distribution differences between positive and negative samples (indicated by high KS scores). Notably, we specifically report the positive rate of top predictions (T30), *i.e.*, the ratio of the top 30 samples with the highest predicted positive probabilities that are true positives. This metric is crucial in predictive tasks involving biological experiments since, in real-world scenarios, the samples are often highly imbalanced, and researchers usually conduct experimental validations on the samples with the highest model prediction confidence. Therefore, a model should provide high confidence for true positive samples. However, it is also important to note that if a model has a high T30 but considerably low accuracy, it may indicate that it occasionally predicts positive samples as negative, risking the omission of pivotal vaccine candidates. To avoid biased evaluations, we advocate assessing models based on their overall performance across multiple metrics rather than overly focusing on a single measure.

We also investigate the influence of protein structure quality on model prediction with an ablation study comparing structures predicted by AlphaFold2 and ESMFold (see Table 9 in Appendix). The performance metrics, detailed in the appendix, conclusively show that AlphaFold2’s predicted structures lead to superior model performance, underscoring the importance of structural accuracy. While ESMFold’s rapid prediction capability is noteworthy, AlphaFold2’s higher precision in structure prediction significantly boosts model performance. This implies that when time is not a critical factor, it is advisable to opt for the more accurate structures provided by AlphaFold2. Unfortunately, the scarcity of available crystal structures limits our ability to contrast their performance with predicted ones, highlighting a need for more high-quality structural data in future studies.

5.2.2 GENERALIZABILITY

We next conduct cross-testing and compare the generalizability of different models (Figure 3). Here, the training and test sets are from different sources. The four web server-based methods are removed as we could not re-train them to avoid data leakage. We focus on AUC and recall to assess the model’s overall learning ability and its capacity to predict true positives. Generally, models trained on the Virus dataset demonstrate higher generalizability compared to those trained on the Bacteria dataset, and those trained on Tumor performed the worst. This may be related to the dataset size, as the number of antigens in Virus is larger than in Bacteria, and both are larger than Tumor.

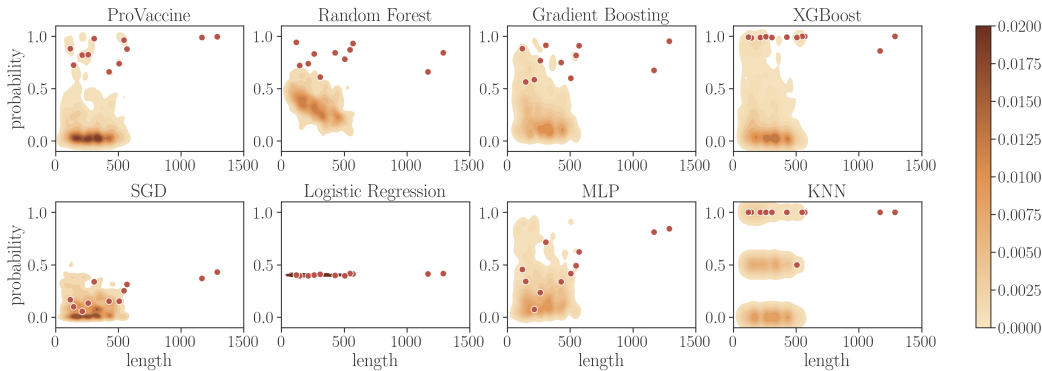


Figure 4: KDE of predicted immunogenicity scores on *Helicobacter pylori* candidates. The 11 experimentally determined immunogen are highlighted by red dots. Only VENUSVACCINE simultaneously identifies all positive samples while providing a reasonable overall distribution.

Moreover, excluding models with prediction issues (*e.g.*, logistic regression, where a 100% recall and a relatively small AUC suggest it trivially predicts all samples as negative), VENUSVACCINE consistently achieved top performance across the six test sets for both metrics. Additionally, the variation in VENUSVACCINE’s performance on the same test set when trained on different datasets was relatively small, which further suggests the model’s robustness in learning intrinsic patterns instead of overfitting to local variations in the data, thereby demonstrating stronger generalizabilities.

5.3 POST-HOC ANALYSIS

This section demonstrates the practical significance of VENUSVACCINE in addressing vaccine design challenges, as a supplement of the quantitative comparisons on well-established datasets and standard measurements. Two significant research subjects are selected to design corresponding ad-hoc analyses, including *Helicobacter pylori* and SARS-CoV-2. Specifically, in the first case, we examine *whether the potential vaccine candidates identified by VENUSVACCINE effectively cover immunogens*, which is frequently assessed by the enrichment of the model’s predictions. In the second case study, we evaluate *whether the ‘likely-effective’ sequences identified by VENUSVACCINE are consistent with experimental evidence, i.e., if the model assigns the top confidence to the experimentally validated vaccine development sequences*.

5.3.1 IDENTIFICATION OF POTENTIAL VACCINE CANDIDATES FOR HELICOBACTER PYLORI

In the first challenge, we use VENUSVACCINE to predict immunogens within the *Helicobacter pylori* proteome. *Helicobacter pylori* is a Group I carcinogen that can persist in the gastric mucosa, causing various gastrointestinal diseases (Malfertheiner et al., 2023). Identifying effective vaccine targets and developing corresponding vaccines can prevent *Helicobacter pylori* infections, thus significantly reducing the incidence of gastric diseases and lowering the rate of gastric cancer (Dos Santos Viana et al., 2021). To this end, we extract the proteome data from the **UniProtKB** database (ver 2024_04) (Consortium, 2019). A total number of 1,858 sequences are extracted after redundancy removal. Although the majority of these 1,858 sequences lack immunogenicity ground truth labels, we refer to the work of Dalsass et al. (2019), which record 11 of these sequences as protective antigens. We use the model trained on **Immuno-Bacteria** for inference since *Helicobacter pylori* is a bacterium. Note that there is no overlap between the training and prediction datasets.

When evaluate the prediction performance, two measurements are considered. (1) *Recall* ((9) in Appendix C) measures the fraction of true protective antigens identified within the set of PVCs predicted by VENUSVACCINE. It is a critical indicator of the model’s ability to accurately retrieve known protective antigens from a larger set of candidates. (2) *Fold-Enrichment* (Dalsass et al., 2019) ((15) in Appendix C) measures the ratio of the true positive rate to the overall prevalence of the condition, characterizing the model’s ability to enrich for positive samples.

To evaluate the performance of the model, we compare the results of VENUSVACCINE with other machine learning approaches. Table 16 compares the recall and enrichment of baseline methods, and Figure 6 visualizes the predicted kernel density estimation (KDE) of immunogenic probabilities for candidates of varying sequence lengths, with the 11 immunogens highlighted in red dots. We exclude the VaxiJen series as they only provide web servers, in which case we could not confirm the absence of data leakage issues in the training data. Among the 1,858 candidates, VENUSVACCINE identifies 123 PVCs. Notably, all the 11 true protective antigens are included within these 123 PVCs. A recall rate of 100% demonstrates VENUSVACCINE’s exceptional ability to extract known protective antigens from the proteome data. In comparison, Random Forest, Gradient Boosting, and XGBoost also achieve a 100% recall rate, while other baseline models do not perform as well. Furthermore, VENUSVACCINE’s fold-enrichment is 15.11, while the highest fold-enrichment among Random Forest, Gradient Boosting, and XGBoost is only 8.81. This indicates that VENUSVACCINE has a stronger predictive ability to enrich for positive samples, which is particularly beneficial when experimental testing is limited to a finite number of samples.

In addition, from this result in Figure 6, three observations could be concluded: (1) The methods in the second row have trouble correctly labeling all the protective antigens. (2) The overall probability distribution found by VENUSVACCINE aligns more closely with reality, with most samples clustering around probabilities close to 0, implying the majority of candidates are non-immunogenic. (3) VENUSVACCINE retains fewer PVCs from the candidates than baseline methods, allowing for the identification of the 11 protective antigens with fewer experimental efforts. The superior performance of VENUSVACCINE could be attributed to its effective integration of multimodal features by the dual-attention mechanism, which captures the complex interplay between sequence, structure, and physicochemical information. In summary, our model exhibits strong robustness and considerable potential to enhance the discovery of immunogenic vaccine targets in addressing the challenges posed by *Helicobacter pylori*.

5.3.2 UNRAVELING SARS-CoV-2 PROTEOME FOR VACCINE TARGET DISCOVERY AND EXPLANATION

In the second challenge, we aim to investigate if our VENUSVACCINE could effectively identify clinically validated vaccine targets from the SARS-CoV-2 proteome. The candidate test set consists of 38 sequences from the SARS-CoV-2 Data Hub (Wheeler et al., 2007). For the ground truth vaccine targets, we refer to the nine vaccines approved by the World Health Organization (Saravanan et al., 2024), five of which use the surface glycoprotein (spike protein) of the novel coronavirus as the vaccine target, while the other four use the whole virus, which also includes the spike protein. The overall prediction performance is in Table 17 in the Appendix, where VENUSVACCINE ranks the most important protein (spike protein) as the top among the 38 proteins. The model also identifies some new potential vaccine targets, *e.g.*, non-structural protein 9 (Nsp9).




Figure 5: Epitope marker on the surface glycoprotein (NCBI ID: YP_009724390.1) by the attention score identifies vaccine targets.

To further validate and understand our predictions, we visualize the prediction results. Figure 5 maps the model’s attention scores to the antigenic epitopes of the sequences and compares them with experimentally obtained epitopes (Shrock et al., 2020; Polyiam et al., 2021; Braun et al., 2024). It shows clearly that the model not only demonstrates excellent predictive performance but also possesses promising interpretability. For clear presentation, we display the chain trace of the spike protein and color the amino acids based on the attention scores from the attention pooling layer, with darker positions indicating higher attention scores. Overall, the model shows higher attention scores around position 400, and the spike protein contains three immunodominant linear B cell epitopes from position 319 to 541. Additionally, there are three regions with deeper colors that overlap with the SARS-CoV-2 spike receptor-binding domain (RBD), including positions 150-164, 291-320, and 439-478. This visualization provides a deeper understanding of the model’s prediction process and enhances the credibility of our approach in identifying potential vaccine targets.

6 RELATED WORK

Developing machine learning models has become a popular approach for many predictive problems in vaccine design. One of the most widely used web servers for identifying potential vaccine candidates is VAXIJEN (Doytchinova et al., 2007), which employs an alignment-free method based on auto cross-covariance (ACC) transformation to identify antigens without relying on sequence alignment. Subsequent methods have introduced additional input features (He et al., 2010) or improved prediction models (Edison et al., 2020; Zhang et al., 2023; Rawal et al., 2022). Compared to other reverse vaccinology tools, these advanced techniques demonstrate superior performance in predicting bacterial protective antigens. Recently, Dimitrov *et al.* explored the immunogenicity of antigens from other sources, such as viruses (Doneva & Dimitrov, 2024) and tumors (Sotirov & Dimitrov, 2024), and evaluated various popular machine learning methods.

These machine learning approaches for predicting immunogenicity have often relied on limited datasets, typically confined to a specific type of pathogen, such as bacteria or viruses, and have been relatively small in scale. For instance, Doytchinova et al. (2007) trained and validated their models using only a few hundred samples. The scarcity of data not only impedes the development of robust predictive models but also restricts the ability to capture the complex immunological responses elicited by diverse antigens. To date, few studies have managed to amalgamate a comprehensive dataset that encompasses antigens from multiple pathogen classes, which is essential for training models capable of predicting immunogenicity with high fidelity across varied biological contexts.

Despite the rapid development of deep learning, the advancement of methodologies for protein property prediction, including immunogenicity prediction, has been relatively slow due to the limited amount of labeled data. Recently, BERT-style PLMs (Lin et al., 2023; Li et al., 2024) have gained increasing attention with an increasing number of prediction tasks fine-tuning models to better fit the specific targets (Schmirler et al., 2023; Li et al., 2023b; Tan et al., 2024b), or attaching additional propagators to process sequence embeddings, such as geometric encoders (Yang et al., 2023; Tan et al., 2023; Zhou et al., 2023; 2024b) or cross-modality aggregators (Tan et al., 2024a;c).

7 CONCLUSION AND DISCUSSION

Despite the rapid advancement of deep learning in providing novel solutions to numerous problems in scientific fields such as protein engineering, the application of advanced models remains largely limited to well-defined engineering problems with relatively mature and abundant data. However, many other issues with significant scientific and practical value have not received sufficient attention, and the methods used for these problems are considerably outdated. In addition to the oversight of important problems due to a lack of interdisciplinary cross-talk between research fields, another major reason is the absence of high-quality datasets and adequate evaluation methods for many specific issues. Preparing and processing datasets for model training and evaluation requires significant effort and domain knowledge, and validating the developed methods also requires feasible evaluation protocols that fit practical requirements. Ideally, these protocols should neither be as disconnected from practical needs as current computational metrics (*e.g.*, perplexity for sequence generation) nor as resource-intensive as wet lab validation processes in terms of time, money, and technical barriers.

As an exploratory work, this study takes immunogenicity prediction in vaccine development as an entry point to investigate how to prepare data, build models, and design evaluation strategies for a scientific problem that seeks powerful deep learning solutions. Although protein representation can be encoded by many pre-trained models, and immunogenicity prediction is inherently one supervised learning task, we emphasize that utilizing deep learning to solve scientific problems should not only focus on algorithm design but also on dataset construction and the practicality of evaluation methods. Here, we demonstrate the overall performance of the model using standard datasets and conventional quantitative evaluation methods, and analyze the model’s reliability in specific applications using post-hoc analysis without conducting entirely new wet lab experiments.

ACKNOWLEDGMENTS

This work was supported by the grants from the National Natural Science Foundation of China (Grant Number 62302291; 12104295), the Computational Biology Key Program of Shanghai Science and Technology Commission (23JS1400600), Shanghai Jiao Tong University Scientific and

Technological Innovation Funds (21X010200843), the Student Innovation Center at Shanghai Jiao Tong University, and Shanghai Artificial Intelligence Laboratory.

REFERENCES

- Jeannette Adu-Bobie, Barbara Capecchi, Davide Serruto, Rino Rappuoli, and Mariagrazia Pizza. Two years into reverse vaccinology. *Vaccine*, 21(7-8):605–610, 2003.
- Ruth Arnon. Overview of vaccine strategies. In *Vaccine Design. Innovative Approaches and Novel Strategies*; Rappuoli, R., Ed.; Caister Academic Press: Norfolk, UK, 2011.
- F. Bagnoli, N. Norais, I. Ferlenghi, M. Scarselli, and R. Rappuoli. *Designing vaccines in the era of genomics*. Vaccine Design: Innovative Approaches and Novel Strategies, 2011.
- Asolina Braun, Louise C Rowntree, Ziyi Huang, Kirti Pandey, Nikolas Thuesen, Chen Li, Jan Petersen, Dene R Littler, Shabana Raji, Thi HO Nguyen, et al. Mapping the immunopeptidome of seven sars-cov-2 antigens across common hla haplotypes. *Nature Communications*, 15(1):7547, 2024.
- Thiago Carvalho, Florian Krammer, and Akiko Iwasaki. The first 12 months of covid-19: a timeline of immunological insights. *Nature reviews. Immunology*, pp. 1–12, 2021.
- Uni Prot Consortium. Uniprot: a worldwide hub of protein knowledge. *Nucleic Acids Research*, (D1):D506–D515, 2019.
- Mattia Dalsass, Alessandro Brozzi, Duccio Medini, and Rino Rappuoli. Comparison of open-source reverse vaccinology programs for bacterial vaccine antigen discovery. *Frontiers in Immunology*, 10, 2019.
- Ivan Dimitrov, Nevena Zaharieva, and Irini Doytchinova. Bacterial immunogenicity prediction by machine learning methods. *Vaccines*, 8(4):709, 2020.
- Nikolet Doneva and Ivan Dimitrov. Viral immunogenicity prediction by machine learning methods. *International Journal of Molecular Sciences*, 25(5), 2024.
- Nikolet Doneva, Irini Doytchinova, and Ivan Dimitrov. Predicting immunogenicity risk in biopharmaceuticals. *Symmetry*, 13(3), 2021.
- Irineu Dos Santos Viana, Maria Luisa Cordeiro Santos, Hanna Santos Marques, Vinicius Lima de Souza Goncalves, Breno Bittencourt de Brito, Filipe Antonio Franca da Silva, Natalia Oliveira e Silva, Fillipe Dantas Pinheiro, Adriano Fernandes Teixeira, Davi Tanajura Costa, et al. Vaccine development against helicobacter pylori: From ideal antigens to the current landscape. *Expert Review of Vaccines*, 20(8):989–999, 2021.
- Irini A Doytchinova, Darren R Flower, Irini A Doytchinova, and Darren R Flower. Vaxijen: a server for prediction of protective antigens, tumour antigens and subunit vaccines. *BMC Bioinformatics*, 2007.
- Ong Edison, Wang Haihe, Mei U Wong, Seetharaman Meenakshi, Valdez Ninotchka, and He Yongqun. Vaxign-ml: Supervised machine learning reverse vaccinology model for improved prediction of bacterial protective antigens. *Bioinformatics*, (10):10, 2020.
- Ahmed Elnaggar, Michael Heinzinger, Christian Dallago, Ghalia Rehawi, Yu Wang, Llion Jones, Tom Gibbs, Tamas Feher, Christoph Angerer, Martin Steinegger, et al. Protrans: Toward understanding the language of life through self-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):7112–7127, 2021.
- Ahmed Elnaggar, Hazem Essam, Wafaa Salah-Eldin, Walid Moustafa, Mohamed Elkerdawy, Charlotte Rochereau, and Burkhard Rost. Ankh: Optimized protein language model unlocks general-purpose modelling. *arXiv:2301.06568*, 2023.
- Noelia Ferruz, Steffen Schmidt, and Birte Höcker. ProtGPT2 is a deep unsupervised language model for protein design. *Nature Communications*, 13(1):4348, 2022.

- Tomas Hayes, Roshan Rao, Halil Akin, Nicholas J Sofroniew, Deniz Oktay, Zeming Lin, Robert Verkuil, Vincent Q Tran, Jonathan Deaton, Marius Wiggert, et al. Simulating 500 million years of evolution with a language model. *bioRxiv*, pp. 2024–07, 2024.
- Yongqun He, Zuoshuang Xiang, and Harry L T Mobley. Vaxign: The first web-based vaccine design program for reverse vaccinology and applications for vaccine development. *Biomed Research International*, 2010(1110-7243):297505, 2010.
- Michael Heinzinger, Konstantin Weissenow, Joaquin Gomez Sanchez, Adrian Henkel, Martin Steinegger, and Burkhard Rost. ProstT5: Bilingual language model for protein sequence and structure. *bioRxiv*, pp. 2023–07, 2023.
- Sven Hellberg, Michael Sjoestroem, Bert Skagerberg, and Svante Wold. Peptide quantitative structure-activity relationships, a multivariate approach. *Journal of Medicinal Chemistry*, 30(7): 1126–1135, 1987.
- Jing Li, Zhongpeng Zhao, Chengzheng Tai, Ting Sun, Lingyun Tan, Xinyu Li, Wei He, Hongjun Li, and Jing Zhang. Virusimmu: a novel ensemble machine learning approach for viral immunogenicity prediction. *bioRxiv*, pp. 2023–11, 2023a.
- Mingchen Li, Liqi Kang, Yi Xiong, Yu Guang Wang, Guisheng Fan, Pan Tan, and Liang Hong. Sesnet: sequence-structure feature-integrated deep learning method for data-efficient protein engineering. *Journal of Cheminformatics*, 15(1):12, 2023b.
- Mingchen Li, Yang Tan, Xinzhu Ma, Bozitao Zhong, Huiqun Yu, Ziyi Zhou, Wanli Ouyang, Bingxin Zhou, Pan Tan, and Liang Hong. ProSST: Protein language modeling with quantized structure and disentangled attention. *Advances in Neural Information Processing Systems*, 37:35700–35726, 2024.
- Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, 2023.
- Ilya Loshchilov, Frank Hutter, et al. Fixing weight decay regularization in adam. *arXiv:1711.05101*, 5, 2017.
- Ali Madani, Ben Krause, Eric R Greene, Subu Subramanian, Benjamin P Mohr, James M Holton, Jose Luis Olmos, Caiming Xiong, Zachary Z Sun, Richard Socher, et al. Large language models generate functional protein sequences across diverse families. *Nature Biotechnology*, 41(8):1099–1106, 2023.
- Peter Malfertheiner, M Constanza Camargo, Emad El-Omar, Jyh-Ming Liou, Richard Peek, Christian Schulz, Stella I Smith, and Sebastian Suerbaum. Helicobacter pylori infection. *Nature Reviews Disease Primers*, 9(1):19, 2023.
- Joshua Meier, Roshan Rao, Robert Verkuil, Jason Liu, Tom Sercu, and Alex Rives. Language models enable zero-shot prediction of the effects of mutations on protein function. *Advances in Neural Information Processing Systems*, 34:29287–29303, 2021.
- Pizza and M. Identification of vaccine candidates against serogroup b meningococcus by whole-genome sequencing. *Science*, 287(5459):1816–20, 2000.
- Kanokporn Polyiam, Waranyoo Phoolcharoen, Namphueng Butkhot, Chanya Srisaowakarn, Arunee Thitithanyanont, Prasert Auewarakul, Tawatchai Hoonsuwan, Marasri Ruengjitchachawalya, Phenjun Mekvichitsaeng, and Yaowaluck Maprang Roshorm. Immunodominant linear b cell epitopes in the spike and membrane proteins of sars-cov-2 identified by immunoinformatics prediction and immunoassay. *Scientific Reports*, 11(1):20383, 2021.
- Roshan M Rao, Jason Liu, Robert Verkuil, Joshua Meier, John Canny, Pieter Abbeel, Tom Sercu, and Alexander Rives. MSA transformer. In *International Conference on Machine Learning*, pp. 8844–8856. PMLR, 2021.

- Kamal Rawal, Robin Sinha, Swarsat Kaushik Nath, P Preeti, Priya Kumari, Srijaanee Gupta, Trapti Sharma, Ulrich Strych, Peter Hotez, and Maria Elena Bottazzi. Vaxi-dl: A web-based deep learning server to identify potential vaccine candidates. *Computers in Biology and Medicine*, 145:105401, 2022.
- Venkatesan Saravanan, Bharath Kumar Chagaleti, and Pavithra Lakshmi NarayananVijay Babu AnandanHaritha ManoharanG.V. AnjanaRamalingam PeramanS. Karthik Raja NamasivayamM. KavisriJesu ArockiarajKathiravan MuthuKumaradossMeivelu Moovendhan. Discovery and development of covid-19 vaccine from laboratory to clinic. *Chemical Biology and Drug Design*, 103(1), 2024.
- Robert Schmirler, Michael Heinzinger, and Burkhard Rost. Fine-tuning protein language models boosts predictions across diverse tasks. *bioRxiv*, pp. 2023–12, 2023.
- Ellen Shrock, Eric Fujimura, Tomasz Kula, Richard T. Timms, and Stephen J. Elledge. Viral epitope profiling of covid-19 patients reveals cross-reactivity and correlates of severity. *Science*, 2020.
- Stanislav Sotirov and Ivan Dimitrov. Application of machine learning algorithms for prediction of tumor t-cell immunogens. *Applied Sciences*, 14(10), 2024.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.
- Jin Su, Chenchen Han, Yuyang Zhou, Junjie Shan, Xibin Zhou, and Fajie Yuan. Saprot: protein language modeling with structure-aware vocabulary. In *The Twelfth International Conference on Learning Representations*, 2023.
- Yang Tan, Bingxin Zhou, Lirong Zheng, Guisheng Fan, and Liang Hong. Semantical and topological protein encoding toward enhanced bioactivity and thermostability. *bioRxiv*, pp. 2023–12, 2023.
- Yang Tan, Mingchen Li, Bingxin Zhou, Bozitao Zhong, Lirong Zheng, Pan Tan, Ziyi Zhou, Huiqun Yu, Guisheng Fan, and Liang Hong. Simple, efficient, and scalable structure-aware adapter boosts protein language models. *Journal of Chemical Information and Modeling*, 2024a.
- Yang Tan, Ruilin Wang, Banghao Wu, Liang Hong, and Bingxin Zhou. Retrieval-enhanced mutation mastery: Augmenting zero-shot prediction of protein language model. *arXiv:2410.21127*, 2024b.
- Yang Tan, Jia Zheng, Liang Hong, and Bingxin Zhou. ProtSolM: Protein solubility prediction with multi-modal features. In *2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 223–232. IEEE, 2024c.
- Michel Van Kempen, Stephanie S Kim, Charlotte Tumescheit, Milot Mirdita, Jeongjae Lee, Cameron LM Gilchrist, Johannes Söding, and Martin Steinegger. Fast and accurate protein structure search with foldseek. *Nature Biotechnology*, pp. 1–4, 2023.
- Mathura S. Venkatarajan and Werner Braun. New quantitative descriptors of amino acids based on multidimensional scaling of a large number of physical–chemical properties. *Molecular Modeling Annual*, 7(12):445–453, 2001.
- David L Wheeler, Tanya Barrett, Dennis A Benson, Stephen H Bryant, Kathi Canese, Vyacheslav Chetvernin, Deanna M Church, Michael Dicuccio, Ron Edgar, and Scott Federhen. Wheeler, d.l. et al. database resources of the national center for biotechnology information. nucleic acids res. 33 database issue: D39-d45. *Nucleic Acids Research*, 35(Database issue):D5–12, 2007.
- Kevin K Yang, Niccolò Zanichelli, and Hugh Yeh. Masked inverse folding with sequence transfer for protein representation learning. *Protein Engineering, Design and Selection*, 36:gza015, 2023.
- Yuhan Zhang, Anthony Huffman, Justin Johnson, and Yongqun He. Vaxign-DL: A deep learning-based method for vaccine design and its evaluation. *bioRxiv*, 2023.
- Bingxin Zhou, Lirong Zheng, Banghao Wu, Yang Tan, Outongyi Lv, Kai Yi, Guisheng Fan, and Liang Hong. Protein engineering with lightweight graph denoising neural networks. *Journal of Chemical Information and Modeling*, 2023.

Bingxin Zhou, Yang Tan, Yutong Hu, Lirong Zheng, Bozitao Zhong, and Liang Hong. Protein engineering in the deep learning era. *mLife*, 3(4):477–491, 2024a.

Bingxin Zhou, Lirong Zheng, Banghao Wu, Kai Yi, Bozitao Zhong, Yang Tan, Qian Liu, Pietro Liò, and Liang Hong. A conditional protein diffusion model generates artificial programmable endonuclease sequences with enhanced activity. *Cell Discovery*, 10(1):95, 2024b.

A DETAILS ON PHYSICOCHEMICAL DESCRIPTORS

To provide protein-level properties, this study follows the convention and employs E-descriptors and Z-descriptors to quantitatively characterize protein sequences.

Venkatarajan & Braun (2001) proposed the E-descriptors in 2001. For each of the 20 naturally occurring amino acids, five numerical values are derived based on the principal component analysis (PCA) of 237 physicochemical properties. The first component E1 strongly correlates with the hydrophobicity of amino acids. E2 provides information about molecular size and steric properties. Components E3 and E5 describe the propensity of amino acids to occur in α -helices and β -strands, respectively. E4 takes into account the partial specific volume, the number of codons, and the relative frequency of amino acids in proteins.

The Z-descriptors, defined by Hellberg et al. (1987), summarize the principal physicochemical properties of amino acids. These descriptors are derived by PCA of a data matrix consisting of 29 molecular descriptors such as molecular weight, pKas, NMR shifts, etc. The first component Z1 reflects the hydrophobicity of amino acids, Z2 their size, and Z3 their polarity. The E and Z descriptors have been widely used for the characterization and classification of proteins and are good predictors of immunogenicity Doytchinova et al. (2007); Rawal et al. (2022).

Table 2: Description of E-descriptors and Z-descriptors.

Name	Description
E1	hydrophobicity of amino acids
E2	molecular size and steric properties
E3	the propensity of amino acids to occur in α -helices
E4	the partial specific volume, the number of codons, and the relative frequency of amino acids in proteins
E5	the propensity of amino acids to occur in β -strands
Z1	the hydrophobicity of amino acids
Z2	size of amino acids
Z3	polarity of amino acids

B BASELINE METHODS

Here we list the sources for running the comparisons in the Experiment section. The machine learning methods are implemented using SKLEARN, the web server prediction tools are obtained from their official websites, and the pre-trained PLMs (for sequence embedding) are sourced from their respective open-source GitHub repositories. Table 3 provides details for accessing these methods.

Table 3: Source of baseline methods.

Name	version	Source
VAXI-DL	-	https://vac.kamalrawal.in/vaxidl/
VAXIJEN2.0	-	http://www.ddg-pharmfac.net/vaxijen/VaxiJen/VaxiJen.html
VAXIJEN3.0	-	https://www.ddg-pharmfac.net/vaxijen3/home/
VIRUSIMMU	-	https://github.com/zhangjbig/VirusImmu
ESM2	t33_650M_UR50D	https://huggingface.co/facebook/esm2_t33_650M_UR50D
Ankh	large	https://huggingface.co/ElnaggarLab/ankh-large
ProtBert	-	uniref&https://huggingface.co/Rostlab/prot_bert

C EVALUATION METRICS

We employ a list of evaluation metrics to provide a comprehensive comparison of different prediction methods.

- AUC-ROC is a performance metric for classification models that measures the ability of the model to distinguish between classes. The ROC curve plots the True Positive Rate (TPR) against the

False Positive Rate (FPR) at various threshold settings. The AUC represents the area under this curve and provides an aggregate measure of performance across all thresholds.

$$\text{AUC} = P(\text{rank}(X_{\text{positive}}) > \text{rank}(X_{\text{negative}})) \quad (6)$$

- Accuracy measures the proportion of correct predictions made by the model out of all predictions.

$$\text{Accuracy} = \frac{\text{True Positives} + \text{True Negatives}}{\text{Total Predictions}} = \frac{TP + TN}{TP + TN + FP + FN} \quad (7)$$

- Precision quantifies the accuracy of positive predictions made by the model. It measures the proportion of true positives among all instances that were predicted as positive. High precision indicates that the model has a low rate of false positives.

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}} = \frac{TP}{TP + FP} \quad (8)$$

- Recall also known as Sensitivity or True Positive Rate, quantifies the ability of the model to identify all relevant instances.

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}} = \frac{TP}{TP + FN} \quad (9)$$

- F1 score is the harmonic mean of Precision and Recall, providing a balance between the two metrics, especially when dealing with imbalanced datasets.

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (10)$$

- MCC is a balanced measure for binary classifications, even if the classes are of different sizes. It considers true and false positives and negatives to produce a coefficient between -1 and +1.

$$\text{MCC} = \frac{(TP \times TN) - (FP \times FN)}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (11)$$

- Top-K accuracy measures the fraction of times the correct label is among the model’s top $K\%$ predictions.

$$\text{Top-K Accuracy} = \frac{\text{Number of correct predictions in top-K}}{\text{Number in top-K}} \quad (12)$$

- Cross-entropy quantifies the difference between two probability distributions: the true distribution and the predicted distribution by the model. y is the true label (1 for positive class, 0 for negative class). p is the predicted probability of the positive class.

$$\text{Cross-Entropy Loss} = -[y \log(p) + (1 - y) \log(1 - p)] \quad (13)$$

- KS statistics measures the maximum difference between the cumulative distribution functions (CDFs) of the positive and negative classes. It’s used to evaluate the discriminatory power of a scoring model.

$$\text{KS} = \max_x |F_{\text{positive}}(x) - F_{\text{negative}}(x)| \quad (14)$$

- Fold-Enrichment measures the ratio of the true positive rate to the overall prevalence of the condition, characterizing the model’s ability to enrich for positive samples.

$$\text{FE} = \frac{TP / (TP + NP)}{(TP + FN) / (TP + FP + TN + FN)} \quad (15)$$

D COMPLETE BASELINE COMPARISON ON IMMUNODB

Tables 4-6 present the detailed performance of all baseline methods, as well as our VENUSVACCINE, on the three benchmark datasets in **ImmunoDB**. The evaluation procedure follows the same methodology as described in the main text. We repeated the assessment on the test set for 10 times, with each time randomly selecting 50% of the samples from the test set, and reported the average and standard deviation across these 10 runs in the tables. Due to limited space, in the main text we only report the average scores of the most important evaluation metrics. Here, we provide a more comprehensive summary, including additional measurements and confidence intervals, to allow for a more thorough comparison of the baseline models.

Table 4: Average performance of baseline methods in **bacteria** immunogenicity prediction tasks from 10 random splits, with the standard deviation reported after \pm .

Model	AUC-ROC	Accuracy	Precision	Recall	F1	MCC	TopK-30	Cross Entropy	KS statistic
RF	87.2±1.8	81.1±1.8	78.9±3.3	62.6±3.2	69.7±2.5	57.1±3.5	77.7±3.7	0.45±0.02	0.60±0.03
GBT	85.8±2.3	80.7±2.0	74.0±4.0	68.8±2.8	71.2±2.9	56.8±4.2	75.9±4.3	0.44±0.03	0.62±0.04
XGBoost	87.1±1.7	80.8±1.5	74.9±2.7	67.7±2.6	71.1±2.1	57.0±2.8	76.1±3.9	0.52±0.06	0.61±0.04
SGD	83.2±2.7	78.6±1.8	74.7±4.0	58.1±3.3	65.3±3.2	51.0±4.1	72.4±4.5	0.61±0.07	0.56±0.05
Logit Reg	77.7±3.0	65.2±2.6	20.0±4.0	0.2±0.5	0.5±0.9	1.8±3.5	67.3±5.9	0.64±0.01	0.47±0.04
MLP	82.4±2.9	78.1±1.5	69.0±3.9	67.3±2.9	68.1±2.8	51.5±3.5	71.6±5.0	0.49±0.03	0.57±0.05
SVM	68.1±2.5	56.7±2.1	42.6±3.7	70.3±3.7	53.0±3.6	18.9±4.4	57.3±4.8	0.64±0.01	0.33±0.04
KNN	82.5±2.4	80.4±1.7	77.2±4.4	61.9±3.6	68.6±3.1	55.3±4.1	73.8±5.8	3.95±0.69	0.52±0.04
Vexi	71.2±1.8	68.1±1.9	55.4±3.2	84.4±2.2	66.9±2.7	41.7±3.5	55.5±4.1	11.51±0.69	0.42±0.04
VaxiJen2.0	78.5±1.5	75.7±1.7	61.0±2.3	88.3±1.8	72.1±1.9	54.6±2.8	62.2±3.0	8.76±0.61	0.57±0.03
VaxiJen3.0	81.4±1.8	83.3±1.5	76.9±3.5	75.1±2.8	76.0±2.7	63.2±3.6	58.6±4.2	3.05±0.55	0.63±0.04
VENUSVACCINE-Ankh	86.6±2.1	82.3±2.0	76.8±5.0	70.4±3.2	73.3±3.0	60.3±4.6	78.5±4.5	0.97±0.10	0.64±0.03
VENUSVACCINE-ESM2	87.6±1.9	80.6±2.0	73.3±3.1	70.2±4.0	71.7±3.3	57.0±4.6	77.0±4.4	0.46±0.05	0.66±0.03
VENUSVACCINE-ProtBert	88.7±1.1	84.5±1.5	83.8±1.6	71.3±3.0	77.0±2.2	65.9±3.0	84.5±1.6	0.51±0.04	0.66±0.02

† The top three are highlighted by **First, Second, Third**.Table 5: Average performance of baseline methods in **virus** immunogenicity prediction tasks from 10 random splits, with the standard deviation reported after \pm .

Model	AUC-ROC	Accuracy	Precision	Recall	F1	MCC	TopK-30	Cross Entropy	KS statistic
RF	96.5±0.3	88.3±0.7	86.4±0.9	90.8±0.8	88.5±0.6	76.7±1.4	98.7±0.8	0.29±0.01	0.80±0.01
GBT	94.7±0.5	86.0±0.7	83.1±1.2	90.2±1.2	86.5±0.7	72.3±1.4	97.3±1.3	0.31±0.01	0.74±0.02
XGBoost	96.5±0.2	89.1±0.4	87.1±1.2	91.5±0.8	89.2±0.6	78.2±0.8	98.8±0.9	0.27±0.01	0.80±0.00
SGD	89.0±1.3	77.3±1.8	85.4±2.4	65.3±2.9	74.0±2.6	56.0±3.7	88.0±2.9	0.52±0.01	0.66±0.02
Logit Reg	85.7±1.4	61.4±1.0	56.3±1.4	99.6±0.2	71.9±1.1	35.6±1.5	82.9±2.8	0.68±0.00	0.61±0.02
MLP	91.2±1.0	85.0±1.2	81.2±1.7	90.8±1.6	85.7±1.3	70.5±2.3	90.6±2.5	0.38±0.02	0.71±0.02
SVM	76.8±1.8	61.6±1.2	58.3±1.4	79.6±2.1	67.3±1.1	25.1±3.1	88.7±1.9	0.71±0.01	0.53±0.03
KNN	90.6±1.0	87.4±0.8	87.4±1.5	87.2±0.9	87.3±0.8	74.8±1.6	86.8±2.3	2.91±0.33	0.75±0.02
Vexi	64.3±1.7	65.3±1.7	61.4±1.7	89.7±1.8	72.9±1.7	33.4±3.5	62.2±2.4	12.49±0.61	0.29±0.03
VaxiJen2.0	82.1±1.0	82.0±1.2	75.1±1.8	95.7±1.0	84.2±1.3	66.6±2.1	74.8±3.4	6.47±0.44	0.64±0.02
VaxiJen3.0	67.5±1.2	68.1±1.8	62.4±2.1	95.3±1.5	75.4±1.8	42.3±3.0	62.5±4.2	11.48±0.64	0.35±0.02
VirusImmu	58.7±2.2	58.8±2.1	56.9±2.5	72.2±2.1	63.6±2.1	18.1±4.4	59.5±4.6	14.86±0.76	0.17±0.04
VENUSVACCINE-Ankh	97.2±0.6	92.2±1.3	91.9±1.5	92.6±2.2	92.3±1.5	84.3±2.7	99.7±0.4	0.46±0.08	0.85±0.03
VENUSVACCINE-ESM2	96.7±0.7	90.3±1.3	89.1±1.5	92.2±1.2	90.6±1.3	80.6±2.7	99.5±0.4	0.38±0.07	0.82±0.02
VENUSVACCINE-ProtBert	96.9±0.4	91.4±0.8	92.8±1.5	89.8±1.0	91.2±0.7	82.8±1.5	97.4±0.6	0.40±0.03	0.84±0.01

† The top three are highlighted by **First, Second, Third**.Table 6: Average performance of baseline methods in **Tumor** immunogenicity prediction tasks from 10 random splits, with the standard deviation reported after \pm .

Model	AUC-ROC	Accuracy	Precision	Recall	F1	MCC	TopK-30	Cross Entropy	KS statistic
RF	74.4±5.3	66.8±4.2	63.6±7.9	37.0±8.4	46.4±8.3	27.0±9.7	63.5±5.2	0.59±0.03	0.42±0.09
GBT	71.3±5.1	65.4±3.2	57.0±4.4	49.8±6.5	53.1±5.3	26.0±6.9	60.0±5.8	0.62±0.06	0.38±0.08
XGBoost	74.2±4.5	69.2±3.6	64.1±5.3	50.2±7.5	56.1±6.4	33.7±7.8	62.2±5.5	0.68±0.08	0.40±0.08
SGD	65.6±3.4	52.6±3.0	45.5±2.8	98.7±1.6	62.2±2.5	29.6±3.8	50.9±7.8	0.77±0.02	0.31±0.06
Logit Reg	60.1±3.2	60.4±2.4	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	45.2±4.8	0.67±0.01	0.25±0.05
MLP	57.3±4.1	60.4±2.4	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	39.1±6.1	0.68±0.00	0.26±0.05
SVM	67.8±4.5	51.4±3.1	43.9±2.6	82.5±6.9	57.3±3.7	15.2±8.9	61.3±6.0	0.64±0.01	0.33±0.08
KNN	56.7±3.3	57.4±3.3	46.5±6.2	44.7±5.2	45.3±4.7	10.8±7.0	49.6±7.8	9.55±1.33	0.12±0.05
VaxiJen2.0	54.3±4.3	54.9±4.2	44.7±6.5	51.4±4.6	47.7±5.3	8.5±8.6	41.3±9.2	16.27±1.53	0.10±0.07
VaxiJen3.0	50.0±0.0	39.0±4.6	39.0±4.6	100.0±0.0	55.9±4.8	0.0±0.0	35.7±4.7	22.00±1.66	0.00±0.00
VENUSVACCINE-Ankh	85.0±2.7	76.9±3.1	65.2±4.1	84.6±1.6	73.5±2.7	55.0±5.2	73.5±3.6	0.48±0.05	0.61±0.05
VENUSVACCINE-ESM2	80.7±4.5	74.0±4.3	66.1±6.9	70.6±6.3	68.2±6.4	46.1±9.2	61.7±9.3	1.37±0.28	0.58±0.06
VENUSVACCINE-ProtBert	79.5±4.2	71.5±4.1	59.3±4.8	78.9±6.7	67.6±4.6	44.7±8.3	68.7±8.4	0.56±0.05	0.54±0.09

† The top three are highlighted by **First, Second, Third**.

E ABLATION STUDY

This section discusses some of the design choices for our VENUSVACCINE. We evaluate the contributions of different inputs to the final results, as well as the advantages of attention pooling compared to simpler global pooling methods. Similar to the previous section, we conducted a full comparison across the three datasets, with results presented in Tables 12-14. The results suggest that encoding sequence information with PLMs is essential for enhanced performance in immunogenicity prediction tasks, and both atomic tokenization and physiochemical descriptors contribute significantly to

Table 7: Average performance of baseline methods in three immunogenicity prediction tasks with the 40% sequence identity cut-off between train and test set.

Model	Bacteria					Virus					Tumor				
	ACC	T30	MCC	F1	KS	ACC	T30	MCC	F1	KS	ACC	T30	MCC	F1	KS
Random Forest	85.5	84.1	68.2	78.8	0.70	89.6	94.3	79.4	88.6	0.81	67.4	67.6	31.3	50.2	0.43
Gradient Boosting	86.4	82.7	70.3	80.8	0.74	78.3	91.4	55.6	77.0	0.74	66.3	65.2	29.3	55.6	0.41
XGBoost	82.0	82.7	61.4	75.5	0.66	90.8	97.1	80.7	89.4	0.86	70.7	69.0	38.8	59.5	0.43
SGD	72.2	74.6	38.0	41.7	0.59	80.0	92.9	61.4	75.4	0.79	55.1	48.6	30.9	64.7	0.32
Logistic Regression	63.5	63.6	0.0	0.0	0.45	59.2	82.9	31.2	68.9	0.73	58.0	48.1	0.0	0.0	0.24
MLP	78.5	75.0	53.5	70.3	0.62	88.8	92.9	76.8	86.9	0.83	58.0	45.2	0.0	0.0	0.27
SVM	54.9	61.4	20.7	56.2	0.38	77.9	84.3	53.7	75.4	0.68	55.4	70.5	20.5	61.0	0.40
KNN	84.3	79.1	65.6	77.2	0.64	91.7	97.1	83.4	90.6	0.84	59.6	50.5	16.1	49.6	0.18
ProVaccine-Ankh	87.6	89.1	73.1	82.2	0.75	87.1	95.7	74.7	84.4	0.90	81.4	84.3	63.3	79.3	0.67
ProVaccine-ESM2	82.3	83.2	62.9	77.2	0.73	92.1	100.0	84.5	91.2	0.94	73.9	62.4	45.7	67.5	0.56
ProVaccine-ProtBert	86.9	88.6	71.7	81.8	0.81	92.9	100.0	85.9	92.0	0.96	73.4	73.3	47.8	70.9	0.56

† The top three are highlighted by **First**, **Second**, **Third**.

Table 8: Average performance of VENUSVACCINE in **Bacteria** immunogenicity prediction tasks from 5 random data split with train/valid/test=8:1:1, with the standard deviation reported after \pm .

Model	AUC-ROC	Accuracy	Precision	Recall	F1	MCC	TopK-30	Cross Entropy	KS statistic
VENUSVACCINE-Ankh	87.3 \pm 2.2	81.3 \pm 1.9	80.1 \pm 5.6	70.8 \pm 8.2	74.7 \pm 2.8	60.7 \pm 3.0	84.9 \pm 1.8	0.85 \pm 0.24	0.63 \pm 0.03
VENUSVACCINE-ESM2	87.4 \pm 1.5	80.3 \pm 1.3	77.7 \pm 3.4	69.8 \pm 3.6	73.4 \pm 2.2	58.1 \pm 2.9	81.4 \pm 3.4	0.63 \pm 0.30	0.62 \pm 0.04
VENUSVACCINE-ProtBert	88.4 \pm 1.8	81.9 \pm 1.8	77.8 \pm 2.6	74.8 \pm 6.2	76.1 \pm 2.4	61.7 \pm 3.9	83.8 \pm 3.2	0.50 \pm 0.10	0.66 \pm 0.05

the representation of the protein. The combination of all three modules achieves optimal prediction performance. The ablation experiments of different protein structure folding methods are shown in Tables 9-11

F ADDITIONAL INFORMATION ON THE ASSESSING DATASET IN CASE STUDY

Tables 16-17 summarize the prediction results on Helicobacter pylori and SARS-COV-2 targets for the two case studies in the post-hoc analysis. The relevant analysis can be found in the main text.

Table 9: Ablation study on different folding methods on **ImmunoDB-Bacteria**.

Fold Type	Model	Accuracy	TopK	MCC	F1	KS-statistic
ESMFold	ProVaccine-Ankh	82.7	78.8	61.8	75.0	0.63
	ProVaccine-ESM2	80.4	77.0	56.4	70.3	0.61
	ProVaccine-ProtBert	82.5	80.1	61.2	74.1	0.66
AlphaFold2	ProVaccine-Ankh	82.3	78.5	60.3	73.3	0.65
	ProVaccine-ESM2	80.7	77.0	57.0	71.7	0.66
	ProVaccine-ProtBert	84.5	84.5	65.9	77.0	0.66

Table 10: Ablation study on different folding methods on **ImmunoDB-Virus**.

Fold Type	Model	Accuracy	TopK	MCC	F1	KS-statistic
ESMFold	ProVaccine-Ankh	91.5	97.5	83.0	91.4	0.85
	ProVaccine-ESM2	89.0	96.2	78.2	89.2	0.80
	ProVaccine-ProtBert	89.7	98.7	79.6	89.4	0.83
AlphaFold2	ProVaccine-Ankh	92.2	99.7	84.3	92.3	0.85
	ProVaccine-ESM2	90.3	99.5	86.6	90.6	0.82
	ProVaccine-ProtBert	91.4	97.4	82.8	91.2	0.84

Table 11: Ablation study on different folding methods on **ImmunoDB-Tumor**.

Fold Type	Model	Accuracy	TopK	MCC	F1	KS-statistic
ESMFold	VENUSVACCINE-Ankh	77.7	77.4	55.8	74.5	0.62
	VENUSVACCINE-ESM2	73.7	67.8	45.5	67.1	0.58
	VENUSVACCINE-ProtBert	70.3	67.8	41.6	67.1	0.51
AlphaFold2	ProVaccine-Ankh	76.9	73.5	55.0	73.6	0.61
	VENUSVACCINE-ESM2	74.0	61.7	46.1	68.2	0.58
	VENUSVACCINE-ProtBert	71.5	68.7	44.8	67.6	0.54

Table 12: Average performance of baseline methods in **Bacteria** immunogenicity prediction tasks from 10 random splits, with the standard deviation reported after \pm .

input module	AUC-ROC	Accuracy	Precision	Recall	F1	MCC	TopK-30	Cross Entropy	KS statistic
Sequence (One-hot)	87.1 \pm 1.0	79.1 \pm 1.2	70.0 \pm 2.6	75.7 \pm 2.7	72.7 \pm 2.1	55.9 \pm 2.8	78.7 \pm 2.8	0.45 \pm 0.02	0.61 \pm 0.02
Ankh									
Sequence (PLM)	86.7 \pm 1.7	81.5 \pm 1.5	78.0 \pm 2.8	70.7 \pm 2.5	74.2 \pm 2.3	60.0 \pm 3.3	63.1 \pm 3.2	0.42\pm0.02	0.84 \pm 0.04
Sequence (PLM)+Structure Token	86.7 \pm 1.9	83.1 \pm 1.5	82.2 \pm 2.2	70.2 \pm 4.1	75.7 \pm 2.7	63.4 \pm 3.3	63.1 \pm 3.6	0.45 \pm 0.04	0.85\pm0.03
Sequence (PLM)+Descriptors	87.1 \pm 1.7	82.2 \pm 1.6	77.1 \pm 3.4	72.4 \pm 3.5	74.6 \pm 2.7	61.1 \pm 3.6	63.4 \pm 3.9	0.42 \pm 0.03	0.81 \pm 0.05
VENUSVACCINE	89.8\pm1.7	86.2\pm0.8	84.5\pm1.3	77.7\pm2.3	80.9\pm1.5	70.3\pm1.7	88.7\pm4.0	0.44 \pm 0.04	0.71 \pm 0.03
ESM									
Sequence (PLM)	88.8 \pm 1.4	83.0 \pm 1.5	78.4 \pm 2.4	72.3 \pm 2.8	75.2 \pm 2.4	62.4 \pm 3.3	68.7 \pm 3.0	0.39\pm0.03	0.80 \pm 0.03
Sequence (PLM)+Structure Token	89.9\pm1.4	84.0 \pm 1.0	81.0\pm2.3	74.8 \pm 3.0	77.8 \pm 2.0	65.4 \pm 2.3	68.8 \pm 2.9	0.41 \pm 0.03	0.85\pm0.03
Sequence (PLM)+Descriptors	88.2 \pm 1.1	82.0 \pm 1.5	76.5 \pm 1.7	75.2 \pm 3.4	75.8 \pm 1.8	61.5 \pm 2.8	67.1 \pm 2.8	0.41 \pm 0.02	0.80 \pm 0.03
VENUSVACCINE	89.8 \pm 1.1	85.0\pm0.5	79.1 \pm 2.3	80.8\pm1.4	80.0\pm1.1	67.9\pm1.3	85.2\pm3.4	0.45 \pm 0.03	0.70 \pm 0.01
ProtBert									
Sequence (PLM)	89.9\pm1.2	86.1\pm1.5	84.6\pm1.9	78.0\pm2.8	81.2\pm2.3	70.3\pm3.4	70.2 \pm 3.4	0.39\pm0.03	0.88\pm0.02
Sequence (PLM)+Structure Token	89.7 \pm 1.9	85.2 \pm 1.5	82.6 \pm 2.8	75.0 \pm 3.9	78.6 \pm 2.7	67.6 \pm 3.6	68.7 \pm 3.1	0.40 \pm 0.05	0.85 \pm 0.04
Sequence (PLM)+Descriptors	88.4 \pm 1.7	82.2 \pm 1.3	76.9 \pm 2.3	74.9 \pm 3.8	75.8 \pm 1.8	61.9 \pm 2.7	66.1 \pm 3.7	0.41 \pm 0.03	0.84 \pm 0.03
VENUSVACCINE	89.7 \pm 1.5	85.8 \pm 2.0	81.8 \pm 3.0	77.2 \pm 2.5	80.2 \pm 2.2	69.2 \pm 3.4	86.5\pm2.5	0.46 \pm 0.06	0.69 \pm 0.03

Table 13: Average performance of baseline methods in **Virus** immunogenicity prediction tasks from 10 random splits, with the standard deviation reported after \pm .

input module	AUC-ROC	Accuracy	Precision	Recall	F1	MCC	TopK-30	Cross Entropy	KS statistic
Sequence (One-hot)	96.8 \pm 0.4	91.5 \pm 0.7	92.1 \pm 1.5	91.2 \pm 0.8	91.6 \pm 0.7	82.9 \pm 1.5	98.0 \pm 1.2	0.26 \pm 0.03	0.84 \pm 0.01
Ankh									
Sequence (PLM)	94.6 \pm 0.9	84.8 \pm 1.5	81.3 \pm 2.3	91.0 \pm 1.3	85.8 \pm 2.5	70.1 \pm 2.9	75.4 \pm 2.5	0.34 \pm 0.02	0.98 \pm 0.01
Sequence (PLM)+Structure Token	97.9\pm0.6	93.7\pm0.6	95.5\pm0.7	92.2 \pm 1.1	93.8\pm0.7	87.5\pm1.2	88.9 \pm 1.8	0.20\pm0.03	0.99\pm0.00
Sequence (PLM)+Descriptors	96.1 \pm 0.7	88.1 \pm 1.0	85.2 \pm 1.6	92.6 \pm 2.0	88.8 \pm 2.1	76.4 \pm 2.0	81.3 \pm 2.1	0.28 \pm 0.02	0.99 \pm 0.01
VENUSVACCINE	97.7 \pm 0.4	92.5 \pm 0.6	91.4 \pm 1.3	95.5\pm1.3	92.8 \pm 0.7	85.2 \pm 1.2	99.5\pm0.4	0.22 \pm 0.02	0.87 \pm 0.02
ESM									
Sequence (PLM)	96.0 \pm 0.8	90.0 \pm 1.1	89.0 \pm 1.7	91.6 \pm 2.1	90.3 \pm 2.1	80.0 \pm 1.9	81.5 \pm 2.0	0.26 \pm 0.03	0.99 \pm 0.01
Sequence (PLM)+Structure Token	97.5 \pm 0.6	93.7 \pm 0.9	95.8\pm0.8	91.7 \pm 1.0	93.7 \pm 0.9	87.4 \pm 1.7	89.1 \pm 1.8	0.29 \pm 0.05	1.00\pm0.01
Sequence (PLM)+Descriptors	96.9 \pm 0.5	91.2 \pm 0.8	91.5 \pm 1.2	90.9 \pm 1.7	91.2 \pm 1.7	82.3 \pm 1.8	84.6 \pm 1.3	0.23\pm0.02	0.99 \pm 0.00
VENUSVACCINE	97.9\pm0.5	94.5\pm0.6	95.7 \pm 0.7	93.6\pm1.0	94.7\pm0.6	89.2\pm1.2	98.9\pm0.5	0.29 \pm 0.03	0.90 \pm 0.01
ProtBert									
Sequence (PLM)	97.2 \pm 0.7	91.4 \pm 1.2	89.3 \pm 1.4	94.2 \pm 2.3	91.7 \pm 2.3	82.9 \pm 2.3	84.8 \pm 2.3	0.23 \pm 0.03	0.99 \pm 0.01
Sequence (PLM)+Structure Token	96.8 \pm 0.5	90.8 \pm 0.9	90.4 \pm 0.9	91.7 \pm 1.5	91.0 \pm 1.7	81.6 \pm 1.9	82.9 \pm 1.7	0.23 \pm 0.02	0.99 \pm 0.01
Sequence (PLM)+Descriptors	96.9 \pm 0.5	91.5 \pm 1.0	89.4 \pm 1.5	94.6\pm2.1	91.9 \pm 2.1	83.0 \pm 2.1	85.3 \pm 2.1	0.24 \pm 0.03	0.99\pm0.01
VENUSVACCINE	97.5\pm0.6	91.6\pm0.9	97.0\pm0.8	93.5 \pm 1.4	92.9\pm0.8	85.6\pm1.5	98.7\pm0.9	0.23\pm0.02	0.88 \pm 0.02

Table 14: Average performance of baseline methods in **Tumor** immunogenicity prediction tasks from 10 random splits, with the standard deviation reported after \pm .

input module	AUC-ROC	Accuracy	Precision	Recall	F1	MCC	TopK-30	Cross Entropy	KS statistic
Sequence (One-hot)	73.1 \pm 4.8	65.5 \pm 4.6	56.4 \pm 6.3	64.5 \pm 6.5	60.0 \pm 5.2	30.3 \pm 9.4	66.5 \pm 6.1	0.90 \pm 0.13	0.39 \pm 0.07
Ankh									
Sequence (PLM)	65.9 \pm 3.3	60.9 \pm 2.0	66.9 \pm 1.0	75.3 \pm 1.0	71.3 \pm 1.2	50.2 \pm 3.3	33.4 \pm 3.3	0.67 \pm 0.00	0.51 \pm 0.06
Sequence (PLM)+Structure Token	83.1 \pm 3.0	75.9 \pm 2.9	66.9 \pm 4.5	80.5 \pm 5.5	72.9 \pm 6.1	52.3 \pm 6.1	60.0 \pm 6.2	0.51 \pm 0.05	0.81\pm0.07
Sequence (PLM)+Descriptors	67.0 \pm 6.3	62.7 \pm 3.2	65.7 \pm 0.0	80.0 \pm 0.0	73.2 \pm 1.2	53.4 \pm 2.3	33.5 \pm 0.0	0.66 \pm 0.01	0.49 \pm 0.09
VENUSVACCINE	85.2\pm1.9	77.7\pm2.1	70.2\pm6.1	83.0\pm3.8	74.5\pm3.3	55.8\pm4.4	61.6\pm2.6	0.48\pm0.03	0.77 \pm 0.03
ESM									
Sequence (PLM)	81.5 \pm 4.3	70.4 \pm 3.6	60.0 \pm 4.7	75.7 \pm 3.4	66.8 \pm 6.8	41.8 \pm 9.1	52.5 \pm 4.0	0.51\pm0.04	0.73 \pm 0.07
Sequence (PLM)+Structure Token	83.4\pm3.4	70.6 \pm 4.2	58.0 \pm 5.5	81.7\pm5.6	67.7\pm5.0	44.3 \pm 8.0	58.6\pm7.6	0.51 \pm 0.05	0.74\pm0.07
Sequence (PLM)+Descriptors	78.4 \pm 2.0	68.6 \pm 2.6	59.8 \pm 5.1	73.9 \pm 4.0	65.9 \pm 3.3	38.0 \pm 4.8	44.3 \pm 3.5	0.55 \pm 0.02	0.72 \pm 0.07
VENUSVACCINE	82.2 \pm 3.5	73.7\pm3.1	67.8\pm3.5	66.8 \pm 5.4	67.1 \pm 4.7	45.5\pm5.9	58.2 \pm 5.6	1.57 \pm 0.02	0.68 \pm 0.07
ProtBert									
Sequence (PLM)	69.7 \pm 3.0	59.9 \pm 3.8	51.8 \pm 6.5	64.7 \pm 5.4	57.3 \pm 5.2	20.9 \pm 7.8	36.5 \pm 4.7	0.61 \pm 0.02	0.61 \pm 0.07
Sequence (PLM)+Structure Token	80.4 \pm 3.0	70.5 \pm 3.3	59.6 \pm 6.0	76.4 \pm 6.1	66.8 \pm 4.0	45.9\pm6.0	55.9 \pm 5.2	0.53\pm0.05	0.67\pm0.06
Sequence (PLM)+Descriptors	75.7 \pm 3.3	69.5 \pm 3.7	58.9 \pm 6.5	76.6 \pm 5.9	66.5 \pm 5.2	40.4 \pm 7.5	46.7 \pm 7.8	0.61 \pm 0.04	0.66 \pm 0.07
VENUSVACCINE	81.0\pm3.2	71.3\pm3.5	59.6\pm5.4	78.1\pm3.9	67.1\pm5.3	43.9 \pm 8.1	67.8\pm7.3	0.58 \pm 0.07	0.52 \pm 0.08

Table 15: Model prediction performance on the first case study of Helicobacter pylori.

Model	Predicted Positive	True Positive	Random Selection	Recall (%)	fold-enrichment
Random Forest	211	11	1.25	100.00	8.81
Gradient Boosting	273	11	1.62	100.00	6.81
XGBoost	287	11	1.70	100.00	6.47
SGD	25	0	0.15	0.00	0.00
Logistic Regression	0	0	0.00	0.00	0.00
MLP	227	4	1.34	36.36	2.98
KNN	284	10	1.68	90.91	5.95
VENUSVACCINE	123	11	0.73	100.00	15.11

Table 16: Detailed likelihood predicted for the 11 determined immunogens under each model.

Protein	Random Forest	Gradient Boosting	XGBoost	SGD	Log Reg	MLP	SVM	KNN	VENUSVACCINE
P69997	0.93	0.910173	0.998961	0.311594	0.412161	0.623967	0.689306	1	0.878434
Q9ZN50	0.87	0.816795	0.997187	0.253575	0.415796	0.492811	0.698149	1	0.96182
Q9ZKE6	0.74	0.586535	0.988803	0.057892	0.394834	0.072825	0.820259	1	0.82105
Q9ZKW5	0.84	0.950726	0.999409	0.429902	0.417487	0.844049	0.72135	1	0.993692
Q9ZKX5	0.78	0.600814	0.986793	0.154398	0.396191	0.41584	0.712237	0.5	0.738437
Q9ZMJ1	0.72	0.564479	0.985005	0.099751	0.396345	0.341039	0.708772	1	0.725007
Q9ZL47	0.83	0.765695	0.994611	0.135861	0.402306	0.233893	0.645672	1	0.82569
Q9ZLT1	0.66	0.673526	0.860356	0.371664	0.412658	0.811522	0.771533	1	0.986235
Q9ZN37	0.84	0.749941	0.99151	0.153825	0.398239	0.338815	0.792457	1	0.659044
Q9ZLJ8	0.61	0.911371	0.989448	0.339596	0.410059	0.715878	0.574296	1	0.978299
POA0R4	0.94	0.880224	0.991948	0.168698	0.401071	0.454899	0.335555	1	0.881635

Table 17: Prediction Details of SARS-COV-2 Targets by VENUSVACCINE.

NCBI ID	Protein Name	predicted label	predicted probability
YP_009724390.1	surface glycoprotein	1	1.000
YP_009742617.1	nsp10	1	1.000
YP_009742616.1	nsp9	1	1.000
YP_009742609.1	nsp2	1	0.999
YP_009742612.1	3C-like proteinase	1	0.998
YP_009725312.1	nsp11	1	0.997
YP_009724396.1	ORF8 protein	1	0.997
YP_009742610.1	nsp3	1	0.997
YP_009724397.2	nucleocapsid phosphoprotein	1	0.995
YP_009742614.1	nsp7	1	0.994
YP_009724395.1	ORF7a protein	1	0.941
YP_009725307.1	RNA-dependent RNA polymerase	1	0.864
YP_009725310.1	endoRNase	1	0.768
YP_009742611.1	nsp4	0	0.494
YP_009724391.1	ORF3a protein	0	0.463
YP_009742608.1	leader protein	0	0.150
YP_009725308.1	helicase	0	0.140
YP_009742613.1	nsp6	0	0.107
YP_009742615.1	nsp8	0	0.047
YP_009725309.1	3'-to-5' exonuclease	0	0.022
YP_009725311.1	2'-O-ribose methyltransferase	0	0.000
YP_009724393.1	membrane glycoprotein	0	0.000
YP_009724392.1	envelope protein	0	0.000
YP_009724394.1	ORF6 protein	0	0.000
YP_009725318.1	ORF7b	0	0.000
YP_009725255.1	ORF10 protein	0	0.000

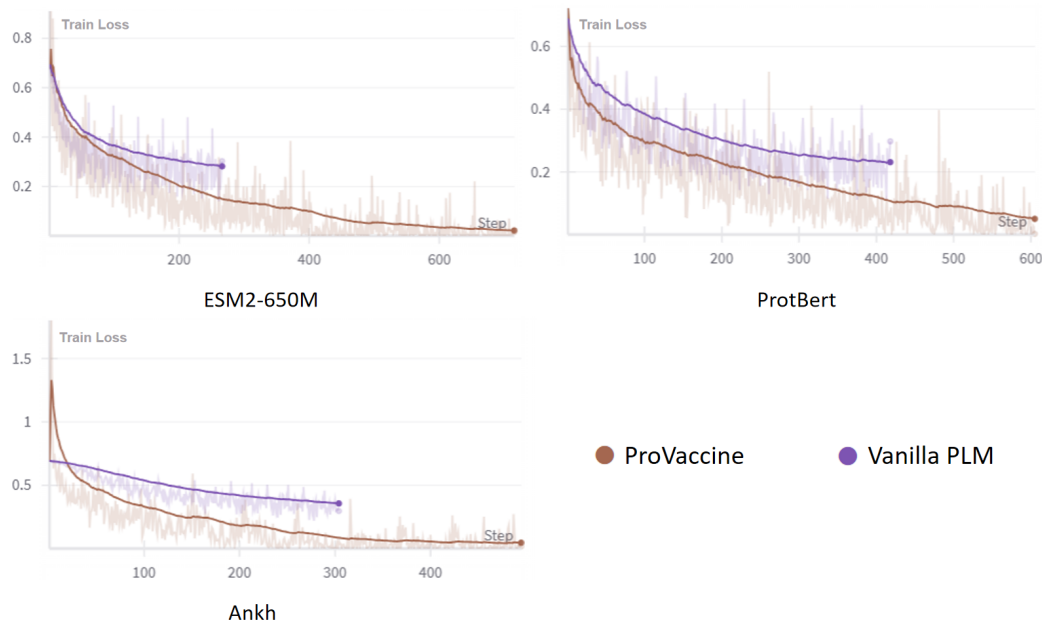


Figure 6: Convergence speed of VENUSVACCINE versus the vanilla PLMs on ImmunoDB-Virus.