

Synthetic Engineering with ImmunoGenAI: Multimodal Deep Learning Integration for Predicting and Optimizing Immune Evasion in Personalized Therapeutic Devices

By Lalithendra Reddy Bhima

AUTHOR BIOGRAPHY

Lalithendra Reddy Bhima is currently a sophomore attending Folsom High School. He is very interested in computer science and mathematics, since elementary school. He has earned numerous prestigious awards at top competitions. He developed a nonprofit that leveraged technology to improve access and literacy in AI, computer science, and STEAM education. Additionally, he enjoys tutoring in STEM and promoting civic education in Sacramento County.

ABSTRACT

ImmunoGenAI characterizes immunogenicity as an emergent property influenced by the intracellular antigen-processing pathway, rather than an isolated function of peptide–HLA binding affinity. It is informed by biological processes including proteasomal cleavage, peptide transport, and HLA presentation, which shape downstream T-cell recognition. ImmunoGenAI optimizes existing therapeutic and vaccine sequences by introducing targeted sequence perturbations that reduce predicted unintended immunogenic peptide formation. These include modifying residues associated with cleavage patterns and HLA binding preferences to lower peptide–HLA stability while preserving structural integrity. By acting across multiple biologically relevant features simultaneously, these changes reduce immunogenicity more effectively than single-objective optimization. This framework is bidirectional: it suppresses immune visibility for therapeutic design and can be inverted for vaccine design to remove unnecessary amino acid substitutions, enabling more reliable control over sequence quality and downstream performance. ImmunoGenAI enables systematic control of immune recognition, outperforming five other industry-standard tools while maintaining structural integrity. ImmunoGenAI utilizes modality-specific neural network encoders, including contrastive and transformer architectures, to project biological inputs into unified latent representations.

FALL 2026

Keywords: *ImmunoGenAI, immunogenicity, immune evasion, personalized therapeutic devices, multimodal deep learning, therapeutic protein engineering, human leukocyte antigen, antigen processing, sequence optimization, structural preservation*

INTRODUCTION

Rationale

One of the most significant challenges in modern biomedical engineering has been the ability to design therapeutic proteins that maintain functional efficacy while minimizing immune recognition (Dellas et al., 2021; Parker et al., 2010). The rapid expansion of immunotherapeutic engineering—valued at an estimated \$4.23 billion market—reflects the growing reliance on personalized biologics in medicine (Coherent Market Insights, 2026). Immune recognition is governed by the highly polymorphic human leukocyte antigen (HLA) system, which comprises over 20,000 alleles (Creary et al., 2019; Robinson et al., 2015). Therapeutic proteins are proteolytically processed into overlapping peptide fragments and presented by major histocompatibility complex (MHC) molecules, where allele-specific binding determines downstream T-cell activation (Calis et al., 2013; Hong et al., 2025; Jurtz et al., 2017; Yewdell, 2006). Although a therapeutic sequence may retain its primary biological function, even minor amino acid substitutions can substantially alter peptide-HLA binding landscapes, increasing immune visibility and reducing clinical durability (Doud & Bloom, 2016; Greaney et al., 2021; Sarkisyan et al., 2016; Sidney et al., 2008).

This clinical limitation creates a particular interest in computational methods capable of predicting and optimizing immune evasion. Immune evasion refers to the ability of engineered sequences to reduce host immune detection through modulation of antigen presentation features (Hamelin et al., 2025; Wec et al., 2021). Quantitatively, predicted HLA binding affinity can shift from strong (<50 nM) to weak (>500 nM) binding thresholds, reducing immune visibility by more than an order of magnitude (Aleksic et al., 2012; Sidney et al., 2008). However, current industry-standard frameworks are predominantly prediction-only systems that evaluate peptide-HLA binding in isolation, or optimization pipelines that introduce local sequence edits without incorporating sequential and structural dependencies (Das et al., 2025; Li et al., 2021; Zinsli et al., 2020). Because immune recognition is nonlinear and depends on high-dimensional sequence-structure interactions, accurately modeling these immunogenic shifts requires architectures capable of learning long-range contextual relationships (Anishchenko et al., 2021; Hayes et al., 2024; Rives et al., 2021). Transformer-based deep learning models, which utilize multi-head self-attention to capture global residue dependencies without vanishing gradient limitations (Vaswani et al., 2017), provide a mechanism for integrating multimodal biological signals within a unified representation space. Yet, existing immune optimization systems rarely fuse structural

FALL 2026

geometry, contextual annotations, allele-specific exposure, and immunopeptidomic data simultaneously (Liang et al., 2024; Zhao & Wang, 2025).

We hypothesize that a multimodal deep learning framework that jointly encodes structural geometry, allele-specific binding distributions, contextual sequence dependencies, and high-dimensional immunopeptidomic features will learn coordinated immune-activation signatures that enable targeted amino acid substitutions to reduce predicted immune recognition while preserving structural stability and functional domain integrity. By training on over 100,000 experimentally validated peptide-HLA entries and evaluating performance across immunogenicity, epitope density, structural preservation, optimization efficiency, and population coverage, this approach aims to establish a scalable and biologically constrained system for synthetic engineering of immune evasion in personalized therapeutic devices (Das et al., 2025; Dellas et al., 2021; Sarkisyan et al., 2016).

BACKGROUND

Synthetic Immunotherapeutic Design

Immunotherapeutic engineering is a rapidly advancing frontier in synthetic biology (\$ 4.23 billion market!), as personalized biologics increasingly reshape modern medicine. In personalized therapeutic design, engineered biologics must maintain their intended biological function while also avoiding unintended immune recognition. This creates a central design challenge because therapeutic proteins are not only evaluated by their primary activity, but also by how the immune system processes and presents their peptide fragments after administration (Coherent Market Insights, 2026).

Immune recognition is governed by the highly polymorphic human leukocyte antigen (HLA) system, where therapeutic proteins are processed into peptide fragments and presented by major histocompatibility complex (MHC) molecules, shaping patient-specific immune responses and downstream T-cell activation (Creary et al., 2019; Robinson et al., 2015; Yewdell, 2006). Because HLA molecules vary across individuals, the same therapeutic sequence can produce different immune-presentation outcomes depending on the patient-specific HLA background and the peptide fragments generated during antigen processing (Creary et al., 2019; Jurtz et al., 2017).

When designing therapeutic sequences, minor amino acid substitutions can significantly alter peptide-HLA binding landscapes, especially since the HLA system comprises over 20,000 alleles. Although the primary biological function may remain intact, small sequence changes can increase antigen presentation, trigger immune activation, and reduce therapeutic efficacy. Therefore, immunotherapeutic sequence design requires

FALL 2026

optimization strategies that reduce unintended immune visibility while preserving the structural and functional properties required for therapeutic performance (Zinsli et al., 2020).

Predicting and Optimizing Immune Evasion

Immune evasion is the ability of pathogens/therapeutic sequences to escape detection and suppression by the host's immune system by reducing immune-triggering features, preserving structural integrity (Hamelin et al., 2025; Wec et al., 2021). In therapeutic engineering, immune evasion is important because unintended immune recognition can reduce treatment durability, increase anti-drug immune responses, and limit the effectiveness of personalized biologics. The goal is not simply to change the sequence, but to reduce immune-triggering features while maintaining the original therapeutic structure and function (Parker et al., 2010; Zinsli et al., 2020).

Accurate prediction of immune evasion is essential prior to sequence modification, as optimization strategies rely on quantitative immune-risk estimation to guide rational amino acid substitutions (Calis et al., 2013; Jurtz et al., 2017). Prediction provides the baseline immune-risk profile of a therapeutic sequence, identifying regions that are more likely to contribute to peptide-HLA binding, antigen presentation, and downstream immune activation. Without accurate prediction, optimization can introduce sequence changes that appear beneficial numerically but fail to preserve biological feasibility (Jurtz et al., 2017; Yewdell, 2006).

By optimizing therapeutic sequences, predicted HLA binding affinity can shift from strong (<50nM) to weak (>500nM) binding thresholds, reducing immune visibility by over an order of magnitude (Aleksic et al., 2012). These threshold changes are important because strong peptide-HLA binding is associated with increased likelihood of antigen presentation, while weaker binding can reduce the probability of immune recognition. However, optimization must remain constrained because excessive or poorly selected substitutions can alter structural topology or disrupt functional domains (Dellas et al., 2021; Sidney et al., 2008).

Current approaches treat prediction and optimization as sequential or loosely coupled steps, lacking unified constraint-aware modeling across immunogenic, structural, and contextual dimensions. This introduces epitope trade-offs, structural drift, and loss of functional robustness, limiting translational feasibility even when predicted immune evasion scores improve. A background problem in this field is therefore the need for optimization methods that evaluate immune risk, structural preservation, and contextual biological constraints together rather than as isolated endpoints (Wec et al., 2021; Zinsli et al., 2020).

Artificial Intelligence

Deep Learning is an A.I. & machine learning subset method that uses neural networks to achieve high levels of accuracy (Li et al., 2021; Rives et al., 2021). In biological sequence modeling, deep learning is useful because immune recognition depends on patterns distributed across amino acid sequences, structural features, and biological context. These patterns may involve long-range residue relationships that are difficult to capture using only local scoring rules or single-objective prediction methods (Rives et al., 2021; Vaswani et al., 2017).

Transformer-based multimodal neural network can interpret long-range sequence dependencies to understand cross-modality immune interactions within therapeutic designs (Liang et al., 2024; Vaswani et al., 2017). Transformer architectures use attention mechanisms to evaluate relationships between sequence positions, allowing the model to identify regions that contribute strongly to immune-risk predictions. In multimodal immunotherapeutic design, this allows sequence information, structural properties, immune labels, and contextual biological features to be represented together in a unified learning framework (Liang et al., 2024; Vaswani et al., 2017).

By integrating multiple biological modalities, artificial intelligence can support immune-evasion prediction and optimization beyond peptide-HLA binding alone. A multimodal framework can learn coordinated immune-activation signatures across genomic, structural, epigenetic, contextual, sequential, and immunological information, allowing targeted amino acid substitutions that reduce predicted immune recognition while maintaining structural stability and functional domain integrity (Das et al., 2025; Kumar et al., 2024).

OBJECTIVES

To address rapidly evolving diseases and the growing demand for personalized biologic therapies, there is a need for precise and scalable design of therapeutic sequences that reduce immune recognition while maintaining structure (Cao et al., 2024; Kumar et al., 2024; Liu et al., 2025; Solá & Griebenow, 2010). Current approaches often evaluate immunogenicity in isolation, failing to account for coordinated biological interactions and leading to unintended immune activation and costly redesign cycles.

This research will develop a tool, ImmunoGenAI, that demonstrates that understanding underlying patterns of immune evasion across coordinated biological dimensions will lead to context-aware immune optimization with structurally stable and functionally viable personalized therapeutic devices.

ImmunoGenAI will be created under the design criteria of a model trained on a variety of independent biological modalities and diverse patient-specific immune-structural conditions. It learns biological (cross-modal immune regulation dynamics) & structural (protein-immune interaction constraints) vs. industry methods. It accepts protein or structure input < 5 minutes per sequence inference/optional manual feature engineering (E2E).

FALL 2026

Primary Endpoints

Increase immune evasion prediction and optimization accuracy, maintain structural stability (RMSD below threshold), reduce false-positive immune activation signals, reduce epitope density, validate performance across independent cohorts and population HLA diversity.

Secondary Endpoints


Maintain cross-modal regulatory coherence, preserve native structural topology, prevent unintended increases in antigen presentation or immune activation.

SOLUTION

ImmunoGenAI utilizes modality-specific neural network encoders, including contrastive and transformer architectures, to project biological inputs into unified latent representations. It formulates the acquisition of underlying immune patterns & grasps evolutionary conserved motifs within therapeutic sequences (Hayes et al., 2024). Additionally, it uses transformers to learn relationships between biologically distant sequence positions and hierarchical interactions across therapeutic sequences. Finally, it utilizes industry standard immunogenicity and structural stability metrics to ensure confidence in immune evasion predictions and optimizations (Jumper et al., 2021).

Figure 1

Overview of user workflow for ImmunoGenAI multimodal deep learning processing pathway.



Note. The pipeline starts with input therapeutic sequence data and advanced multimodal contextualization, including secondary structure, binding partner, HLA genotype, cell type, methylation level, patient age, clinical outcome, and design weights. The prediction/optimization stage applies risk scoring, constraint weighting, and iterative redesign to generate optimized constructs with structural RMSD and functional preservation while back translating to DNA/RNA. The model creation workflow includes data preprocessing and planning, data encoding and feature creation, transformer-based latent embedding alignment, fusion concatenation, and calibration and constraint integration with regularization.

IMPLEMENTATION

Model training dataset

For ImmunoGenAI, we used key datasets from the National Institutes of Health (NIH), Immune Epitope Database (IEDB), RCSB Protein Data Bank (PDB), AlphaFold, TCGA, GEO, and the Pan-Cancer Atlas (Jumper et al., 2021). These sources formed the raw biomedical dataset foundation for the project and were used to create curated multimodal datasets containing 33.28 billion rows and over 2TB of biological data.

We turned raw biomedical data into a clean dataset through a structured compilation process. First, we performed human-only sequence filtering by retaining only human-derived sequences, specifically

Homo-Sapiens sequences. This ensured that the model was trained on biologically relevant human data rather than mixed-species inputs that could introduce noise into immune prediction and therapeutic sequence optimization.

After filtering, we removed redundant proteins and sequences shorter than 50 amino acids. This step eliminated fragments, stabilized training inputs, and reduced noise through the preprocessing script. We then performed dataset harmonization by converting heterogeneous structures to JSON so that different types of biomedical data could be stored and processed in a consistent format across the pipeline.


Because the datasets exceeded conventional memory limits at over 2TB, we indexed large datasets for efficient streaming access. The major roadblock was that the datasets were too large for conventional processing, so we used an engineering solution of chunk processing to cluster data/sequences. This allowed the dataset to be processed in manageable sections while still preserving access to the full multimodal dataset.

The final compilation workflow moved from raw biomedical datasets into human-only sequence filtering, dataset harmonization, redundancy and fragment removal, and feature extraction. Feature extraction included CIBERSORTx and pVACseq, allowing the cleaned dataset to support downstream immune profiling, neoantigen prediction, and multimodal ImmunoGenAI model development (Kumar et al., 2024; Zhang et al., 2025).

ENCODING

Figure 2

The Encoding Procedure of the ImmunoGenAI Integration.



FALL 2026

Note. Panel A shows a series of each numeric expression of a feature matrix and conversion into a vectorized layout for each encoder. Panel B analyzes the multimodal transformer-based architecture characterized by each encoder.

Next, sequences were transformed into numerical representations using multimodal encoding. Biological features—including sequence information, structural properties, and immunological context—were vectorized and processed using transformer-based neural network encoders to capture relationships between amino acid positions and immune-relevant patterns.

Multimodal Feature Vectorization transforms six heterogeneous datasets into structured numerical feature vectors, allowing different biological modalities to be represented in a unified machine-learning compatible format. Genomic, Epigenetic, Structural, Contextual, Immune, and Sequence features were represented through a Feature Matrix and Vector Extraction process.

Sequence Tokenization converts biological sequences into discrete token representations by mapping each residue to a fixed vocabulary of the 20 canonical amino acids, allowing protein sequences and immune alleles to be processed by transformer-based neural encoders.

Through the application of Representation Learning, the encoded biological modalities were processed through Learned Embedding, Multi-Head Self-Attention, Transformer (x3), Global Attention Pooling, Layer Norm, and Modality Latent Vector. This encoding procedure allowed sequence information, structural properties, and immunological context to be represented in a unified machine-learning compatible format while capturing relationships between amino acid positions and immune-relevant patterns.


Fusion Architecture and Prediction

A fusion model was then trained to integrate these modalities and predict immunogenicity scores. Training was performed using a train/dev/test split, with hyperparameter tuning and regularization techniques applied to ensure generalization and stability. Inputted sequence would process under each individual layer, including the 6 encoders and the Fusion network to create the final Immune evasion score from 0-1. Predicted score is evaluated from specific constraint engineering features that let ImmunoGenAI be tailored to the patient.

OPTIMIZATION

Figure 3

The optimization procedure layout of ImmunoGenAI translating and gradient-guided generation for therapeutic sequences that are more effective in controlling immune evasion sustainability.



Note. Panel A utilizes gradient sensitivity to identify mutation-prone positions influencing immunogenicity across sequences including the canonical amino acid representations. Panel B describes the mutation landscape which evaluates sequence changes impacting immune evasion score distribution.

Gradient-guided attribution reveals how sensitive the model's prediction is to each amino-acid position.

Iterative mutation search explores candidate variants across sequence positions relative to the initial input sequence (up to 100 variants). Model scoring guides selection toward higher predicted immune-evasion sequences.

More possible mutations result in a “more complex” search space at the expense of computation. Fine-tune mutation constraints to balance exploration & convergence while preventing overfitting & underfitting.

STATISTICAL ANALYSIS


We conducted an in-depth statistical data analysis to quantify and verify the performance of my tool. A broad, robust test set of 258 therapeutic sequences + a benchmark set of 20 well-characterized proteins and structurally validated immunogenic motifs with established immunogenic profiles were used.

RESULTS

We use multiple previously established metrics such as the immunogenicity score, epitope density, distribution of HLA coverage, structural preservation, optimization efficiency, and modality ablation study.

Figure 4

ImmunoGenAI significantly improves immune evasion scores



Note. ImmunoGenAI significantly improves immune evasion scores when compared to original sequences on a test subset of $n=278$. ImmunoGenAI increased immune evasion scores from $\mu = 0.520$ to $\mu = 0.839$ (~61% improvement), with highly significant statistical difference ($p = 5.64 \times 10^{-55}$). The y-axis is the immunogenicity score on a scale of 0-1. The open points outside the ImmunoGenAI Box and Whisker Plot are outliers that are beyond 1.5 times the interquartile range. The horizontal divisions present in each box (from top to bottom) are the upper whisker, 3rd quartile, median, 1st quartile, and lower whisker.

Immune evasion score

As noted in previous studies, immune evasion can preserve structural and functional integrity of therapeutic efficacy. Figure 4 demonstrates that ImmunoGenAI increased immune evasion from 0.520 ± 0.276

FALL 2026


to 0.839 ± 0.046 on the 278-sequence test dataset. Extrapolating this shift suggests a substantial improvement in predicted immune modulation compared to the original sequences with a benchmark improvement of 43.1% and an average improvement of 61.27% across the full test dataset. This shows that ImmunoGenAI's results were statistically significant ($p < 0.0001$) in predicted immunogenicity scores while maintaining ideal values for other metrics.

Secondary endpoints

The secondary endpoints were evaluated across the benchmark dataset sequences ($n=20$). ImmunoGenAI maintains native structural topology and preserves conformational integrity throughout optimization. Structural deviation is minimal, with an average RMSD of $0.48 \text{ \AA}$ observed across the benchmark set ($n=20$). This indicates negligible backbone displacement and confirms that immune modulation is achieved without destabilizing tertiary architecture or disrupting broader multimodal biological consistency. These metrics compare ImmunoGenAI immunogenicity prediction approaches to 5 other industry standard tools. The results display that ImmunoGenAI achieved the highest median immune evasion score and most consistent optimization across all evaluated sequences.

Figure 5

ImmunoGenAI Significantly Improves the Normalized Immune Evasion Scores




FALL 2026

Note. ImmunoGenAI significantly improves the normalized immune evasion scores where the lighter shaded boxes indicate the original sequence and the darker shaded boxes indicate the optimized sequences being re-predicted from each tool. Box and Whisker Plot (n=20) comparing immune evasion with legend indicating the color for each prediction sequence.

Figure 6

ImmunoGenAI and other Immune Evasion Predictors on Epitope Density




Note. ImmunoGenAI and other immune evasion predictors on the epitope density of the simulated regions where light indicates the original sequence and the dark indicates the optimized sequence. Box and whisker plot (n=20) comparing the epitope density of each original and optimized method. ImmunoGenAI produced the lowest median epitope density, indicating reduced predicted immunogenicity compared with existing methods.

Immune-evasion score was used to evaluate the predicted ability of optimized therapeutic sequences to reduce unintended immune recognition. Higher immune-evasion values indicate stronger predicted reduction in immune-triggering features. Epitope density was used as a complementary metric to quantify the relative burden of predicted immunogenic peptide regions after optimization. Lower epitope density indicates fewer predicted immune-visible regions within the optimized sequence set. The mean immune-evasion score for each method is depicted in Figure 5, and the calculated epitope density for each method is depicted in Figure 6.

FALL 2026

Figure 7




Note. Structural preservation and RMSD comparison across optimization methods in Box and whisker plot (n=20 benchmark proteins). Light-shaded boxplots represent TM-score, where higher values indicate greater structural similarity, and dark-shaded boxplots represent RMSD, where lower values indicate less structural deviation. ImmunoGenAI achieved the highest median TM-score and the lowest RMSD compared to the other tools, while all other tools were assessed using their internal prediction algorithms with their corresponding optimization methods as indicated in the legend.

Structural preservation was evaluated to determine whether sequence optimization maintained the original therapeutic protein framework. Template Modeling Score (TM-score) was used to quantify structural similarity between original and optimized protein structures, while Root Mean Square Deviation (RMSD) was used to measure the average distance between corresponding backbone $C\alpha$ atoms. Higher TM-score values and lower RMSD values indicate stronger preservation of native structural topology. The structural preservation and RMSD comparison for each optimization method is depicted in Figure 7.


Figure 8

ImmunoGenAI Technique Yields Significantly Improved Optimization Efficiency



Note. ImmunoGenAI technique yields significantly improved optimization efficiency compared to the other tools, while all other tools were assessed using their internal prediction algorithms with their corresponding optimization methods as indicated in the legend. Optimization efficiency was measured as Δ immunogenicity per mutation, where more negative values indicate a greater reduction in predicted immunogenicity for each amino acid substitution introduced. Box and whisker plot (n=20) comparing the optimization efficiency with legend indicating each optimization method derived from the primary prediction tool.

Figure 9



Note. ImmunoGenAI technique yields improved normalized population coverage across diverse set of HLA allele coverage. Higher normalized coverage values indicate broader predicted HLA allele representation across population-level immune presentation contexts. Box and whisker plot (n=20) comparing population coverage with legend indicative of each prediction method with respective prediction tools.

Optimization efficiency was evaluated by comparing the change in predicted immunogenicity relative to the number of sequence modifications introduced (Doud & Bloom, 2016; Parker et al., 2010). This metric was used to determine whether immune-evasion improvements required extensive mutation or could be achieved through targeted amino acid substitutions. Population HLA coverage was evaluated to estimate how broadly the optimized sequences maintained predicted performance across diverse HLA presentation contexts. The optimization efficiency for each method is depicted in Figure 8, and the estimated population-level HLA coverage is depicted in Figure 9.

FALL 2026

Table 1

The Evaluation Dataset Included 258 Therapeutic Protein Sequences, and 20 Benchmark Proteins with Validated Immunogenic Profiles


P01308	Insulin
P01241	Somatotropin
P01588	Erythropoietin
P01563	Interferon alpha-2
P01375	Tumor necrosis factor
P05013	Interferon alpha-6
P60568	Interleukin-2
P04141	Granulocyte-macrophage colony-stimulating factor
P15692	Vascular endothelial growth factor A, long form
P09919	Granulocyte colony-stimulating factor
P02763	Alpha-1-acid glycoprotein 1
P02745	Complement C1q subcomponent subunit A
P02647	Apolipoprotein A-I
P02765	Alpha-2-HS-glycoprotein
P04004	Vitronectin
P01009	Alpha-1-antitrypsin
P69905	Hemoglobin subunit alpha
P68871	Hemoglobin subunit beta
P00918	Carbonic anhydrase 2
P61626	Lysozyme C

Note. The evaluation dataset included 258 therapeutic protein sequences from the Therapeutic Targets Database (TTD) and 20 benchmark proteins with validated immunogenic profiles, enabling robust statistical assessment and generalization beyond standard curated datasets.

FALL 2026

Figure 10

Structural Topology Comparison of Benchmark Proteins



Note.


Depicted are structural topology comparisons of benchmark proteins (Original [Red] vs. ImmunoGenAI Optimized [Blue]). Each of the benchmark sets was simulated in 3D using the TM-align tool. ImmunoGenAI preserves protein structural integrity during optimization, as shown by strong overlap between original and optimized structures across multiple therapeutic proteins.

Sequence Evaluation

ImmunoGenAI does not introduce indiscriminate sequence alterations. Although targeted amino acid substitutions are introduced to modulate immunogenicity, epitope architecture and structural integrity are largely preserved. Across the 278-sequence test dataset, the model maintained ~88% peptide-HLA (Jaccard's coefficient) identity while achieving substantial immune modulation.

Figure 11

ImmunoGenAI across different sequence evaluation metrics.



Note. Panel A illustrates the distribution of backbone RMSD scores across the benchmark set (n=20), with the dashed vertical red line indicating the median RMSD score of 0.48 Å. Panel B depicts the population shift in predicted immune evasion scores between original and ImmunoGenAI-optimized sequences across the benchmark set (n=20), where the optimized sequence distribution shifts toward higher immune evasion scores. Panel C compares baseline immunogenicity scores to ImmunoGenAI-optimized scores across all therapeutic proteins, showing that underlying pattern and context are preserved while overall immune evasion scores increase. Improvements are not uniform across therapeutic proteins, indicating that optimization respects intrinsic immunogenic architecture. The simulation was evaluated from a full test set (n=278). Panel D shows a

FALL 2026

leave-one-modality-out ablation study, where removal of each encoder decreases model performance compared to the full multimodal model, supporting the purpose of incorporating all six encoders. The y-axis represents the mean predicted immunogenicity score for each outcome.

OPTIMIZATION OUTCOMES

An important observation is that optimization effects were not uniform across all therapeutic proteins, suggesting that immunogenicity is highly context-dependent and influenced by intrinsic sequence architecture. The ablation study further highlights that each modality contributes uniquely to overall performance, confirming that multimodal integration is critical for balanced optimization.

A mild trade-off between mutation count and immune evasion was observed; however, most sequences required only 1–3 mutations, demonstrating that significant improvements can be achieved efficiently without excessive sequence modification. Additionally, due to the scale and heterogeneity of biological datasets, processing pipelines were adapted to ensure complete and unbiased analysis across all sequences.

Overall, ImmunoGenAI advances current approaches by enabling constraint-aware, sequence-level control of immune recognition. Unlike existing methods that optimize single objectives, this framework identifies key sequence regions that drive immunogenic responses and selectively modifies them to minimize unintended immunogenic effects while improving therapeutic and vaccine sequence efficacy (Zhang et al., 2025; Zhao & Wang, 2025).

CONCLUSION

In this paper, we introduce ImmunoGenAI, a multimodal deep learning tool that uses prediction and optimization algorithms to improve the immune evasion of a therapeutic sequence. We find that deep learning is particularly effective in enabling precise, sequence-level control of immune recognition by identifying and modifying key regions that drive immunogenic responses. By reducing unintended immunogenic peptide formation while preserving structural integrity, ImmunoGenAI improves therapeutic reliability without requiring extensive sequence modification.

This research demonstrates that sequences learned by a multimodal transformer neural network undergo selection more aligned with human immune presentation patterns, thereby increasing immune evasion beyond the industry standard. ImmunoGenAI significantly enhances immune compatibility in synthetic immunotherapeutic design for the same structural framework, development time, and manufacturing resources — enabling safer, longer-lasting personalized therapies which help save lives.

ImmunoGenAI exceeded the claim and met all of the engineering endpoints by significantly increasing immune evasion and reducing unintended antigen presentation while preserving structural stability. ImmunoGenAI outperforms 5 other immune prediction tools and their corresponding optimization methods, including industry-standard IEDB immunogenicity prediction scores.

ImmunoGenAI can deliver a 61.27% relative improvement in immune evasion within the same structural framework, improving immunotherapeutic design efficiency. Therapy cost can be reduced due to the increase in efficiency. Case study conducted on Monoclonal Antibody (mAb) Therapies & other benchmark sequences (Lim et al., 2022; Solá & Griebenow, 2010). ImmunoGenAI can be applied to therapeutic biologics developed for diseases such as COVID-19, Cancer, autoimmune disorders, and genetic enzyme replacement therapies (Cao et al., 2024; Greaney et al., 2021; Liu et al., 2025).

ImmunoGenAI can be applied across a wide range of biomedical domains, including cancer immunotherapies, infectious disease vaccines (e.g., COVID-19), autoimmune disorder treatments, and enzyme replacement therapies. In cancer immunotherapy and infectious disease vaccines, ImmunoGenAI can modulate antigen-derived peptide repertoires by reducing high-affinity HLA-binding epitopes or enhancing desired immunodominant regions, thereby refining T-cell activation profiles and improving specificity of immune targeting. In autoimmune and enzyme replacement therapies, it can selectively disrupt self-reactive epitope formation to minimize aberrant immune activation and anti-drug antibody responses, while preserving protein function and physiological activity across diverse therapeutic contexts. By redesigning existing biologics and guiding the development of new ones, it offers a pathway toward more precise and personalized treatments.

ImmunoGenAI decreases the likelihood of unintended immune recognition as opposed to industry-standard optimization techniques. The emergence of compensatory epitope hotspots and anti-drug antibody formation is mitigated through coordinated, context-aware multimodal optimization. This is paramount because some therapeutic proteins may inherently contain localized immunogenic regions that require coordinated modulation to prevent compensatory immune amplification during optimization (Fig. 11).

Beyond improving biological performance, this approach has the potential to reduce development costs and accelerate therapeutic design by minimizing trial-and-error optimization. Overall, ImmunoGenAI represents a new paradigm in immunotherapeutic engineering, where targeted, constraint-aware sequence optimization can systematically reduce unintended immune recognition while enhancing treatment efficacy.

Future work will extend ImmunoGenAI to optimize additional regulatory regions, including untranslated regions (UTRs), to further refine control over antigen expression and immune presentation. ImmunoGenAI can be scaled up even more by adding more parallel computing and training for a longer period

FALL 2026

of time. Developing an ImmunoGenAI research API will allow researchers to submit HTTP requests to predict and optimize therapeutic protein sequences, improving accessibility to the tool. Applying transfer learning on ImmunoGenAI with additional population-scale HLA datasets and disease-specific immunogenicity data can enhance adaptability across therapeutic domains. Industry collaboration and licensing opportunities, along with lab testing of ImmunoGenAI, can improve confidence and provide external validation. Together, these results position ImmunoGenAI as a scalable tool for frontier immunotherapeutic design.

ACKNOWLEDGEMENTS

I would like to thank my advisor who has a PhD in Colombia University, for guiding and supporting me in this research. I would also like to thank the dataset availability I utilized from the public sources: National Institutes of Health (NIH), Immune Epitope Database (IEDB), RCSB Protein Data Bank (PDB), AlphaFold, TCGA, GEO, and the Pan-Cancer Atlas.

AVAILABILITY AND REQUIREMENTS

Project name: ImmunoGenAI: Multimodal deep learning integration for predicting and optimizing immune evasion in personalized therapeutic devices.

Project home page: <https://github.com/lalithbhima/ImmunoGenAI>; Operating system(s): Platform independent Programming language: Python Other requirements: Python 3.9.4 or newer License: MIT Any restrictions to use by non-academics: license needed.

REFERENCES

- Aleksic, M., Liddy, N., Molloy, P. E., Pumphrey, N., Vuidepot, A., Chang, K. M., & Jakobsen, B. K. (2012). Different affinity windows for virus and cancer-specific T-cell receptors: Implications for therapeutic strategies. *European Journal of Immunology*, 42(12), 3174–3179. <https://doi.org/10.1002/eji.201242606>
- Anishchenko, I., Pellock, S. J., Chidyausiku, T. M., Ramelot, T. A., Ovchinnikov, S., Hao, J., Bafna, K., Norn, C., Kang, A., Bera, A. K., DiMaio, F., Carter, L., Chow, C. M., Montelione, G. T., & Baker, D. (2021). De novo protein design by deep network hallucination. *Nature*, 600(7889), 547–552. <https://doi.org/10.1038/s41586-021-04184-w>
- Calis, J. J., Maybeno, M., Greenbaum, J. A., Weiskopf, D., De Silva, A. D., Sette, A., Keşmir, C., & Peters, B. (2013). Properties of MHC class I presented peptides that enhance immunogenicity. *PLOS Computational Biology*, 9(10), e1003266. <https://doi.org/10.1371/journal.pcbi.1003266>
- Cao, Y., Yan, W., Yi, W., Yin, Q., & Li, Y. (2024). Bioengineered therapeutic systems for improving antitumor immunity. *National Science Review*, 12(1), nwae404. <https://doi.org/10.1093/nsr/nwae404>
- Coherent Market Insights. (2026). Synthetic biology market analysis & forecast: 2026–2033. <https://www.coherentmarketinsights.com/market-insight/synthetic-biology-market-112>
- Creary, L. E., Gangavarapu, S., Mallempati, K. C., Montero-Martín, G., Caillier, S. J., Santaniello, A., Hollenbach, J. A., Oksenberg, J. R., & Fernández-Viña, M. A. (2019). Next-generation sequencing reveals new information about HLA allele and haplotype diversity in a large European American population. *Human Immunology*, 80(10), 807–822. <https://doi.org/10.1016/j.humimm.2019.07.275>
- Das, D., Bhaduri, S., Pramanick, A., & Mitra, P. (2025). DeepPROTECTNeo: A deep learning-based personalized and RV-guided optimization tool for TCR epitope interaction using context-aware transformers. *bioRxiv*. <https://doi.org/10.1101/2025.01.04.631301>
- Dellas, N., Liu, J., Botham, R. C., & Huisman, G. W. (2021). Adapting protein sequences for optimized therapeutic efficacy. *Current Opinion in Chemical Biology*, 64, 38–47. <https://doi.org/10.1016/j.cbpa.2021.03.005>
- Doud, M. B., & Bloom, J. D. (2016). Accurate measurement of the effects of all amino-acid mutations on influenza hemagglutinin. *Viruses*, 8(6), 155. <https://doi.org/10.3390/v8060155>
- Greaney, A. J., Loes, A. N., Crawford, K. H. D., Starr, T. N., Malone, K. D., Chu, H. Y., & Bloom, J. D. (2021). Comprehensive mapping of mutations in the SARS-CoV-2 receptor-binding domain that affect recognition by polyclonal human plasma antibodies. *Cell Host & Microbe*, 29(3), 463–476.e6. <https://doi.org/10.1016/j.chom.2021.02.003>
- Hamelin, D. J., Scicluna, M., Saadie, I., Mostefai, F., Grenier, J. C., Baron, C., Caron, E., & Hussin, J. G. (2025). Predicting pathogen evolution and immune evasion in the age of artificial intelligence.

-
- Computational and Structural Biotechnology Journal, 27, 1370–1382.
<https://doi.org/10.1016/j.csbj.2025.03.044>
- Hayes, T., Rao, R., Akin, H., Sofroniew, N. J., Oktay, D., Lin, Z., Verkuijl, R., Tran, V. Q., Deaton, J., Wiggert, M., et al. (2024). Simulating 500 million years of evolution with a language model. *bioRxiv*.
<https://doi.org/10.1101/2024.07.01.600583>
- Hong, K. L., Gao, B., Stalder, L., Ehling, R. A., Horvath, M., Wehrle, S., Foster, J., Junet, V., Kucharczky, J., Lallée, S., Diringer, M. C., Yang, Q., Vazquez-Lombardi, R., & Reddy, S. I. (2025). Predicting TCR antigen specificity at proteome-scale with synthetic immune cells and machine learning. *bioRxiv*.
<https://doi.org/10.1101/2025.07.23.666264>
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., Back, T., Petersen, S., Reiman, D., Clancy, E., Zielinski, M., Steinegger, M., Pacholska, M., Berghammer, T., Bodenstein, S., Silver, D., Vinyals, O., Senior, A. W., Kavukcuoglu, K., Kohli, P., & Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583–589.
<https://doi.org/10.1038/s41586-021-03819-2>
- Jurtz, V., Paul, S., Andreatta, M., Marcatili, P., Peters, B., & Nielsen, M. (2017). NetMHCpan-4.0: Improved peptide-MHC class I interaction predictions integrating eluted ligand and peptide binding affinity data. *Journal of Immunology*, 199(9), 3360–3368. <https://doi.org/10.4049/jimmunol.1700893>
- Kumar, A., Dixit, S., Srinivasan, K., M, D., & Vincent, P. M. D. R. (2024). Personalized cancer vaccine design using AI-powered technologies. *Frontiers in Immunology*, 15, 1357217.
<https://doi.org/10.3389/fimmu.2024.1357217>
- Li, G., Iyer, B., Prasath, V. B. S., Ni, Y., & Salomonis, N. (2021). DeepImmuno: Deep learning-empowered prediction and generation of immunogenic peptides for T-cell immunity. *Briefings in Bioinformatics*, 22(6), bbab160. <https://doi.org/10.1093/bib/bbab160>
- Liang, B., Wang, N., Huang, H., Hou, Y., & Qiang, X. (2024). VIEPred: Predicting viral immune evasion with multimodal representation learning. *Journal of Physics: Conference Series*, 2868, 012029.
<https://doi.org/10.1088/1742-6596/2868/1/012029>
- Lim, Y. W., Adler, A. S., & Johnson, D. S. (2022). Predicting antibody binders and generating synthetic antibodies using deep learning. *mAbs*, 14(1), 2069075.
<https://doi.org/10.1080/19420862.2022.2069075>
- Liu, S., Zhong, Z., Tu, Q., & Jin, M. (2025). Synthetic biology approaches to enhance cancer immune responses. *Frontiers in Immunology*, 16, 1698158. <https://doi.org/10.3389/fimmu.2025.1698158>
- Parker, A. S., Zheng, W., Griswold, K. E., & Bailey-Kellogg, C. (2010). Optimization algorithms for functional deimmunization of therapeutic proteins. *BMC Bioinformatics*, 11, 180.
<https://doi.org/10.1186/1471-2105-11-180>

FALL 2026

-
- Rives, A., Meier, J., Sercu, T., Goyal, S., Lin, Z., Liu, J., Guo, D., Ott, M., Zitnick, C. L., Ma, J., & Fergus, R. (2021). Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15), e2016239118. <https://doi.org/10.1073/pnas.2016239118>
- Robinson, J., Halliwell, J. A., Hayhurst, J. D., Flicek, P., Parham, P., & Marsh, S. G. (2015). The IPD and IMGT/HLA database: Allele variant databases. *Nucleic Acids Research*, 43(D1), D423–D431. <https://doi.org/10.1093/nar/gku1161>
- Sarkisyan, K. S., Bolotin, D. A., Meer, M. V., Usmanova, D. R., Mishin, A. S., Sharonov, G. V., Ivankov, D. N., Bozhanova, N. G., Baranov, M. S., Soylemez, O., Bogatyreva, N. S., Vlasov, P. K., Egorov, E. S., Logacheva, M. D., Kondrashov, A. S., Chudakov, D. M., Putintseva, E. V., Mamedov, I. Z., Tawfik, D. S., Lukyanov, K. A., & Kondrashov, F. A. (2016). Local fitness landscape of the green fluorescent protein. *Nature*, 533(7603), 397–401. <https://doi.org/10.1038/nature17995>
- Sidney, J., Assarsson, E., Moore, C., Ngo, S., Pinilla, C., Sette, A., & Peters, B. (2008). Quantitative peptide binding motifs for 19 human and mouse MHC class I molecules derived using positional scanning combinatorial peptide libraries. *Immunome Research*, 4, 2. <https://doi.org/10.1186/1745-7580-4-2>
- Solá, R. J., & Griebenow, K. (2010). Glycosylation of therapeutic proteins: An effective strategy to optimize efficacy. *BioDrugs*, 24(1), 9–21. <https://doi.org/10.2165/11530550-000000000-00000>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* (pp. 6000–6010). Curran Associates Inc. <https://arxiv.org/abs/1706.03762>
- Wec, A. Z., Lin, K. S., Kwasniewski, J. C., Sinai, S., Gerold, J., & Kelsic, E. D. (2021). Overcoming immunological challenges limiting capsid-mediated gene therapy with machine learning. *Frontiers in Immunology*, 12, 674021. <https://doi.org/10.3389/fimmu.2021.674021>
- Yewdell, J. W. (2006). Confronting complexity: Real-world immunodominance in antiviral CD8+ T cell responses. *Immunity*, 25(4), 533–543. <https://doi.org/10.1016/j.immuni.2006.09.005>
- Zhang, G., Fu, Y., Buratto, D., Chan, K. C., Wang, H., Huang, Y., Huang, L., & Zhou, R. (2025). Reengineering the antigen optimization process for superior neoantigen vaccine design. <https://doi.org/10.1101/2025.08.03.668363>
- Zhao, Y., & Wang, H. (2025). Artificial intelligence-driven circRNA vaccine development: Multimodal collaborative optimization and a new paradigm for biomedical applications. *Briefings in Bioinformatics*, 26(3), bbaf263. <https://doi.org/10.1093/bib/bbaf263>
- Zinsli, L. V., Stierlin, N., Loessner, M. J., & Schmelcher, M. (2020). Deimmunization of protein therapeutics—Recent advances in experimental and computational epitope prediction and deletion. *Computational and Structural Biotechnology Journal*, 19, 315–329. <https://doi.org/10.1016/j.csbj.2020.12.024>