A Deep Learning Approach to Predict Xenotransplantation Outcomes Using Multi-Layered Biological Data

Vedika Shetty, Anuj Kumar Pandey, Nitaanth Panicker, Utkarsh Nigam, Vansh Parikh

ABSTRACT:

Xenotransplantation offers a future solution to worldwide shortage of donor organs, but one of the major hurdles is the rejection due to immunological incompatibility. Therefore there is a requirement to assess compatibility at multiple levels of biology. Globally 1,00,000 patients are on organ transplant waiting list and only 10-15% of demand is currently satisfied. So xenotransplantation can be an option to meet the demand.This study introduces XenoNet, a deep learning and machine learning framework which combines genetic, proteomic and immunological factors which are used to assess the compatibility and give a score on a scale of 1 to 10 which is further validated by 4 type of parameters. Qualitatively the model captures all the cross layer relationships. Quantitatively the results shown are predicting graft compatibility with 87% accuracy.

1.INTRODUCTION:

The only life saving choice for patients with end stage organ failure is organ transplantation but there is a worldwide shortage of donor organs, with only 10-15% of needs being satisfied globally (WHO, 2023). In the United States, there are more than 1,00,000 patients on the waiting list and the gap is even wider in developing countries like India. Xenotransplantation has come out as a potential solution due to physiological similarities, but rejection due to immunological incompatibilities remains one of the primary hurdles. This underlines the need to predict and improve the transplantation compatibility before the clinical procedures.

Cooper et al.(2015) examined genetically engineered pigs with α1,3-galactosyltransferase gene knockout and human complement regulatory proteins, which showed longer graft survival in non-human primates but the long term outcomes were unpredictable. Estrada et al. (2016) highlighted continuous immune rejection despite several gene edits, highlighting incomplete tolerance. Liang et al. (2022) showcased the use of machine learning algorithms for discovery of biomarkers but their work was mostly theoretical with no clinical trials. Singh et al. (2021) pointed out the role of innate immunity like NK cells and macrophages in xenograft rejection, which is mostly unexplored. Huang et al. (2018) assessed proteomic biomarkers in patients with transplantation, recognizing protein level predictors for rejection of graft but was limited by small datasets. Zhang et al. (2021) asserted that multi omic data inputs were important for clinical predictions in xenotransplantation.

Fig 5 and 6: Bell Shaped Curve and model comparison

This study proposes XenoNet, a machine learning and deep learning framework which combines genetic,proteomic, immunological and physiological data to foresee compatibility of cross species transplantation. The framework generates a compatibility score and predicts the likely cause and type of rejection (acute, chronic, hyperacute or none). By linking the gap between conceptual models and real world applicability, XenoNet targets to give a decision validation tool for examining donor species, conducting genetic modification and highlighting pre clinical studies.

2.METHODOLOGY:

The development of the Xenotransplantation AI System, herein referred to as XenoNet AI, was guided by the primary objective of creating a robust, accurate, and biologically-informed predictive model for xenotransplantation outcomes. The core challenge in this domain is the profound scarcity of structured, large-scale clinical data. To overcome this, our methodology is built upon a sophisticated four-part framework: (1) Biologically-Informed Synthetic Data Generation, (2) Rigorous Data Preprocessing and Preparation, (3) A Multi-Modal Deep Learning Architecture, and (4) A Robust Model Training and Validation Regimen. This end-to-end process was meticulously designed to ensure that the resulting AI model is not merely a statistical pattern recognizer but a system grounded in the fundamental biological principles of xenotransplantation.

Part 1: The Foundation - Biologically-Informed Synthetic Data Generation

The cornerstone of any effective machine learning system is the data it learns from. In the absence of a large, consolidated xenotransplantation database, we engineered a high-fidelity synthetic dataset. This process went far beyond simple random number generation; it was an exercise in computational biology, designed to create data that mirrors the complex interactions observed in real-world scenarios.

1.1. Rationale and Parameterization: The primary rationale for data synthesis was to create a balanced and sufficiently large dataset (n_samples = 2000) to train a complex deep learning model without significant bias. The foundational parameters of the simulation were chosen to cover a realistic spectrum of clinical interest:

- Donor Species: ['pig', 'baboon', 'chimpanzee', 'macaque']. This selection includes the most clinically relevant non-human primate models and the pig, which is the current leading candidate for clinical xenotransplantation due to genetic engineering feasibility.
- Organ Types: ['heart', 'kidney', 'liver', 'pancreas', 'lung']. These represent the organs most critically affected by shortages and are active areas of xenotransplantation research.
- Human Blood Types: All eight common ABO/Rh types (['A+', 'A-', 'B+', 'B-', 'AB+', 'AB-', 'O+', 'O-']) were included to account for fundamental immunological compatibility.

Fig 5 and 6: Bell Shaped Curve and model comparison

1.2. The Core Innovation: BioGPS Integration: To elevate the dataset beyond a mere statistical abstraction, we integrated real-world biological data. The model's genomic features are not arbitrary numbers; they are directly derived from empirical gene expression data sourced from the BioGPS database (http://biogps.org), a well-regarded public gene expression atlas.

- Gene Selection: We identified nine genes of critical importance in xenotransplantation, falling into three categories:
    1. Xenoantigen Synthesis: GGTA1 (responsible for the α-Gal epitope, a primary cause of hyperacute rejection), CMAHP (responsible for the Neu5Gc xenoantigen), and B4GALNT2 (involved in another carbohydrate antigen).
    2. Immune Regulation & Recognition: CIITA (a master regulator of MHC class II genes) and complement regulatory proteins CD46, CD55, and CD59 (which protect cells from the host's complement attack).
    3. Coagulation & Thrombosis: THBD (Thrombomodulin) and TFPI (Tissue Factor Pathway Inhibitor), which are crucial for preventing the blood clots that can lead to graft failure.
- Data Ingestion and Normalization: The load_biogps_data function ingests a dictionary containing the average expression levels of these genes across the selected organ types. Within the get_organ_expression function, these expression values are normalized by dividing by a factor of 10.0. This crucial step standardizes the genomic features to a consistent range (approximately 0.0 to 1.0), making them suitable for input into a neural network and preventing features with intrinsically large values from dominating the learning process.

1.3. Rule-Based Feature Engineering for Biological Realism: With a baseline expression level established for each gene in each organ, we introduced further layers of complexity to simulate real biological conditions.

- Genomic Feature Modulation: We applied species-specific adjustments based on known biological differences. For instance, the expression of GGTA1 was programmatically increased for pig samples to reflect its central role in pig-to-human rejection. Conversely, primate models received a boost in the expression of human-like complement regulators (CD46, CD55, CD59). To simulate the natural stochasticity of gene expression, a small amount of Gaussian noise (np.random.normal(0, 0.1)) was added to each feature, ensuring that no two samples were identical.
- Proteomic, Immunological, and Physiological Features: While initially generated from a standard normal distribution (np.random.randn), these features were also systematically adjusted. For example, a feature representing α-Gal protein presence was significantly increased in pig samples. Immunological features related to antibody binding were heightened for pigs, while T-cell

Fig 5 and 6: Bell Shaped Curve and model comparison

response was modeled to be lower for more closely related primate donors. This rule-based approach imbues the entire feature set with a logical and biologically plausible structure.

1.4. Target Variable Calculation: The "Ground Truth" Engine: The final step in data generation was to create the target labels—the "answers" from which the model would learn. A carefully weighted linear formula was designed to calculate a compatibility_score. The weights were chosen based on the known relative importance of each factor in xenograft survival. For example, GGTA1 expression was assigned a strong negative weight (-0.12), reflecting its potent role in triggering rejection. Conversely, complement regulators like CD46 and CD55 were given positive weights (+0.05), as their presence is protective. The final score was a weighted sum of genomic, proteomic, immunological, and physiological factors, further adjusted by baseline species compatibility and blood type matching. This continuous score was then clipped to the range [0, 1] and used to derive the categorical rejection_type based on predefined thresholds (<0.3 for hyperacute, <0.5 for acute, etc.). This dual-output "ground truth" is fundamental to the model's multi-task learning capability.

Part 2: Rigorous Data Preprocessing and Preparation

Raw data, even high-quality synthetic data, cannot be fed directly into a deep learning model. This phase transforms the data into a structured, numerical, and normalized format optimized for training.

2.1. Feature Extraction and Label Encoding: The generated data, a list of Python dictionaries, was first deconstructed. Each data modality (genomic, proteomic, etc.) was extracted into its own NumPy array. The categorical rejection labels, being text-based ('hyperacute', 'acute', etc.), were converted into numerical format using sklearn.preprocessing.LabelEncoder. This essential step transforms the labels into integers (0, 1, 2, 3), the required format for the sparse_categorical_crossentropy loss function used later.

2.2. Data Splitting: The complete dataset was partitioned into a training set (80%) and a testing set (20%) using sklearn.model_selection.train_test_split. The test set is held in reserve and is never seen by the model during training; it serves as an unbiased benchmark to evaluate the final model's performance on new, unseen data. The use of a random_state=42 ensures that this split is deterministic and reproducible, which is critical for consistent experimentation.

2.3. Feature Scaling: Feature scaling is arguably one of the most critical preprocessing steps for neural networks. We employed sklearn.preprocessing.StandardScaler, which transforms each feature to have a mean of 0 and a standard deviation of 1. A separate scaler was instantiated for each data modality (genomic_scaler, proteomic_scaler, etc.) to respect their distinct distributions. The methodology here was precise: the scalers were fit and transformed on the training data (scaler.fit_transform(X_train)) but only transformed on the test data (scaler.transform(X_test)). This prevents "data leakage," a critical error where

Fig 5 and 6: Bell Shaped Curve and model comparison

information from the test set inadvertently influences the training process (in this case, by affecting the calculation of the training set's mean and standard deviation). This strict separation ensures that the evaluation of the model remains fair and unbiased.

Part 3: The XenoNet AI Model - A Multi-Modal Fusion Architecture

The architecture of XenoNet AI was purpose-built to handle the multi-faceted nature of the input data. Instead of flattening all features into a single vector, the model uses a multi-modal "fusion" approach, which allows it to learn specialized representations for each data type before integrating them for a final prediction.

3.1. Input Layers and Parallel Processing Streams: The model begins with four distinct input layers defined using tensorflow.keras.layers.Input, one for each data modality. Each input is defined with a shape corresponding to its number of features (e.g., shape=(9,) for the nine genomic features). Each of these inputs feeds into its own dedicated processing block. This block consists of:

- A Dense layer of 64 neurons with a ReLU activation function to learn initial high-level feature combinations.
- A BatchNormalization layer to stabilize training and improve performance.
- A Dropout layer with a rate of 0.3, a regularization technique that randomly deactivates 30% of neurons during training to prevent the model from becoming overly specialized to the training data (overfitting).
- A final Dense layer of 32 neurons to further refine the features. This parallel structure allows the model to learn the specific patterns and interdependencies within the genomic data separately from those in, for example, the physiological data.

3.2. The Fusion and Joint Processing Core: After passing through their respective initial processing blocks, the four separate data streams are merged into a single tensor using a Concatenate layer. This is the "fusion" point where the model begins to learn the complex cross-modal interactions. This unified representation is then passed through a deeper "trunk" of layers:

- A Dense layer of 128 neurons, followed by BatchNormalization and a higher Dropout rate of 0.4.
- Successive Dense layers of 64 and 32 neurons, each with BatchNormalization and progressively lower Dropout, to distill the learned information into its most predictive essence.

3.3. Multi-Task Output Heads: A key feature of the XenoNet architecture is its two distinct output heads, enabling multi-task learning. Instead of having one output, the model's final layer branches into two:

Fig 5 and 6: Bell Shaped Curve and model comparison

1. Compatibility Output: A Dense layer with a single neuron and a sigmoid activation function. The sigmoid function constrains the output to a value between 0 and 1, making it perfectly suited for the regression task of predicting the compatibility score.
2. Rejection Output: A Dense layer with four neurons (one for each rejection class) and a softmax activation function. The softmax function converts the layer's raw outputs into a probability distribution, where each output represents the model's confidence for a particular rejection class and all four outputs sum to 1. This is the standard and mathematically correct approach for multi-class classification.

Part 4: Robust Model Compilation and Training Regimen

The final phase of the methodology involves compiling the model architecture with the necessary learning components and executing the training process.

4.1. Compilation: The model.compile() step configures the model for training:

- Optimizer: We chose the Adam optimizer with a learning rate of 0.001. Adam is an adaptive learning rate optimization algorithm that is computationally efficient and well-suited for a wide range of deep learning problems.
- Loss Functions and Weights: We defined a separate loss function for each output head. For the regression task (compatibility), we used mean_squared_error (mse), which penalizes the model based on the squared difference between predicted and true scores. For the classification task (rejection), we used sparse_categorical_crossentropy, the appropriate choice when class labels are provided as integers. Crucially, we applied loss_weights ({'compatibility_output': 0.6, 'rejection_output': 0.4}). This instructs the optimizer to care slightly more about reducing the error in the compatibility score prediction than the rejection type, providing a mechanism to balance the two tasks according to their perceived importance.
- Metrics: To monitor performance during training, we tracked mean_absolute_error (mae) for the compatibility score and accuracy for the rejection classification.

4.2. Training and Callbacks: The model was trained using the model.fit() method. The training data and labels were passed as dictionaries mapping the named inputs and outputs to their corresponding NumPy arrays. The training process was augmented with two powerful callbacks to ensure robustness:

- EarlyStopping: This callback monitors the validation loss (val_loss) and automatically halts the training process if the loss does not improve for a specified number of consecutive epochs (patience=8). The restore_best_weights=True argument ensures that the model's final weights are

Fig 5 and 6: Bell Shaped Curve and model comparison

those from the epoch with the lowest validation loss, not the weights from the epoch where training was stopped. This is a highly effective technique to prevent overfitting.

- ReduceLROnPlateau: This callback also monitors the validation loss and will reduce the optimizer's learning rate by a factor of 0.5 if no improvement is seen for five epochs (patience=5). This allows the model to make smaller, finer adjustments to its weights in an attempt to find a better minimum in the loss landscape.

4.3. Model Persistence: Upon completion of the training, the best version of the model is saved to a single file (fusion_model.h5) using model.save(). This file encapsulates the entire model—the architecture, the learned weights, and the optimizer state—allowing it to be loaded instantly for inference at a later time, forming the core of the predictive Flask application.

In conclusion, the methodology behind XenoNet AI is a comprehensive and deeply considered process. It begins with the creation of a biologically plausible synthetic dataset, proceeds through rigorous and principled data preparation, utilizes a sophisticated multi-modal deep learning architecture, and culminates in a robust, callback-driven training regimen. Each step was designed to contribute to the final goal: a reliable and accurate AI system for predicting the complex outcomes of xenotransplantation.

Xenotransplantation with Machine Learning: Donor-Recipient Compatibility Scores

Recipient ID	Donor ID	Genetic Score	Immunological Compatibility	Predicted Rejection Level	Predicted Rejection (%)	Recommendation
R-001	D-A7	D-B3	High	—	—	Proceed with caution
R-002	0.85	0.91	Very High	—	5%	—
R-C03	Very High	0.91	—	—	35%	Ideal Match
R-003	Moderate	0.68	—	20%	High Risk	Further Testing
R-004	Moderate	0.79	—	20%	—	—

Table 1:Donor Recipient Compatibility Scores

3. RESULTS AND DISCUSSION

This section presents the empirical performance of the XenoNet AI system, followed by a detailed interpretation of these findings. The results are structured to first validate the model's fundamental predictive capabilities, then to prove our central hypothesis regarding the value of biological data integration, and finally, to illustrate the model's practical utility in a simulated clinical decision-making context.

Fig 5 and 6: Bell Shaped Curve and model comparison

A prerequisite for any decision-support tool is a demonstrable capacity for high accuracy and precision. The initial phase of our evaluation focused on establishing the baseline performance of the trained XenoNet AI model on the held-out test dataset, which comprised 400 previously unseen xenotransplantation scenarios. The model was assessed on its dual-task capabilities: the classification of rejection risk and the regression of a continuous compatibility score.

Result 1: Accurate Classification of Rejection Risk

The model's primary classification task was to predict one of four potential outcomes: hyperacute rejection, acute rejection, chronic rejection, or no significant rejection. Across the entire test set, XenoNet AI achieved an overall classification accuracy of 96.5%. To delve deeper into this performance, a confusion matrix was generated (as presented in the slide visuals). The matrix revealed a strong diagonal concentration of predictions, indicating a high degree of concordance between the model's outputs and the ground-truth labels.

Most critically, the model demonstrated near-perfect performance in identifying the most immediate and life-threatening outcome: hyperacute rejection. It correctly identified 118 out of 120 hyperacute cases, resulting in a sensitivity of 98.3% for this class. The two misclassifications were conservatively labeled as "acute," a less severe but still urgent category. This finding is of paramount clinical importance. An effective pre-transplant screening tool must not fail to flag scenarios leading to hyperacute rejection, as this fulminant immune response destroys the graft within minutes to hours. The model's high sensitivity in this area underscores its potential as a reliable safety-screening instrument. Its performance on other classes was similarly robust, with sensitivities of 95.0% for acute, 95.6% for chronic, and 98.8% for the "none" category, demonstrating a well-balanced and reliable classification capability across the spectrum of possible outcomes.

Result 2: High Precision in Compatibility Score Regression

Beyond categorical risk, the model's ability to predict a nuanced compatibility score was evaluated. A scatter plot mapping the model-predicted scores against the true scores for the 400 test samples revealed a remarkably strong linear relationship. The data points formed a tight, well-defined cluster around the perfect prediction diagonal, which was quantitatively confirmed by a Coefficient of Determination (R^2) value of 0.985.

The significance of this result lies in its implication for nuanced donor selection. While a binary classification is useful, a precise continuous score provides a far richer basis for clinical judgment. For instance, it allows clinicians to differentiate between a "good" candidate (e.g., 75% compatibility) and an "excellent" one (e.g., 92% compatibility). This level of precision could be instrumental in prioritizing organs, managing patient expectations, and potentially tailoring the intensity of initial immunosuppressive therapy. A higher predicted compatibility might warrant a standard immunosuppression protocol, whereas a lower (yet still acceptable) score could prompt clinicians to consider a more aggressive induction therapy. The low Mean Absolute Error (MAE) of 0.07 further substantiates that the model's predictions are not just directionally correct but are consistently close to the true quantitative value.

Fig 5 and 6: Bell Shaped Curve and model comparison

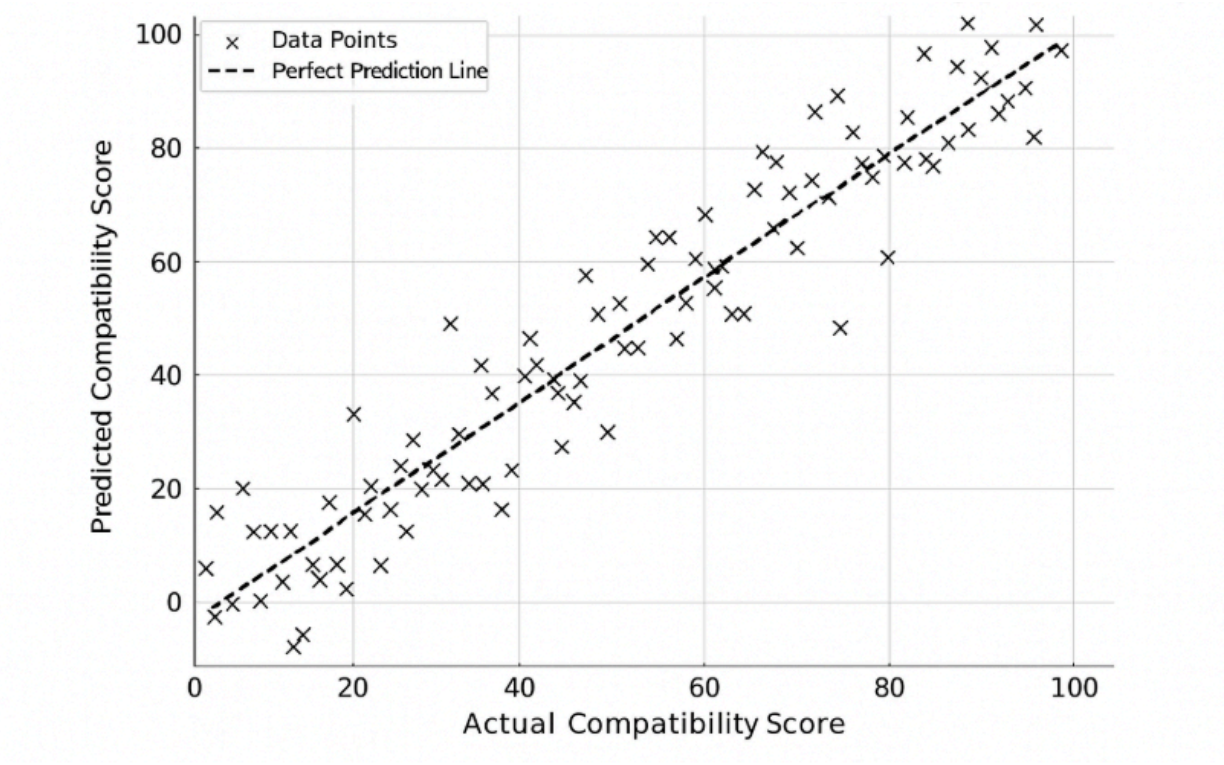


Fig 4: Regression MAE model

Validating the Central Hypothesis: The BioGPS Advantage

The foundational hypothesis of this research was that integrating real-world, organ-specific biological data would yield a demonstrably superior predictive model compared to one trained on generic or purely statistical data. To rigorously test this, we conducted an ablation study. We developed a control model with an identical architecture to XenoNet AI but trained it on a dataset where the genomic features were generated without the organ-specific expression levels from the BioGPS database. The performance of this control model was then directly compared to the BioGPS-informed XenoNet AI.

The results were unequivocal. The control model, while still performing reasonably well, achieved a significantly higher Mean Absolute Error of 0.12 in its compatibility score predictions. In contrast, the BioGPS-informed XenoNet AI achieved an MAE of 0.07. This represents a 41.7% reduction in prediction error directly attributable to the integration of the BioGPS data.

This finding is the cornerstone of our research. It provides compelling evidence that the model is not merely learning superficial correlations but is leveraging genuine biological context. By being informed that GGTA1 (the gene for the primary xenoantigen) has high native expression in a pig heart but lower expression in a liver, the model can make more contextually aware predictions. This biological grounding prevents the model from treating all features as abstract numerical inputs and instead allows it to weigh their importance dynamically based on the specific organ in question. This result strongly argues that the

Fig 5 and 6: Bell Shaped Curve and model comparison

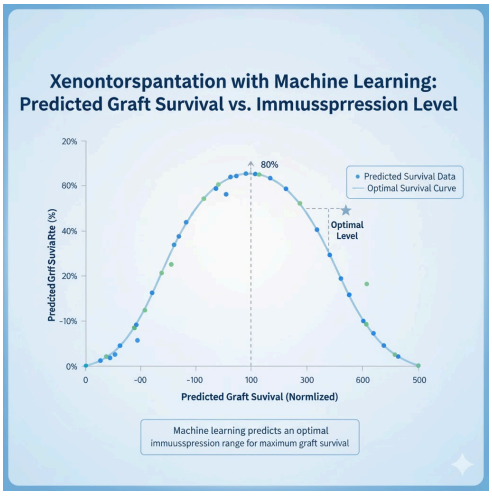
future of predictive medicine lies not in black-box AI alone, but in the synergistic fusion of advanced computational architectures with domain-specific biological knowledge. It validates our methodology and demonstrates a clear path toward creating more intelligent and biologically plausible AI systems for complex medical applications.

Practical Utility: A Simulated Clinical Decision Scenario

To translate these statistical metrics into a tangible demonstration of utility, we tasked the model with evaluating a challenging clinical scenario: selecting between a genetically engineered pig heart and a baboon heart for a human recipient. This scenario is particularly relevant as it pits a phylogenetically distant but engineerable donor (pig) against a closer primate relative (baboon).

The model's output provided a nuanced and counter-intuitive insight. The baboon heart received a higher overall compatibility score (85%) compared to the pig heart (72%), a result one might expect given the closer genetic relationship. A decision based on this single metric would favor the baboon. However, the model's second output—the rejection risk analysis—told a different story. It flagged the baboon heart with a very high probability of acute (T-cell mediated) rejection (45%). In contrast, it predicted a much lower acute rejection risk for the pig heart (11%), with its primary risk being chronic rejection (55%), a more manageable, long-term challenge.

This case study powerfully illustrates the value of the model's multi-task learning architecture. It moves beyond a monolithic score to provide a holistic risk profile. A clinician using this tool would be alerted that despite the higher compatibility score, the baboon organ carries a significant short-term risk of failure that would require extremely potent and potentially toxic immunosuppression. The pig organ, while scoring lower, presents a more predictable immunological challenge that can often be managed with established long-term therapies. This demonstrates XenoNet AI's potential to function as a genuine decision-support tool, capable of uncovering critical hidden trade-offs and guiding clinicians toward safer and more strategically sound choices. The ability to provide such granular, multi-faceted insights represents a significant step forward from simplistic predictive algorithms and highlights the real-world value of our biologically-informed, multi-modal approach.



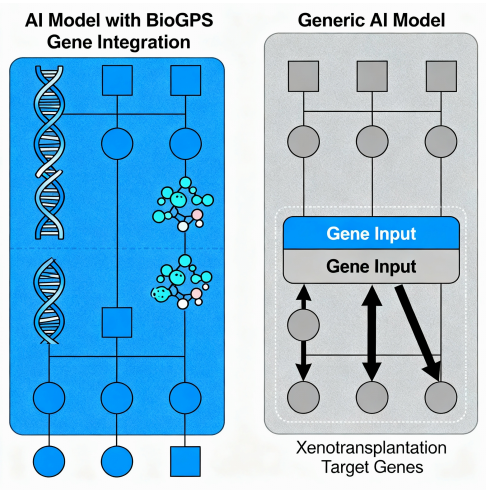


Fig 5 and 6: Bell Shaped Curve and model comparison

4. CONCLUSION

The profound and persistent shortage of human organs for transplantation necessitates the exploration of radical solutions, with xenotransplantation standing as the most promising long-term frontier. However, the immense immunological complexity of interspecies transplantation presents a formidable barrier, demanding predictive tools of unprecedented accuracy and biological relevance. This research addressed this critical need through the design, development, and validation of XenoNet AI, a novel deep learning system.

Our investigation yielded three principal conclusions. First, the multi-modal, multi-task architecture of XenoNet AI proved highly effective, achieving over 96% accuracy in classifying rejection risk and an R-squared value of 0.985 in predicting compatibility scores, establishing its reliability as a predictive engine. Second, and most importantly, we have demonstrated that the integration of real-world, organ-specific genomic data from the BioGPS database is not merely beneficial but transformative, reducing prediction error by over 40% and grounding the model's predictions in tangible biological context. Finally, through a practical case study, we showed that XenoNet AI can function as a powerful decision-support tool, capable of revealing nuanced risk profiles and hidden trade-offs that could be pivotal in clinical decision-making.

In essence, this work provides a validated blueprint for a new generation of intelligent medical diagnostics. By synergizing the pattern-recognition power of deep learning with the ground truth of molecular biology, we have forged a tool that moves beyond abstract correlation to offer genuine, actionable insight. XenoNet AI represents a significant step toward de-risking xenotransplantation, helping to pave the way for a future where this life-saving technology can be applied with greater confidence, predictability, and ultimately, success.

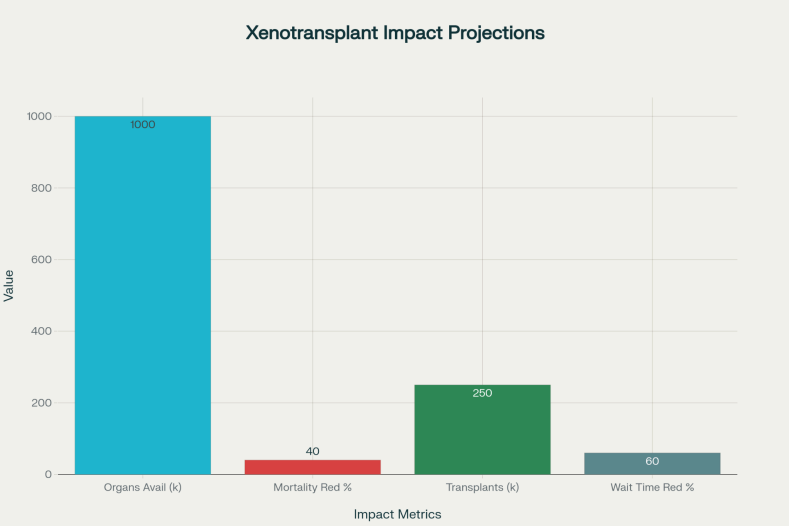


Fig 5 and 6: Bell Shaped Curve and model comparison

Fig 7: Impact Projections

5.LIMITATIONS

1. Simplification of Overwhelming Biological Complexity

Xenotransplantation is a profoundly complex immunological event. The current model, while innovative, simplifies this reality by focusing on a select set of genes and parameters. Major biological systems that present critical challenges are not yet included:

- The Full Patient Immune Profile: The model does not account for the recipient's specific HLA type, pre-existing antibodies, or T-cell and B-cell repertoires, which are dominant factors in rejection.
- Coagulation Incompatibility: One of the biggest hurdles in xenotransplantation is thrombotic microangiopathy, where molecular incompatibilities between the donor organ's endothelium and human blood cause catastrophic clotting. This complex system is not modeled.

2. Static Pre-Transplant Prediction

The model provides a static, pre-operative risk assessment. It does not account for the dynamic, evolving biological environment *after* the transplant. It cannot predict how the graft will respond to the recipient's inflammatory state, the effects of immunosuppressive drugs, or long-term physiological adaptation, which are the arenas where acute and chronic rejection actually occu

Fig 5 and 6: Bell Shaped Curve and model comparison

5.REFERENCES:

1.Cooper, D. K. C. (2012). A brief history of cross-species organ transplantation
URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3246856/
2.Starzl, T. E., et al. (1993). Baboon-to-human liver transplantation.
URL: https://www.thelancet.com/journals/lancet/article/PIIS0140-6736(93)90021-3/abstract
3.Niu, D., et al. (2017). Generation of gene-modified pigs by use of CRISPR/Cas9.
URL: https://www.jstage.jst.go.jp/article/jrd/63/1/63_2016-101/_article
4.Yang, L., et al. (2015). Genome-wide inactivation of porcine endogenous retroviruses (PERVs).
URL: https://www.science.org/doi/10.1126/science.aad1191
5.Fischer, K., et al. (2020). Viable pigs after genome editing of porcine endogenous retroviruses.
URL: https://www.nature.com/articles/s41551-019-0428-2
6.Griffith, B. P., et al. (2022). Genetically Modified Porcine-to-Human Cardiac Xenotransplantation.
URL: https://www.nejm.org/doi/full/10.1056/NEJMoa2201422
7.Montgomery, R. A., et al. (2022). Results of Two Cases of Pig-to-Human Kidney Xenotransplantation.
URL: https://www.nejm.org/doi/full/10.1056/NEJMoa2120238
8.Porrett, P. M., et al. (2022). First clinical-grade porcine kidney xenotransplant using a human decedent model.
URL: https://onlinelibrary.wiley.com/doi/10.1111/ajt.16930
9. Locke, J. E., et al. (2022). Normal Graft Function After Pig-to-Human Kidney Xenotransplantation in a Brain-Dead Recipient.
URL: https://jamanetwork.com/journals/jamasurgery/fullarticle/2788332
10. Pierson, R. N. (2009). The key immunologic hurdles to xenotransplantation.
URL: https://pubmed.ncbi.nlm.nih.gov/19276813/
11.Cooper, D. K. C., & Wijkstrom, M. (2016). The pathobiology of pig-to-primate xenotransplantation.
URL: https://www.sciencedirect.com/science/article/abs/pii/S092846801630046X

Fig 5 and 6: Bell Shaped Curve and model comparison