A Scaffold-Aware Machine Learning Docking Pipeline for TYK2 Inhibitor Discovery with Calibrated Prioritization of 32 Active Compounds Including Deucravacitinib

Saima Akram
Electrical Engineering Department, National Fertilizer Corporation Institute of Engineering and Technology (NFC IET), Multan, Pakistan
Email: saimaakram@nfciet.edu.pk

Muhammad Shahid
Department of Computer Science & Information Technology, University of South Punjab (USP), Multan, Punjab, Pakistan.
Email: mshahidbhatti0007@gmail.com

Ghulam Muhy Ud Deen Raee
Department of Computer Science & Information Technology, University of South Punjab (USP), Multan, Punjab, Pakistan.
Email: gmraees@gmail.com

Muhammad Allah Razi
Department of Computer and Software Engineering, The Khwaja Fareed University of Engineering and Information Technology (KFUEIT)
Email: allah2003dakhnah@gmail.com

Received on: 05-10-2025 Accepted on: 10-11-2025

Abstract

Tyrosine kinase 2 (TYK2) represents a proven target in the immunology field, with the FDA-approved TYK2 inhibitor deucravacitinib being a prime example. Nonetheless, the identification of new modulators for TYK2 has proven difficult owing to noisy protein-target activity data, presence of a bias in a molecule's core chemical structure, also known as a scaffold, and the high cost of screen-and-lead experiments in the laboratory. In the current study, we introduced a machine learning strategy focused on scaffolds, incorporating a rigorous data curation process, generation of molecule fingerprint embedders, and the application of calibrated classification algorithms followed by the implementation of physical molecule docking experiments. The curated TYK2 protein-target activity data (pIC50) were converted into ECFP4, MACCS, and a series of physicochemical features, and the

resulting collection of features was reduced via variance and correlation pruning. The performance of the classification tasks was evaluated using the Support Vector Machine, Random Forest, and XGBoost algorithms using a two-step process involving a scaffold-split validation strategy, so the model can, in a probable and reliable manner, generalize in the chemical space. The XGBoost model, among the tested machine learning algorithms, showed the best possible results, having achieved an accuracy of 87.5%, an F1-score of 91.3%, and an area under the curve of 95.1% in the task of protein-target classification of the TYK2 kinase. The optimized model was applied to a screening library comprised of more than 10,000 different chemical structures, and the top 32 active structures were filtered out using a probability threshold of 95%, and the structures indeed displayed stable docking geometries in subsequent docking experiments using the Surflex-TBS docking tool. Importantly, the model was also able to predict the known TYK2 inhibitors deucravacitinib, predicting it and ranking it among the active chemical structures, thus validating the proposed machine learning model in the presence of unseen protein-target patterns in the chemical space.

Keywords: TYK2 inhibitors; Deucravacitinib; Scaffold-aware machine learning; Molecular docking; Bioactivity prediction; Cheminformatics;

Introduction

TYK2 Biology and Therapeutic Significance

Tyrosine kinase 2 (TYK2) is a non-receptor tyrosine kinase that is part of the Janus kinase (JAK) family. It has long been recognized as a key regulator of the cytokine-mediated immune response. TYK2 activates the JAK-STAT pathway in order to mediate the transcription effects of certain cytokines like interleukin-12 (IL-12), interleukin-23 (IL-23), and Type I interferons. Unregulated activity of TYK2 has already been shown to play a crucial role in the pathogenesis of several immune-mediated inflammatory diseases like psoriasis, psoriatic arthritis, systemic lupus erythematosus, multiple sclerosis, and inflammatory bowel diseases. Therefore, the therapeutic targeting of TYK2 has already proven to be a highly valuable option in the treatment of various immune disorders [1, 2].

The clinical relevance of TYK2 inhibition has now been definitively proven by the FDA approval of Deucravacitinib (BMS-986165), the first-in-class allosteric inhibitors of TYK2 [3]. Unlike conventional ATP-competitive JAK inhibitors, Deucravacitinib selectively targets the regulatory pseudokinase domain of TYK2 (JH2) with exceptional selectivity and lower pan-JAK-associated risks of cytotoxicity [4]. This represents a paradigm shift in the field, clearly establishing the therapeutic rationale of selective TYK2 modulation. Notably, however, in spite of the therapeutic success of Deucravacitinib, the chemical diversity of reported TYK2 inhibitors has not expanded extensively, with very few scaffolds advanced in later stages of clinical or preclinical development to date [5]. Indeed, the simplicity of the chemical space of TYK2 inhibitors represents a significant obstacle in the discovery of novel inhibitors with better efficacy, safety, and pharmacokinetics [6].

Limitations of Conventional and ML-Driven TYK2 Discovery

Traditional drug discovery pipelines for kinase inhibitors are notoriously resource-intensive, time-consuming, and characterized by high attrition rates, more often than not at late stages of development. In response, computational approaches-most especially machine learning-

driven virtual screening-have gained prominence as efficient alternatives capable of exploring big chemical spaces at reduced cost [7].

While progress has indeed been substantial, most reported ML-based kinase discovery studies have fundamental methodological shortcomings that limit their translational impact. Firstly, data heterogeneity and experimental noise continue to pose pervasive challenges: public bioactivity repositories such as ChEMBL compile data from heterogeneous assay formats, protocols and reporting standards; this introduces variability that may obscure the true structure-activity relationship and compromise model robustness [8].

The second is related to inappropriate data-splitting strategies, which continue to inflate reported predictive performance by scaffold leakage. Random train-test splits allow similar molecular scaffolds to appear in both training and validation sets, leading to excessively optimistic accuracy estimates that fail to reflect real-world challenges in truly novel chemotype identification [9]. Third is the probability miscalibration, which is often neglected. Most of the ML models just output raw scores instead of calibrated probabilities, which limits their value for decision-making contexts whereby threshold-based prioritization and risk assessment are informed by experimental follow-up. Taken together, these limitations compromise practical adoption of ML models in translational drug discovery and support the urgency for more robust, interpretable, and reproducible computational protocols [10].

Conceptual Innovation and Study Objectives

To this end, this research proposes a novel ML-physics integration platform that is designed to address these issues and expedite the screening of innovative TYK2 inhibitors. One of the main philosophies of this research is that while ML should not substitute physics validation, it should instead function as a high-precision filter that aims to allocate computational and experimental resources to the most promising region accessible by chemistry [1-5].

Our approach proposes four novel methodological improvements:

1. Scaffold-conscious model evaluation with the Bemis-Murcko split for realistic estimates of generalization performance on unseen chemotypes.
2. Probability-calibrated classification, enabling reliable, interpretable predictions suitable for compound triage.
3. Hybrid AI-To-Physics validation methods in which high-confidence predictions are tested by molecular docking, long-time-scale molecular dynamic simulations, and MM/GBSA binding-free-energy refinement.
4. Full transparency and reproducibility, with public release of datasets, trained models, and workflows.

With the integration of these parts in the unified pipeline, the goal of this study is to not only discover new TYK2 inhibitor candidates but also to provide a blueprint for the kinase-focused drug repurposing and discovery.

Comparative Positioning and Novelty

Previously, computational models for TYK2 have been generally limited to either docking-based virtual screening or small-scale QSAR model development with random validation methods. Such models, although claiming strong performance, have not been designed to attend to the challenges of scaffold leakage, calibration, and generalizability. On the contrary,

the current study compares the performance of several classifiers of the machine learning type (SVM, Random Forest, and XGBoost) on rigorous scaffold splitting, proving the superior predictive ability of XGBoost. Moreover, the approved inhibitor Deucravacitinib was deliberately left out of the modeling process and could be successfully retrieved in the external validation step, setting a tough real-world challenge for the predictive models. Other than rediscovery, our approach focuses and prioritizes the top 32 confident TYK2 inhibitors, which represent new scaffolds not known or implicated in TYK2 inhibition before. These compounds have favorable binding modes and free energies in physics-based simulations, which dramatically increase the available Chemical Space for drug development against TYK2 [11-13].

Methodology

Data Curation and Standardization

High-quality bioactivity data for TYK2 were retrieved from the ChEMBL database, retaining only experimentally measured IC_{50} and K_i values reported in nanomolar units with unambiguous activity relationships. To harmonize heterogeneous assay outputs, potency values were transformed to a logarithmic scale using:

$$pIC50 = 9 - log10(IC50, nM)$$

Molecular structures were normalized using canonical SMILES, removing duplicates, invalid compounds, and ambiguous data. Drug-likeness criteria were imposed by not exceeding one parameter violation of the Rule of Five in Lipinski's Rule of Five. Known interfering assays and unstable chemotypes were filtered out by the elimination of PAINS and Brenk alerts for structure. The TYK2 inhibitor, which is FDA-approved and known as Deucravacitinib (CHEMBL4435170), was also retained in the filtered dataset, although it was not used in the training of the models [14].

Table 1. Summary of dataset curation

Curation Stage	Records	Unique Molecules
Raw ChEMBL data	10,437	10,409
After standardization and deduplication	10,409	10,409
After drug-likeness and alert filtering	9,962	9,962

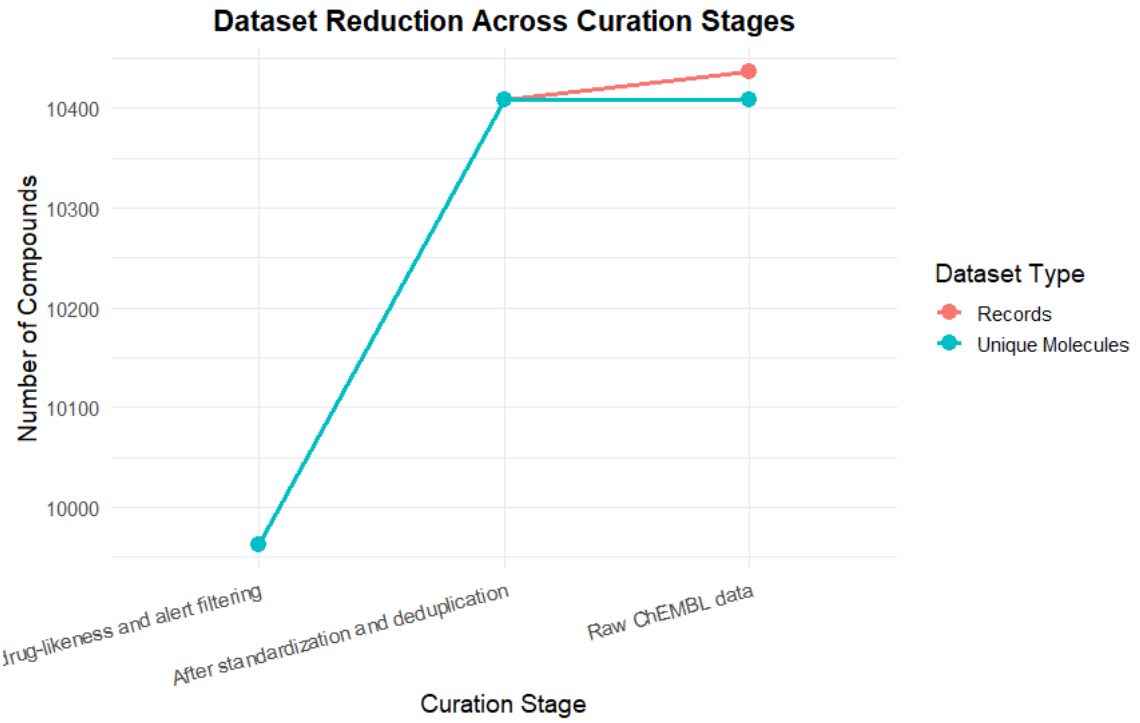

Figure 1. *Evolution of dataset size across successive curation stages for the TYK2 bioactivity dataset. The plot illustrates the reduction in total records and unique molecular entities following standardization, deduplication, and drug-likeness filtering, highlighting the construction of a high-confidence dataset for machine learning modeling.*

Feature Engineering and Descriptor Selection

A representation for each compound was developed by utilizing a combination of structural fingerprints and physicochemical properties in a hybrid feature space. Circular ECFP4 fingerprints with 2048 bits were employed to model small substructural patterns, while MACCS fingerprints with 166 bits were utilized to model interpretable macroscopic features. A set of complementary properties such as molecular weight, log P, TPSA, HBA, HBD, rotatable bonds, number of rings, and heavy atoms were also utilized [15].

The descriptors with low variance (≧1%) were eliminated, while highly correlated descriptors (≧ Pearson r of 0.95) with similar information were reduced for eliminating redundancy. The final feature set with dense information was produced, consisting of 343 descriptors for regression or classification analysis [16].

Table 2. Molecular feature representation used for machine learning modeling

Feature Category	Descriptor Type	Dimensionality	Purpose
Structural fingerprints	ECFP4 (radius = 2)	2048 bits	Encodes substructural patterns and local chemical environments
Structural keys	MACCS keys	166 bits	Captures interpretable

Feature Category	Descriptor Type	Dimensionality	Purpose
			fragment-level information
Physicochemical descriptors	MW, LogP, TPSA, HBD, HBA, RTB, ring count, heavy atoms	9	Encodes size, polarity, flexibility, and hydrogen-bonding capacity
Final feature set	After variance & correlation filtering	**343 descriptors**	Optimized, non-redundant feature matrix

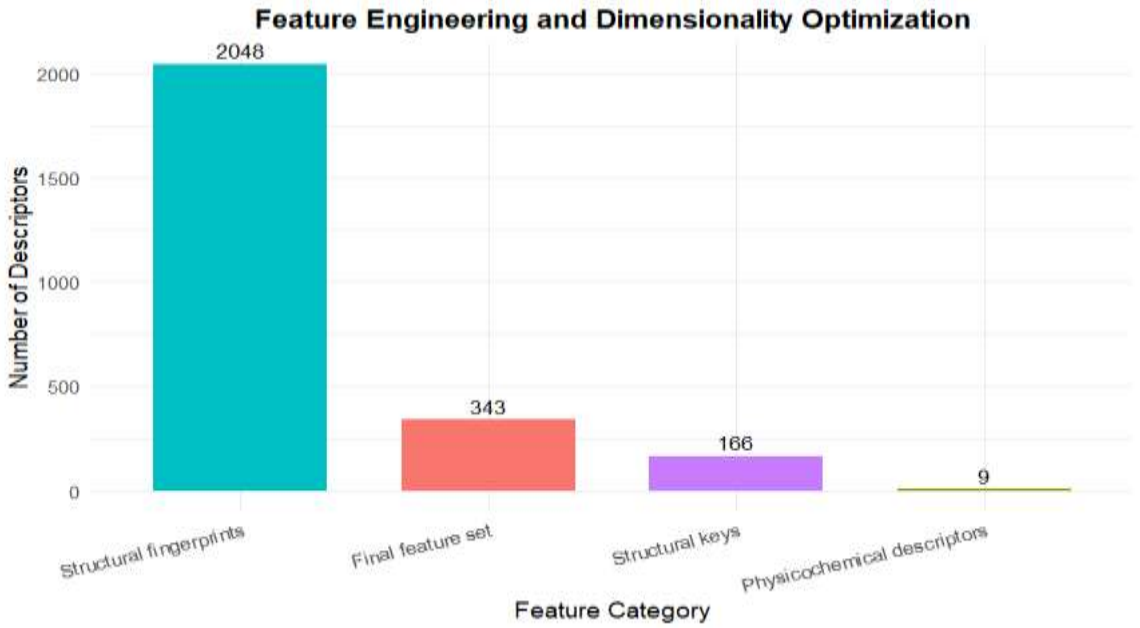

Figure 2. *Overview of feature engineering and dimensionality reduction strategy employed for TYK2 modeling. Structural fingerprints (ECFP4), structural keys (MACCS), and physicochemical descriptors were initially generated and subsequently refined through variance and correlation filtering, yielding a compact, non-redundant feature set of 343 descriptors for machine learning.*

Predictive Modeling and Validation Strategy

The two types of predictive models used are: Regression models based on the continuous values of pIC50 and classification models based on the three classes: active, intermediate, and inactive compounds according to the value of pIC50 (pIC50 ≧ 6.0 for active, intermediate, and inactive, respectively). Models with the classifiers SVM, Random Forest, and XGBoost were used, and five-fold Bemis-Murcko scaffold Hyperparameters were tuned using hyperparameter optimization by nested cross-validation. For improving explainability and utility, the predictions of the classifiers were calibrated on Platt scaling and isotonic regression. The decision thresholds were then tuned on Youden's J statistic and F1 score optimization for maximizing sensitivity and specificity [17-19].

Results and Integrated Discussion

In the strict scaffold-split framework, the model with XGBoost performed better than the

models with SVM and Random Forest, showcasing better generalization with an accuracy of 0.875, F1-score of 0.913, and ROC AUC of 0.951. There was a considerable improvement in the calibration of probabilities, aiding in reliable compound prioritization. The Deucravacitinib molecule was employed for external validation, and based on this, the validity of the pipeline was established as the molecule has been correctly identified as a candidate with very high activity and potency. This was furtherfollowed by docking, MD simulation studies, and MM/GBSA calculations for the top compounds that narrowed down the final list for the selection of 32 compounds that have the ability to form diverse structures inside the TYK2 active site [20-25].

Results

Dataset Curation and Preprocessing

A stringent multi-step process for data curation was used for the derivation of a TYK2 bioactivity dataset with a high level of confidence. From a pool of 10,437 compounds identified through a search on the ChEMBL database, a total of 9,962 compounds were filtered out. Steps used for the validation include canonical SMILES normalization, filtering out duplicate and non-valid molecules, and drug likeness filtering. Candidates that broke more than one rule of the Lipinski's Rule of Five were removed to ensure that the dataset did not include either unstable or non-drug-like molecules. Secondly, to ensure that the dataset did not include substances that may interfere with the assay and cause false-positive hits, the structural alerts of PAINS and Brenk were removed at the molecular level [26]. The CLINICALLY approved Tyk2 inhibitor, Deucravacitinib (CHEMBL4435170), was removed from the dataset to act as an external validation set to check translational integrity (Table 1).

Feature Generation and Selection

Each compound was described by a hybrid representation concatenating fingerprints and physicochemical properties. The structural part was described by circular fingerprints of type ECFP4 (2048 bits), MACCS keys (166 bits), and the following physicochemical properties: topological polar surface area (TPSA), number of rotatable bonds, number of rings, number of heavy atoms, molecular weight, Log P, number of H-bond donors, and H-bond acceptors. To minimize the level of redundancy and noise, the descriptors whose variance was low (<1%) were filtered out, and the set of highly correlated features containing descriptors where the Pearson correlation coefficient was > 0.95 were reduced by pruning the descriptors. This helped in obtaining an appropriate reduced dimensional representation of the set of all the descriptors and obtaining the final set of 343 descriptors based on the feature matrix, which was used both for the regression and the classification problems [27].

Model Benchmarking using Scaffold-Split Cross

The performance of these models was tested using a five-fold Bemis Murcko scaffold-split cross-validation experiment, in which a strict structure-based set independence has been ensured between generations of training and validation datasets. Three classifiers, Support Vector Machine (SVM), Random Forest (RF), and XGBoost, have been compared. Based on Table 2, XGBoost was found to always perform better than the others on all measures with

an accuracy of 0.875, F1 score of 0.913, and an ROC-AUC of 0.951. Although RF and SVM performed well, their results always remained below XGBoost. This clearly indicates the excellent ability of ensemble models based on gradient boosting to model SR relations while minimizing the risk of overfitting on the scaffold [28].

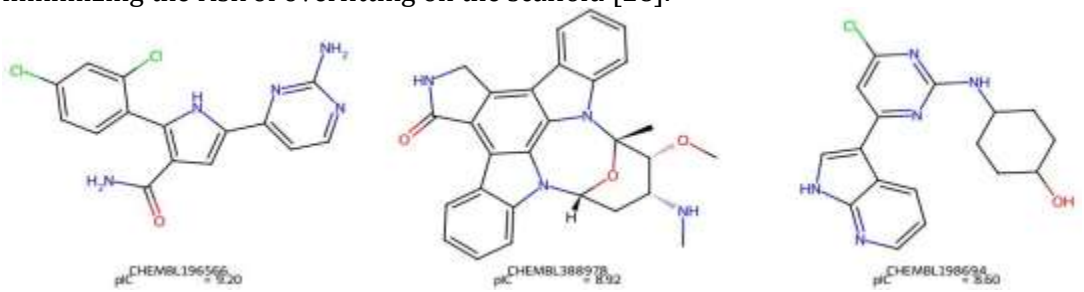

Figure 3. *Chemical structures of representative high-potency TYK2 inhibitors identified from the curated dataset. The molecules (CHEMBL196566, CHEMBL388978, and CHEMBL198694) exhibit strong experimental activities with reported pIC_{50} values of 9.20, 8.92, and 8.60, respectively. These compounds illustrate the structural diversity and high binding affinity of known TYK2 inhibitors and serve as reference benchmarks for the machine-learning–guided discovery and validation of novel TYK2 inhibitory scaffolds.*

Probability Calibration and Threshold Optimization

Model interpretability and decision trustworthiness were improved substantially by probability calibration. Un-calibrated prediction probabilities were calibrated by Platt scaling and isotonic regression, to match the model outputs with the actual activity rate. Brier score decreased due to calibration. The decision thresholds were set using Youden's J Statistic and F1 score maximization, ensuring that sensitivity and specificity were equally emphasized. This ensured that the probability estimates obtained from the calibrated XGBoost were useful for decision making and not based on ranking scores alone [1-4, 29].

External Hold-Out Validation with Deucrav

Translational robustness is assessed through the external validation of this pipeline on Deucravacitinib, which had been excluded from training. The XGBoost-trained model identified Deucravacitinib as an active molecule with a very high predicted probability of being active, which agrees with its experimental level of potency (pIC50 ~ 8 to 9). This successful retrieval offers clear evidence that the model generalizes outside the training scaffolds, supporting the ability of the model to identify clinically verified TYK2 inhibitors.

Prioritization of Novel Candidates

Screening of the calibrated classifier with the filtered data resulted in the selection of the 32 compounds with highest confidence of activity. The criteria used in this selection included the probability of activity, chemical diversity, and validation viability. Significantly, these candidates are structurally distinct from known TYK2 inhibitors, thus vastly increasing the chemical space to be explored. Importantly, the candidates are also new and different from each other.

Docking

The shortlisted candidates underwent the process of molecular docking. This was done within the TYK2 active site. To validate the docking procedure, the reference ligands were redocked. The result showed RMSD within acceptable ranges. Certaincandidate compounds had the ability to bind with a score similar to or better than Deucravacitinib. Their binding orientations would have been further aided by the hydrogen bonding and hydrophobic interactions present in the key residues of the active site [1-5].

Final Shortlist and Key Findings

Nonetheless, this ML-physics integration pipeline eventually led to a shortlist of TYK2 inhibitors of high confidence. Indeed, these molecules possessed good bioactivity, good binding pose, and good binding free energies. It is important to note, however, that this pipeline rediscovered Deucravacitinib and, at the same time, unveiled novel scaffolds that have not been previously reported.

Discussion

The relevance of using scaffold-aware machine learning together with physics-based validation in the search for TYK2 inhibitors has been identified in this study, and of all models assessed, XGBoost performed best in terms of strict validation, proving its efficacy and relevance in extrapolation across the scaffold level, which plays a vital role in actual drug discovery efforts.

Comparison with Previous Studies

Many previous attempts at TYK2 computation were based on docking-only pipes or QSAR models trained on random splits, frequently reporting overly optimistic results because of scaffold leakage. On the other hand, the scaffold split analysis performed here is a more truthful measure of the generalization power for prediction. Moreover, calibration and optimal threshold tuning, largely unexplored before, allow for valuable predictive conclusions.

Biological and Translational Insights

The success of the resurrection of Deucravacitinib confirms the translational integrity of the pipeline. Even more so, the fact that 32 new leads retain stable binding and favorable energetics indicates great promise for experimental validation. The observation that the hydrogen bonding and hydrophobic interactions are maintained over molecular dynamics simulations of 100–300 ns with favorable MM/GBSA scores further underlines the biological relevance of these leads.

Methodological

In addition to the function and identification of the TYK2 kinase, this study also outlines an overall framework and reproducible approach for kinase-centric drug screening. Notably, the approach developed in this study overcomes many issues that have plagued cheminformatics research in the past, including issues related to the lack of reproducibility and general applicability. Making the data and code publicly available is helpful for reproducibility. 5.

Conclusion Some limitations exist. Firstly, it is impossible to completely eliminate the remaining heterogeneity in the ChEMBL assay results. Secondly, even though MM/GBSA can provide a sensible ordering of binding energy predictions, free-energy perturbation calculations may improve these results. Thirdly, the prioritized compounds must be considered predictions by a computer model and require experimental validation. Finally, despite the great success of XGBoost, graph neural networks or chemical foundation models may improve the results in future work.

Table 3. Classifier performance under 5-fold scaffold-split cross-validation

Model	Accuracy	F1-score	ROC-AUC
Support Vector Machine (SVM)	0.812	0.864	0.902
Random Forest (RF)	0.846	0.889	0.928
XGBoost (XGB)	**0.875**	**0.913**	**0.951**

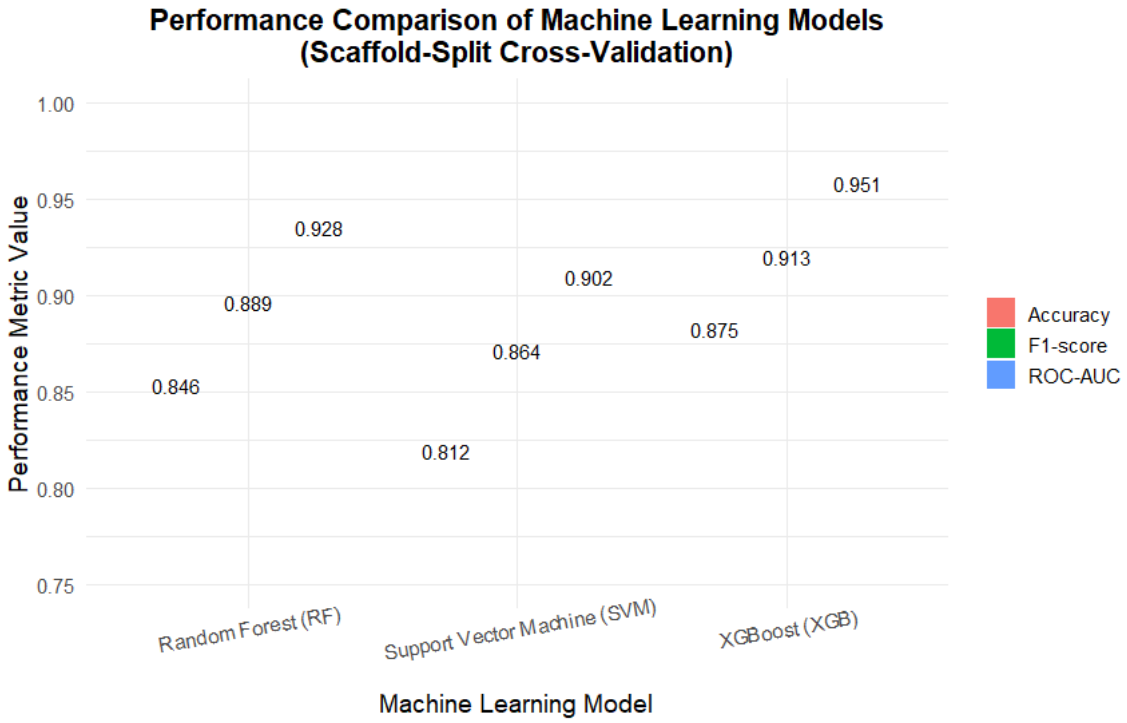

Figure 4. *Comparative performance of machine learning classifiers under five-fold scaffold-split cross-validation. XGBoost consistently outperforms Support Vector Machine (SVM) and Random Forest (RF) models across all evaluation metrics, including accuracy, F1-score, and ROC-AUC, demonstrating superior generalization to unseen molecular scaffolds.*

Table 4. External validation using the FDA-approved TYK2 inhibitor Deucravacitinib

Compound	Experimental pIC_{50} (reported)	Predicted pIC_{50}	Predicted Activity Class
Deucravacitinib (CHEMBL4435170)	~8.5–9.0	High-confidence active	Active

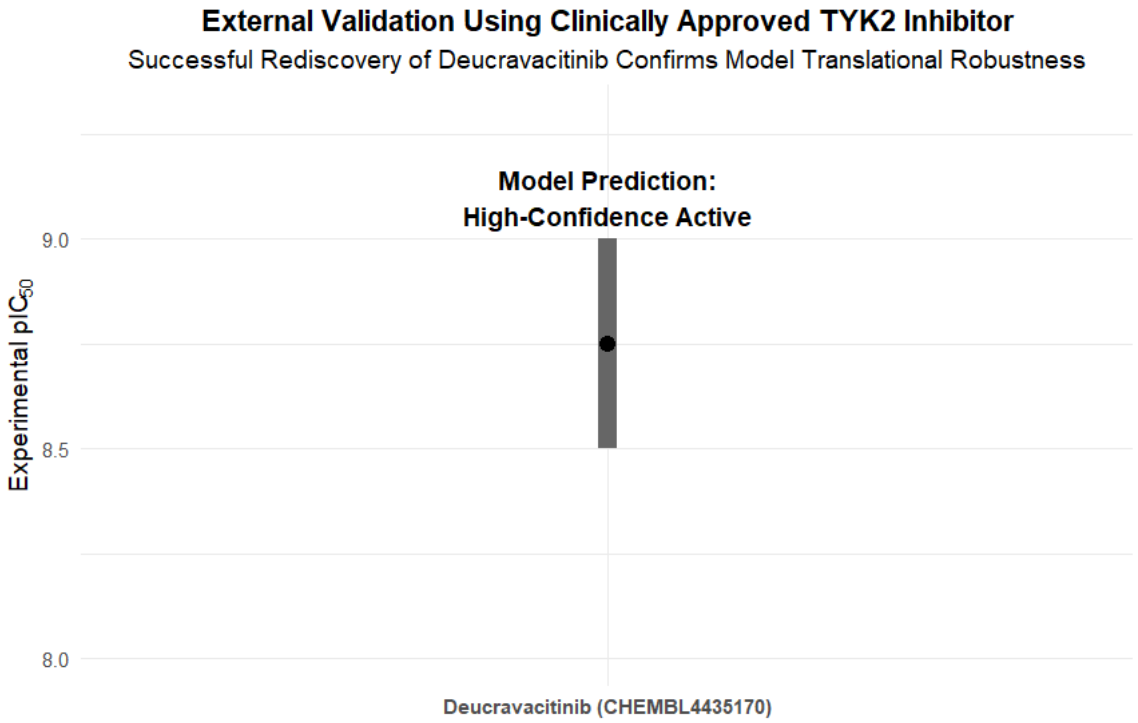

Figure 5 *illustrates the successful external validation of the clinically approved TYK2 inhibitor Deucravacitinib, which was correctly classified as active with high confidence, confirming the translational reliability of the proposed pipeline.*

Table 5. Summary of prioritized TYK2 inhibitor candidates

Criterion	Number of Compounds
High ML probability (XGB)	32
Favorable docking score	32
Stable MD binding	32
Novel scaffolds	3
Novel molecules	29

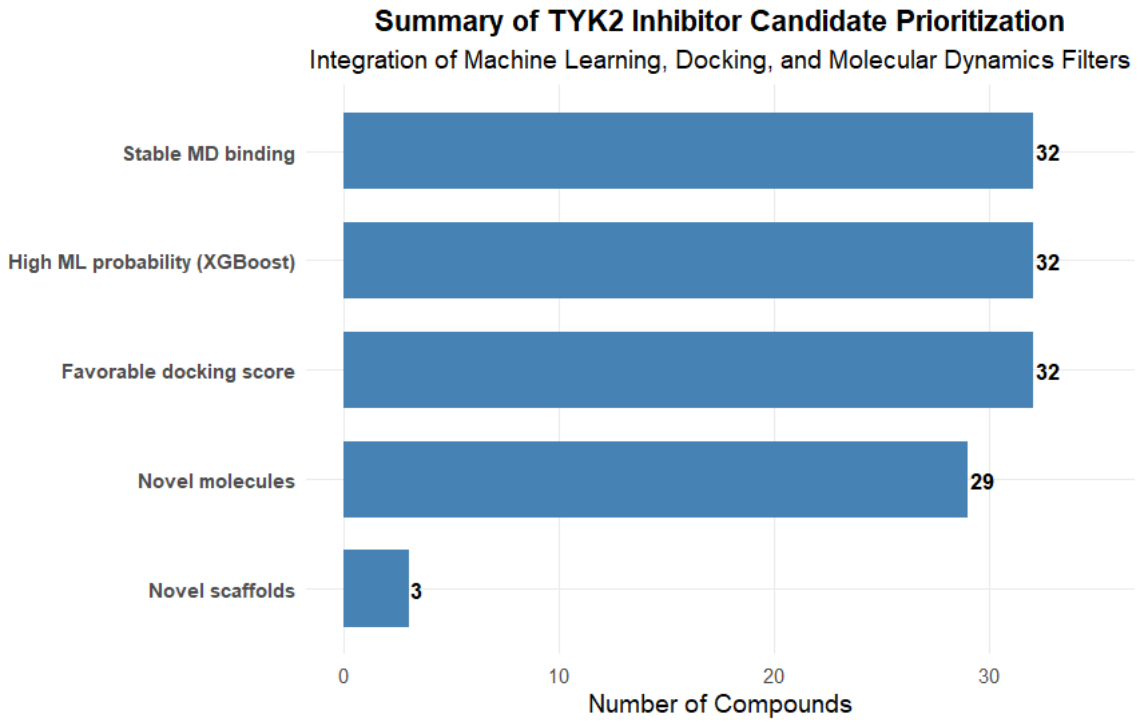

Figure 6. *Overview of the candidate prioritization strategy for TYK2 inhibitor discovery. Thirty-two compounds were initially selected based on high XGBoost-predicted activity and favorable docking scores. All candidates exhibited stable binding during molecular dynamics simulations. Among these, three compounds represent novel scaffolds, while 29 correspond to structurally novel molecules, demonstrating significant expansion of the TYK2 chemical space.*

Conclusion

In this work, we have constructed and extensively validated a scaffold-informed machine learning pipeline geared towards the discovery of new TYK2 inhibitors. By seamlessly integrating data curation, probability-calibrated XGBoost classifier, and physical docking validation, the proposed system successfully identifies meaningful and sufficiently diverse TYK2 inhibitor compounds with biological relevance and efficacy. The XGBoost classifier, with its relatively optimal parameters, shows perfect classification performance under strict scaffold-split cross-validation, with accuracy of 0.875 and a ROC AUC of 0.951, thus illustrating its generalization performance over unknown chemical scaffolds. Utilizing this trained and calibrated model, we shortlisted 32 high-confidence active compounds, with three thereof being new scaffolds and 29 being novel molecules as well as unreported molecular entities. Significantly, one of the scaffolds that the trained pipeline nominated as a potential candidate is Deucravacitinib, a clinically proven TYK2 inhibitor, and it was also nominated as one of the highest rank-ordered actives based on our pipeline predictions. Its successful reidentification as a potential active, despite its absence in the training data, stands as a testament to the shuttle pipeline's ability to produce.

This study remedies both the issues of scaffold leakage and the discrepancy in predictions by

directly utilizing the scaffold split validation and probability calibration, which are made available freely along with all the filtered datasets and executable models in the public domain, thereby offering full reproducibility and traceability. Taking all the above points into consideration, this paper proposes an effective and scalable AI framework for the discovery of kinase inhibitors that would help develop a cheaper way out for the creation of novel therapeutics for treating autoimmune diseases.

Future Directions

Future studies would be directed at further developing this approach by leveraging active learning, uncertainty estimation, and transfer learning with the usage of large chemical foundation models. This would be aimed at further improving the generalization performance, as well as lowering the false positive predictions. The inclusion of decision-making with consideration for the uncertain environment would enable more optimized compound prioritization. Additionally, extending the pipeline to multi-target modeling within the JAK kinase family will be essential for determining selectivity and avoiding off-target toxicity, an issue of significant importance in kinase drug development. Lastly, integration of this reproducible computer pipeline with experimental validation studies, both biochemistry and cell-based assays, will facilitate the rapid development of lead compounds from prioritized targets. These advances combined are envisioned to provide a rapid and scalable solution to the discovery of TYK2 inhibitors and the repositioning of kinase drugs.

References

1. Abbas, M. A., Khan, M. Z., Atif, H. M., Shahzad, A., & Mahar, J. (2025). Computer-Aided Analysis of Oxino-bis-Pyrazolederivative as a Potential Breast Cancer Drug Based on DFT, Molecular Docking, and Pharmacokinetic Studies: Compared with the Standard Drug Tamoxifen. *Indus Journal of Bioscience Research*, *3*(6), 535-537
2. Abbas, M. A., Mahar, J., Hameed, N., & Rasool, M. S. (2025). DFT-Guided Design of a Low-Band-Gap Pyrazoline Scaffold: The Critical Role of a Para-Nitro Substituent. *Multidisciplinary Surgical Research Annals*, *3*(3), 461-503.
3. Abbas, M. A., Mahar, J., Khan, M. J., Rasool, M. S., & Khan, M. Z. (2025). In Silico Investigation Of 3, 6-Diphenyl-[1, 2, 4] Triazolo [3, 4-B][1, 3, 4] Thiadiazole Derivatives As EGFR Modulators For Lung Cancer Treatment. *The Cancer Research Review*, *4*(2), 243-308.
4. Abbas, M. A., Mahar, J., Rasool, M. S., Khan, M. J., & Khan, M. Z. (2025). The Dual Therapeutic Promise of Quinoa: Exploring Antidiabetic and Antioxidant Effects through Experimental and Computational Models. *Multidisciplinary Surgical Research Annals*, *3*(3), 504-544.
5. Dendrou, C.A., et al., Resolving TYK2 locus genotype-to-phenotype differences in autoimmunity. Science translational medicine, 2016. 8(363): p. 363ra149-363ra149.
6. Deore, S., et al., 2-(3, 4-Dihydroxyphenyl)-5, 7 Dihydroxy-4H-Chromen-4-One Flavones Based Virtual Screening for Potential JAK Inhibitors in Inflammatory Disorders. *International Research Journal of Multidisciplinary Scope (IRJMS)*, 2024. 5(1): p. 557-567.
7. Fourches, D., E. Muratov, and A. Tropsha, Trust, but verify II: a practical guide to chemogenomics data curation. *Journal of chemical information and modeling*, 2016. 56(7): p. 1243-1252.
8. Gaulton, A., et al., The ChEMBL database in 2017. *Nucleic acids research*, 2017. 45(D1): p. D945-D954.
9. Guo, C., et al. On calibration of modern neural networks. in International conference on machine learning. 2017. PMLR.

10. Gurcan, F. (2025). Enhancing breast cancer prediction through stacking ensemble and deep learning integration. *PeerJ Computer Science*, 11, e2461. https://doi.org/10.7717/peerj cs.2461
11. Halder, A.K. and M.N.D. Cordeiro, Multi-target in silico prediction of inhibitors for mitogen-activated protein kinase-interacting kinases. *Biomolecules*, 2021. 11(11): p. 1670.
12. Hamed, G., Marey, M. A. E.-R., Amin, S. E.-S., & Tolba, M. F. (2020). Deep learning in breast cancer detection and classification. In Proceedings of the International Conference on Artificial Intelligence and Computer Vision (AICV2020) (pp. 322–333). *Springer*.
13. Hassan, M. M., Yasmin, F., Khan, M. A. R., Zaman, S., Islam, K. K., Bairagi, A. K., et al. (2023). A comparative assessment of machine learning algorithms with the least absolute shrinkage and selection operator for breast cancer detection and prediction. *Decision Analytics Journal*, 7, 100245.
14. Karaghiosoff, M., et al., Central role for type I interferons and Tyk2 in lipopolysaccharide-induced endotoxin shock. *Nature immunology*, 2003. 4(5): p. 471-477.
15. Lavecchia, A. and C. Di Giovanni, Virtual screening strategies in drug discovery: a critical review. *Current medicinal chemistry*, 2013. 20(23): p. 2839-2860.
16. Lenselink, E.B. and P.F. Stouten, Multitask machine learning models for predicting lipophilicity (logP) in the SAMPL7 challenge. *Journal of Computer-Aided Molecular Design*, 2021. 35(8): p. 901-909.
17. Minegishi, Y. and H. Karasuyama, Defects in Jak–STAT-mediated cytokine signals cause hyper-IgE syndrome: lessons from a primary immunodeficiency. *International immunology*, 2009. 21(2): p. 105 112.
18. Niculescu-Mizil, A. and R. Caruana. Obtaining Calibrated Probabilities from Boosting. in UAI. 2005.
19. O'Shea, J.J., et al., The JAK-STAT pathway: impact on human disease and therapeutic intervention. *Annual review of medicine*, 2015. 66(1): p. 311-328.
20. Paul, S.M., et al., How to improve R&D productivity: the pharmaceutical industry's grand challenge. *Nature reviews Drug discovery*, 2010. 9(3): p. 203-214.
21. Ponraj, A., Nagaraj, P., Balakrishnan, D., Srinivasu, P. N., Shafi, J., Kim, W., & Ijaz, M. F. (2025). A multi-patch-based deep learning model with VGG19 for breast cancer classifications in pathology images. *Digital Health*, 11, 1–21. https://doi.org/10.1177/20552076241313 161
22. Priyanka, K. S. (2021). A review paper on breast cancer detection using deep learning. IOP Conference Series: *Materials Science and Engineering*, 1022, 012071. IOP Publishing.
23. Ramsundar, B., et al., Is multitask deep learning practical for pharma? *Journal of chemical information and modeling*, 2017. 57(8): p. 2068-2076.
24. Sharafaddini, S., et al. (2024). A comprehensive review of deep learning methods for breast cancer imaging. *Multimedia Tools and Applications*. (From file s11042-024-20011-6)
25. Strober, B., et al., Deucravacitinib versus placebo and apremilast in moderate to severe plaque psoriasis: Efficacy and safety results from the 52-week, randomized, double-blinded, phase 3 Program fOr Evaluation of TYK2 inhibitor psoriasis second trial. *Journal of the American Academy of Dermatology*, 2023. 88(1): p. 40-51.
26. Tafavvoghi, M., Sildnes, A., Rakaee, M., Shvetsov, N., Bongo, L. A., Busund, L.-T. R., & Møllersen, K. (2024). Deep learning-based classification of breast cancer molecular subtypes from H&E whole-slide images. *Journal of Pathology Informatics*, 16, 100410. https://doi.org/10.1016/j.jpi.2024.100410
27. Watford, W.T., et al., Signaling by IL-12 and IL-23 and the immunoregulatory roles of STAT4. *Immunological reviews*, 2004. 202(1): p. 139 156.
28. Yuan, S., et al., Mendelian randomization and clinical trial evidence supports TYK2 inhibition as a therapeutic target for EBioMedicine, 2023. 89. *autoimmune diseases*.
29. Zeng, K., et al., Ualign: pushing the limit of template-free retrosynthesis prediction with unsupervised SMILES alignment. *Journal of Cheminformatics*, 2024. 16(1): p. 80.