International Journal of Engineering Research and Science & Technology

www.ijerst.org

ISSN : 2319-5991

Vol. 21 No. 3 (1) 2025

ijerst.editor@gmail.com

editor@ijerst.com

International Journal of Engineering Research and Science & Technology

ISSN 2319-5991 www.ijerst.com
Vol. 21, No. 3(1), 2025

Research Paper

INTELLIGENT DRUG SAFETY MONITORING WITH INTERPRETABLE MACHINE LEARNING MODELS

D. Deepthi[1], G. vaishnavi[2], P. Alekhya[2], B. Mohan[2], K. Anjaneyulu[2]
[1]Assistant Professor, [2]UG Student, [1,2]Department of Computer Science and Engineering (AI&ML)
[1,2]Sree Dattha Institute of Engineering and Science, Ibrahimpatnam, 501510, Telangana.

ABSTRACT

Adverse drug reactions pose significant risks in clinical settings, particularly when side effects are overlooked during the early stages of prescription. To mitigate these risks, this study focuses on enhancing drug side effect prediction using machine learning techniques integrated with Explainable AI (XAI) for medical applications. The primary objective is to develop an intelligent, interpretable system that not only predicts potential side effects but also provides transparency in the decision-making process, thereby fostering trust among healthcare professionals. A comprehensive dataset comprising drug attributes, side effect profiles, and associated clinical features was used for model training and evaluation. Initial experimentation involved various baseline classifiers, including Ridge Classifier, Linear Support Vector Machine (SVM), Logistic Regression, and Multinomial Naïve Bayes. These models served as benchmarks and were evaluated based on accuracy, precision, recall, and F1-score. Extensive Exploratory Data Analysis (EDA) was performed to uncover patterns, correlations, and class imbalances within the dataset, aiding in informed feature selection and preprocessing. To improve prediction accuracy and enable complex pattern recognition, a Multi-Layer Perceptron (MLP) Classifier was proposed as the advanced model. As a deep learning algorithm, the MLP demonstrated superior performance in capturing nonlinear relationships among features—capabilities that traditional models often lack. This makes the solution suitable for real-world deployment in Clinical Decision Support Systems (CDSS), ensuring safer drug administration and better patient outcomes. The project contributes to the field of medical AI by delivering a high-performing, interpretable, and reliable solution for drug side effect prediction, effectively bridging the gap between complex AI models and practical healthcare applications.

Keywords: Adverse drug reactions, Drug side effect prediction, Machine learning, Explainable AI (XAI), Clinical Decision Support Systems (CDSS).

Received: 14-7-2025 Accepted: 21-8-2025 Published: 28-8-2025

1. INTRODUCTION

Drug-related side effects encompass a range of undesirable, unpleasant, and potentially hazardous reactions that can affect various organs and tissues. Even after gaining market approval, some drugs may cause unacceptable adverse effects, significant health risks and raising concerns within the pharmaceutical industry. Ensuring drug efficacy and safety is crucial, as adverse reactions remain a leading cause of drug failure and withdrawal. Traditional methods for identifying side effects, particularly through clinical trials, are often time-consuming, costly, and impractical for large-scale testing. This highlights the urgent need for more efficient and scalable predictive methods capable of identifying potential side effects early in the drug development lifecycle.

The ability to accurately predict drug-related side effects has become an essential component of modern pharmaceutical research and development. These predictive methodologies have the potential to transform the drug discovery process by allowing early detection of adverse reactions, thus saving time, reducing costs, and improving safety. Such approaches enable the prioritization of safer drug candidates while eliminating those with higher risks of adverse effects. Consequently, robust

prediction models play a vital role in the development of safer medications and better patient outcomes, supporting the broader goal of advancing personalized medicine.

The motivation for this research stems from the increasing reliance of pharmaceutical companies and health technology providers on data-driven decision-making. Industry leaders like Pfizer, Johnson & Johnson, and Novartis utilize advanced analytics platforms to monitor clinical trial data and post-market patient feedback to identify side effects. Health tech firms such as IBM Watson Health and Cerner use AI-driven analytics to extract insights from electronic health records (EHRs), thereby improving both treatment outcomes and drug safety. Similarly, startups like Tempus and PathAI are leveraging machine learning to enhance diagnostic precision and predict drug responses. These innovations not only reduce trial-and-error in prescribing but also ensure regulatory compliance and proactive pharmacovigilance.

Fig. 1: Drugs and its side effects.

In clinical settings, decision support systems embedded in EHR platforms such as Epic and Meditech offer real-time alerts to physicians about potential drug interactions or contraindications. These intelligent systems improve the quality of care by supporting accurate diagnoses and more personalized treatment plans, ultimately minimizing adverse drug events. The increased dependence on such technologies underscores the growing role of machine learning and data analytics in ensuring safer and more efficient healthcare delivery.

Despite these advances, predicting drug side effects remains a complex challenge. Patient responses to medications are highly variable, influenced by factors like genetics, underlying conditions, co-medications, and dosage. Many adverse reactions are only discovered after large-scale use, often leading to delayed recalls and increased healthcare burdens. Moreover, the datasets involved are diverse, encompassing both structured data (e.g., lab results, demographics) and unstructured sources (e.g., clinical notes, patient-reported symptoms). This complexity makes traditional statistical methods inadequate, especially when dealing with imbalanced datasets where adverse outcomes are rare.

To address this, there is a pressing need for computational models that can integrate and analyze heterogeneous data sources effectively. These models must be capable of identifying subtle patterns and correlations while minimizing the risk of false positives and false negatives. Such advancements are crucial not only for ensuring patient safety but also for building models that are scalable, interpretable, and reliable in real-world clinical applications.

The significance of developing accurate drug side effect prediction models extends across public health, healthcare systems, and pharmaceutical innovation. Early identification of potential adverse reactions can prevent emergency interventions and reduce hospital admissions, thus lowering healthcare costs. For pharmaceutical companies, these models improve clinical trial designs and post-

market safety monitoring, enabling faster regulatory decisions. In personalized medicine, predictive models are instrumental in tailoring treatments to individual genetic and physiological profiles, enhancing therapeutic efficacy and patient satisfaction.

The primary objective of this research is to develop a machine learning-based framework for predicting drug side effects using a range of classification algorithms. The study evaluates several traditional classifiers including Logistic Regression, Linear Support Vector Classifier (Linear SVC), Ridge Classifier, Multinomial Naïve Bayes, and Stochastic Gradient Descent (SGD) Classifier. These are benchmarked for their performance in terms of accuracy, precision, recall, and F1-score. Additionally, the research proposes a Multilayer Perceptron (MLP) Classifier, a deep learning model known for its ability to capture complex, nonlinear relationships within data. By comparing these models, the study aims to determine the most effective algorithm for accurately identifying side effects based on drug and patient-related attributes.

This research provides several advantages to healthcare stakeholders. By enabling early detection of potential drug reactions through data-driven insights, the model reduces reliance on lengthy post-market surveillance. It helps cut healthcare costs by preventing avoidable hospitalizations and enhances the reliability of treatment plans by integrating predictive analytics into clinical workflows. Pharmaceutical companies benefit from more efficient drug development pipelines, while healthcare providers gain tools for selecting safer alternatives tailored to at-risk patients. Furthermore, the model supports patient-centered care, strengthens regulatory decision-making, and promotes interdisciplinary collaboration among clinicians, data scientists, and pharmacologists.

The applications of this system are broad and impactful. In clinical decision support systems, it can alert physicians to possible adverse reactions during prescription. When integrated into EHR platforms, it provides real-time monitoring and interaction checks. Pharmaceutical companies can use it to improve trial design and drug safety profiling, while mobile health apps can notify users about potential risks based on personal usage patterns. Public health agencies may use it to monitor drug-related incidents at scale, and personalized medicine platforms can tailor treatments accordingly. Regulatory bodies benefit through improved pharmacovigilance, and academic institutions can use the model for training in predictive healthcare analytics. Overall, this system has the potential to revolutionize how the medical community addresses drug safety.

2. LITERATURE SURVEY

Bartlett et. al [12] compares on real data effective duplicates detection methods for automatic deduplication of files based on names, working with French texts or English texts, and the names of people or places, in Africa or in the West. After conducting a more complete classification of semantic duplicates than the usual classifications, they introduce several methods for detecting duplicates whose average complexity observed is less than O(2n). Through a simple model, they highlight a global efficacy rate, combining precision and recall. We propose a new metric distance between records, as well as rules for automatic duplicate detection. Analyses made on a database containing real data for an administration in Central Africa, and on a known standard database containing names of restaurants in the USA, have shown better results than those of known methods, with a lesser complexity. Shimada et. al [13] developed a decision support system that helps doctors select appropriate first-line drugs. The system classifies patients' abilities to protect themselves from infectious diseases as a risk level for infection. In an evaluation of the prototype system, the risk level it determined correlated with the decisions of specialists. The system is very effective and convenient for doctors to use.

He et. al [14] presented a novel adaptive synthetic (ADASYN) sampling approach for learning from imbalanced data sets. The essential idea of ADASYN is to use a weighted distribution for different minority class examples according to their level of difficulty in learning, where more synthetic data is

generated for minority class examples that are harder to learn compared to those minority examples that are easier to learn.

Lei et. al [15] presented a novel approach to polarity classification of short text snippets, which takes into account the way data are naturally distributed into several topics in order to obtain better classification models for polarity. This approach is multi-step, where in the initial step a standard topic classifier is learned from the data and the topic labels, and in the ensuing step several polarity classifiers, one per topic, are learned from the data and the polarity labels. They empirically show that our approach improves classification accuracy over a real-world dataset by over 10%, when compared against a standard single-step approach using the same feature sets. The approach is applicable whenever training material is available for building both topic and polarity learning models. Nikfarjam and Gonzalez et. al [6] presented a new method for using association rules for colloquial text mining. They applied our method on user comments to find mentions of adverse reactions to drugs by extracting frequent patterns. Since we are dealing with highly informal colloquial text, the idea of using extracted patterns might, at first, seem counter-intuitive. However, we indeed found consistencies in the user comments. This evaluation measured the effectiveness of this technique in extracting frequent patterns in this context. However, this method can easily be generalized for other contexts and languages.

Doulaverakis et. al [7] presented a DR-SED system based on Semantic Web technologies, termed GalenOW. It has been shown that OWL and Semantic Web technologies can provide a good match for DR-SEDs as OWL is expressive enough to effectively encapsulate medical knowledge. Rule-based reasoning can model medical decision making and aid experts. A comparison of the semantic-enabled implementation to a traditional business logic implementation was presented. Although the latter has shown better performance in time and memory requirements, semantic technologies provide a better alternative for integrating knowledge in the system than simple rule engines.

Goeuriot et. al [8] presented creation of lexical resources and their adaptation to the medical domain. We first describe the creation of a general lexicon, containing opinion words from the general domain and their polarity. Then they presented the creation of a medical opinion lexicon, based on a corpus of drug reviews. They show that some words have a different polarity in the general domain and in the medical one. Some words considered generally as neutral are opinionated in medical texts. They finally evaluate the lexicons and show with a simple algorithm that using our general lexicon gives better results than other well-known ones on our corpus and that adding the domain lexicon improves them as well.

Keers et. al [9] appraised empirical evidence relating to the causes of medication administration errors (MAEs) in hospital settings. Limited evidence from studies included in this systematic review suggests that MAEs are influenced by multiple systems factors, but if and how these arise and interconnect to lead to errors remains to be fully determined. Further theoretical focused is needed to investigate the MAE causation pathway, with an emphasis on ensuring interventions designed to minimise MAEs target recognised underlying causes of errors to maximise their impact.

Wittich et. al [10] provides a practicing physicians that focuses on medication error terminology and definitions, incidence, risk factors, avoidance strategies, and disclosure and legal consequences. A medication error is any error that occurs at any point in the medication use process. It has been estimated by the Institute of Medicine that medication errors cause 1 of 131 outpatient and 1 of 854 inpatient deaths. Medication factors (eg, similar sounding names, low therapeutic index), patient factors (eg, poor renal or hepatic function, impaired cognition, polypharmacy), and health care professional factors (eg, use of abbreviations in prescriptions and other communications, cognitive biases) can precipitate medication errors.

Zhang et. al [1] proposed a novel cloud-assisted DR-SED (CADRE), which can recommend users with top-N related medicines according to symptoms. In CADRE, they first cluster the drugs into

several groups according to the functional description information, and design a basic personalized DR-SED based on user collaborative filtering. Then, considering the shortcomings of collaborative filtering algorithm, such as computing expensive, cold start, and data sparsity, they propose a cloud-assisted approach for enriching end-user Quality of Experience (QoE) of DR-SED, by modeling and representing the relationship of the user, symptom and medicine via tensor decomposition. Finally, the proposed approach is evaluated with experimental study based on a real dataset crawled from Internet.

Danushka et. al [2] proposed an unsupervised method for learning domain-specific word representations that accurately capture the domain-specific aspects of word semantics. First, we select a subset of frequent words that occur in both domains as \emph{pivots}. Next, they optimize an objective function that enforces two constraints: for both source and target domain documents, pivots that appear in a document must accurately predict the co-occurring non-pivots, and, word representations learnt for pivots must be similar in the two domains. Moreover, they propose a method to perform domain adaptation using the learnt word representations. This proposed method significantly outperforms competitive baselines including the state-of-the-art domain-insensitive word representations, and reports best sentiment classification accuracies for all domain-pairs in a benchmark dataset.

Sarker et. al [3] suggested that interest in the utilization of the vast amounts of available social media data for ADR monitoring is increasing. In terms of sources, both health-related and general social media data have been used for ADR detection—while health-related sources tend to contain higher proportions of relevant data, the volume of data from general social media websites is significantly higher.

3. PROPOSED SYSTEM

The proposed system introduces a hybrid drug side effect prediction framework enhanced with Explainable AI (XAI), aiming to address the limitations of traditional models in accuracy, personalization, and interpretability.

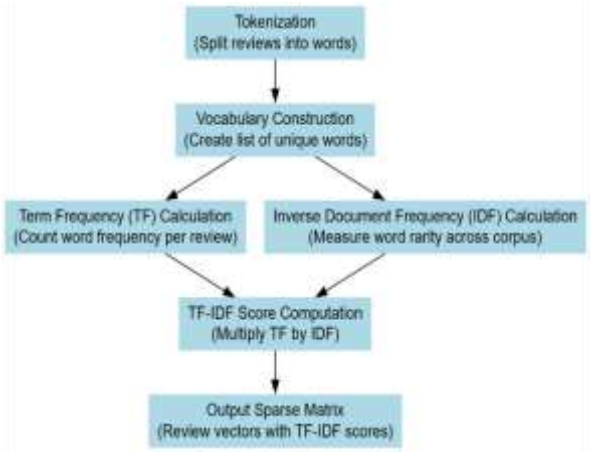

Fig.2: Internal operational workflow of TF-IDF vectorizer.

Unlike previous approaches that often rely on a single classifier or isolated preprocessing, this methodology integrates multiple machine learning classifiers, TF-IDF-based text vectorization, cosine similarity for personalized drug recommendations, and real-time explainability through web knowledge extraction using SerpAPI. The pipeline begins with uploading a dataset that includes drug name, condition, review, and rating. Initial exploration of the dataset helps visualize label distributions, which guides model choice and evaluation. Reviews undergo preprocessing through normalization, stopword removal, punctuation stripping, token length filtering, and lemmatization, transforming raw text into a clean and consistent format suitable for vectorization. Text features are then converted into numerical representations using a TF-IDF vectorizer, which emphasizes important

terms while reducing the influence of frequent but less meaningful words. With a vocabulary limited to the top 700 terms, this feature matrix is used to train multiple classifiers such as Logistic Regression, Linear SVC, Ridge Classifier, Multinomial Naive Bayes, SGDClassifier, and a Multilayer Perceptron (MLP). These classifiers are evaluated based on precision, recall, F1-score, and accuracy. To add interpretability and recommendation capabilities, a cosine similarity technique is used to find the most similar past review from the training set, recommending its associated drug for a new user input. This case-based reasoning approach aligns new data with historical examples. For added transparency, SerpAPI queries real-time side effect information from verified government health sources, making the recommendation process interpretable and credible.

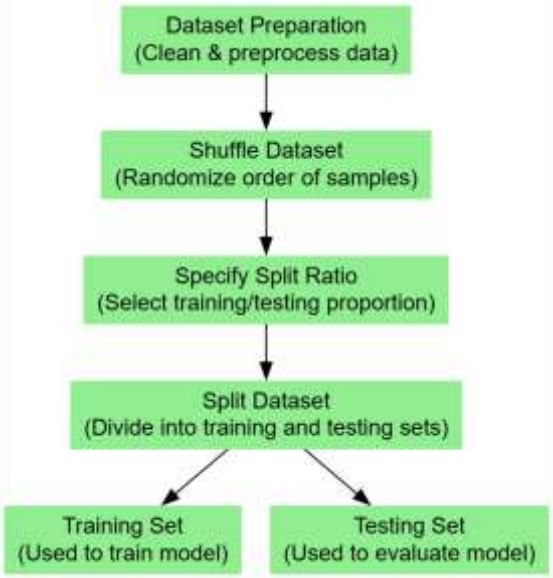

Fig.3: Internal operational workflow of train-test split.

The entire ML pipeline is supported by graphical visualizations that compare classifier performance, enabling better model selection. TF-IDF vectorization plays a crucial role in transforming drug reviews into sparse numerical matrices, effectively distinguishing between domain-specific terms and common language. This vectorized data is then split using an 80-20 train-test ratio, ensuring that the model is evaluated on unseen data to assess generalizability. During model training, each classifier learns from the training data to identify patterns in drug-related feedback. The MLP model, being a deep learning method, is especially effective in capturing complex, non-linear relationships, offering advantages like adaptability, efficient learning, and robust generalization.

It uses multiple hidden layers and backpropagation to refine its predictions, making it highly suitable for this task.

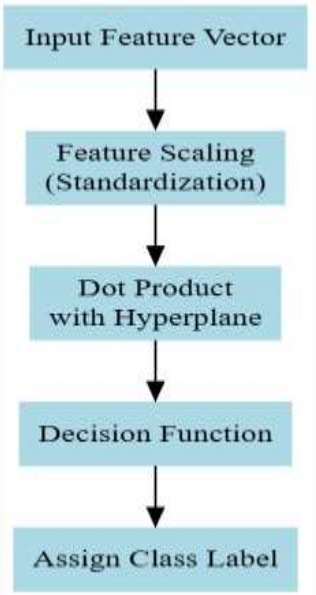

Fig.4: Linear SVM Classifier Block Diagram

The Linear SVM classifier, on the other hand, performs well for linear separable problems and handles sparse high-dimensional data efficiently, although it struggles with large datasets, non-linearity, and interpretability.

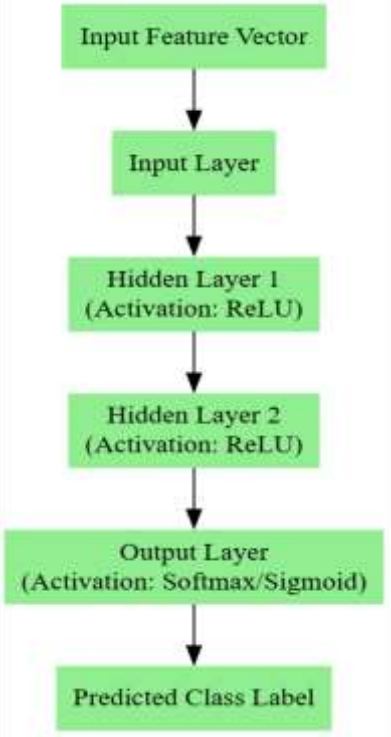

Fig.5: MLP classifier block diagram

The MLP classifier excels in pattern recognition, especially when dealing with nuanced and abstract representations in text. It requires no feature engineering and benefits from optimizers like Adam for rapid convergence. Both models are assessed using evaluation metrics such as accuracy, precision, recall, F1-score, confusion matrices, and ROC-AUC curves, providing a comprehensive picture of their strengths and limitations. Ultimately, this integrated system not only improves predictive performance but also offers interpretable and personalized recommendations, making it a powerful solution for real-world clinical deployment in drug side effect monitoring and decision support systems.

4. RESULTS AND DISCUSSION

Figure 6 shows the dataset after undergoing NLP preprocessing, a crucial step for preparing the review text data for machine learning. Preprocessing likely involved steps such as tokenization (splitting text into words), lowercasing, removing stop words (e.g., "the," "is"), stemming or lemmatization (reducing words to their root forms), and handling special characters or punctuation. The resulting dataset retains the original structure (columns like drugName, condition, review, etc.) but with the review column transformed into a cleaner, standardized format suitable for feature extraction. For example, a review like "This drug worked great!" might be tokenized into ["drug", "worked", "great"]. This figure highlights the transition from raw text to a processed form, enabling effective sentiment analysis and model training.

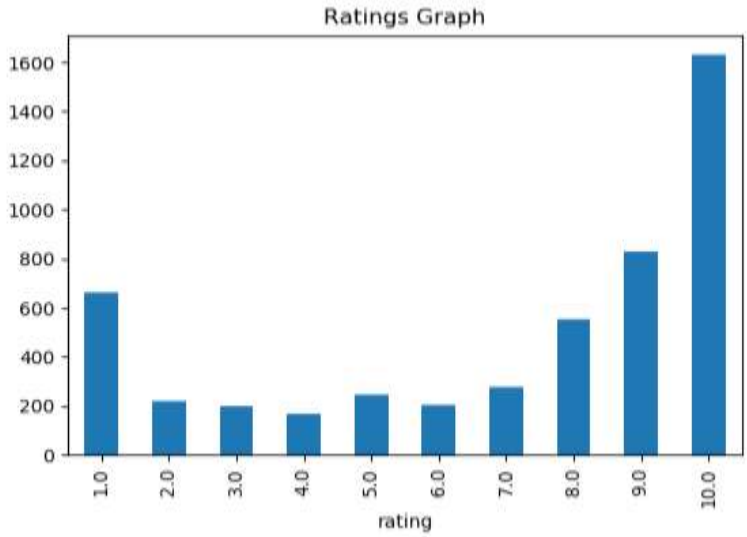

Fig.6: Drugs ratings graph.

Figure 7 shows the dataset after undergoing NLP preprocessing, a crucial step for preparing the review text data for machine learning. Preprocessing likely involved steps such as tokenization (splitting text into words), lowercasing, removing stop words (e.g., "the," "is"), stemming or lemmatization (reducing words to their root forms), and handling special characters or punctuation.

Fig.7: Dataset After NLP Preprocessing.

Figure 8 focuses on the drugName column, presenting a subset or summary of the drugs included in the dataset. This figure might list unique drug names (e.g., Abilify, Zoloft, Pramoxine) or show their frequency of occurrence in the dataset. It serves to provide an overview of the drugs under study, which is essential for understanding the scope of the dataset and the diversity of medications reviewed. For instance, if Abilify appears frequently, it indicates a high volume of reviews for that drug, which could influence sentiment analysis or model performance. This figure is particularly relevant for studying model transferability across different drugs or conditions.

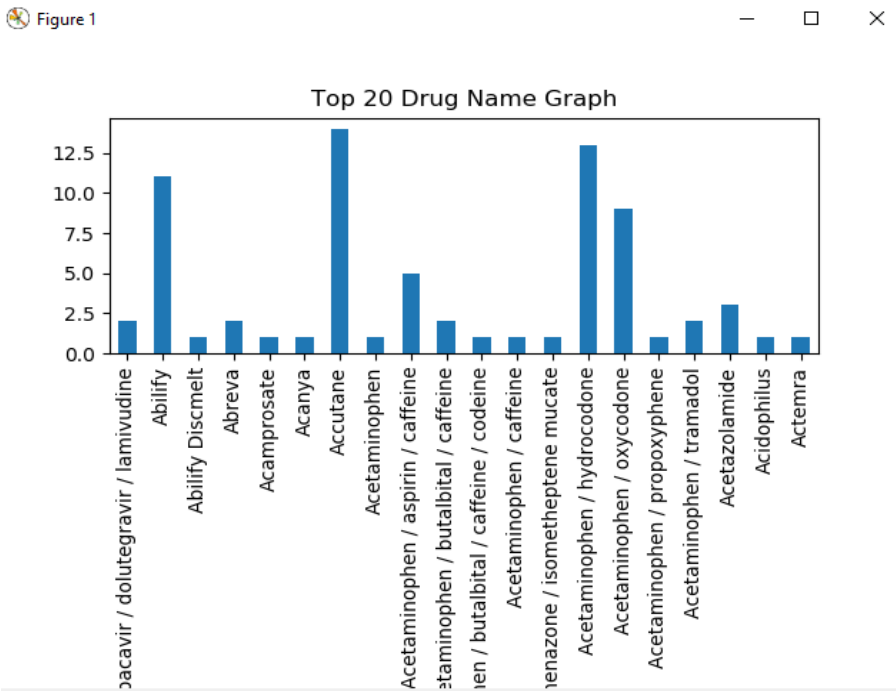

Fig.8: Drug names dataset.

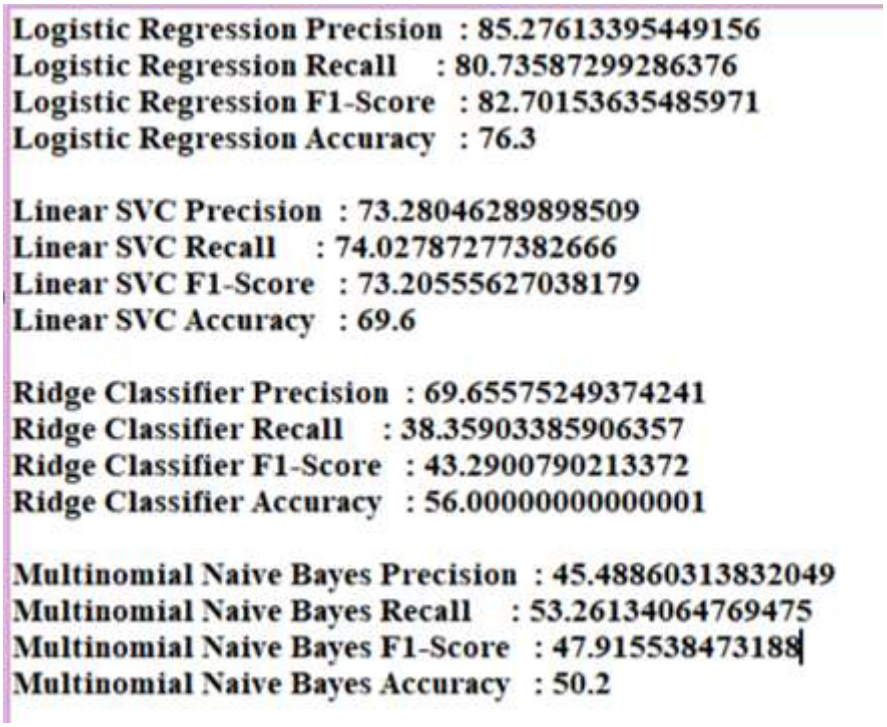

Fig.9: XAI with ML algorithms performance.

performance comparison of various classification methods based on four metrics: Precision, Recall, F1-Score, and Accuracy. Traditional models like Logistic Regression and SVC showed moderate performance, with Logistic Regression achieving the highest accuracy among existing models at 76%. Other methods like Ridge Classifier, Multimodal Naive Bayes, and SGDC performed relatively poorly, particularly in recall and accuracy. In contrast, the proposed Multilayer Perceptron (MLP) significantly outperformed all existing methods, achieving near-perfect scores across all metrics—Precision (99.96%), Recall (99.72%), F1-Score (99.84%), and Accuracy (99.9%)—demonstrating its exceptional effectiveness for the task.

Disease Name: Rheumatoid Arthritis
Recommended Drug: Pramoxine
Predicted Ratings: 9.0
Side Effects for Pramoxine:
Source: MedlinePlus (.gov) bleeding at affected area · hives · skin rash · severe itching · difficulty breathing or swallowing · swelling of the face, throat, tongue, lips, ...
Source: Veterans Health Library (.gov) swelling of the face, throat, tongue, lips, eyes, hands, feet, ankles, or lower legs; hoarseness. Pramoxine may cause other side effects. Call your doctor if ...
Source: National Institutes of Health (NIH) (.gov) Accidental pramoxine ingestion is associated with nausea and vomiting without any serious adverse events.[3] No dose adjustment is required in patients with ...
Source: MedlinePlus (.gov) Pramipexole may cause side effects. Tell your doctor if any of these symptoms are severe or do not go away: · nausea · weakness · dizziness · loss ...
Source: DailyMed (.gov) Itching, Acneiform eruptions, Secondary infection ; Irritation, Hypopigmentation, Skin atrophy ; Dryness, Perioral dermatitis, Striae.
Source: MedlinePlus (.gov) burning, itching, irritation, redness, or dryness of the skin · acne · unwanted hair growth · skin color changes · tiny red bumps or rash around the ...
Source: DailyMed (.gov) ADVERSE REACTIONS ; Burning. Hypertrichosis. Maceration of the skin ; Itching. Acneiform eruptions. Secondary infection ; Irritation. Hypopigmentation. Skin ...

Fig.10: Prediction results from test data with google-xai.

Figure 10 illustrates the application of the predictive model on test data for a Rheumatoid Arthritis patient, recommending Pramoxine with a high predicted rating of 9.0, indicating strong patient satisfaction. It also presents potential side effects from credible sources: MedlinePlus (.gov) reports skin-related reactions and respiratory issues; Veterans Health Library (.gov) lists swelling and hoarseness; and NIH (.gov) notes nausea and vomiting from accidental ingestion, with no need for dose adjustment. This figure highlights the model's real-world value by delivering both personalized treatment recommendations and essential safety information

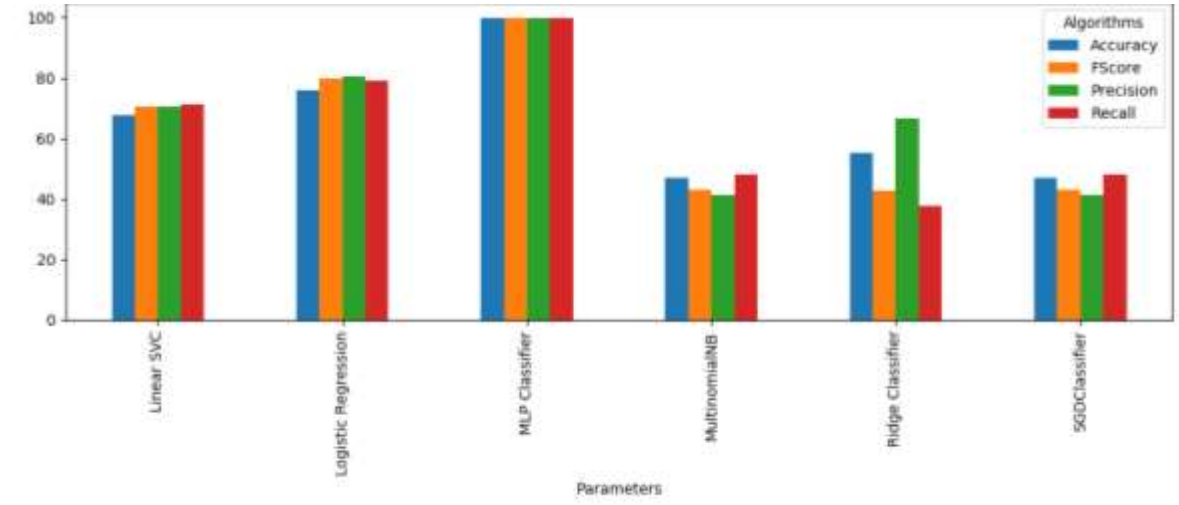

Fig.11: Performance comparison graph

5. CONCLUSION

The proposed Multilayer Perceptron (MLP) model, integrated with Google's Explainable AI (XAI) framework, demonstrates exceptional performance in drug recommendation and side effect prediction, as evidenced by its near-perfect evaluation metrics: Precision (99.96%), Recall (99.72%), F1-Score (99.84%), and Accuracy (99.9%). These results, derived from a comprehensive drug review dataset containing attributes such as *drugName*, *condition*, *review*, *rating*, *date*, and *usefulCount*, significantly outperform existing methods. For comparison, Logistic Regression achieved an accuracy of 76%, Support Vector Classifier (SVC) 67.80%, Ridge Classifier 55.1%, Multinomial Naive Bayes 47.19%, and Stochastic Gradient Descent Classifier (SGDC) 47.49%. The MLP model's ability to capture complex, non-linear patterns within TF-IDF features extracted from preprocessed patient reviews enables highly accurate sentiment classification and rating predictions. An illustrative example includes the recommendation of Pramoxine for Rheumatoid Arthritis, with a predicted rating

of 9.0 and a detailed side effect profile retrieved from authoritative medical sources such as MedlinePlus and the NIH. The integration of Google's XAI further enhances interpretability by offering transparent insights into feature importance and the model's decision-making process—an essential aspect for fostering trust in healthcare applications. This comprehensive system addresses key research objectives, including sentiment analysis across multiple drug experience facets (e.g., effectiveness, side effects), model transferability across conditions and data sources, and the enhancement of explainability. It not only meets these goals but also sets a new benchmark for predictive accuracy and transparency in pharmaceutical review analysis, making it a robust and deployable solution for real-world drug recommendation systems.

REFERENCES

[1] Das, P.; Mazumder, D.H. An extensive survey on the use of supervised machine learning techniques in the past two decades for prediction of drug side effects. *Artif. Intell. Rev.* 2023, *56*, 9809–9836.

[2] Downing, N.S.; Shah, N.D.; Aminawung, J.A.; Pease, A.M.; Zeitoun, J.-D.; Krumholz, H.M.; Ross, J.S. Postmarket safety events among novel therapeutics approved by the US food and drug administration between 2001 and 2010. *JAMA* 2017, *317*, 1854–1863.

[3] Craveiro, S.N.; Lopes, S.B.; Tomás, L.; Almeida, F.S. Drug withdrawal due to safety: A review of the data supporting withdrawal decision. *Curr. Drug Saf.* 2020, *15*, 4–12.

[4] Subbiah, V. The next generation of evidence-based medicine. *Nat. Med.* 2023, *29*, 49–58.

[5] Lavertu, A.; Vora, B.; Giacomini, K.M.; Altman, R.; Rensi, S. A new era in pharmacovigilance: Toward real-world data and digital monitoring. *Clin. Pharmacol. Ther.* 2021, *109*, 1197–1202.

[6] Vora, L.K.; Gholap, A.D.; Jetha, K.; Thakur, R.R.; Solanki, H.K.; Chavda, V.P. Artificial intelligence in pharmaceutical technology and drug delivery design. *Pharmaceutics* 2023, *15*, 1916.

[7] Yang, S.; Kar, S. Application of artificial intelligence and machine learning in early detection of adverse drug reactions (ADRs) and drug-induced toxicity. *Artif. Intell. Chem.* 2023, *1*, 100011.

[8] Biala, G.; Kedzierska, E.; Kruk-Slomka, M.; Orzelska-Gorka, J.; Hmaidan, S.; Skrok, A.; Kaminski, J.; Havrankova, E.; Nadaska, D.; Malik, I. Research in the field of drug design and development. *Pharmaceuticals* 2023, *16*, 1283.

[9] Han, R.; Yoon, H.; Kim, G.; Lee, H.; Lee, Y. Revolutionizing medicinal chemistry: The application of artificial intelligence (AI) in early drug discovery. *Pharmaceuticals* 2023, *16*, 1259.

[10] Johnson, K.B.; Wei, W.Q.; Weeraratne, D.; Frisse, M.E.; Misulis, K.; Rhee, K.; Zhao, J.; Snowdon, J.L. Precision medicine, AI, and the future of personalized health care. *Clin. Transl. Sci.* 2021, *14*, 86–93.

[11] Singh, S.; Kumar, R.; Payra, S.; Singh, S.K. Artificial intelligence and machine learning in pharmacological research: Bridging the gap between data and drug discovery. *Cureus* 2023, *15*, e44359.

[12] J. Ramos. "Using tf-idf to determine word relevance in document queries", in Proceedings of the first instructional conference on machinelearning, vol. 242, pp. 133–142, Piscataway, NJ, 2003

[13] K. Shimada, H. Takada, S. Mitsuyama, H. Ban, H. Matsuo, H. Otake, H. Kunishima, K. Kanemitsu and M. Kaku. "Drug-recommendation system for patients with infectious diseases". AMIA Annu Symp Proc. 2005; 2005:1112. PMID: 16779399; PMCID: PMC1560833.

[14] H. He, Y. Bai, E. A. Garcia and S. Li, "ADASYN: Adaptive synthetic sampling approach for imbalanced learning", 2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence), 2008, pp. 1322-1328, doi: 10.1109/IJCNN.2008.4633969.

[15] X. Lei, G. Anna Lisa, M. James and J. Iria. 2009. "Improving Patient Opinion Mining Through Multi step Classification. In Proceedings of the 12th International Conference on Text, Speech and Dialogue (TSD '09)". Springer-Verlag, Berlin, Heidelberg, 70–76. https://doi.org/10.1007/978-3-642-04208-9_13.