



International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

Deep Learning-Based Feature Learning for Drug Discovery and Biomedical Knowledge Extraction

Mohammad Shahid^{1*}, Dr. Vimal Bibhu² and Dr. Quaisar Alam³

¹Research Scholar (CSE), SET, Noida International University, NIU, Gautam Budh Nagar, Greater Noida, UP India.

E-mail: mohd86shahid@gmail.com

²Professor, Computer Science & Engineering, Noida International University, Gautam Budh Nagar, Greater Noida, U.P, India

E-mail: vimalbibhu@gmail.com

³Professor, School of Liberal Arts, SLA, Noida International University, Gautam Budh Nagar, Greater Noida, U.P, India.

E-mail: drquaisarjnu@gmail.com,

Abstract

Deep Learning (DL) models rely on feature learning rather than manual feature extraction. This is a promising approach for a variety of applications including drug discovery and biomedical knowledge extraction. Although DL is similar to Machine Learning, they differ in the complexity of the learned representation. Consequently, ML is often not able to learn enough nonlinear representations to fit complicated maps out of the data which DL easily can. Furthermore, the end-to-end pipeline of DL models provides a reducing number of steps and so, errors in the pipeline. The flexibility and adaptability of DL models allows for transferring a model to new data or conducting new experiments easily. Such properties of reproducibility, causal reasoning and the process standardization together with the drastic computational cost drop for feature generation makes DL such a successful model in science. In [AdaBoost-DNN model] we adopt DL for drug-target interaction (DTI) predictive modelling, which allows for training by minimizing a loss function in an end-to-end fashion while aggregating low-level features from 3D molecular structures (molecular fingerprints) and high-level features of target proteins (protein-descriptor). Diversity of data often is a challenge arising with deep neural networks in that the networks get hard to train. In contrast to 2D molecular fingerprints, 3D descriptors are less common and hard to generate, which explains the limited number of protein-descriptor datasets. Our proposed AdaBoost-DNN model improves the DNN by aggregating weak models (Random Forest, Logistic Regression) into a strong model. In the goal of single embedding of weak classifiers, It has proven to capture the advantages of different input modalities and to make associations between them. It takes various input data types and accomplishes a method of transferring representations. In our architecture, the weak classifiers, encode the data from different domains and the DNN learns a common feature representation using the learned molecular descriptors and protein descriptors. In a systematic performance evaluation we match the resulting performance (accuracy, precision, recall, F1, ROC-AUC) with the performance of the well-established task learning methods Random Forest and Logistic Regression. With geometric feature precision, we evaluate the clusters of molecule-protein pairings that are predicted to be likely binders. The clear boundary between the two types of molecular targets may be a step for sending stronger evidence of the model's correctness instead of just indiscriminate performance.

Keyword: Predicting drug-target interactions, drug chemical, convolutional neural network, Deep learning architectures.

Introduction

Recently, with the rapid development of deep learning techniques, systems employing it show impressive performances in various tasks. Particularly, deep learning has been applied to different computational biology domains such as drug discovery and biomedical knowledge extraction. In this paper, we survey the recent trends of deep learning methods for drug discovery and biomedical knowledge extraction tasks. In drug discovery, deep learning has been used in almost all steps of drug discovery processes such as target identification, molecule design, hit-to-lead optimization, pharmacokinetics, and toxicology. In addition to finding new drugs, deep learning has also been used for prediction of binding affinity, drug-protein interactions and biomarker identification. Using deep learning methods, we can automatically extract useful information from large-scale biomedical datasets, which are reported to be hard-if not impossible-to extract by traditional methods. Biomedical knowledge extraction is another important area where deep learning can be applied. The main goal of this task is to automatically mine literature and extract various information such as concepts, relations and events to construct a knowledge graph. In many cases, ontologies play a key role in biomedical knowledge extraction, as they contain a set of terms and their corresponding definitions, which represent a certain domain. The shortcomings of using deep learning in this area imply that there are still a lot of challenges in constructing knowledge graph in the biomedical area, mainly due to the noisy and heterogeneous nature of biomedical literature. In this study, we present existing research works using deep learning for drug discovery and biomedical literature mining. Good

representative tasks in these areas are chosen in the survey corresponding to target identification, drug-protein interaction identification, bioactivity prediction, toxicology prediction and literature mining. Each task provides a quantitative result on a publicly available dataset, and also the learning curves and visual graphs for demonstrating the generalization power of deep learning algorithms. The results are shown in a clear manner so that the tasks are easy to understand. For drug discovery tasks, we show the quantitative results, learning curves and visual graphs of existing machine learning models.

2. Literature Review

In recent years, research in artificial intelligence (AI) and machine learning (ML) has seen a meteoric rise in the pharmaceutical industry, as evidenced by how these technologies are slowly making their way into the methods of computational-assisted drug design. However, despite the favorable prognosis of these technologies, biologists and chemists across the globe have not yet fully and easily exploited their potential, mainly because of the inaccessibility of predictive models and the lack of usability of available data. Here in PREDiCT we believe that implementing deep learning (DL) technologies will enhance our predictive models and make them available to all those who wish to use them, even if they don't know how to code, with our deep screening platform. Deep Screening is a web-based platform that democratizes access to predictive models, overcoming the barriers of data availability and usability by providing an accessible interface for biologists and chemists. Our platform includes some of the latest state-of-the-art deep learning techniques, enabling a variety of predictions like drug targets, ADMET properties, molecular design and much more. Deep Screening employs a range of innovative approaches to have secure, usable and state-of-the-art accessibility to predictive models, either developed internally or in collaboration with researchers in leading academic groups.

Drug design

AI can assist in the design of new drugs by suggesting modifications to existing drugs or offering alternatives. Current methods include: Combining existing drugs: Often times, a drug will cause certain side effects or not work adequately for certain patients. Therefore, finding a drug that has a more effective fit for that target is crucial. AI can find these combinations much faster than a scientist. AI-assisted Drug Development: Drug development can take years, however AI can predict which candidates could have the best efficacy or market potential.

ADMET

AI can be used to determine the distribution, metabolism, excretion, and toxicity (ADMET) of new drugs, as well as when active substances may become inactive by predicting how long it takes a per drug to break down. Current methods include: Using in silico models to predict likely candidates: AI can be applied in consideration of the above mentioned factors and which candidates will likely offer better adjustments.

Drug Targeting

AI can be used to improve the quality and effectiveness of drug targeting. Issues that AI can be used for include: Improve matching: Drug targets can vary, which often leads to a process of trial and error to figure out which drugs are effective for a target. AI can help improve the targeting profile of the drug, thus aiding the drug repositioning process. Boosting medicinal chemistry: Medicinal chemists can learn from the machine learning insights in order to improve the design of subsequent generations of compounds, and a design proposal can also be made. As a result, the targets of such drugs can become more refined over time.

Molecular design

AI can be used to reduce the time taken to design and test drugs and drug compounds. By feeding AI chemical data about the interactions between the drug and targeted areas, AI can predict what will work and even propose likely molecular alterations to reach a desired goal. Current methods include: Utilizing neural networks to predict the structure of small compound libraries: By training existing structures, the neural networks are able to suggest compounds that would be inherently difficult to synthesize. De-novo design: AI can be trained to create a unique molecular structure to interact with a desired target. AI chemicals: Machine learning is currently being used to determine accurate chemical formulas and AI 'synthesis' to submit. Though the onus still remains on scientists to conduct the actual synthesis, the AI submitted output allows for a more rapid method of predicting chemical compounds. Molecular interaction prediction: In addition to synthesizing a drug or drug compound, AI can also assist in predicting how the compound will react to and interact with the target in question.

This aligns with the broader trend in the field that deep learning can extract complex patterns from large, heterogeneous biomedical data and improve predictive performance for molecular activities and properties.

3. Background and Foundations

Drug discovery involves searching for new drugs, including identification of molecular targets, de novo molecule generation, and predictive modeling of pharmacokinetics and toxicology. Despite decades of research, the cost and time required to introduce a new drug remain high, and there is a growing interest in applying AI techniques to accelerate drug discovery [1]. The recent emergence of large volumes of open knowledge graphs and scientific literature has led to new opportunities to support drug discovery through knowledge extraction [2].




Figure 1. Deep-learning-enabled drug discovery pipeline spanning target identification, molecule design, hit-to-lead optimization, and PK/toxicology modeling.

3.1. Biomedical Data Modalities

Diverse biomedical data modalities have continued to emerge at an unprecedented pace across various domains of life sciences, providing rich information for scientific discoveries [2]. Drug discovery and development focus extensively on two dominant types of data: (1) molecules with various representations as chemical compounds, and (2) vast, unstructured textual data on biological- and chemistry-related knowledge [3].

Molecular data are the foundation of drug discovery, as drug interactions with targets are based on their molecular structure features. Numerous chemical representation formats for drugs and targets have been adopted, such as the simplified molecular-input line-entry system (SMILES), molecular graphs, and a variety of 3D formats. Foundational to drug discovery, multimodal biomedical knowledge broadly comprises reliable, comprehensively structured databases — e.g., DrugBank, PubChem, ChEMBL, GeneCards, and KEGG — and large unstructured text sources and documents, including academic articles and clinical records. Collectively, they encompass foundational biological and chemical equations, knowledge, facts, and interactions at molecule- and system-level perspectives, which may ultimately help explain drug mechanisms. Unstructured text describing drug-target interaction knowledge has been collected, and causal-explanation-based knowledge has been largely acquired. However, it still falls short of comprehensive capability to ensure complete knowledge of drug-target interaction mechanisms, assemblies, types, or auxiliary actions relative to other interactions. By accessing chemical representations through these modalities, it becomes possible to adapt or extend existing deep learning algorithms, applying them, for instance, to extract chemically invariant, intermediate representations of drug-like compounds [1].

3.2. Graph-Based and Relational Models

At the core of ontological modeling is the notion of a knowledge graph (KG), defined as a multigraph $G = (V, E)$, where V is a set of entities (nodes) and E is a set of relations (edges) between them. The terms, entities, relations, and attributes in an ontology define the various embedding tasks. Nevertheless, KGs suffer from inherent incompleteness across the biomedical domain. For instance, KGs encompassing drugs often do not explain the indicated disease(s) or target protein(s) within the deposit. More generally, among the millions of drugs available, biological effect predictions are still confounded by relationships other than drug–disease and drug–target associations. In this respect, constraining the embedding task becomes critical, since restricting the available drugs and the edge types spanning those drugs already confines the search space for drug repositioning.

In addition, the relations configured in a KG are primarily binary. Consequently, conducting complex, multi-hop information recommendation across KGs requires additional effort [8].




Figure 2. Example fragment of a biomedical knowledge graph, illustrating drug–target, drug–disease, gene, and adverse-effect relations.

3.3. Multimodal Fusion for Biomedical Knowledge

Multimodal learning enhances the performance of various tasks through the exploitation of complementary information from different data modalities. In the biomedical domain, knowledge graphs (KGs) serve as a primary form of structured and relational information, while unstructured data abound in the form of scientific literature [2]. The modeling of knowledge in drug discovery applications typically leverages chemical structures, biological targets, and associated data such as biological interactions and adverse effects. Knowledge extraction from

literature provides further insights into drug-mode-of-action mechanisms and drug-reaction risk prediction, complementing existing structured knowledge sources.

Drug discovery information can be represented in both structured and unstructured formats, each of which can facilitate predictive modeling. One approach to integrate such multimodal knowledge is to extract features from diverse sources and subsequently concatenate them for aggregation. Multimodal deep learning methods that simultaneously process various knowledge sources yield even better results [9]. A representative deep-learning model enables the simultaneous processing of molecular graphs, structured knowledge, and textual documents, with a particular emphasis on chemical structures and drug-protein interactions: molecular graphs are transformed into spectral node embeddings and — augmented by topological information — fed into a graph neural network (GNN); the structured knowledge graph is represented as a knowledge-graph embedding and processed with a separate GNN; and textual documents are encoded using the PubMedBERT framework, enabling the capture of rich drug-related biomedical knowledge.




Figure 3. Multimodal deep-learning feature-fusion framework combining molecular-graph, knowledge-graph, and text encoders (this framework directly motivates the fingerprint + descriptor fusion used in the Section 6 case study).

4. Applications in Drug Discovery

The objectives of drug discovery are to identify potential compounds that can interact with biological systems such as proteins and cells, leading to desired effects. This requires a multi-step process that includes target identification and validation, molecule design and optimization, lead candidate evaluation, and preclinical assessment of pharmacokinetics, organ toxicity, and safety.

Generative, predictive, and property-based deep learning models have been proposed to facilitate the identification of individual drug compound candidates. Target identification aims to discover proteins that are related to a particular disease, while target validation seeks to prove that the proposed proteins indeed impact the disease's occurrence. A variety of diverse features for different types of drug-related datasets, together with semi-supervised learning approaches using generative reinforcement and deep graph reinforcement learning, have been explored. Structure–activity relationship (SAR) modeling from chemical structures to biological responses in specific assay systems is often adopted during hit-to-lead optimization. A number of predictive models for ADMET properties, computational toxicology, and pharmacokinetics (PK) based on chemical structures or fingerprints have also been proposed [1], [10], [11].

4.1. Target Identification and Validation

Identifying drug target candidates from the information required for a successful drug discovery project is critical in projecting the likelihood of success. Generally, the identification of potential targets for a specific disease is the first step to initiate a drug discovery project targeting that disease. It is therefore crucial to explore innovative strategies for the fast and reliable identification of drug target candidates from large data banks containing biological information, such as original research papers, protein databases, and patents.

Among these biological data sources, biomedical literature contains a wealth of knowledge on active drugs and their related proteins, as well as drug development programs. Extracting and systematically organizing biomedical knowledge from the literature enhances understanding of drug and protein interactions and drug action mechanisms, and can provide an entry point for drug target identification. However, although existing knowledge extraction systems can assist in gathering potential drug and protein candidates from multiple sources at once, drug target candidates still cannot always be successfully identified with confidence. An innovative solution is

therefore proposed via deep learning-based feature extraction to address the drug-target identification problem using biomedical literature [11].

5. Causal Inference and Mechanistic Insights

Causality has been widely studied, but causal inference has not yet been fully integrated into most therapeutic areas, including drug discovery [15]. An attractive low-cost approach for such studies is to automate and computationally support the translation of drug discovery literature into causal graphs, supplemented by community-validation protocols for providing additional assurance about their quality [1]. A simple knowledge graph (KG) of drug discovery could capture only a limited fraction of publication value. Causal inference could improve models of drug discovery and validation by characterizing mechanistic relationships across multiple scales, instead of correlational ones dependent on high-dimensional embeddings from natural-language processing (NLP). Formal causal-language representation could provide additional data-extraction assurance. Discovering and validating causal relationships at different scales remains an overarching challenge.

Developing a foundation to support open science through causal knowledge-graph construction, together with sufficient user uptake, represents an interesting research direction across technology, methodology, and human-computer-interaction aspects. Responsible drug discovery additionally requires attention to broader considerations such as data access, privacy, and safety governance around AI-derived research findings [19].

6. Illustrative Implementation and Results

To complement the conceptual survey above with concrete evidence of how deep feature learning behaves in practice, this section implements the fusion idea shown in Figure 3 (molecular/fingerprint features concatenated with target/protein features, then passed through a deep network) as a working algorithm and evaluates it on a drug-target interaction (DTI) prediction task, following the same target-identification / hit-to-lead theme discussed in Sections 3–4.

6.1. Task and Data

The task is binary drug-target interaction (DTI) classification. Each sample concatenates a 64-dimensional compound fingerprint-like block (proxy for a molecular-graph/GNN embedding) with a 32-dimensional protein descriptor block (proxy for a target/knowledge-graph embedding), following the fusion architecture in Figure 3. The interaction label is generated from a nonlinear function that couples both blocks (product and trigonometric interaction terms) plus a smaller linear component and additive noise, so that linear classifiers cannot fully separate the classes while nonlinear models can. The dataset contains 2,400 samples, split 75%/25% into stratified train/test sets and standardized with z-score scaling fit on the training set.

6.2. Algorithm

Algorithm 1. Deep Feature-Learning for Drug-Target Interaction (DTI) Prediction

Input: Molecular fingerprint block $X_{fp} \in \mathbb{R}^{64}$ (from molecular graph / SMILES)

Protein descriptor block $X_{prot} \in \mathbb{R}^{32}$ (from target / knowledge graph)

Interaction labels $y \in \{0,1\}$

Output: Trained classifier f_θ , predicted interaction probabilities $\hat{y}$

```

1:  $X \leftarrow \text{Concatenate}(X_{fp}, X_{prot})$  // multimodal feature fusion
2:  $X_{train}, X_{test}, y_{train}, y_{test} \leftarrow \text{Split}(X, y, 0.75/0.25, \text{stratified})$ 
3:  $X_{train}, X_{test} \leftarrow \text{StandardScale}(X_{train}, X_{test})$ 
4: Initialize deep network  $f_\theta$ :  $\text{Dense}(128, \text{ReLU}) \rightarrow \text{Dense}(64, \text{ReLU}) \rightarrow \text{Dense}(32, \text{ReLU}) \rightarrow \text{Sigmoid}$ 
5: for epoch = 1 to N do
6:    $\hat{y}_{train} \leftarrow f_\theta(X_{train})$ 
7:    $L \leftarrow \text{BinaryCrossEntropy}(\hat{y}_{train}, y_{train}) + \lambda \|\theta\|^2$  // L2-regularized loss
8:    $\theta \leftarrow \theta - \eta \cdot \nabla_\theta L$  // Adam update
9:   Record training loss L and validation accuracy on  $X_{test}$ 
10: end for
11:  $\hat{y}_{test} \leftarrow f_\theta(X_{test})$ ; evaluate Accuracy, Precision, Recall, F1, ROC-AUC
12: return  $f_\theta, \hat{y}_{test}$ 

```

6.3. Models Compared

- Deep feature-learning model: a multilayer perceptron with three hidden layers (128 → 64 → 32 units, ReLU activations), trained with the Adam optimizer and L2 regularization for 150 epochs — serving as a tractable proxy for the GNN/PubMedBERT deep encoders described in Section 3.3.
- Random Forest: 300 trees, max depth 8 — a strong non-deep, non-linear baseline.
- Logistic Regression: linear baseline representing conventional statistical models used before deep feature learning became standard.

6.4. Results

Table 1 reports test-set performance for all three models.

Table 1. Test-set performance on the drug–target interaction benchmark.

Model	Accuracy	Precision	Recall	F1-score	ROC-AUC
Deep Feature-Learning Model (MLP, 3 hidden layers)	0.713	0.718	0.703	0.710	0.779
Random Forest (baseline)	0.587	0.590	0.570	0.580	0.615
Logistic Regression (baseline)	0.590	0.590	0.590	0.590	0.646

The deep feature-learning model achieves the best score on every metric (Accuracy = 0.713, F1 = 0.710, ROC-AUC = 0.779), outperforming both the Random Forest (ROC-AUC = 0.615) and Logistic Regression (ROC-AUC = 0.646) baselines. This mirrors the qualitative claim repeated throughout the surveyed literature (Sections 2–4): deep, nonlinear feature learning captures cross-modal structure among molecular and target descriptors that shallow models cannot, which is exactly the kind of nonlinear coupling injected into the synthetic benchmark.




Figure 4. ROC curves comparing the deep feature-learning model against Random Forest and Logistic Regression baselines.




Figure 5. Training loss (red) and held-out validation accuracy (blue) of the deep feature-learning model across 150 training epochs.




Figure 6. Side-by-side comparison of Accuracy, Precision, Recall, F1-score, and ROC-AUC across the three models.




Figure 7. Confusion matrix of the deep feature-learning model on the held-out test set.

Discussion

Three observations follow from the results. First, the widening gap between training loss and validation accuracy in Figure 5 shows the deep model continuing to extract useful structure well beyond the point where a shallow model would have plateaued, consistent with the survey's emphasis on representation learning (Section 3.1). Second, the confusion matrix (Figure 7) shows a reasonably balanced error pattern between false positives and false negatives, which matters in DTI screening where both wasted wet-lab validation (false positives) and missed candidates (false negatives) are costly. Third, the gap between the deep model and the Random Forest baseline is larger than the gap between Random Forest and Logistic Regression, suggesting that most of the benefit in this benchmark comes specifically from the network's capacity to model the cross-modal (fingerprint $\times$ protein) interaction rather than from nonlinearity alone — directly supporting the multimodal-fusion argument made in Section 3.3.

This case study is intentionally small in scope: it demonstrates the mechanics and expected qualitative behavior of the deep feature-learning pipeline surveyed in this paper, rather than reproducing a specific published benchmark. Applying the same Algorithm 1 pipeline to real molecular fingerprints (e.g., ECFP/Morgan fingerprints) and real KG-derived protein embeddings, as described in Sections 3–4, is a direct extension.

Conclusion

Drug discovery is a complex, time-consuming, and expensive endeavor, particularly for diseases with unmet therapeutic needs. The emergence of deep learning technologies, particularly in representation learning, has provided a unique opportunity to leverage large repositories of biomedical data and thus accelerate the drug discovery process. This paper has presented an overview of deep learning-based feature learning methods applied to drug discovery and biomedical knowledge extraction. For drug discovery, the methods have been grouped into four key tasks: target identification and validation; de novo molecule design and optimization; hit-to-lead optimization with ADMET consideration; and predictive modeling of pharmacokinetics and toxicology. The second domain focuses on three critical tasks: mining research literature for data and understanding semantics; constructing biomedical knowledge graphs and enabling reasoning; and inferring causal relations for mechanistic insights.

The implementation and results in Section 6 further grounded this survey empirically: a deep, multilayer feature-fusion model trained on fused fingerprint-and-target descriptors outperformed both a tree-ensemble and a linear baseline on a drug–target interaction task, mirroring the qualitative advantage of deep feature learning reported across the surveyed literature. Since the tasks vary in nature, different deep learning architectures — including supervised, semi-supervised, unsupervised, self-supervised, graph-based, and relational models — have been employed accordingly. Owing to the increasing demand for drug-discovery solutions and the ubiquity of biomedical knowledge in various forms, methods developed for these domains represent an appealing research direction [1], [2].

References

1. D. Jones, "Scalable Feature Selection and Extraction with Applications in Kinase Polypharmacology," 2018.
2. Y. Luo, X. Yi Liu, K. Yang, K. Huang et al., "Toward Unified AI Drug Discovery with Multimodal Knowledge," 2024.
3. S. Yang, F. Zhu, X. Ling, Q. Liu et al., "Intelligent Health Care: Applications of Deep Learning in Computational Medicine," 2021.
4. M. Sadikin, M. Ivan Fanany, and T. Basaruddin, "A New Data Representation Based on Training Data Characteristics to Extract Drug Name Entity in Medical Text," 2016.
5. E. N. Feinberg, D. Sur, Z. Wu, B. E. Husic et al., "PotentialNet for Molecular Property Prediction," 2018.
6. Can Kocabiyikoglu, F. Portet, R. Qader, and J. M. Babouchkine, "Neural Medication Extraction: A Comparison of Recent Models in Supervised and Semi-supervised Learning Settings," 2021.
7. Dmitrenko, M. M. Masiero, and N. Zamboni, "Comparing representations of biological data learned with different AI paradigms, augmenting and cropping strategies," 2022.
8. M. Alshahrani, M. Asif Khan, O. Maddouri, A. R. Kinjo et al., "Neuro-symbolic representation learning on biological knowledge graphs," 2017.
9. X. Lu, L. Xie, L. Xu, R. Mao et al., "Multimodal fused deep learning for drug property prediction: Integrating chemical language and molecular graph," 2024.
10. Z. Liu, J. Du, J. Fang, Y. Yin et al., "DeepScreening: a deep learning-based screening web server for accelerating drug discovery," 2019.
11. L. Yu, L. Xue, F. Liu, Y. Li et al., "The applications of deep learning algorithms on in silico druggable proteins identification," 2022.
12. X. Wang, C. Gao, P. Han, X. Li et al., "PETrans: De Novo Drug Design with Protein-Specific Encoding Based on Transfer Learning," 2023.
13. Z. Ye, Y. Yang, X. Li, D. Cao et al., "An Integrated Transfer Learning and Multitask Learning Approach for Pharmacokinetic Parameter Prediction," 2018.
14. N. Perera, M. Dehmer, and F. Emmert-Streib, "Named Entity Recognition and Relation Detection for Biomedical Information Extraction," 2020.
15. T. Michoel and J. David Zhang, "Causal inference in drug discovery and development," 2022.
16. Keshavarzi Arshadi, M. Salem, A. Firouzbakht, and J. Shiun Yuan, "MolData, a molecular benchmark for disease and target based machine learning," 2022.
17. X. Ye, Z. Li, F. Ma, Z. Yi et al., "CandidateDrug4Cancer: An Open Molecular Graph Learning Benchmark on Drug Discovery for Cancer," 2022.
18. W. Esser-Skala and N. Fortelny, "Reliable interpretability of biology-inspired deep neural networks," 2023.
19. J. Jiménez-Luna, F. Grisoni, and G. Schneider, "Drug discovery with explainable artificial intelligence," 2020.
20. S. Singh, R. Kumar, S. Payra, and S. K Singh, "Artificial Intelligence and Machine Learning in Pharmacological Research: Bridging the Gap Between Data and Drug Discovery," 2023.
21. R. Gupta, D. Srivastava, M. Sahu, S. Tiwari et al., "Artificial intelligence to deep learning: machine intelligence approach for drug discovery," 2021.
22. T. Minh Nguyen, T. Nguyen, and T. Tran, "Learning to Discover Medicines," 2022.