

Received 28 May 2026

Accepted 26 June 2026

DOI: 10.48308/CMCMA.5.1.15

AMS Subject Classification: 68T07; 65K10

A Particle Swarm Optimization-Based Deep Learning Framework for Drug–Target Interaction Prediction

Raziyeh Masumshah^a and Changiz Eslahchi^a

Accurate prediction of drug-target interactions (DTIs) plays a central role in computational drug discovery and drug repositioning. Although deep learning models have substantially improved DTI prediction performance, most existing approaches rely exclusively on gradient-based optimization, which may suffer from unstable convergence or suboptimal solutions in sparse and high-dimensional biological datasets. In this study, we propose PSO-DTI, a swarm intelligence-enhanced neural framework that integrates similarity-based relational feature transformation with Particle Swarm Optimization (PSO) for adaptive neural weight refinement. Drugs and proteins are transformed into similarity-profile representations using Jaccard and cosine similarity measures, respectively. A neural classifier is then optimized using a validation-guided PSO-based selective weight adoption strategy. Experiments on benchmark *C. elegans* and Human datasets demonstrate consistent performance improvements over representative state-of-the-art models across AUROC, AUPRC, Accuracy, F1-score, and MCC metrics. Ablation studies confirm the contribution of similarity-based preprocessing and PSO-enhanced optimization. Case analyses further support the biological plausibility of top-ranked predictions. These findings suggest that integrating global search heuristics into neural training pipelines can enhance robustness and generalization in DTI prediction. Copyright © 2026 Shahid Beheshti University.

Keywords: Drug-Target Interaction; Deep Learning; Particle Swarm Optimization; Similarity-Based Representation Learning.

1. Introduction

Identifying reliable drug-target interactions (DTIs) is a fundamental step in drug discovery and therapeutic development. Traditional experimental identification methods are costly and time-consuming, motivating the development of computational approaches capable of efficiently predicting potential interactions. Machine learning and deep learning techniques have significantly improved predictive performance by learning latent representations from chemical and biological data [1, 2]. Beyond DTI modeling, computational pharmacology also addresses drug-drug interaction (DDI) prediction and drug-disease association (DDA) inference. These related tasks highlight the importance of robust feature representation and optimization strategies in biomedical interaction modeling. Insights from these domains motivate the exploration of alternative training paradigms for DTI prediction [3, 4, 5]. From a computational perspective, predicting drug-target interactions can be formulated as a link prediction problem, where the objective is to infer potential associations between drugs and target proteins based on available data [6]. Over the last decade, a wide range of techniques have been proposed, ranging from similarity-based matrix factorization to advanced deep learning models. Traditional methods often rely on handcrafted features, whereas modern deep learning and graph representation learning techniques are capable of extracting latent features from complex biological data [7, 8, 9].

Recent studies have pushed the boundaries of DTI modeling. For example, Graph-DTI employs heterogeneous network embedding to capture complex topological relationships between drugs and targets [8], while attention-based fusion models like FusionDTI demonstrate enhanced learning of fine-grained structural information [9]. Knowledge graph-based approaches further integrate rich semantic context into DTI prediction [10], and cross-domain fusion frameworks show promising generalization

^a Department of Computer and Data Sciences, Faculty of Mathematical Sciences, Shahid Beheshti University, 1983969411, Tehran, Iran.

* Correspondence to: C. Eslahchi. Email: ch-eslahchi@sbu.ac.ir

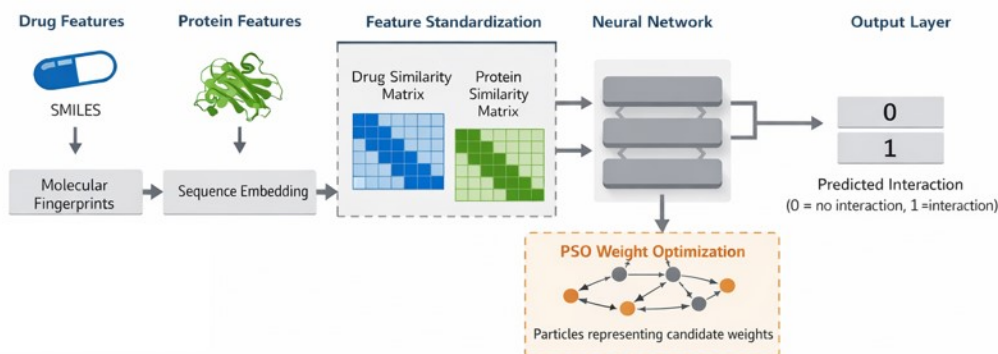


Figure 1. Overview of the PSO-DTI framework for drug-target interaction prediction. Drug features (SMILES) and protein sequences are transformed into standardized similarity profiles. These representations are processed by a neural network, whose weights are iteratively refined using Particle Swarm Optimization (PSO). The model outputs a binary prediction (0 = no interaction, 1 = interaction) for each drug-protein pair.

to new drug and target entities [11]. Despite these advances, significant challenges remain, including data sparsity, model interpretability, and the intrinsic nonconvexity of learning models applied to high-dimensional biological data [12].

Most existing DTI prediction frameworks use gradient-based optimization techniques (e.g., backpropagation) for training deep models. While these methods are effective in many settings, they can be sensitive to initialization, prone to local minima, and may exhibit unstable convergence when trained on sparse or noisy biological datasets. This limitation motivates the exploration of alternative optimization strategies that can complement or enhance gradient descent techniques. Particle Swarm Optimization (PSO), introduced by Kennedy and Eberhart [13], is a population-based metaheuristic that simulates the social behavior of particles moving through a search space. PSO has been successfully applied to a variety of neural network optimization tasks due to its simplicity, global search capability, and good balance between exploration and exploitation [14]. In biomedical contexts, swarm intelligence methods have demonstrated value in optimizing model parameters where gradient descent alone may struggle. However, the direct integration of PSO into neural network weight optimization for DTI prediction remains underexplored. Most neural DTI models rely entirely on gradient-based training without leveraging global search heuristics such as PSO, which can enhance exploration of the parameter space and reduce the risk of premature convergence to suboptimal solutions.

In this work, we introduce a PSO-enhanced neural optimization framework tailored for Drug-Target Interaction prediction. Our approach utilizes Particle Swarm Optimization to iteratively explore candidate weight configurations based on validation loss performance. When PSO finds a better configuration, the model updates to these weights; otherwise, it retains its previous parameters, offering a selective adaptive training mechanism that stabilizes convergence. To assess the proposed framework, comprehensive experiments are performed on benchmark DTI datasets, benchmarking our method against standard training strategies and representative state-of-the-art models. Experimental results demonstrate that integrating PSO into neural training improves predictive accuracy and training robustness.

2. Methodology

The objective of the drug-target interaction prediction task is to determine whether a given drug interacts with a specific target protein. Given sets of drugs and target proteins, the task is formulated as a binary classification problem, where each drug-protein pair is mapped to a label indicating the presence (1) or absence (0) of an interaction. This formulation allows the model to learn patterns from known interactions and generalize to unseen pairs. An overview of the architecture is illustrated in Fig. 1.

2.1. Datasets

To evaluate the effectiveness and generalization capability of the proposed framework, experiments were conducted on two benchmark drug-target interaction (DTI) datasets, namely the *C. elegans* and Human datasets. These datasets are commonly used in computational DTI studies and provide structured interaction information between drug compounds and their corresponding protein targets. The *C. elegans* dataset comprises 1,767 drug compounds and 1,876 protein targets, resulting in 7,786 experimentally validated drug-target interactions. The Human dataset includes 2,726 drugs and 2,001 human protein targets, with a total of 6,728 confirmed interactions. The two datasets differ in biological complexity and interaction density, thereby enabling a comprehensive evaluation of the proposed model under distinct experimental settings [15, 16, 17]. The inclusion of these two datasets allows the proposed model to be evaluated under different biological contexts and interaction distributions.

The variation in the number of drugs, targets, and known interactions between the two datasets provides a comprehensive setting to assess the stability, scalability, and predictive capability of the proposed approach.

2.2. Feature Representation

Each drug is encoded using its molecular fingerprint representation derived from the SMILES strings. Specifically, 881-bit PubChem molecular fingerprints are generated using the RDKit cheminformatics toolkit. These fingerprints capture structural and chemical substructures in a fixed-length binary vector of dimension 881. Such representations encode the presence or absence of predefined molecular fragments, providing a compact yet informative description of molecular topology and functional groups.

For target proteins, sequence-based descriptors are derived directly from the amino acid sequences (SEQ). Each amino acid residue is mapped to a unique integer index, transforming the protein sequence into a numerical sequence representation. To ensure dimensional consistency across all protein instances and to facilitate stable model training, protein sequences are padded or truncated to a fixed maximum length of 1200, following established practices in previous DTI studies. Consequently, each protein is represented as a 1200-dimensional sequence-based feature vector.

Formally,

- $X_D \in \mathbb{R}^{n_d \times 881}$ denotes the drug feature matrix, where n_d represents the number of drugs.
- $X_T \in \mathbb{R}^{n_t \times 1200}$ denotes the protein feature matrix, where n_t represents the number of target proteins.

Given the heterogeneous biological origins of drugs and proteins, these representations serve as standardized numerical embeddings that enable unified downstream learning within the predictive framework.

2.3. Similarity-Based Representation

While fixed-length feature vectors encode intrinsic properties of drugs and proteins, they do not explicitly capture relational dependencies among entities. To incorporate such structural relationships, similarity-based transformations are applied, generating interaction-aware representations within each modality.

2.3.1 Drug Similarity

Each drug is represented by an 881-dimensional binary fingerprint vector. To quantify pairwise structural similarity between drugs, Jaccard similarity is employed, which is particularly suitable for binary feature spaces.

The similarity between drugs d_i and d_j is defined as:

$$\text{Sim}_D(i, j) = \frac{|F_i \cap F_j|}{|F_i \cup F_j|}, \quad (1)$$

where F_i and F_j denote the binary fingerprint sets of drugs d_i and d_j , respectively.

This yields a drug similarity matrix

$$S_D \in \mathbb{R}^{n_d \times n_d}. \quad (2)$$

Thus, each drug is re-encoded by its similarity profile relative to all other drugs, capturing global structural relationships within the drug space.

2.3.2 Protein Similarity

Each protein is represented by a 1200-dimensional sequence-derived feature vector. To measure similarity between protein targets in the continuous embedding space, cosine similarity is utilized, as it effectively captures angular similarity independent of vector magnitude.

The similarity between proteins t_i and t_j is defined as:

$$\text{Sim}_T(i, j) = \frac{E_i \cdot E_j}{\|E_i\| \|E_j\|}, \quad (3)$$

where E_i and E_j denote the feature vectors of proteins t_i and t_j , respectively.

This results in a protein similarity matrix

$$S_T \in \mathbb{R}^{n_t \times n_t}. \quad (4)$$

Accordingly, each protein is represented through its relational similarity vector within the protein embedding space, enriching the representation with contextual structural information.

2.4. PSO-Enhanced Neural Optimization Framework

After constructing similarity-based representations, the drug-target interaction prediction task is formulated as a supervised binary classification problem using a neural network model.

Let

$$Z_{ij} = [S_D(i, :), S_T(j, :)] \quad (5)$$

denote the concatenated similarity-based feature vector of drug d_i and protein t_j , where $S_D(i, :)$ represents the i -th row of the drug similarity matrix and $S_T(j, :)$ represents the j -th row of the protein similarity matrix.

The neural network, parameterized by W , outputs a predicted interaction probability \hat{y}_{ij} for each drug-protein pair. The model is trained by minimizing the binary cross-entropy loss:

$$\mathcal{L}(W) = - \sum_{i,j} [y_{ij} \log(\hat{y}_{ij}) + (1 - y_{ij}) \log(1 - \hat{y}_{ij})], \quad (6)$$

where $y_{ij} \in \{0, 1\}$ denotes the ground-truth interaction label.

2.4.1 PSO-Based Weight Optimization

Instead of relying exclusively on conventional gradient-based optimization, Particle Swarm Optimization (PSO) is integrated to enhance the exploration of the neural network parameter space and mitigate the risk of convergence to local minima.

In this study, PSO is applied as a refinement step for the neural network parameters. The neural network is first trained using standard gradient-based optimization. After this initial training stage, the learned network weights serve as the initial positions for the particles in the PSO search space. Particle velocities are initialized with small random values sampled from a uniform distribution. Given that the model utilizes a single drug representation and a single protein representation, the optimization process is handled by a single particle. The PSO parameters are set as follows: inertia weight $\omega = 0.7$, and cognitive and social acceleration coefficients $c_1 = c_2 = 1.5$. The optimization then proceeds for 10 iterations.

The network weights, which represent the parameters to be optimized, can be collectively denoted as a vector for particle k within the swarm at a given iteration. Let W_k represent this vector of trainable parameters:

$$W_k = [w_1, w_2, \dots, w_m], \quad (7)$$

where m is the total number of trainable parameters in the network, and k is the index for the particle.

At iteration t , the velocity $v_k(t+1)$ and position $W_k(t+1)$ of each particle are updated according to the following equations:

$$v_k(t+1) = \omega v_k(t) + c_1 r_1 (pbest_k - W_k(t)) + c_2 r_2 (gbest - W_k(t)), \quad (8)$$

$$W_k(t+1) = W_k(t) + v_k(t+1), \quad (9)$$

Key components in these updates include:

- $pbest_k$: The best position (optimal weight vector) found so far by particle k .
- $gbest$: The best position found so far by any particle in the entire swarm.
- ω : The inertia weight, balancing exploration and exploitation.
- c_1 and c_2 : The acceleration coefficients, influencing how strongly particles are attracted to their personal best and the global best positions.
- r_1 and r_2 : Random numbers drawn uniformly from $[0, 1]$, introducing stochasticity into the search.

Through iterative cooperation and information sharing among particles, PSO searches for the optimal network weight configuration that minimizes the training objective.

2.4.2 Selective Weight Adoption Mechanism

To promote stable convergence and prevent performance fluctuations during the hybrid optimization process, an adaptive weight selection mechanism is applied at each iteration.

Specifically, after the PSO-based update, the candidate weight vector is evaluated on a validation set. If the updated parameters lead to a reduction in the validation loss, the new weight configuration is accepted and integrated into the network; otherwise, the previous weights are retained.

Formally, let $\mathcal{L}_{val}(W)$ denote the validation loss. The update rule is defined as:

$$W(t+1) = \begin{cases} W_{PSO}(t+1), & \text{if } \mathcal{L}_{val}(W_{PSO}(t+1)) < \mathcal{L}_{val}(W(t)), \\ W(t), & \text{otherwise.} \end{cases} \quad (10)$$

This selective acceptance strategy acts as a safeguard against detrimental PSO perturbations, enforcing monotonic improvement on the validation set and significantly enhancing the overall stability of the optimization process.

3. Results

3.1. Evaluation Metrics

To comprehensively evaluate the performance of the proposed PSO-DTI framework, five widely adopted evaluation metrics were employed: area under the receiver operating characteristic curve (AUROC), area under the precision–recall curve (AUPRC), Accuracy, F1-score, and Matthews correlation coefficient (MCC).

Among these metrics, AUROC and AUPRC are particularly important for drug–target interaction prediction due to the inherent class imbalance commonly observed in biological interaction datasets. While Accuracy may be biased toward the majority class, AUPRC provides a more informative assessment of predictive performance on positive interaction samples. In addition, MCC offers a balanced evaluation by considering all elements of the confusion matrix, including true positives, true negatives, false positives, and false negatives.

Beyond standard performance reporting, the training stability of the proposed model and the statistical significance of performance improvements over baseline methods were systematically analyzed to ensure robustness and reproducibility of the results. To ensure a fair comparison, all methods were trained and evaluated under identical hardware and software conditions, using an NVIDIA RTX 3080 GPU, a batch size of 128, and 50 training epochs.

3.2. Baseline Methods

To assess the effectiveness of the proposed PSO-DTI model, experiments were conducted on the *C. elegans* and Human benchmark datasets. For a fair comparison, all baseline methods were reimplemented and trained using the same similarity-transformed input features as the proposed model. Furthermore, to ensure a fair and consistent performance evaluation, a 5-fold cross-validation protocol was adopted on each dataset. Specifically, the entire drugtarget interaction dataset, including both known interacting pairs (positive samples) and unknown or non-interacting pairs (negative samples), was randomly partitioned into five folds at the interaction-pair level. In each iteration, four folds were used for model development, while the remaining fold was reserved for testing. Within the training portion, a validation subset was further utilized to guide the PSO-based selective weight adoption mechanism. This evaluation protocol was applied consistently across all compared methods. In total, eight representative baseline methods were selected, each employing recent deep learning architectures for drug–target interaction prediction:

- **CPI-GNN**: Employs graph neural networks for drug molecular graph encoding and convolutional neural networks for protein sequence modeling [15].
- **BACPI**: Integrates CNN-based feature extraction with dual attention mechanisms to capture compound–protein interaction patterns [18].
- **CPGL**: Utilizes graph attention networks combined with LSTM modules to enhance robustness in feature learning [19].
- **BINDTI**: Encodes molecular graphs and protein representations using advanced attention-based fusion mechanisms [20].
- **FOTF-CPI**: Enhances transformer architectures through optimized fragmentation strategies for interaction prediction [21].
- **CAT-DTI**: Applies CNN- and transformer-based encoders with cross-attention mechanisms for interaction modeling [22].
- **DO-GMA**: Combines CNN and GCN encoders with gated multi-head attention and bilinear feature fusion [23].
- **CAMF-DTI**: Incorporates coordinate attention and multi-scale feature fusion to achieve refined representation learning [24].

All methods were evaluated using the same experimental protocol and performance metrics, including AUROC, AUPRC, Accuracy, F1-score, and MCC, to ensure a fair and comprehensive comparison.

3.3. Experimental Results

3.3.1 Performance on the *C. elegans* Dataset

As summarized in Table 1, PSO-DTI achieved the highest performance across all evaluation metrics on the *C. elegans* dataset. Compared with the strongest baseline (DO-GMA), PSO-DTI improved AUROC and AUPRC by 0.2%, Accuracy and F1-score by 0.4%, and MCC by 0.7%, demonstrating enhanced stability in binary classification. Paired *t*-tests confirmed that all improvements were statistically significant ($p < 0.0001$), highlighting the robustness of the proposed method.

These improvements indicate that integrating particle swarm optimization into neural network weight training effectively enhances global exploration capability, mitigates convergence to local minima, and results in improved generalization.

To further analyze the practical applicability of the proposed method, we also evaluated its computational efficiency on the *C. elegans* dataset. As shown in Table 3, the runtime of PSO-DTI was compared with DO-GMA, which is one of the strongest baseline models and incorporates multiple computationally intensive components, including CNN, attention mechanisms, GCN, and MLP layers. The results show that PSO-DTI requires only 7 minutes for training, while DO-GMA requires approximately 40 minutes under the same experimental conditions. In terms of inference efficiency, PSO-DTI needs 0.04 seconds per sample, whereas DO-GMA requires 0.18 seconds per sample. These results indicate that PSO-DTI achieves superior predictive performance while maintaining significantly lower computational cost.

3.3.2 Performance on the Human Dataset

On the Human dataset, PSO-DTI similarly outperformed all baseline models. Relative to DO-GMA, PSO-DTI improved AUROC by 0.3%, AUPRC by 0.2%, Accuracy and F1-score by 0.5%, and MCC by 1.2%, confirming superior robustness in balanced interaction prediction. Paired *t*-tests again showed that these gains were statistically significant ($p < 0.0001$), as shown in Table 2.

Overall, these results confirm that the swarm-based optimization strategy provides consistent and stable improvements across heterogeneous DTI datasets, demonstrating the effectiveness and generalizability of the PSO-DTI model.

In addition to predictive performance, we also analyzed the computational cost of the proposed model on the Human dataset. As shown in Table 4, PSO-DTI was compared with DO-GMA under identical hardware and training settings. The results show that PSO-DTI requires approximately 8 minutes for training, whereas DO-GMA requires about 46 minutes. Furthermore, the evaluation time of PSO-DTI is 0.05 seconds per sample, compared with 0.18 seconds per sample for DO-GMA. These results further confirm that PSO-DTI not only improves prediction accuracy but also offers a more efficient computational framework for drug–target interaction prediction.

Table 1. Performance comparison on the *C. elegans* dataset.

Method	AUROC		AUPRC		Accuracy		F1-score		MCC	
CPI-GNN	0.986	0.0015	0.986	0.0014	0.949	0.0021	0.918	0.0023	0.829	0.0025
BACPI	0.986	0.0014	0.986	0.0013	0.949	0.0020	0.948	0.0019	0.933	0.0018
CPGL	0.986	0.0016	0.986	0.0015	0.928	0.0024	0.928	0.0022	0.853	0.0026
BINDTI	0.982	0.0018	0.983	0.0017	0.966	0.0019	0.966	0.0018	0.932	0.0017
FOTF-CPI	0.990	0.0013	0.990	0.0012	0.966	0.0018	0.966	0.0017	0.932	0.0016
CAT-DTI	0.983	0.0017	0.986	0.0015	0.967	0.0019	0.964	0.0020	0.932	0.0018
DO-GMA	0.993	0.0008	0.993	0.0008	0.974	0.0016	0.973	0.0017	0.948	0.0015
CAMF-DTI	0.987	0.0015	0.984	0.0016	0.948	0.0022	0.949	0.0021	0.895	0.0024
PSO-DTI (Ours)	0.995	0.0012	0.995	0.0012	0.978	0.0030	0.977	0.0020	0.955	0.0030

Table 2. Performance comparison on the Human dataset.

Method	AUROC		AUPRC		Accuracy		F1-score		MCC	
CPI-GNN	0.967	0.0016	0.966	0.0015	0.907	0.0023	0.906	0.0022	0.834	0.0024
BACPI	0.967	0.0015	0.967	0.0014	0.905	0.0024	0.907	0.0023	0.835	0.0023
CPGL	0.968	0.0016	0.967	0.0015	0.902	0.0025	0.904	0.0024	0.832	0.0026
BINDTI	0.981	0.0013	0.976	0.0014	0.940	0.0019	0.938	0.0020	0.879	0.0019
FOTF-CPI	0.983	0.0012	0.980	0.0013	0.941	0.0018	0.932	0.0019	0.881	0.0018
CAT-DTI	0.982	0.0013	0.969	0.0015	0.942	0.0018	0.944	0.0019	0.886	0.0018
DO-GMA	0.986	0.0010	0.984	0.0008	0.950	0.0016	0.951	0.0017	0.900	0.0015
CAMF-DTI	0.987	0.0012	0.984	0.0013	0.948	0.0020	0.949	0.0021	0.895	0.0022
PSO-DTI (Ours)	0.989	0.0012	0.986	0.0012	0.955	0.0030	0.956	0.0020	0.912	0.0030

Table 3. Runtime comparison of PSO-DTI and DO-GMA on the *C. elegans* dataset

Method	Training Time (min)	Evaluation Time (sec/sample)
PSO-DTI	7	0.04
DO-GMA	40	0.18

3.3.3 Training Stability

The training stability of PSODTI was evaluated by monitoring the AUROC score over 50 training epochs, as illustrated in Figures 2 and 3. The model demonstrates rapid and steady convergence, with the AUROC increasing from 0.781 at the first epoch to 0.996 at epoch 50. Notably, PSODTI reaches nearoptimal performance around epoch 40, after which only minor fluctuations

Table 4. Runtime comparison of PSO-DTI and DO-GMA on the Human dataset

Method	Training Time (min)	Evaluation Time (sec/sample)
PSO-DTI	8	0.05
DO-GMA	46	0.18

are observed. This behavior indicates that the model maintains stable learning dynamics during the later stages of training. The observed convergence pattern suggests that integrating particle swarm optimization contributes to a more robust optimization process. By complementing conventional gradient-based learning, PSO helps guide the model toward more suitable regions of the parameter space, thereby reducing sensitivity to initialization and decreasing the likelihood of convergence to suboptimal local minima. To further examine the contribution of the PSO component, a neural network trained solely using gradient descent (BaseNN) was included in the ablation study (Section 3.4), representing a gradient-only optimization scenario without PSO. The comparative results presented in Tables 5 and 6 show that PSODTI consistently outperforms this baseline model, indicating that the PSOenhanced optimization strategy contributes to improved predictive performance and more effective model training.

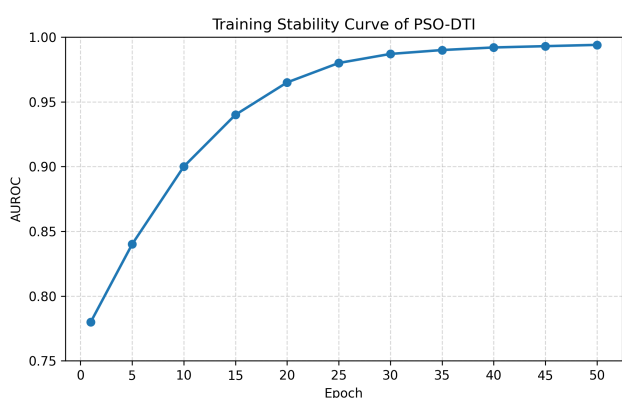


Figure 2. Training stability of PSO-DTI on the *C. elegans* dataset. AUROC increases steadily and reaches near-optimal performance by epoch 40.

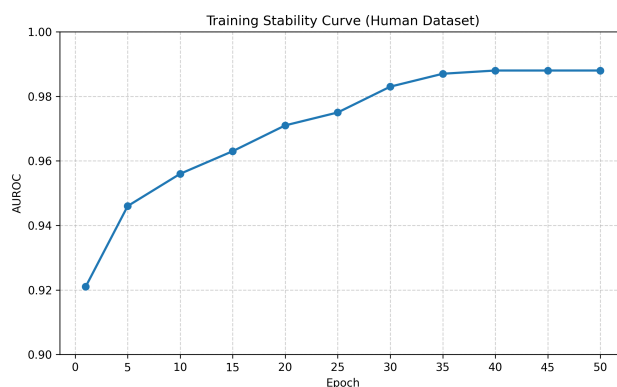


Figure 3. Training stability of PSO-DTI on the Human dataset. AUROC increases steadily during training and reaches near-optimal performance by epoch 40, demonstrating stable convergence.

3.4. Ablation Study

To investigate the contribution of individual components within the proposed PSO-DTI framework, ablation experiments were conducted on both the *C. elegans* and Human datasets. Three model variants were evaluated: (i) a base neural network trained solely using gradient descent (Base-NN), (ii) a neural network incorporating similarity-based standardized inputs without PSO optimization (NN + Similarity), and (iii) the complete PSO-DTI framework integrating both similarity construction and PSO-enhanced weight optimization.

On the *C. elegans* dataset, incorporating similarity-based feature standardization consistently improved predictive performance over the base neural model, highlighting the effectiveness of similarity transformation in enhancing representation consistency across heterogeneous drug and target inputs. Further integrating PSO-based weight optimization resulted in the highest performance across all evaluation metrics. This improvement indicates that the swarm-based optimization mechanism significantly enhances global search capability, mitigates convergence to local minima, and stabilizes the learning process. Detailed results for this dataset are summarized in Table 5.

A similar trend was observed on the Human dataset. While similarity-based preprocessing led to noticeable gains in predictive accuracy and robustness, the most substantial improvements were achieved after incorporating PSO. The complete PSO-DTI model consistently outperformed both the base neural network and the similarity-only variant. These results demonstrate that PSO not only refines learned representations but also improves decision boundary stability and reduces misclassification across both positive and negative interaction classes. The corresponding results for the Human dataset are reported in Table 6. Overall, these findings confirm that the combined use of similarity-based feature standardization and PSO-enhanced optimization yields consistent and robust performance gains across heterogeneous DTI datasets.

3.5. Case Study

To further evaluate the biological relevance and practical applicability of the proposed PSO-DTI framework, case studies were conducted on both the Human and *C. elegans* datasets, focusing on top-ranked novel drug-target interaction predictions.

Table 5. Ablation study results on the *C. elegans* dataset.

Model Variant	AUROC	AUPRC	Accuracy	F1-score	MCC
Base-NN	0.988	0.987	0.969	0.968	0.939
NN + Similarity	0.991	0.990	0.972	0.971	0.944
PSO-DTI	0.995	0.995	0.978	0.977	0.955

Table 6. Ablation study results on the Human dataset.

Model Variant	AUROC	AUPRC	Accuracy	F1-score	MCC
Base-NN	0.982	0.979	0.944	0.943	0.889
NN + Similarity	0.986	0.983	0.949	0.948	0.897
PSO-DTI	0.989	0.986	0.955	0.956	0.912

These analyses aim to assess the model's ability to identify pharmacologically meaningful and biologically plausible molecular relationships beyond the interactions observed during training.

3.5.1 *C. elegans* Dataset

A similar validation strategy was applied to the *C. elegans* dataset. As in the Human dataset, all candidate interactions evaluated in this analysis were excluded from the training data to ensure an unbiased assessment of novel predictions. Several top-ranked drug–target pairs were associated with proteins involved in essential metabolic and regulatory pathways, highlighting the model's ability to generalize across species.

For example, PSO-DTI predicted a high-confidence interaction between Imatinib and ABL-1, the *C. elegans* ortholog of the human ABL kinase. Given the conserved role of ABL signaling in cellular regulation, this prediction is biologically plausible and illustrates the model's capacity to capture functionally conserved drug–target interactions.

Another significant prediction involved Gefitinib and LET-23, the *C. elegans* homolog of the human epidermal growth factor receptor (EGFR). As Gefitinib is a known EGFR inhibitor, the identification of this cross-species interaction further supports the ability of PSO-DTI to recognize conserved molecular interaction patterns.

3.5.2 Human Dataset

Following training on known interactions, all unknown drug–protein pairs in the Human dataset were ranked according to their predicted interaction probabilities. The top twenty highest-confidence predictions were selected for biological validation using literature evidence and publicly available resources, including DrugBank and PubMed. To prevent information leakage, all evaluated drug–target pairs were strictly excluded from the training set during model development. Consequently, the ranked interactions represent previously unobserved associations from the model's perspective.

Among the top-ranked predictions, the interaction between Metformin and PRKAA1 (AMPK) was identified with high confidence. Extensive studies have demonstrated that Metformin activates AMPK signaling pathways, thereby regulating glucose metabolism and cellular energy homeostasis. The successful prioritization of this interaction indicates that PSO-DTI can recover pharmacologically relevant relationships that are well supported by existing biological evidence.

Another notable prediction involved Sorafenib and KDR (VEGFR2). Sorafenib is a well-characterized multi-kinase inhibitor with established activity against VEGFR2 in cancer therapy. The model's ability to rank this interaction among the top candidates further demonstrates its effectiveness in identifying therapeutically significant kinase-related drug–target associations.

3.5.3 Implications of Case Study Findings

Collectively, the presented case studies suggest that PSO-DTI is capable of prioritizing biologically plausible drug–target interactions that extend beyond the explicitly known training data. The high-confidence predictions reflect underlying pharmacological relevance and evolutionary conservation, as supported by independent literature evidence.

While the integration of swarm-based optimization is likely to contribute to more stable ranking behavior and improved global search capability, further controlled experiments would be required to isolate its specific impact on ranking consistency. Overall, these findings highlight the potential of PSO-DTI as a computational screening tool for prioritizing candidate drug–target pairs for subsequent experimental validation, thereby supporting translational research and cross-species interaction studies.

It should be noted that, although literature-based validation supports the biological plausibility of several top-ranked predictions, experimental verification remains necessary to confirm novel interactions. The case studies presented here therefore serve as evidence of biological consistency rather than definitive experimental confirmation.

4. Discussion

The results obtained on both the Human and *C. elegans* datasets indicate that incorporating Particle Swarm Optimization (PSO) into neural network training yields consistent performance gains in drug–target interaction (DTI) prediction. Although the numerical improvements over strong baseline models are moderate, they are observed systematically across multiple evaluation metrics, including AUROC, AUPRC, Accuracy, F1-score, and MCC. This consistency suggests that the contribution of PSO is not confined to a single performance dimension, but rather reflects broader improvements in model robustness and decision boundary stability.

From an optimization perspective, conventional deep learning approaches for DTI prediction rely primarily on gradient-based methods, which can be sensitive to parameter initialization and may exhibit unstable convergence behavior in high-dimensional and sparse biological settings. By introducing a swarm-based search mechanism, the proposed framework expands exploration of the parameter space beyond local gradient trajectories. Rather than replacing gradient descent entirely, PSO functions as a complementary global search heuristic that can reduce the likelihood of premature convergence to suboptimal solutions. The selective weight adoption mechanism further stabilizes training by accepting only those parameter updates that improve validation performance, thereby mitigating detrimental oscillations during optimization.

Another key aspect of the proposed framework lies in the similarity-based relational representation of drugs and targets. Transforming raw molecular fingerprints and protein sequence descriptors into similarity profiles enables the model to incorporate global structural context within each modality. This relational embedding strategy effectively positions each entity within its neighborhood of biologically related counterparts, potentially enhancing representation consistency across heterogeneous feature spaces. The ablation study results demonstrate that both similarity construction and PSO-based optimization independently contribute to predictive performance, while their combination produces the most pronounced improvements.

The conducted case studies further support the biological plausibility of high-confidence predictions generated by PSO-DTI. Several top-ranked drug–target interactions were consistent with well-established pharmacological mechanisms reported in the literature. Importantly, these interactions were excluded from the training process, indicating genuine model generalization rather than memorization. While literature-based validation provides supportive evidence of biological consistency, experimental validation remains necessary to confirm novel predictions in real-world settings.

Despite these encouraging results, several limitations should be acknowledged. First, evaluation was conducted on two benchmark datasets which, although widely adopted, may not fully capture the scale and diversity of real-world DTI networks. Validation across additional organisms, larger interaction graphs, and temporally split datasets would provide a more comprehensive assessment of generalizability. Second, incorporating PSO introduces additional computational overhead compared with purely gradient-based optimization. Although training times remained manageable in the present study, scalability to substantially larger datasets warrants further investigation. Third, the current framework does not incorporate three-dimensional protein structural information, graph-based molecular representations, or multi-omics data, all of which could potentially enhance representational richness and predictive accuracy.

Future research directions include exploring hybrid optimization strategies that more tightly integrate PSO with adaptive gradient-based methods, enabling efficient global exploration followed by refined local optimization. Additionally, incorporating knowledge graph representations and structural biology information may further improve contextual reasoning in DTI prediction. Finally, systematic analysis of convergence behavior and training stability across multiple random seeds would provide deeper insight into the dynamics of swarm-enhanced neural optimization.

Overall, this study demonstrates that integrating global search heuristics into deep learning pipelines offers a viable and biologically informed strategy for improving robustness in computational drug–target interaction prediction. The proposed PSO-DTI framework provides a flexible foundation for future extensions toward more scalable and structurally enriched biomedical interaction modeling systems.

5. Conclusion

In this work, we introduced PSO-DTI, a swarm-enhanced neural framework for drug–target interaction prediction that integrates similarity-based relational representations with Particle Swarm Optimization–guided weight refinement. By combining structured similarity embedding with global search heuristics, the proposed approach offers an alternative training paradigm for neural DTI models beyond purely gradient-driven optimization.

Experimental evaluations on benchmark Human and *C. elegans* datasets demonstrated that PSO-DTI achieves consistent improvements in predictive performance and ranking quality compared with representative deep learning baselines. Ablation analyses confirmed that both similarity transformation and PSO-based optimization contribute meaningfully to model behavior, while their combination yields the strongest results. Furthermore, literature-supported case studies indicated that high-confidence

predictions produced by PSO-DTI are aligned with known pharmacological mechanisms, highlighting the framework's potential utility as a computational screening tool.

Rather than positioning PSO as a replacement for gradient descent, this study emphasizes the value of hybrid optimization perspectives in biomedical machine learning. The findings suggest that incorporating global search strategies into neural training can enhance robustness when modeling sparse and heterogeneous biological interaction data.

Future extensions toward structurally enriched representations, larger-scale datasets, and coordinated hybrid optimization schedules may further improve scalability and generalization performance. A potential direction for future research is extending the proposed framework to support additional feature representations, such as structure-based or graph-based encodings of drugs and proteins, which may enable evaluation on a wider range of DTI datasets. Overall, PSO-DTI contributes to the ongoing development of more stable, interpretable, and biologically informed computational approaches for prioritizing candidate drug–target interactions.

Acknowledgments

The authors gratefully acknowledge the valuable comments and careful evaluation provided by the anonymous reviewers, which helped to enhance the quality of this work.

References

1. D.B. Searls. Using bioinformatics in gene and drug discovery. *Drug Discovery Today*, 5(4):135–143, 2000.
2. R. Chen, X. Liu, S. Jin, J. Lin, and J. Liu. Machine learning for drug–target interaction prediction. *Molecules*, 23(9):2208, 2018.
3. R. Masumshah, R. Aghdam, and C. Eslahchi. A neural network-based method for polypharmacy side effects prediction. *BMC Bioinformatics*, 22(1):1–17, 2021.
4. R. Masumshah and C. Eslahchi. DPSP: A multimodal deep learning framework for polypharmacy side effects prediction. *Bioinformatics Advances*, 3(1):vbad110, 2023.
5. R. Masumshah and C. Eslahchi. PSO-FeatureFusion: A general framework for fusing heterogeneous features via particle swarm optimization. *Bioinformatics Advances*, 2025.
6. Y. Yamanishi, M. Araki, A. Gutteridge, W. Honda, and M. Kanehisa. Prediction of drug–target interaction networks from the integration of chemical and genomic spaces. *Bioinformatics*, 24(13):i232–i240, 2008.
7. H. Öztürk, A. Özgür, and E. Ozkirimli. DeepDTA: Deep drug–target binding affinity prediction. *Bioinformatics*, 34(17):i821–i829, 2018.
8. X. Qu, G. Du, J. Hu, and Y. Cai. Graph-DTI: A new model for drug–target interaction prediction based on heterogeneous network graph embedding. *Current Computer-Aided Drug Design*, 20(6):1013–1024, 2024.
9. Z. Meng, Z. Meng, K. Yuan, and I. Ounis. FusionDTI: Fine-grained binding discovery with token-level fusion for drug–target interaction. *arXiv preprint*, 2024.
10. J. Liu, Y. Lu, S. Guan, T. Jiang, Y. Ding, Q. Fu, Z. Cui, and H. Wu. Drug–target interaction prediction by combining transformer and graph neural networks. *Current Bioinformatics*, 19(4):316–326, 2024.
11. X. Li, H. Yang, K. Hu, R. Wu, L. Liu, and R. Su. CDI-DTI: A strong cross-domain interpretable drug–target interaction prediction framework based on multi-strategy fusion. *arXiv preprint arXiv:2510.19520*, 2025.
12. B. Ahmad, K. Ouahada, and H. Hamam. Machine learning for drug–target interaction prediction: A comprehensive review of models, challenges, and computational strategies. *Computational and Structural Biotechnology Journal*, 2026.
13. J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceedings of ICNN'95 – International Conference on Neural Networks*, volume 4, pages 1942–1948. IEEE, 1995.
14. R. Poli, J. Kennedy, and T. Blackwell. Particle swarm optimization: An overview. *Swarm Intelligence*, 1(1):33–57, 2007.
15. M. Tsubaki, K. Tomii, and J. Sese. Compound–protein interaction prediction with end-to-end learning of neural networks for graphs and sequences. *Bioinformatics*, 35(2):309–318, 2019.
16. L. Chen, X. Tan, D. Wang, F. Zhong, X. Liu, T. Yang, X. Luo, K. Chen, H. Jiang, and M. Zheng. TransformerCPI: Improving compound–protein interaction prediction by sequence-based deep learning with self-attention mechanism. *Bioinformatics*, 36:4406–4414, 2020.
17. P. Li, Y. Li, C.-Y. Hsieh, S. Zhang, X. Liu, H. Liu, S. Song, and X. Yao. TrimNet: learning molecular representation from triplet messages for drug–target interaction prediction. *Briefings in Bioinformatics*, 22(4):bbaa266, 2021.
18. M. Li, Z. Lu, Y. Wu, and Y. Li. BACPI: A bi-directional attention neural network for compound–protein interaction and binding affinity prediction. *Bioinformatics*, 38(7):1995–2002, 2022.
19. M. Zhao, M. Yuan, Y. Yang, and S.X. Xu. CPGL: Prediction of compound–protein interaction by integrating graph attention network with long short-term memory neural network. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 20(3):1935–1942, 2022.
20. L. Peng, X. Liu, L. Yang, L. Liu, Z. Bai, M. Chen, X. Lu, and L. Nie. BINDTI: A bi-directional intention network for drug–target interaction identification based on attention mechanisms. *IEEE Journal of Biomedical and Health Informatics*, 29(3):1602–1612, 2024.
21. Z. Yin, Y. Chen, Y. Hao, S. Pandiyan, J. Shao, and L. Wang. FOTF-CPI: A compound–protein interaction prediction transformer based on the fusion of optimal transport fragments. *iScience*, 27(1), 2024.
22. Y. Qian, J. Wu, and Q. Zhang. CAT-CPI: Combining CNN and transformer to learn compound image features for predicting compound–protein interactions. *Frontiers in Molecular Biosciences*, 9:963912, 2022.

23. L. Peng, J. Mao, G. Huang, G. Han, X. Liu, W. Liao, G. Tian, and J. Yang. DO-GMA: An end-to-end drug–target interaction identification framework with a depthwise overparameterized convolutional network and the gated multihead attention mechanism. *Journal of Chemical Information and Modeling*, 65(3):1318–1337, 2025.
24. J. Mi, C. Li, D. Jiang, and J. Wan. CAMF-DTI: Enhancing drug–target interaction prediction via coordinate attention and multi-scale feature fusion. *Current Issues in Molecular Biology*, 47(11):964, 2025.