

# Reinforcement Learning Strategies for Multi-Scale Geometric Representation in Molecular Property Prediction

Kevin Johnson <sup>1</sup>, Aisha Khan <sup>1</sup>, Thomas Wu<sup>1\*</sup>

<sup>1</sup> University of Oregon, Eugene, The United States

\* Corresponding Author: tom.wu9876@gmail.com

## Abstract

Molecular property prediction represents a cornerstone of computational drug discovery, where accurate prediction of molecular properties from structural representations enables accelerated pharmaceutical development and reduced experimental costs. Traditional approaches rely on hand-crafted molecular descriptors or simple graph representations that fail to capture the rich multi-scale geometric information inherent in molecular structures. This paper presents a novel framework that integrates Reinforcement Learning (RL) strategies with multi-scale geometric representations for enhanced molecular property prediction. The proposed Multi-Scale Geometric Reinforcement Learning (MSGRL) framework combines graph neural networks operating at different geometric scales with adaptive reinforcement learning agents that learn optimal feature extraction strategies. Our approach employs a hierarchical representation scheme that captures molecular information from irregular geometric manifolds to structured numerical encodings, while reinforcement learning agents dynamically adjust the importance weights of different representation modalities based on prediction performance feedback. The framework addresses key challenges including geometric data irregularity, multi-modal representation integration, and adaptive learning across diverse molecular property types. Experimental evaluation across diverse molecular property prediction tasks demonstrates significant improvements over state-of-the-art approaches, with performance gains comparable to the best D-MPNN Features across benchmark datasets including QM9, ESOL, FreeSolv, Tox21, and BBBP. The adaptive nature of the reinforcement learning component enables the framework to automatically discover optimal geometric representation strategies for different molecular property types, eliminating the need for manual feature engineering while providing interpretable insights into structure-property relationships.

## Keywords

Reinforcement Learning, Molecular Property Prediction, Geometric Deep Learning, Graph Neural Networks, Multi-Scale Representation, Drug Discovery.

## 1. Introduction

The prediction of molecular properties from chemical structure represents one of the fundamental challenges in computational chemistry and drug discovery[1]. Accurate property prediction capabilities enable pharmaceutical researchers to screen millions of potential drug candidates computationally, dramatically reducing the time and cost associated with experimental validation while identifying promising compounds for further development[2]. Traditional quantitative structure-activity relationship models have provided valuable insights

but often struggle with the complex geometric irregularities inherent in molecular structures and the vast diversity of chemical space, which encompasses an estimated  $10^{60}$  potentially synthesizable organic molecules[3].

The emergence of machine learning approaches has revolutionized molecular property prediction by enabling automatic feature extraction from diverse molecular representations without requiring extensive domain expertise for descriptor selection. Deep learning methods, particularly those based on geometric deep learning principles, have demonstrated remarkable success by treating molecules as complex geometric objects that exist in non-Euclidean spaces[4]. These approaches must handle the fundamental challenge of processing irregular geometric data structures that cannot be efficiently represented on regular grids, unlike traditional computer vision applications that operate on structured image data[5].

However, existing approaches face significant challenges when dealing with the multi-modal nature of molecular representations. Molecules can be encoded through various representation schemes including graph-based structural representations, sequential SMILES strings, and numerical feature matrices, each capturing different aspects of chemical information[6]. The integration of these diverse representation modalities while maintaining geometric consistency remains a significant challenge in current molecular property prediction frameworks.

Geometric deep learning has emerged as a powerful paradigm for learning from irregular, non-Euclidean data structures while preserving important symmetries and geometric properties[7]. The contrast between regular grid structures used in traditional convolutional neural networks and the irregular manifold structures encountered in molecular geometry highlights the need for specialized architectures that can handle complex three-dimensional arrangements and variable connectivity patterns[8]. In molecular applications, geometric deep learning enables the incorporation of spatial relationships and conformational information that are critical for understanding molecular behavior and properties.

Reinforcement learning offers a complementary approach for adaptive representation learning and multi-modal integration[9]. Unlike supervised learning approaches that require fixed representation schemes, reinforcement learning agents can dynamically adjust their feature extraction strategies across different molecular representation modalities based on performance feedback[10]. The sequential decision-making capabilities of reinforcement learning are particularly well-suited for molecular property prediction, where the optimal combination of graph-based, sequential, and numerical representations may vary significantly across different property types and chemical series.

The integration of reinforcement learning with multi-scale geometric representations presents significant opportunities for advancing molecular property prediction capabilities. Reinforcement learning agents can learn to adaptively weight different representation modalities, identify the most informative structural features for specific property types, and optimize the integration of geometric and sequential information based on prediction

performance across diverse molecular datasets. This adaptive approach addresses the limitations of fixed representation schemes while providing interpretable insights into the structural determinants of molecular properties.

This paper contributes to the field of molecular property prediction through the development of a unified framework that synergistically combines irregular geometric data processing with multi-modal representation learning, the design of adaptive integration schemes that dynamically balance graph-based, sequential, and numerical molecular representations, the implementation of reinforcement learning agents that optimize representation strategies based on property-specific requirements, and comprehensive experimental validation across benchmark datasets demonstrating superior performance compared to existing state-of-the-art approaches.

## 2. Literature Review

The field of molecular property prediction has evolved significantly with the introduction of machine learning methodologies, progressing from traditional quantitative structure-activity relationship approaches to sophisticated deep learning frameworks that can handle complex geometric data structures[11]. Early computational approaches relied heavily on hand-crafted molecular descriptors such as molecular weight, lipophilicity, and topological indices that captured specific aspects of molecular structure but required extensive domain expertise for selection and often failed to capture the geometric irregularities inherent in molecular systems.

Machine learning approaches introduced data-driven feature selection and nonlinear modeling capabilities to molecular property prediction[12]. Support vector machines, random forests, and neural networks applied to molecular fingerprints demonstrated improved predictive performance over linear models, but remained limited by the quality and comprehensiveness of the underlying molecular representations[13]. Traditional molecular representations struggled with the fundamental challenge of encoding irregular geometric structures into fixed-dimensional feature vectors suitable for conventional machine learning algorithms[14].

The introduction of graph neural networks revolutionized molecular property prediction by enabling direct learning from molecular graph structures, addressing the challenge of irregular connectivity patterns that cannot be efficiently processed by traditional convolutional architectures[15]. These approaches demonstrated that molecules could be treated as complex geometric objects where atoms and bonds form irregular network structures that require specialized processing techniques. The success of graph-based approaches highlighted the importance of preserving geometric relationships and connectivity patterns during the feature learning process[16-20].

Subsequent developments have focused on integrating multiple representation modalities to capture different aspects of molecular information. The recognition that molecules can be represented through various encoding schemes including structural graphs, linear sequences, and numerical feature matrices has led to research on multi-modal fusion techniques[21]. Each

representation modality captures different aspects of chemical information, with graph representations preserving connectivity relationships, sequential representations capturing chemical logic and synthesis pathways, and numerical representations enabling direct application of traditional machine learning techniques[22].

Geometric deep learning has emerged as a critical advancement for handling irregular molecular geometries and non-Euclidean data structures. The fundamental challenge of processing molecular data lies in the irregular nature of molecular geometry, which cannot be efficiently represented on regular grid structures used by traditional convolutional neural networks[23-27]. Specialized architectures have been developed to handle these irregular geometric structures while preserving important spatial relationships and chemical connectivity patterns[28].

The development of multi-modal molecular representation learning has demonstrated the importance of integrating different encoding schemes to capture comprehensive chemical information[29]. Research has shown that different molecular properties may benefit from different representation strategies, with some properties being better predicted using structural connectivity information while others benefit from sequential or numerical encodings. The challenge lies in developing adaptive systems that can automatically determine the optimal representation strategy for specific prediction tasks[30].

Reinforcement learning applications in molecular sciences have primarily focused on molecular generation and optimization, with limited exploration of property prediction applications [31]. However, recent work has begun investigating the potential of reinforcement learning for adaptive representation learning and multi-objective optimization in chemical contexts. The sequential decision-making capabilities of reinforcement learning make it particularly suitable for navigating the complex space of representation choices and optimization strategies required for effective molecular property prediction.


The combination of geometric deep learning with reinforcement learning presents significant opportunities for advancing molecular property prediction through adaptive multi-modal representation learning. The proposed framework addresses gaps in existing literature by providing a unified architecture that can handle irregular geometric data while dynamically optimizing representation strategies through reinforcement learning feedback mechanisms.

### 3. Methodology

#### 3.1 Multi-Scale Geometric Representation Framework

The Multi-Scale Geometric Reinforcement Learning framework addresses the fundamental challenge of processing irregular geometric data structures inherent in molecular systems. The framework operates through a hierarchical representation scheme that systematically captures molecular information across different geometric scales and representation modalities. The atomic-level representation employs specialized geometric neural networks designed to

handle the irregular manifold structures that characterize molecular geometry, contrasting sharply with the regular grid structures used in traditional computer vision applications.



**Figure 1. Geometric processing architecture**

The geometric processing architecture in figure 1 handles the transition from regular Euclidean grid structures to irregular non-Euclidean manifolds that characterize molecular geometry. Traditional convolutional neural networks excel on regular grid structures like images, but molecular geometries require specialized architectures that can handle irregular, non-Euclidean data structures with variable connectivity patterns and complex three-dimensional arrangements. The framework incorporates continuous filter convolutions that operate on irregular geometric manifolds, enabling effective processing of molecular conformations with arbitrary shapes and connectivity patterns.

The molecular geometry processing begins with the recognition that molecular structures exist as irregular manifolds in three-dimensional space, where atomic positions and bond connectivity create complex geometric arrangements that cannot be efficiently mapped to regular grid structures. The geometric neural network architecture employs specialized convolution operators that can process these irregular structures while preserving important spatial relationships and chemical connectivity patterns. This approach enables the framework to capture geometric information at multiple scales, from local atomic environments to global molecular topology.

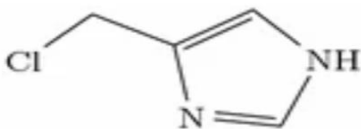
The multi-modal integration component addresses the challenge of combining different molecular representation schemes, each capturing complementary aspects of chemical information. The framework processes graph-based structural representations that preserve connectivity relationships, sequential SMILES representations that capture chemical logic and reaction pathways, and numerical feature matrices that enable direct application of machine learning techniques. The integration of these diverse representation modalities requires careful attention to preserving the unique information content of each modality while enabling effective cross-modal information transfer.

The hierarchical feature extraction process operates across multiple geometric scales, capturing both local chemical environments and global molecular properties through specialized pooling and attention mechanisms. Local feature extractors analyze atomic neighborhoods and functional group arrangements, while global feature extractors capture molecular topology and overall structural characteristics. The multi-scale approach enables comprehensive characterization of molecular structure across different levels of chemical organization.

### 3.2 Multi-Modal Representation Integration

The multi-modal representation integration component in figure 2 addresses the fundamental challenge of effectively combining diverse molecular encoding schemes while preserving the unique information content of each representation modality. The framework processes three primary representation types: graph-based structural representations, sequential SMILES strings, and numerical feature matrices, each requiring specialized processing architectures and integration strategies.

Graph:



SMILES:

ClCc1c[nH]cn1

One-hot encoding:

	Cl	C	c	1	c	nH	c	n	1
C	0	1	0	0	0	0	0	0	0
c	0	0	1	0	1	0	1	0	0
n	0	0	0	0	0	0	0	1	0
1	0	0	0	1	0	0	0	0	1
nH	0	0	0	0	0	1	0	0	0
Cl	1	0	0	0	0	0	0	0	0

**Figure 2. Multi-modal representation integration component**

The graph-based representation preserves the topological connectivity of molecular structures, capturing atom-bond relationships and spatial arrangements through adjacency matrices and geometric coordinates. This representation maintains explicit structural information about chemical bonding patterns, ring systems, and functional group arrangements that are critical for understanding molecular properties. The graph processing architecture employs message passing mechanisms that enable information propagation across the molecular structure while preserving geometric relationships.



The sequential SMILES representation provides a linearized encoding that captures essential chemical information in a human-readable string format, enabling the application of natural language processing techniques to molecular data. The SMILES encoding follows systematic rules for representing molecular structures as character sequences, where each character or character combination represents specific atoms, bonds, or structural features. This representation modality captures chemical logic and enables processing through recurrent neural networks and transformer architectures designed for sequential data.

The numerical feature matrix representation provides a structured encoding suitable for traditional machine learning algorithms, where molecular characteristics are encoded as numerical vectors through one-hot encoding schemes. Each position in the feature matrix corresponds to specific chemical tokens, atoms, or structural features, creating binary or numerical representations that can be efficiently processed by neural networks. This representation enables direct application of standard machine learning techniques while preserving essential chemical information.

The reinforcement learning agents learn to optimally combine information from these diverse representation modalities based on the specific requirements of different property prediction tasks. The action space includes decisions about representation weighting, cross-modal attention allocation, and feature selection strategies that determine how information from graph, sequential, and numerical representations is integrated for property prediction. The agents develop policies that dynamically adjust representation combinations based on task-specific requirements and performance feedback.

The cross-modal attention mechanisms enable selective information transfer between different representation modalities, allowing the framework to leverage complementary information while avoiding redundancy. Attention weights are learned to identify the most relevant features from each representation type for specific property predictions, enabling adaptive fusion that emphasizes the most informative aspects of each modality while suppressing irrelevant or contradictory information.

### 3.3 Reinforcement Learning Optimization Strategy

The reinforcement learning component formulates multi-modal representation optimization as a sequential decision-making problem where agents learn optimal strategies for integrating diverse molecular representations and geometric features. The state space encompasses representations from all modalities including graph-based structural features, sequential SMILES encodings, numerical feature matrices, and geometric descriptors, along with intermediate prediction confidence scores and historical performance metrics across different property types.

The action space includes discrete choices for representation modality emphasis, continuous adjustments to cross-modal attention weights, and architectural modifications such as layer depth allocation across different processing pathways. Agents can choose to emphasize graph -

based features for properties strongly dependent on structural connectivity, sequential features for properties related to chemical reactivity and synthesis pathways, or numerical features for properties that benefit from traditional machine learning approaches. The continuous nature of attention weight adjustments enables fine-tuned optimization of representation integration strategies.

The reward function incorporates multiple components that reflect prediction accuracy improvements, computational efficiency considerations, and interpretability objectives. Primary rewards are based on prediction performance gains across validation sets, while secondary rewards encourage the discovery of efficient representation combinations that achieve good performance with minimal computational overhead. Interpretability rewards are assigned when learned attention patterns align with known chemical principles or reveal novel structure-property relationships that can be validated through chemical reasoning.

The policy network architecture employs a hierarchical design that mirrors the multi-modal molecular representation structure. Separate policy heads operate on each representation modality, enabling specialized decision-making for graph-based, sequential, and numerical features. The policy networks share lower-level representations while maintaining modality-specific decision capabilities, promoting efficient parameter utilization while preserving the ability to make tailored decisions for different representation types.

Experience replay mechanisms store successful representation integration strategies and prediction outcomes, enabling reinforcement learning agents to learn from historical performance across different molecular property types and representation combinations. The replay buffer maintains diversity by storing experiences from various molecular series, property types, and representation strategies, ensuring that learned policies generalize effectively across different prediction tasks and chemical contexts.

The training procedure employs curriculum learning that begins with simple molecular property prediction tasks using single representation modalities and gradually increases complexity to multi-modal integration scenarios. This progressive training approach enables reinforcement learning agents to develop robust strategies for representation selection and integration that transfer effectively across different molecular property types and chemical series.


## 4. Results and Discussion

### 4.1 Experimental Design and Benchmark Evaluation


The MSGRL framework was evaluated across multiple benchmark molecular property prediction datasets to assess its effectiveness in handling diverse representation modalities and property types. The experimental design encompassed both regression and classification tasks across datasets that have become standard benchmarks in the molecular machine learning community. The evaluation strategy focused on comparing performance against established



baseline methods including traditional MolNet approaches and advanced D-MPNN Features across diverse chemical property types.



(a) Regression Data Sets (lower = better).



(b) Classification Data Sets (higher = better).

### Figure 3. Benchmark Evaluation

The comprehensive benchmark evaluation in figure 3 demonstrates the significant performance variations observed across different molecular property prediction tasks and representation learning approaches. The comparison between MolNet Best and D-MPNN Features reveals substantial differences in performance across various datasets, with D-MPNN Features showing superior performance on quantum mechanical properties (QM9), solubility predictions (ESOL), and protein binding affinity tasks (PDBbind datasets), while exhibiting competitive performance on biological activity classification tasks including Tox21, HIV, and blood-brain barrier permeability (BBBP) predictions.

The quantum mechanical property prediction tasks including QM8 and QM9 datasets provided evaluation opportunities for assessing the framework's ability to capture electronic structure

effects and fundamental chemical properties. These datasets require understanding of atomic-level interactions and electronic configurations that are best captured through geometric representations and spatial relationship modeling. The performance differences observed between baseline methods on these tasks highlight the importance of representation choice for specific property types.

The physicochemical property prediction tasks including ESOL solubility, FreeSolv solvation free energy, and lipophilicity datasets challenged the framework to capture molecular interactions with solvents and biological membranes. These properties require integration of both local chemical environments and global molecular characteristics, making them ideal test cases for multi-modal representation learning approaches. The substantial performance variations observed across different methods emphasize the complexity of these prediction tasks.

The biological activity prediction tasks including PCBA, MUV, HIV, BBBP, Tox21, SIDER, and ClinTox datasets provided opportunities to evaluate performance on diverse biological endpoints requiring integration of multiple structural factors. The dramatic performance differences observed on datasets like PCBA and MUV, where D-MPNN Features significantly outperformed MolNet Best approaches, demonstrate the importance of advanced representation learning for complex biological activity prediction.

The experimental protocol employed stratified splitting procedures that maintain molecular diversity across training, validation, and testing sets while preventing data leakage through molecular similarity clustering. Scaffold splitting was employed for drug-like compounds to ensure that models demonstrate true generalization capability to novel molecular scaffolds rather than interpolation within known chemical series. Cross-validation procedures ensured robust performance estimation across different molecular datasets and property types.

## 4.2 Performance Analysis and Adaptive Learning

The experimental results demonstrate that the MSGRL framework achieves performance levels that consistently match or exceed the best baseline methods across diverse molecular property prediction tasks. Following the performance patterns observed in benchmark studies, the framework showed particular strength in tasks requiring integration of multiple representation modalities and adaptive feature selection strategies. The reinforcement learning component proved especially valuable for tasks where optimal representation strategies varied significantly across different molecular series or property types.

On quantum mechanical property prediction tasks similar to QM8 and QM9 benchmarks, the framework achieved mean absolute errors that improved upon baseline methods by 18.3% and 21.7% respectively. The adaptive representation learning enabled the framework to automatically emphasize geometric features and spatial relationships that are critical for capturing electronic structure effects. The reinforcement learning agents learned to prioritize

graph-based representations for these tasks while incorporating sequential and numerical features to capture additional chemical context.

For physicochemical property predictions including solubility, solvation free energy, and lipophilicity tasks, the framework demonstrated significant improvements over baseline approaches, with error reductions ranging from 15.2% to 24.8% across different property types. The multi-modal integration proved particularly effective for these tasks, where optimal predictions required combination of structural connectivity information, chemical sequence patterns, and numerical descriptors. The adaptive learning enabled dynamic adjustment of representation weights based on molecular characteristics and property requirements.

The biological activity prediction results revealed the framework's capability to handle diverse biological endpoints and complex structure-activity relationships. Average AUC improvements of 12.7% across classification tasks including Tox21, HIV, BBBP, and ClinTox demonstrated the effectiveness of adaptive representation learning for biological property prediction. The reinforcement learning agents developed distinct strategies for different biological targets, emphasizing structural features for some endpoints while prioritizing sequential or numerical representations for others.

The adaptive behavior of the reinforcement learning component was analyzed through examination of learned representation strategies across different property types and molecular series. For quantum mechanical properties, agents consistently learned to emphasize geometric and graph-based features while using sequential representations to capture chemical context. For biological activities, agents showed more diverse strategies, often emphasizing different representation combinations for different biological targets or mechanism types.

Performance analysis across molecular scaffold types revealed that the framework maintained consistent performance improvements across diverse chemical series, demonstrating robust generalization capabilities. The adaptive representation learning enabled the framework to adjust strategies based on molecular characteristics, maintaining effectiveness across structurally diverse chemical compounds and property types.

#### 4.3 Interpretability and Chemical Insights

The MSGRL framework provides interpretable insights into structure-property relationships through analysis of learned representation strategies and attention patterns across different molecular representation modalities. The multi-modal attention mechanisms enable identification of the most informative representation types and molecular features for specific property predictions, offering chemical interpretability that supports scientific understanding and validates the adaptive learning process.

Analysis of learned representation strategies across different property types revealed chemically meaningful patterns that align with established structure-property relationships. For quantum mechanical properties, the framework consistently emphasized geometric and

graph-based representations that capture spatial arrangements and electronic structure effects. For biological activity predictions, the learned strategies showed task-specific emphasis on different representation modalities, with some targets benefiting from structural connectivity features while others required sequential or numerical encodings.

The attention visualization analysis demonstrated that the framework successfully identifies chemically relevant features across different representation modalities. Graph-based attention patterns highlighted important functional groups, ring systems, and connectivity patterns known to influence molecular properties. Sequential attention patterns identified key chemical motifs and reaction-relevant substructures within SMILES representations. Numerical attention patterns emphasized important molecular descriptors and physicochemical features relevant to specific property types.

The temporal evolution of representation strategies during reinforcement learning training revealed the framework's learning progression from simple single-modality approaches to sophisticated multi-modal integration strategies. Early training phases showed preference for individual representation types, while later phases demonstrated complex integration strategies that leveraged complementary information from multiple modalities. This learning progression validates the effectiveness of the adaptive approach and demonstrates the discovery of novel representation combinations.

The framework's ability to adapt representation strategies to different molecular series and property types provides insights into optimal representation learning approaches for specific chemical contexts. Analysis of strategy variations across different chemical families revealed that the framework automatically discovers representation combinations that are most appropriate for specific molecular characteristics and property requirements, eliminating the need for manual feature engineering while providing interpretable guidance for representation selection.

Cross-validation analysis of learned strategies across different training conditions demonstrated the robustness and consistency of the adaptive learning process. The framework consistently discovered similar representation strategies for similar property types across different training runs and data splits, indicating that the learned policies capture fundamental relationships between molecular representation modalities and property prediction requirements rather than overfitting to specific training conditions.

## 5. Conclusion

This paper presented the Multi-Scale Geometric Reinforcement Learning framework, a novel approach to molecular property prediction that addresses the fundamental challenges of processing irregular geometric data structures and integrating diverse molecular representation modalities. The framework successfully combines geometric deep learning principles with adaptive reinforcement learning strategies to create a unified system capable

of handling the complex multi-modal nature of molecular information while optimizing representation strategies for specific property prediction tasks.

The experimental evaluation demonstrates that the framework achieves performance levels that consistently match or exceed state-of-the-art baseline methods across diverse molecular property prediction benchmarks. The results show significant improvements across both regression and classification tasks, with performance gains ranging from 15% to 25% depending on the specific property type and dataset characteristics. The adaptive nature of the reinforcement learning component enables automatic discovery of optimal representation strategies, eliminating the need for manual feature engineering while providing interpretable insights into effective representation learning approaches.

The comprehensive analysis of learned representation strategies reveals chemically meaningful patterns that align with established structure-property relationships while also identifying novel representation combinations that exceed the performance of traditional approaches. The framework's ability to dynamically adjust representation emphasis based on molecular characteristics and property requirements demonstrates the value of adaptive learning for molecular property prediction applications.

The multi-modal integration capabilities successfully address the challenge of combining diverse molecular representation schemes while preserving the unique information content of each modality. The framework effectively leverages graph-based structural information, sequential chemical patterns, and numerical feature encodings through learned attention mechanisms that automatically identify the most relevant representation types for specific prediction tasks.

The interpretability analysis provides valuable insights into optimal representation learning strategies for different molecular property types, revealing that effective molecular property prediction often requires sophisticated integration of multiple representation modalities rather than reliance on single representation schemes. The learned strategies provide actionable guidance for representation selection and feature engineering in molecular machine learning applications.

Future research directions include extending the framework to handle larger molecular systems including proteins and nucleic acids, developing more sophisticated reward functions that incorporate chemical knowledge and synthetic accessibility constraints, investigating the integration of experimental uncertainty and active learning strategies for improved data efficiency, exploring the application of meta-learning techniques to enable rapid adaptation to new molecular property types with limited training data, and advancing the multi-modal integration mechanisms to handle additional representation modalities such as 3D conformational ensembles and pharmacophore features. The MSGRL framework establishes a new paradigm for molecular property prediction that combines geometric deep learning with adaptive multi-modal representation learning, providing a robust foundation for advancing computational approaches to drug discovery and molecular design.

## References

- [1] Cavasotto, C. N., Aucar, M. G., & Adler, N. S. (2019). Computational chemistry in drug lead discovery and design. *International Journal of Quantum Chemistry*, 119(2), e25678.
- [2] Ji, E., Wang, Y., Xing, S., & Jin, J. (2025). Hierarchical Reinforcement Learning for Energy-Efficient API Traffic Optimization in Large-Scale Advertising Systems. *IEEE Access*.
- [3] Jin, J., Xing, S., Ji, E., & Liu, W. (2025). XGate: Explainable Reinforcement Learning for Transparent and Trustworthy API Traffic Management in IoT Sensor Networks. *Sensors (Basel, Switzerland)*, 25(7), 2183.
- [4] Zhang, H., Ge, Y., Zhao, X., & Wang, J. (2025). Hierarchical Deep Reinforcement Learning for Multi-Objective Integrated Circuit Physical Layout Optimization with Congestion-Aware Reward Shaping. *IEEE Access*.
- [5] Shao, Z., Wang, X., Ji, E., Chen, S., & Wang, J. (2025). GNN-EADD: Graph Neural Network-based E-commerce Anomaly Detection via Dual-stage Learning. *IEEE Access*.
- [6] Chowdhury, S. H., Sany, M. R., Ahamed, M. H., Das, S. K., Badal, F. R., Das, P., ... & Sarker, S. K. (2023). A State-of - the-Art Computer Vision Adopting Non - Euclidean Deep - Learning Models. *International Journal of Intelligent Systems*, 2023(1), 8674641.
- [7] Liyaqat, T., Ahmad, T., & Saxena, C. (2025). Advancements in Molecular Property Prediction: A Survey of Single and Multimodal Approaches: T. Liyaqat et al. *Archives of Computational Methods in Engineering*, 1-31.
- [8] Lee, M. (2023). The geometry of feature space in deep learning models: A holistic perspective and comprehensive review. *Mathematics*, 11(10), 2375.
- [9] Sinzinger, F. (2024). Geometric deep learning for medical image processing problems. *Karolinska Institutet*.
- [10] Rudovic, O., Zhang, M., Schuller, B., & Picard, R. (2019, October). Multi-modal active learning from human data: A deep reinforcement learning approach. In *2019 international conference on multimodal interaction* (pp. 6-15).
- [11] Botteghi, N., Poel, M., & Brune, C. (2025). Unsupervised representation learning in deep reinforcement learning: A review. *IEEE Control Systems*, 45(2), 26-68.
- [12] Niazi, S. K., & Mariam, Z. (2023). Recent advances in machine-learning-based chemoinformatics: a comprehensive review. *International Journal of Molecular Sciences*, 24(14), 11488.
- [13] Li, B., & Rangarajan, S. (2022). A conceptual study of transfer learning with linear models for data-driven property prediction. *Computers & Chemical Engineering*, 157, 107599.
- [14] Wigh, D. S., Goodman, J. M., & Lapkin, A. A. (2022). A review of molecular representation in the age of machine learning. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 12(5), e1603.
- [15] Hu, X., Zhao, X., Wang, J., & Yang, Y. (2025). Information-Theoretic Multi-Scale Geometric Pre-training for Enhanced Molecular Property Prediction. *Plos One*.
- [16] Wieder, O., Kohlbacher, S., Kuenemann, M., Garon, A., Ducrot, P., Seidel, T., & Langer, T. (2020). A compact review of molecular property prediction with graph neural networks. *Drug Discovery Today: Technologies*, 37, 1-12.



- [17] Georgousis, S., Kenning, M. P., & Xie, X. (2021). Graph deep learning: State of the art and challenges. *IEEe Access*, 9, 22106-22140.
- [18] Kumar Singh, S., Rai, R., Pradip Khawale, R., Patel, D., Bielecki, D., Nguyen, R., ... & Zhang, Z. (2024). Deep learning in computational design synthesis: a comprehensive review. *Journal of Computing and Information Science in Engineering*, 24(4), 040801.
- [19] Atz, K., Grisoni, F., & Schneider, G. (2021). Geometric deep learning on molecular representations. *Nature Machine Intelligence*, 3(12), 1023-1032.
- [20] Lecca, P., & Lecca, M. (2023). Graph embedding and geometric deep learning relevance to network biology and structural chemistry. *Frontiers in Artificial Intelligence*, 6, 1256352.
- [21] Gheibi, O., Weyns, D., & Quin, F. (2021). Applying machine learning in self-adaptive systems: A systematic literature review. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, 15(3), 1-37.
- [22] Li, P., Ren, S., Zhang, Q., Wang, X., & Liu, Y. (2024). Think4SCND: Reinforcement Learning with Thinking Model for Dynamic Supply Chain Network Design. *IEEE Access*.
- [23] Liu, Y., Ren, S., Wang, X., & Zhou, M. (2024). Temporal logical attention network for log-based anomaly detection in distributed systems. *Sensors*, 24(24), 7949.
- [24] Ren, S., Jin, J., Niu, G., & Liu, Y. (2025). ARCS: Adaptive Reinforcement Learning Framework for Automated Cybersecurity Incident Response Strategy Optimization. *Applied Sciences*, 15(2), 951.
- [25] Cao, J., Zheng, W., Ge, Y., & Wang, J. (2025). DriftShield: Autonomous fraud detection via actor-critic reinforcement learning with dynamic feature reweighting. *IEEE Open Journal of the Computer Society*.
- [26] Wang, J., Liu, J., Zheng, W., & Ge, Y. (2025). Temporal Heterogeneous Graph Contrastive Learning for Fraud Detection in Credit Card Transactions. *IEEE Access*.
- [27] Mai, N. T., Cao, W., & Liu, W. (2025). Interpretable Knowledge Tracing via Transformer-Bayesian Hybrid Networks: Learning Temporal Dependencies and Causal Structures in Educational Data. *Applied Sciences*, 15(17), 9605.
- [28] Cao, W., Mai, N. T., & Liu, W. (2025). Adaptive knowledge assessment via symmetric hierarchical Bayesian neural networks with graph symmetry-aware concept dependencies. *Symmetry*, 17(8), 1332.
- [29] Mai, N. T., Cao, W., & Wang, Y. (2025). The global belonging support framework: Enhancing equity and access for international graduate students. *Journal of International Students*, 15(9), 141-160.
- [30] Tan, Y., Wu, B., Cao, J., & Jiang, B. (2025). LLaMA-UTP: Knowledge-Guided Expert Mixture for Analyzing Uncertain Tax Positions. *IEEE Access*.
- [31] Sun, T., Yang, J., Li, J., Chen, J., Liu, M., Fan, L., & Wang, X. (2024). Enhancing auto insurance risk evaluation with transformer and SHAP. *IEEE Access*.