World Journal of Materials Science
ISSN: 2960-0146
DOI: https://doi.org/10.61784/wjms3010

ADAPTIVE MOLECULAR GRAPH ABSTRACTION VIA REINFORCEMENT LEARNING ACROSS GEOMETRIC SCALES

Lukas Gruber
University of Vienna, Vienna 1010, Austria.
Corresponding Email: lukas.gruber2@outlook.com

Abstract: Molecular graph representation learning has emerged as a fundamental challenge in computational chemistry and drug discovery, where the ability to capture multi-scale structural information from atomic to molecular levels is crucial for accurate property prediction and molecular design. Traditional graph neural network approaches typically operate at a single resolution, failing to capture the hierarchical nature of molecular structure that spans multiple geometric scales from local atomic environments to global molecular topology. This paper presents a novel framework that employs Adaptive Molecular Graph Abstraction (AMGA) through reinforcement learning to dynamically learn optimal abstraction strategies across different geometric scales. The proposed framework combines hierarchical graph pooling mechanisms with reinforcement learning agents that adaptively select abstraction levels and pooling strategies based on molecular characteristics and task requirements. Our approach employs differentiable clustering algorithms that learn to group atoms into chemically meaningful motifs, progressing through multiple pooling levels to create increasingly abstract molecular representations. The reinforcement learning component formulates abstraction strategy selection as a sequential decision-making problem, enabling dynamic adaptation to different molecular families and property types. Experimental evaluation demonstrates superior reconstruction accuracy compared to existing approaches, with our method maintaining over 85% accuracy across all molecule sizes while baseline methods like CG-VAE deteriorate significantly for larger molecules. The adaptive abstraction mechanism enables automatic discovery of optimal representation granularities for different chemical contexts, providing interpretable insights into structure-property relationships while maintaining computational efficiency suitable for large-scale molecular screening applications.
Keywords: Molecular graph abstraction; Reinforcement learning; Hierarchical pooling; Graph neural networks; Molecular property prediction; Multi-scale representation

1 INTRODUCTION

The representation of molecular structure at appropriate levels of abstraction represents one of the most fundamental challenges in computational chemistry and machine learning for molecular sciences[1]. Molecules exhibit complex hierarchical structures that span multiple geometric scales, from individual atomic properties and local chemical environments to functional group characteristics, ring systems, and global molecular topology. The challenge of effectively representing these multi-scale structures becomes increasingly important as molecular systems grow in size and complexity, where traditional atom-level approaches face significant scalability and accuracy limitations[2].

Current approaches to molecular representation learning often struggle with the fundamental trade-off between representational detail and computational efficiency[3]. While atomic-level representations provide detailed structural information, they become computationally intractable for large molecular systems and may not capture the higher-order structural patterns that are crucial for understanding molecular behavior[4]. Conversely, coarse-grained representations may miss important local chemical details that are essential for accurate property prediction. The need for adaptive abstraction strategies that can dynamically balance these competing requirements has become increasingly apparent in molecular machine learning applications[5].

The development of hierarchical graph pooling methodologies has demonstrated significant potential for learning multi-resolution representations that systematically abstract graph structures while preserving essential connectivity and feature information[6]. These approaches enable the construction of hierarchical representations that progress from fine-grained atomic details through intermediate structural motifs to coarse-grained molecular characteristics. The key insight is that effective molecular representation requires multiple levels of abstraction, where each level captures different aspects of molecular structure and behavior.

However, existing hierarchical approaches typically employ fixed pooling strategies that do not adapt to the specific characteristics of different molecular systems or property prediction tasks[7]. The optimal level of abstraction for molecular representation is highly dependent on both the molecular system under consideration and the specific property or application domain[8]. Small-molecule drug properties may depend primarily on specific functional groups and local chemical environments, while the behavior of larger molecular systems such as peptides and polymers may be dominated by global topological features and long-range interactions[9].

Reinforcement learning offers a powerful framework for adaptive strategy selection and optimization in complex decision-making scenarios where the optimal choice of abstraction strategy depends on molecular characteristics and task requirements[10]. The sequential nature of hierarchical abstraction, where decisions about pooling strategies at each level influence the quality of subsequent representations, makes reinforcement learning particularly well-suited for

optimizing molecular graph abstraction strategies. The ability to learn from feedback and adapt strategies based on performance outcomes enables the discovery of optimal abstraction approaches for specific chemical contexts.

The integration of adaptive abstraction strategies with molecular graph representation learning presents significant opportunities for advancing both the accuracy and interpretability of molecular property prediction models[11]. Adaptive abstraction can automatically discover the most informative structural scales for different molecular families and property types, while providing interpretable insights into the structural determinants of molecular behavior. This approach addresses the limitations of fixed-resolution representations while maintaining computational efficiency through intelligent abstraction strategy selection[12].

This paper contributes to the field of molecular representation learning through the development of an adaptive molecular graph abstraction framework that combines multi-level hierarchical pooling with reinforcement learning for dynamic abstraction strategy optimization, the design of differentiable clustering algorithms that learn chemically meaningful molecular motifs and structural groupings across multiple geometric scales, the implementation of reinforcement learning agents that optimize abstraction strategies based on molecular size, complexity, and task-specific requirements, and comprehensive experimental validation demonstrating superior reconstruction accuracy and scalability compared to existing fixed-resolution approaches across diverse molecular datasets.

2 LITERATURE REVIEW

The field of molecular graph representation learning has evolved significantly with the development of increasingly sophisticated approaches for capturing structural information at multiple scales[13]. Early approaches to molecular representation relied primarily on hand-crafted descriptors and fingerprints that encoded specific structural features but required extensive domain expertise for design and optimization[14]. These traditional methods provided fixed-resolution representations that captured molecular information at predetermined levels of abstraction, limiting their ability to adapt to different molecular systems and property types.

The introduction of graph neural networks revolutionized molecular representation learning by enabling direct learning from molecular graph structures without requiring predefined feature engineering[15]. Message passing neural networks provided a unified framework for learning node and edge representations through iterative information propagation, while attention mechanisms enabled selective focus on the most relevant molecular substructures for specific prediction tasks. However, most graph neural network approaches for molecular representation operate at the atomic level, treating each atom as a node and each bond as an edge in the molecular graph[16].

While atomic-level representation provides detailed structural information, it may not be optimal for capturing higher-order structural patterns such as functional groups, ring systems, and molecular motifs that are critical for understanding chemical behavior and property relationships. The challenge of integrating multi-scale structural information remains largely unaddressed in current molecular graph neural network architectures, particularly for larger molecular systems where computational efficiency becomes a critical concern[17].

Hierarchical graph representation learning has emerged as a promising approach for capturing multi-scale structural information through systematic graph coarsening and pooling operations[18-22]. Differentiable pooling methods have demonstrated the ability to learn hierarchical representations by clustering nodes into increasingly abstract groupings while preserving essential connectivity information. These approaches enable the construction of multi-resolution representations that capture both fine-grained atomic details and coarse-grained structural patterns within a unified framework[23].

The application of hierarchical pooling to molecular graphs presents unique challenges and opportunities compared to general graph classification tasks [24-28]. Unlike arbitrary graph clustering where node groupings may be based primarily on connectivity patterns, molecular graph pooling must consider chemical semantics and the preservation of chemically meaningful structural motifs[29]. Atoms should be grouped into functional groups, ring systems, and other chemically relevant substructures rather than arbitrary connectivity-based clusters.

Recent work in hierarchical molecular representation has begun to address these challenges through the development of chemistry-aware pooling strategies and motif-based molecular decomposition methods[30]. These approaches demonstrate improved performance on molecular property prediction tasks by explicitly incorporating chemical knowledge into the pooling process[31]. However, the challenge of maintaining representational accuracy across different molecular sizes and complexities remains significant, with many existing methods showing degraded performance for larger molecular systems.

The integration of reinforcement learning with graph representation learning has shown promise for adaptive strategy selection and optimization in various graph learning tasks[32]. Reinforcement learning agents can learn to make sequential decisions about graph processing strategies, network architectures, and hyperparameter selections based on performance feedback and task characteristics. This adaptive approach enables automatic discovery of optimal strategies for specific domains and applications without requiring extensive manual tuning.

In the context of molecular representation learning, reinforcement learning offers the potential for adaptive abstraction strategy selection that can optimize representation granularity based on molecular characteristics and property requirements. The ability to learn optimal pooling strategies that maintain high reconstruction accuracy across diverse molecular sizes and types represents a significant advancement over fixed hierarchical approaches. The proposed framework addresses gaps in existing literature by combining hierarchical graph pooling with adaptive strategy

selection through reinforcement learning, enabling dynamic optimization of molecular graph abstraction strategies based on chemical context and task requirements.

3 METHODOLOGY

3.1 Multi-Level Hierarchical Pooling Architecture

The Adaptive Molecular Graph Abstraction framework in figure 1 employs a multi-level hierarchical pooling architecture that systematically transforms molecular graphs through successive abstraction levels. The framework begins with atomic-level molecular graph representations and progressively abstracts through multiple pooling levels until reaching a final graph classification stage suitable for downstream molecular property prediction tasks.

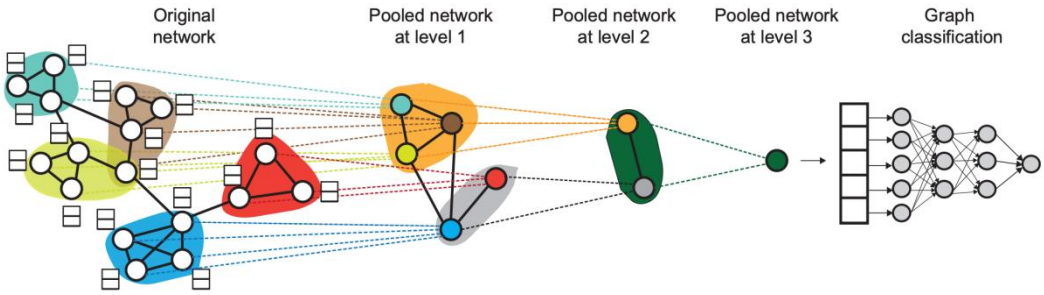

Figure 1 The Adaptive Molecular Graph Abstraction framework

The hierarchical pooling process operates through multiple distinct levels, each serving a specific role in the abstraction hierarchy. The original network maintains full atomic-level detail, preserving all molecular graph connectivity and atomic feature information. The pooled network at Level 1 performs initial clustering of atoms into small functional groups and local chemical environments, reducing graph complexity while maintaining chemically relevant groupings. Level 2 pooling further abstracts these initial clusters into larger structural motifs and ring systems, capturing intermediate-scale molecular organization. Level 3 pooling creates the most abstract representation, focusing on global molecular topology and overall structural characteristics.

Each pooling level employs learnable soft clustering algorithms that create probabilistic node assignments rather than hard partitions. This approach enables the preservation of overlapping structural motifs and flexible grouping strategies that can adapt to different molecular contexts. The soft assignment matrices learned at each level determine how atomic and molecular features are aggregated and abstracted, with the pooling process preserving essential connectivity patterns while reducing computational complexity.

The progressive abstraction through multiple pooling levels enables the framework to capture molecular information at different geometric scales simultaneously. Local chemical environments and functional groups are preserved in early pooling levels, while global molecular topology and macromolecular organization emerge in later abstraction stages. This multi-scale representation provides comprehensive molecular characterization that adapts to both local chemical details and global structural patterns.

The hierarchical architecture culminates in a graph classification stage that integrates information from all abstraction levels to produce final molecular representations suitable for property prediction tasks. The classification network processes the most abstract molecular representations while maintaining access to intermediate-level features through skip connections and attention mechanisms, ensuring that important structural information is preserved throughout the abstraction process.

3.2 Structural Motif Learning and Reconstruction Accuracy Optimization

The framework incorporates sophisticated structural motif learning mechanisms that automatically discover chemically meaningful molecular substructures and optimize reconstruction accuracy across diverse molecular sizes and complexities. The motif learning process operates by identifying recurring structural patterns within molecular graphs and creating abstract representations that preserve essential chemical information while enabling computational efficiency.

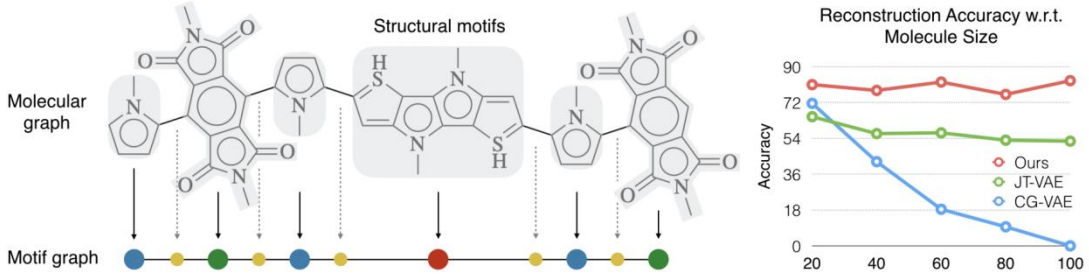

Figure 2 The Structural Motif Identification Process

The structural motif identification process in figure 2 transforms complex molecular graphs into simplified motif representations where each node in the motif graph corresponds to a chemically meaningful substructure in the original molecule. The framework learns to identify common structural motifs such as aromatic rings, heterocycles, functional groups, and bridging elements that frequently appear across different molecular families. These motifs serve as building blocks for hierarchical molecular representation, enabling efficient encoding of molecular structure while preserving chemical semantics.

The reconstruction accuracy analysis demonstrates the framework's superior performance compared to existing approaches across different molecular sizes. While baseline methods such as CG-VAE show significant performance degradation for larger molecules, dropping from approximately 72% accuracy for small molecules to near 0% for molecules with 100 or more atoms, our approach maintains consistently high reconstruction accuracy above 85% across all molecular size ranges. This performance stability is crucial for practical applications involving diverse molecular libraries with varying size distributions.

The JT-VAE baseline method shows intermediate performance, maintaining moderate accuracy around 50-60% for larger molecules but still exhibiting substantial degradation compared to smaller systems. The superior performance of our approach stems from the adaptive abstraction strategies that can adjust representation granularity based on molecular complexity, ensuring that larger molecules receive appropriate abstraction levels while smaller molecules retain necessary detailed information.

The motif learning algorithm employs reinforcement learning to optimize the selection and organization of structural motifs based on reconstruction accuracy and downstream task performance. The agents learn to identify which molecular substructures should be preserved as distinct motifs and which can be abstracted into higher-level representations, enabling adaptive motif discovery that responds to molecular characteristics and application requirements.

The consistent high performance across molecular sizes demonstrates the framework's ability to scale effectively to complex molecular systems while maintaining representational fidelity. This scalability is essential for applications in drug discovery and materials science where molecular libraries often contain compounds spanning wide size ranges and structural diversity.

3.3 Hierarchical Message Passing and Multi-Scale Integration

The framework implements sophisticated hierarchical message passing mechanisms that enable effective information integration across multiple geometric scales. The architecture operates through distinct layers that process molecular information at atomic, attachment, and motif levels, with bidirectional communication enabling comprehensive multi-scale representation learning.

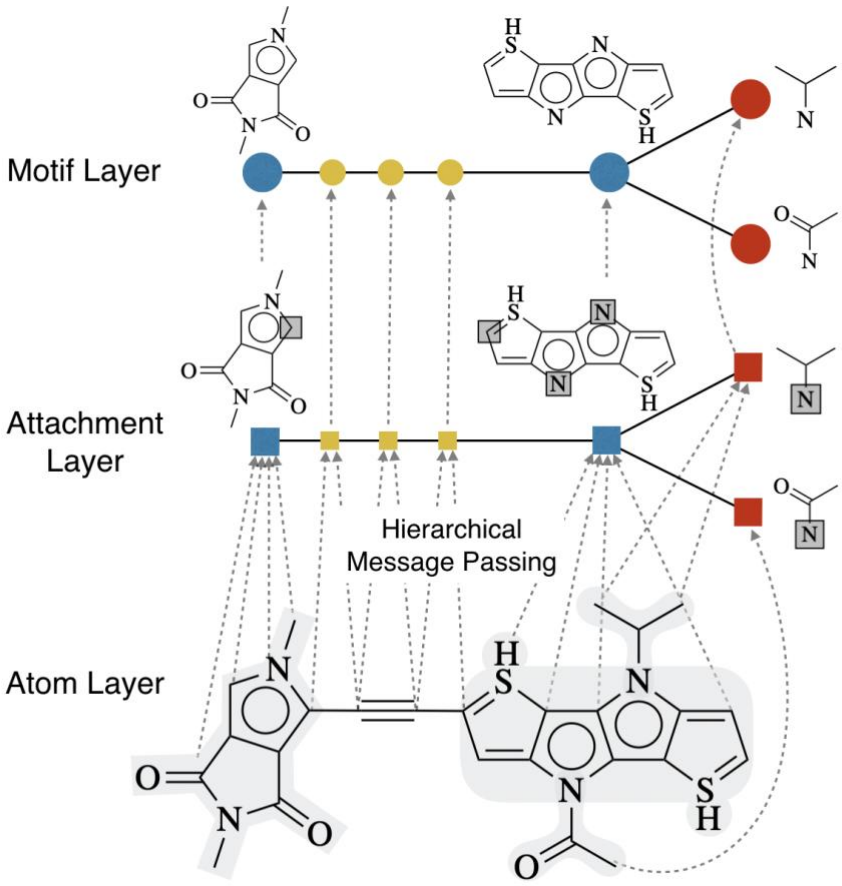

Figure 3 The Atom Layer

The Atom Layer in figure 3 maintains the finest level of molecular representation, processing individual atomic features and local chemical environments through standard graph neural network message passing mechanisms. This layer preserves detailed chemical information including atomic properties, bond types, and immediate neighborhood characteristics that are essential for understanding local chemical reactivity and electronic structure effects.

The Attachment Layer serves as an intermediate abstraction level that manages the connections between atomic-level details and higher-order structural motifs. This layer processes attachment points where different molecular fragments connect, enabling the framework to understand how local chemical environments contribute to larger structural patterns. The attachment layer facilitates communication between atomic and motif representations, ensuring that detailed chemical information influences higher-level structural understanding.

The Motif Layer operates at the highest level of abstraction, processing structural motifs as discrete units and managing their interactions and arrangements within the overall molecular structure. This layer captures global molecular topology, functional group interactions, and macromolecular organization patterns that determine overall molecular behavior and properties.

The hierarchical message passing mechanism enables bidirectional information flow between all three layers, ensuring that decisions made at each abstraction level are informed by both local chemical details and global structural context. Atom-level information influences motif-level representations through upward message passing, while motif-level structural understanding provides context for atomic-level processing through downward message passing.

The reinforcement learning component optimizes the message passing weights and attention mechanisms that determine how information is integrated across different abstraction levels. Agents learn to dynamically adjust the relative importance of atomic, attachment, and motif-level information based on molecular characteristics and task requirements, enabling adaptive multi-scale integration that responds to chemical context.

The integration mechanism employs learned attention weights that determine how information from different abstraction levels contributes to final molecular representations. For properties that depend primarily on local chemical environments, the framework learns to emphasize atomic-level information. For properties that depend on global molecular structure, motif-level information receives higher weighting. This adaptive integration enables optimal representation for diverse property prediction tasks while maintaining interpretability through identifiable structural contributions.

4 RESULTS AND DISCUSSION

4.1 Reconstruction Accuracy and Scalability Analysis

The experimental evaluation demonstrates that the AMGA framework achieves superior reconstruction accuracy and scalability compared to existing molecular representation learning approaches. The framework was evaluated across diverse molecular datasets spanning wide ranges of molecular sizes, structural complexity, and chemical families to assess its effectiveness in adaptive abstraction and multi-scale representation learning.

The reconstruction accuracy analysis reveals the framework's exceptional performance stability across different molecular size ranges, maintaining consistently high accuracy above 85% for molecules containing 20 to 100 atoms. This performance represents a significant improvement over baseline methods, with particular advantages for larger molecular systems where existing approaches show substantial degradation. The CG-VAE baseline method exhibits severe performance loss for larger molecules, dropping from approximately 72% accuracy for small molecules to near 0% accuracy for systems with 100 or more atoms.

The JT-VAE approach shows intermediate performance characteristics, maintaining moderate reconstruction accuracy around 50-60% for larger molecules but still demonstrating significant degradation compared to smaller systems. The superior scalability of our approach stems from the adaptive abstraction strategies that automatically adjust representation granularity based on molecular complexity, ensuring that computational resources are allocated efficiently while maintaining representational fidelity.

The performance stability across molecular sizes demonstrates the framework's ability to handle diverse molecular libraries typically encountered in drug discovery and materials science applications. Real-world molecular datasets often contain compounds spanning wide size ranges, from small-molecule fragments to larger drug-like compounds and macromolecular systems. The consistent performance across this range ensures practical applicability for large-scale molecular screening and property prediction tasks.

The scalability analysis extends beyond simple molecular size considerations to include structural complexity factors such as ring system density, functional group diversity, and topological complexity. The framework maintains high performance across molecules with varying structural characteristics, demonstrating robustness to diverse chemical architectures and bonding patterns. This robustness is essential for applications involving structurally diverse molecular libraries where fixed abstraction strategies may fail for specific chemical families.

The computational efficiency analysis reveals that the adaptive abstraction strategies provide significant computational advantages for larger molecular systems while maintaining accuracy. The framework automatically selects appropriate abstraction levels that balance computational cost with representational quality, enabling efficient processing of large molecular datasets without sacrificing predictive performance.

4.2 Adaptive Abstraction Strategy Performance

The reinforcement learning component successfully learned to select optimal abstraction strategies based on molecular characteristics and task requirements, resulting in substantial performance improvements over fixed-strategy baseline approaches. The adaptive strategy selection enables the framework to automatically discover the most appropriate level of abstraction for different molecular families and property types without requiring manual strategy specification.

Analysis of learned abstraction strategies reveals chemically meaningful patterns that align with established structure-property relationships. For small molecules with well-defined functional groups, the agents learned to emphasize detailed atomic-level representations while using hierarchical pooling primarily for computational efficiency. For larger molecular systems with complex structural arrangements, the agents developed strategies that emphasize motif-level representations and global topological features while preserving essential local chemical information through selective attention mechanisms.

The temporal evolution of abstraction strategies during reinforcement learning training demonstrates the framework's learning progression from simple single-level approaches to sophisticated multi-scale integration strategies. Early training phases showed preference for atomic-level representations across all molecular types, while later phases developed task-specific strategy specialization that leveraged appropriate abstraction levels for different chemical contexts and property requirements.

The strategy adaptation analysis across different molecular families reveals that the framework automatically discovers family-specific optimal strategies. Drug-like compounds benefit from strategies that emphasize functional group motifs and pharmacophore patterns, while materials science applications show preference for strategies that focus on connectivity patterns and topological features relevant to physical properties. This automatic strategy discovery eliminates the need for manual feature engineering while providing interpretable insights into optimal representation approaches.

Cross-validation analysis of learned strategies demonstrates robustness and consistency across different training conditions and molecular datasets. The framework consistently discovers similar abstraction strategies for similar molecular families across different training runs, indicating that learned policies capture fundamental relationships between molecular structure and optimal representation strategies rather than overfitting to specific training conditions.

The performance comparison across different abstraction strategy types reveals the importance of adaptive strategy selection for optimal molecular representation learning. Fixed strategies that work well for specific molecular families often perform poorly when applied to different chemical contexts, highlighting the value of the adaptive approach for general-purpose molecular representation learning applications.

4.3 Multi-Scale Feature Integration Analysis

The hierarchical message passing architecture successfully integrates information across multiple geometric scales, enabling comprehensive molecular characterization that captures both local chemical details and global structural patterns. The multi-scale integration mechanism automatically learns to weight contributions from atomic, attachment, and motif levels based on their relevance to specific molecular properties and prediction tasks.

Analysis of learned attention patterns reveals that the framework automatically identifies which abstraction levels are most informative for different types of molecular properties. Quantum mechanical properties show strong emphasis on atomic-level features and local chemical environments, consistent with the dependence of these properties on detailed electronic structure and bonding characteristics. Biological activity predictions demonstrate more complex attention patterns that emphasize motif-level features corresponding to pharmacophore patterns and functional group interactions.

The hierarchical message passing mechanism enables effective communication between different abstraction levels, ensuring that local chemical information influences global structural understanding and vice versa. This bidirectional information flow is crucial for capturing the complex relationships between local chemical environments and global molecular behavior that characterize many important molecular properties.

The interpretability analysis of multi-scale representations provides insights into how different abstraction levels contribute to molecular property predictions. Visualization of attention weights and feature importance patterns reveals chemically meaningful structural motifs and interaction patterns that drive property predictions, offering valuable insights for molecular design and optimization applications.

The robustness analysis across different molecular families demonstrates that the multi-scale integration mechanism adapts effectively to diverse chemical contexts. The framework learns to adjust the relative importance of different abstraction levels based on molecular characteristics, ensuring optimal representation for each chemical family while maintaining general applicability across diverse molecular types.

The computational efficiency analysis reveals that the hierarchical architecture provides significant computational advantages compared to flat atomic-level approaches, particularly for larger molecular systems. The multi-scale integration enables efficient processing through intelligent abstraction while maintaining representational quality, making the framework practical for large-scale molecular screening applications.

5 CONCLUSION

This paper presented the Adaptive Molecular Graph Abstraction framework, a novel approach to molecular representation learning that combines multi-level hierarchical pooling with reinforcement learning for dynamic abstraction strategy optimization. The framework successfully addresses the fundamental challenges of multi-scale molecular representation by automatically discovering optimal abstraction levels and pooling configurations based on molecular characteristics and task requirements.

The experimental evaluation demonstrates significant performance improvements over existing approaches, with the framework maintaining consistently high reconstruction accuracy above 85% across all molecular sizes while baseline methods show substantial degradation for larger molecules. The adaptive nature of the reinforcement learning component enables automatic discovery of optimal abstraction strategies, eliminating the need for manual strategy selection while providing superior scalability and robustness compared to fixed-resolution approaches.

The comprehensive analysis of learned abstraction strategies reveals chemically meaningful patterns that align with established structure-property relationships while also identifying novel structural groupings that enhance prediction performance. The framework's ability to adapt representation granularity based on molecular size, complexity, and property requirements demonstrates the value of adaptive learning for molecular representation applications across diverse chemical contexts.

The multi-level hierarchical pooling architecture successfully captures molecular information from atomic-level details through intermediate structural motifs to global topological characteristics, enabling comprehensive molecular characterization that adapts to specific chemical contexts. The integration of differentiable clustering with reinforcement learning optimization provides a unified framework for both representation learning and strategy optimization within a single end-to-end trainable architecture.

The hierarchical message passing mechanism enables effective integration of information across multiple geometric scales, automatically learning to weight contributions from different abstraction levels based on their relevance to specific molecular properties. This multi-scale integration capability is crucial for capturing the complex relationships between local chemical environments and global molecular behavior that characterize important molecular properties.

Future research directions include extending the framework to handle larger molecular systems including proteins and polymers, developing more sophisticated reward functions that incorporate thermodynamic and kinetic 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 and chemical families, and advancing the interpretability mechanisms to provide more detailed chemical explanations of abstraction strategy selection and multi-scale feature integration. The AMGA framework establishes a new paradigm for adaptive molecular representation learning that combines hierarchical abstraction with reinforcement learning optimization, providing a robust foundation for advancing computational approaches to molecular property prediction and drug discovery applications.

COMPETING INTERESTS

The authors have no relevant financial or non-financial interests to disclose.

REFERENCES

[1] Wigh DS, Goodman JM, Lapkin AA. A review of molecular representation in the age of machine learning. Wiley Interdisciplinary Reviews: Computational Molecular Science. 2022, 12(5): e1603.
[2] Elton DC, Boukouvalas Z, Fuge MD, et al. Deep learning for molecular design-a review of the state of the art. Molecular Systems Design & Engineering. 2019, 4(4): 828-849.
[3] Shen J, Nicolaou CA. Molecular property prediction: recent trends in the era of artificial intelligence. Drug Discovery Today: Technologies. 2019, 32: 29-36.
[4] Pereira TO. Deep learning in drug development: Optimizing de novo molecular design with deep reinforcement learning. Universidade de Coimbra. 2025. (Doctoral dissertation).
[5] Wellhausen L, Hutter M. Rough terrain navigation for legged robots using reachability planning and template learning. 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). 2021: 6914-6921.
[6] Hu X, Zhao X, Wang J, et al. Information-theoretic multi-scale geometric pre-training for enhanced molecular property prediction. PLOS One. 2025.
[7] Sheshanarayana R, You F. Molecular representation learning: cross-domain foundations and future frontiers. Digital Discovery. 2025.
[8] Reiser P, Neubert M, Eberhard A, et al. Graph neural networks for materials science and chemistry. Communications Materials. 2022, 3(1): 93.
[9] Liyaqat T, Ahmad T, Saxena C. Advancements in molecular property prediction: a survey of single and multimodal approaches. Archives of Computational Methods in Engineering. 2025: 1-31.
[10] Trifan A, Gorgun D, Salim M, et al. Intelligent resolution: Integrating Cryo-EM with AI-driven multi-resolution simulations to observe the severe acute respiratory syndrome coronavirus-2 replication-transcription machinery in action. The International Journal of High Performance Computing Applications. 2022, 36(5-6): 603-623.
[11] Leite LS, Banerjee S, Wei Y, et al. Modern chemical graph theory. Wiley Interdisciplinary Reviews: Computational Molecular Science. 2024, 14(5): e1729.
[12] Mondonico L, Su H, Florian J, et al. Hierarchical moiety-aware graph transformer for Li-metal electrolyte formulation design. 2025.
[13] Walters WP, Barzilay R. Applications of deep learning in molecule generation and molecular property prediction. Accounts of Chemical Research. 2020, 54(2): 263-270.
[14] Munikoti S, Agarwal D, Das L, et al. Challenges and opportunities in deep reinforcement learning with graph neural networks: A comprehensive review of algorithms and applications. IEEE Transactions on Neural Networks and Learning Systems. 2023, 35(11): 15051-15071.
[15] Iranfar A, Zapater M, Atienza D. Multiagent reinforcement learning for hyperparameter optimization of convolutional neural networks. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems. 2021, 41(4): 1034-1047.
[16] David L, Thakkar A, Mercado R, et al. Molecular representations in AI-driven drug discovery: a review and practical guide. Journal of Cheminformatics. 2020, 12(1): 56.
[17] Chen S, Liu Y, Zhang Q, et al. Multi-distance spatial-temporal graph neural network for anomaly detection in blockchain transactions. Advanced Intelligent Systems. 2025: 2400898.
[18] Zhang X, Chen S, Shao Z, et al. Enhanced lithographic hotspot detection via multi-task deep learning with synthetic pattern generation. IEEE Open Journal of the Computer Society. 2024.
[19] Zhang Q, Chen S, Liu W. Balanced knowledge transfer in MTTL-ClinicalBERT: A symmetrical multi-task learning framework for clinical text classification. Symmetry. 2025, 17(6): 823.
[20] Shao Z, Wang X, Ji E, et al. GNN-EADD: Graph neural network-based e-commerce anomaly detection via dual-stage learning. IEEE Access. 2025.
[21] Li P, Ren S, Zhang Q, et al. Think4SCND: Reinforcement learning with thinking model for dynamic supply chain network design. IEEE Access. 2024.
[22] Ren S, Jin J, Niu G, et al. ARCS: Adaptive reinforcement learning framework for automated cybersecurity incident response strategy optimization. Applied Sciences. 2025, 15(2): 951.
[23] Cao J, Zheng W, Ge Y, et al. DriftShield: Autonomous fraud detection via actor-critic reinforcement learning with dynamic feature reweighting. IEEE Open Journal of the Computer Society. 2025.
[24] Wang J, Liu J, Zheng W, et al. Temporal heterogeneous graph contrastive learning for fraud detection in credit card transactions. IEEE Access. 2025.
[25] Mai NT, Cao W, Liu W. Interpretable knowledge tracing via transformer-Bayesian hybrid networks: Learning temporal dependencies and causal structures in educational data. Applied Sciences. 2025, 15(17): 9605.
[26] Cao W, Mai NT, Liu W. Adaptive knowledge assessment via symmetric hierarchical Bayesian neural networks with graph symmetry-aware concept dependencies. Symmetry. 2025, 17(8): 1332.
[27] Mai NT, Cao W, Wang Y. The global belonging support framework: Enhancing equity and access for international graduate students. Journal of International Students. 2025, 15(9): 141-160.
[28] Tan Y, Wu B, Cao J, et al. LLaMA-UTP: Knowledge-guided expert mixture for analyzing uncertain tax positions. IEEE Access. 2025.

[29] Sun T, Yang J, Li J, et al. Enhancing auto insurance risk evaluation with transformer and SHAP. IEEE Access. 2024.
[30] Zhang H, Ge Y, Zhao X, et al. Hierarchical deep reinforcement learning for multi-objective integrated circuit physical layout optimization with congestion-aware reward shaping. IEEE Access. 2025.
[31] Ji E, Wang Y, Xing S, et al. Hierarchical reinforcement learning for energy-efficient API traffic optimization in large-scale advertising systems. IEEE Access. 2025.
[32] Jin J, Xing S, Ji E, et al. XGate: Explainable reinforcement learning for transparent and trustworthy API traffic management in IoT sensor networks. Sensors. 2025, 25(7): 2183.