



MOLECULAR GRAPH NEURAL NETWORKS IN DRUG DISCOVERY: A NARRATIVE REVIEW

George Papadopoulos^{1*}, Eleni Georgiou¹

1. Department of Computational Pharmaceutical Sciences, Faculty of Medicine, National and Kapodistrian University of Athens, Athens, Greece.

ARTICLE INFO

Received:

09 July 2025

Received in revised form:

21 November 2025

Accepted:

23 November 2025

Available online:

28 December 2025

Keywords: Molecular graph neural networks, Drug discovery, Molecular property prediction, Graph representation learning, Equivariant neural networks, Molecular generation

ABSTRACT

Drugs are molecules, and molecules are naturally represented as graphs, with atoms as nodes and bonds as edges, making graph neural networks an intuitive computational framework for medicinal chemistry. The modern era began around 2017, when message-passing neural networks demonstrated that molecular representations could be learned directly from graph structures rather than relying on handcrafted features, marking a conceptual shift from encoding chemistry to enabling models to discover task-relevant chemical abstractions. Since then, the field has expanded to include attention-based models, geometric networks, equivariant architectures, and self-supervised molecular foundation models, which now impact property prediction, molecular design, synthesis planning, protein–ligand modeling, and drug repurposing. Despite this rapid progress, molecular graph learning still faces challenges such as data sparsity, noisy labels, uncertain generalization, and limited interpretability, issues that are particularly critical in drug discovery, where confident extrapolation is often more important than retrospective benchmark performance. Emerging directions focus on tighter integration with protein structure models, prospective validation in real discovery programs, and autonomous design–make–test–analyze systems, suggesting that the next phase will reward models that combine chemical insight, uncertainty awareness, and practical deployability. This narrative review traces the intellectual trajectory of molecular graph neural networks, linking early innovations in message passing to contemporary foundation models and highlighting their current and future roles in computational drug discovery.

This is an **open-access** article distributed under the terms of the [Creative Commons Attribution-Non Commercial-Share Alike 4.0 License](https://creativecommons.org/licenses/by-nc-sa/4.0/), which allows others to remix, and build upon the work non commercially.

To Cite This Article: Papadopoulos G, Georgiou E. Molecular Graph Neural Networks in Drug Discovery: A Narrative Review. *Pharmacophore*. 2025;16(6):34-44. <https://doi.org/10.51847/8W12OUixI5>

Introduction

Molecules have long invited graph-based reasoning because medicinal chemists naturally think in terms of atoms, bonds, substituents, scaffolds, and local chemical environments. The message-passing neural network framework crystallized this intuition by treating molecular representation as a learnable exchange of information across bonded neighborhoods [1]. In this view, the model does not merely receive a molecule as a fixed string or descriptor vector; it progressively constructs a chemical representation through local interactions. This framing resonated strongly with drug discovery because potency, selectivity, toxicity, and physicochemical behavior often emerge from relationships among molecular substructures rather than from isolated atom counts.

Before graph neural networks became prominent, cheminformatics relied heavily on fingerprints, descriptors, and similarity measures that encoded decades of expert knowledge but also fixed the representational vocabulary in advance. MoleculeNet helped expose both the value and the limitations of these traditional representations by placing graph-based learning beside established molecular machine-learning baselines across diverse prediction tasks [2]. The conceptual leap of GNNs was not simply better feature engineering, but the replacement of static molecular encodings with differentiable operations that could be optimized for specific endpoints. Later comparison studies sharpened this lesson by showing that graph-based representations could be powerful, but that their advantage depended on dataset size, chemical diversity, endpoint quality, and validation strategy [3].

This review follows the emergence of molecular graph neural networks as an intellectual journey rather than as a systematic inventory. Early successes in quantum chemistry, property prediction, and graph generation created the sense that molecular graphs might become a general substrate for AI-driven chemistry [4, 5]. Subsequent work expanded that substrate into attention, equivariance, self-supervision, synthesis planning, and protein–ligand modeling. The aim here is to connect these developments into a coherent story about how graph learning changed the questions computational drug discovery could ask.

Corresponding Author: George Papadopoulos; Department of Computational Pharmaceutical Sciences, Faculty of Medicine, National and Kapodistrian University of Athens, Athens, Greece. E-mail: george.papadopoulos@gmail.com.

The structure of the review is roughly chronological but intentionally thematic. It begins with the formative pre-2020 period, when message passing, SchNet, and early benchmarks established the field's foundations [1, 2, 4, 6]. It then follows the proliferation of architectures, the scaling of molecular datasets, and the migration of GNNs beyond property prediction into generation, synthesis, structural biology, and repurposing. The later sections turn toward trust, deployment, limitations, and emerging paradigms, because the most important question today is no longer whether molecular GNNs can learn chemistry, but how reliably they can support decisions in drug discovery.

The Rise of Graph Neural Networks for Molecules (Pre-2020)

The MPNN Framework and Its Immediate Successors

The modern molecular GNN story began with the message-passing neural network framework, which unified several graph-learning ideas under a simple cycle of neighborhood aggregation, state updating, and readout [1]. This abstraction mattered because it gave chemists and machine-learning researchers a shared language for comparing models that had previously appeared architecturally distinct. Almost immediately, the framework encouraged adaptations that encoded bond features, atom environments, and molecular-level pooling in ways that aligned with cheminformatics practice. By the time directed message-passing models were analyzed for property prediction, the field had begun to understand learned molecular representations not as a novelty but as a serious alternative to engineered descriptors [7].

SchNet and the Incorporation of 3D Geometry

SchNet represented a second turning point because it treated molecules not only as topological graphs but also as physical systems embedded in three-dimensional space [4]. Its continuous-filter convolutions allowed interatomic distances to shape learned representations, making the model especially influential for quantum-mechanical and energy-related properties. This was an important conceptual expansion: molecular graphs could now carry geometry, not merely connectivity. Later geometric deep-learning perspectives framed this move as part of a broader transition from two-dimensional chemical graphs toward representations that respect molecular shape, symmetry, and spatial interaction [8].

Graph Attention and the First ADMET Benchmarks

As molecular graph learning matured, attention mechanisms offered a way to let models assign different importance to neighboring atoms and bonds during representation learning. Graph attention ideas became particularly visible in drug discovery through models such as AttentiveFP, which used attention to improve molecular property prediction and provide chemically suggestive atom-level weighting [9]. Benchmarks such as MoleculeNet made these developments consequential by forcing models to compete across toxicity, solubility, quantum, and bioactivity tasks rather than on isolated examples [2]. This early benchmarking culture helped establish ADMET and property prediction as the proving ground on which molecular GNNs first earned credibility.

The Proliferation of Architectures (2020–2022)

Message Passing Refinements

Around 2020, the field shifted from proving that message passing could work to asking which details of message passing mattered most. Directed molecular graph approaches emphasized the difference between atom-centered and bond-centered propagation, while models such as PotentialNet highlighted how staged interaction modeling could better reflect molecular environments relevant to property prediction [6]. Studies of learned molecular representations further showed that architecture, featurization, and data splitting could change conclusions about model superiority [7]. This period was less glamorous than the first wave, but it was scientifically important because it transformed GNN design from enthusiasm into craft.

Equivariant Networks and Tensor Field Networks

The next major architectural wave came from physics-inspired symmetry. $E(n)$ -equivariant graph neural networks showed how molecular representations could transform consistently under rotations, translations, and reflections, a property essential when learning from three-dimensional structures [10]. Equivariant message passing extended this idea toward tensorial molecular properties and spectra, reinforcing the view that molecular deep learning should respect the geometry of physical space [11]. These models reframed drug design as a problem where chemical identity and spatial configuration must be learned together rather than treated as separate preprocessing steps.

Graph Transformers

Graph transformers emerged as molecular modeling absorbed lessons from natural language processing, especially the power of self-attention and pre-training. Self-supervised graph transformer models on large molecular corpora showed that attention could be used not merely for local neighborhood weighting but for richer global representation learning [12]. This was a conceptual bridge between small, task-specific molecular predictors and the later idea of molecular foundation models. The transformer turn also changed the field's vocabulary, shifting attention from handcrafted graph operators toward scalable representation learning over large chemical spaces.

Table 1 consolidates the review's historical argument by showing how each major molecular GNN phase expanded the representational language available for drug-discovery modeling.

Table 1. Conceptual Evolution of Molecular Graph Neural Networks in Drug Discovery

Evolutionary phase	Core representational idea	Representative model families or concepts	Drug-discovery function enabled	Main conceptual gain	Persistent weakness
Early molecular graph learning, 2017–2019	Molecules are treated as atom–bond graphs in which local neighborhoods exchange information	Message-passing neural networks, directed message passing, molecular graph convolution	Property prediction, early ADMET prediction, toxicity estimation, benchmark comparison	Replaced fixed descriptors with learnable molecular representations	Performance depends strongly on dataset quality, splitting strategy, and endpoint definition
Geometry-aware molecular modeling	Molecular graphs are extended with spatial distances, conformers, and 3D molecular structure	SchNet, continuous-filter networks, geometry-enhanced GNNs	Quantum properties, shape-sensitive physicochemical prediction, structure-informed screening	Added physical and spatial realism beyond 2D topology	Increased computational burden and sensitivity to conformer generation
Attention-based molecular GNNs	Neighboring atoms, bonds, or substructures receive variable learned importance	Graph attention networks, AttentiveFP-like models, feature-wise attention	ADMET prediction, atom-level saliency, chemically suggestive model inspection	Improved flexibility and introduced visually interpretable atom weighting	Attention weights may appear chemically meaningful without proving causal explanation
Equivariant and symmetry-aware architectures	Molecular representations transform consistently under rotation, translation, and reflection	E(n)-equivariant GNNs, tensor field networks, equivariant message passing	Protein–ligand modeling, 3D property prediction, molecular interaction learning	Aligned learning with physical symmetries of molecular systems	Technically complex and more difficult to reproduce or deploy
Graph transformers and self-supervised learning	Large molecular corpora are used to pre-train reusable molecular representations	Graph transformers, contrastive molecular learning, masked graph modeling	Transfer learning, few-shot prediction, general-purpose molecular representation	Shifted the field from task-specific predictors toward reusable chemical intelligence	Transfer is unreliable when pre-training chemistry differs from downstream discovery chemistry
Multi-modal and biological-context GNNs	Ligand graphs are combined with proteins, disease networks, omics, assays, or biomedical knowledge graphs	Drug–target affinity models, protein–ligand graph models, biomedical knowledge-graph GNNs	Binding affinity prediction, target identification, repurposing, mechanism-aware prioritization	Moved GNNs from isolated molecular prediction toward biological decision context	Requires heterogeneous data integration, careful validation, and domain-specific interpretability
Generative and synthesis-aware graph models	Molecular graphs are generated, optimized, and linked to synthetic feasibility	Junction-tree VAEs, graph generative models, reinforcement learning, retrosynthesis models	De novo design, compound optimization, virtual library exploration, synthesis planning	Turned graph learning from molecular reading into molecular proposal and optimization	Generated molecules may be unrealistic, difficult to synthesize, or over-optimized to proxy objectives
Foundation-model and autonomous-loop era, 2023–2025	Molecular graph learning becomes part of broader reusable, adaptive, and closed-loop discovery systems	Molecular foundation models, graph–language systems, active learning, design–make–test–analyze loops	Prospective discovery, experiment prioritization, adaptive compound design, knowledge-guided decision support	Connects graph learning to real discovery workflows and human–AI experimentation	Translational value still depends on prospective validation, uncertainty estimation, reproducibility, and governance

3D-Aware Models and Conformer Ensembles

The rise of 3D-aware molecular GNNs reflected a growing recognition that a two-dimensional graph is often an incomplete description of a drug-like molecule. Geometry-enhanced molecular representation learning explicitly incorporated spatial information to improve property prediction and made conformational context part of the learned representation [13]. This development mattered for drug discovery because ligand binding, permeability, selectivity, and reactivity often depend on shape and orientation rather than graph topology alone. The field increasingly began to treat conformers not as inconvenient variants but as evidence that molecular identity is dynamic.

Multi-Modal Fusion

Although molecular GNNs began with small-molecule graphs, drug discovery quickly pushed them toward multi-modal settings. GraphDTA, for example, paired molecular graph representations with protein sequence information to predict drug–target binding affinity, illustrating how ligand graphs could be fused with biological context [14]. Broader graph machine-learning discussions in drug discovery also emphasized that real projects combine chemical structures with assays, omics, disease networks, images, and prior knowledge [15]. Multi-modal fusion therefore became a natural response to a central limitation of pure molecular graphs: a drug’s meaning depends on the biological system in which it acts.

Scaling to Large-Scale Molecular Datasets

Self-Supervised Pre-Training

Self-supervised pre-training changed the ambition of molecular GNNs by suggesting that models could learn broadly useful chemical representations before seeing a specific assay. Strategies such as molecular contrastive learning used alternative views of the same molecule to encourage robust graph representations [16]. Large-scale graph transformers similarly explored how molecular corpora could support pre-trained models that transfer across downstream endpoints [12]. This move echoed a broader pattern in AI: once architectures stabilized, scale and pre-training became the route toward generality.

Transfer Learning and Few-Shot Molecular Prediction

Transfer learning became attractive because most drug discovery endpoints are data-poor, especially for new targets, rare mechanisms, or early-stage projects. Pre-trained molecular graph representations promised to reduce dependence on large labeled assay collections by importing chemical knowledge learned from broader compound corpora [16]. The same motivation animated compact reviews and application studies that framed GNNs as tools for extracting reusable structure–property information across tasks [17]. Yet the promise was never automatic, because transfer depends on whether the pre-training chemistry resembles the downstream chemical and biological question.

Benchmarking at Scale (MoleculeNet, TDC, and Beyond)

Benchmarking played a formative role in molecular graph learning because it created shared tasks, common datasets, and a vocabulary for progress. MoleculeNet was especially influential, placing molecular GNNs in direct comparison with fingerprints and other machine-learning baselines across many endpoints [2]. Later evaluation work complicated the story by emphasizing robustness, out-of-distribution behavior, and the difficulty of interpreting benchmark gains as evidence of real-world drug discovery value [18]. The field’s benchmark culture therefore evolved from celebrating leaderboard movement toward asking whether evaluation protocols resemble prospective chemical decision-making.

Beyond Property Prediction: GNNs for Generation and Synthesis

Molecular De Novo Design with Graph Generative Models

Once GNNs had proven useful for prediction, the field naturally asked whether graphs could also be generated. Junction tree variational autoencoders decomposed molecules into chemically meaningful substructures, offering a generative route that respected molecular validity more directly than atom-by-atom construction [5]. Reinforcement-learning approaches then connected molecular generation to optimization, making the design process resemble iterative exploration of chemical space under property-driven objectives [19]. This transition from prediction to generation marked a philosophical shift: GNNs were no longer only reading molecules, but helping imagine new ones.

Retrosynthesis and Reaction Prediction

Synthesis planning brought graph learning into contact with one of medicinal chemistry’s most practical questions: whether a proposed molecule can be made. Machine-learning approaches to computer-aided synthesis planning showed how reaction prediction and retrosynthetic reasoning could be formalized as learnable transformations over molecular structures [20]. In this setting, graphs offered a natural way to represent changes in bonds, atoms, and reaction centers. The appeal was not merely automation, but the possibility of linking design models with synthetic feasibility so that generated molecules would be imaginative yet actionable.

Multi-Step Synthesis Planning

Multi-step synthesis planning required GNNs to move beyond single reactions and participate in search over possible routes. Reviews of machine learning in synthesis planning emphasized that reaction predictors become most useful when coupled to heuristic search, template selection, and route scoring [20]. This integration changed the role of molecular graph models from endpoint predictors to components in larger decision systems. In drug discovery, that distinction matters because a model that proposes a potent compound but cannot connect it to a feasible route may create more burden than value.

The Interface with DNA-Encoded Libraries and Chemical Space Exploration

Ultra-large chemical spaces, including DNA-encoded libraries, made molecular representation learning especially attractive because exhaustive experimental exploration is impossible. Graph networks for molecular design helped frame chemical space navigation as a representation problem in which molecules can be embedded, compared, optimized, and prioritized [21]. The same logic supported virtual screening applications where graph-based models could triage vast libraries before experimental testing. At this interface, GNNs became less a replacement for chemistry than a mapmaking tool for regions of chemical space too large for human intuition alone.

Structural Biology and GNNs at the Interface

Protein–Ligand Interaction Prediction

Protein–ligand interaction prediction became one of the clearest signs that molecular GNNs were moving beyond isolated compound property prediction. GraphDTA represented ligands as molecular graphs while pairing them with protein sequence features, showing how graph learning could help model binding affinity in a biologically grounded setting [14]. MGraphDTA

extended this direction with multiscale graph modeling and explainability, reflecting the field's desire to represent both local chemical contacts and broader interaction patterns [22]. Later binding-affinity models such as GraphscoreDTA reinforced the view that protein–ligand prediction is not merely a molecular property task, but an interface problem where ligand chemistry, protein context, and structural complementarity must be learned together [23].

Protein Structure Representations

As structural biology became increasingly computational, graph representations of proteins began to complement ligand graph models. Graph machine-learning perspectives in drug discovery emphasized that proteins can be represented through residues, contact maps, pockets, surfaces, or interaction networks, each encoding a different biological abstraction [15]. Geometric deep learning then gave the field a vocabulary for learning over molecular and macromolecular structures while respecting spatial relationships [8]. The growing connection between ligand graphs and protein structural graphs made drug discovery feel less like separate modeling of “compound” and “target” and more like representation learning over interacting molecular systems.

Drug Repurposing and Target Identification

Drug repurposing brought GNNs into the broader biomedical network, where drugs, diseases, proteins, pathways, and adverse events can be modeled as connected entities. Graph convolutional approaches to polypharmacy side effects demonstrated how relational structure could reveal clinically relevant drug interactions in large biomedical graphs [24]. This work helped expand the imagination of molecular GNNs beyond atom–bond graphs toward heterogeneous knowledge graphs and interactomes. In this setting, the promise of graph learning was not just to predict a molecule's properties, but to reposition it within a web of biological mechanisms and therapeutic opportunities. **Figure 1** illustrates how molecular graph neural networks extend beyond atom–bond property prediction by embedding candidate drugs within heterogeneous biomedical networks of diseases, proteins, pathways, adverse events, and therapeutic opportunities.




Figure 1. From Molecular Graph Prediction to Biomedical Network Repositioning for Drug Repurposing

Interpretability and Trust in Molecular Graph Models

Attention as an Explanation

Attention mechanisms entered molecular graph learning partly as a route to better prediction, but they were quickly interpreted as a possible window into chemical reasoning. AttentiveFP made this tension visible because its atom-level attention weights could be mapped back onto molecular structures, inviting comparison with pharmacophoric intuition [9]. Yet attention is not automatically explanation, and its apparent chemical plausibility can obscure whether the model is identifying causal motifs or merely correlational shortcuts. The field therefore learned to treat attention maps as useful hypotheses rather than definitive mechanistic accounts.

Shapley Values and Subgraph Attribution

More formal attribution methods attempted to move beyond visually appealing attention maps toward systematic explanations of graph predictions. Hierarchical informative GNNs with feature-wise attention reflected this broader effort to identify which substructures and features drive molecular predictions [25]. Explainable drug–target affinity models similarly tried to connect model outputs to interpretable chemical and biological interaction patterns [22]. The enduring challenge is validation: an attribution is only useful in medicinal chemistry if it helps a scientist decide what to synthesize, modify, or distrust.

The Tension between Accuracy and Explainability

In practice, the most accurate molecular graph model is not always the most useful one. Comparison studies between descriptor-based and graph-based models showed that performance advantages can be context-dependent, which keeps simpler representations relevant when interpretability, robustness, or speed matters [3]. Compact reviews of molecular property prediction with GNNs likewise warned that model complexity must be justified by the demands of the task and the quality of

the available data [17]. This tension has become central to adoption: drug discovery teams need models that are not only accurate in aggregate but also credible at the level of individual chemical decisions.

Integration Into the Drug Discovery Workflow

Virtual Screening and Hit Prioritization

Virtual screening became one of the most visible application areas for molecular GNNs because it offered a direct route from prediction to experimental prioritization. The discovery of antibacterial candidates using deep learning showed how learned molecular representations could reshape screening cascades by identifying compounds that might be overlooked by conventional similarity-based approaches [26]. Graph attention models for molecular property prediction further strengthened the idea that GNNs could triage large compound collections by learning subtle structure–activity patterns [9]. The most important lesson from these examples was not that GNNs replace assays, but that they can change which molecules reach the assay in the first place.

Lead Optimisation and Multi-Parameter Scoring

Lead optimisation is inherently multi-objective, requiring teams to balance potency, selectivity, permeability, solubility, metabolism, toxicity, and synthetic feasibility. Molecular property prediction studies using learned graph representations helped make this balancing act computationally tractable by enabling models to share information across related endpoints [7]. ADMET-focused GNN work, including graph attention approaches, made the technology especially relevant to the recurring medicinal chemistry problem of improving one property without damaging another [9]. In this role, GNNs function less like final arbiters and more like navigational instruments for multi-parameter chemical design.

Deployment in Pharmaceutical Industry Settings

Industry deployment forced molecular GNNs to confront problems that are less visible in academic benchmarks: assay drift, inconsistent metadata, proprietary chemical spaces, and integration with legacy cheminformatics systems. Reviews of graph machine learning in drug discovery described this transition from promising algorithms to workflow components, emphasizing the importance of data governance, model monitoring, and human expert review [15]. Comparative studies also reminded practitioners that graph models must earn their place against strong baselines, not against outdated or poorly tuned alternatives [3]. The result has been a more mature view in which GNNs are valuable when embedded within decision workflows rather than treated as stand-alone engines of discovery.

Barriers to Adoption

The barriers to adoption are not only technical, but institutional. Data silos restrict the diversity and scale of training sets, while inconsistent assay protocols make it difficult to learn transferable chemical–biological relationships [18]. Computational cost and specialized expertise can further slow deployment, particularly when teams cannot easily reproduce published architectures or adapt them to internal data. These barriers help explain why the field’s narrative has shifted from model invention toward validation, reproducibility, and operational trust.

Current Limitations and Unfulfilled Promises

Poor Generalization to Novel Chemistry

The central limitation of molecular GNNs is not whether they can interpolate within familiar chemistry, but whether they can extrapolate to new scaffolds, mechanisms, and chemical series. Robustness studies have emphasized that out-of-distribution evaluation can sharply change conclusions about model utility, especially when test compounds differ meaningfully from the training set [18]. This matters profoundly in drug discovery, where the molecules of greatest interest are often precisely those that depart from known chemical neighborhoods. The field is therefore learning that benchmark success must be interpreted through the lens of chemical distance and prospective relevance.

Sensitivity to Training Data Noise and Bias

Drug discovery data are noisy because assays vary, labels are context-dependent, and inactive compounds may simply not have been tested under the right conditions. Molecular representation studies showed early that model conclusions depend strongly on dataset construction, splitting strategy, and endpoint definition [7]. ADMET and bioactivity prediction are especially vulnerable because sparse positives, publication bias, and measurement variability can make a learned graph model appear more certain than the biology permits. Better uncertainty estimation, careful curation, and honest validation are therefore as important as architectural novelty.

Computational Cost and Environmental Impact

The scaling of molecular GNNs has brought practical costs in hardware, energy, and engineering complexity. Graph neural networks for chemistry and materials science have become increasingly sophisticated, but that sophistication often comes with heavier training pipelines and greater dependence on specialized infrastructure [27]. Equivariant and 3D-aware models add particular computational demands because they process spatial information and sometimes conformational ensembles [10, 13]. The field now faces a familiar AI dilemma: the models that best capture molecular reality may be the hardest to train, reproduce,

and deploy sustainably. **Figure 2** shows how increasing molecular realism in graph neural networks can improve chemical representation while also raising hardware, energy, reproducibility, and deployment burdens.




Figure 2. Computational Sustainability Trade-Offs in Advanced Molecular Graph Neural Networks

The Reproducibility Deficit

Reproducibility remains a persistent weakness in molecular deep learning. Benchmarking papers such as MoleculeNet created a shared foundation, but later work showed that comparisons can still be distorted by hyperparameter choices, splitting protocols, featurization details, and incomplete reporting [2, 18]. This problem becomes more severe as models combine graph architectures, geometry, pre-training, and task-specific fine-tuning. For molecular GNNs to become dependable scientific instruments, reproducibility must be treated as part of model design rather than as an afterthought.

Table 2 translates the review’s limitations and future outlook into a readiness framework for judging when molecular GNNs can support real drug-discovery decisions rather than only retrospective model evaluation.

Table 2. Translational Readiness Framework for Molecular GNNs in Drug Discovery

Translational domain	What molecular GNNs contribute	Evidence needed before practical adoption	Key failure mode	Readiness-enhancing strategy	Human decision role
Molecular property and ADMET prediction	Learns structure–property relationships directly from atom–bond and feature interactions	Scaffold-split validation, temporal validation, external assay validation, comparison against strong descriptor and fingerprint baselines	Retrospective benchmark success fails to generalize to new scaffolds or assay settings	Use uncertainty calibration, chemical-domain applicability estimates, and prospective compound testing	Medicinal chemists judge whether predicted trade-offs are chemically actionable
Virtual screening and hit prioritization	Reorders compound libraries using learned molecular similarity beyond conventional fingerprints	Prospective hit-rate improvement, diversity of selected molecules, comparison with docking and similarity workflows	Model enriches familiar chemistry while missing novel or unusual active scaffolds	Combine GNN scoring with diversity filters, assay-aware uncertainty, and orthogonal screening evidence	Screening teams decide which predictions justify experimental budget
Lead optimization and multi-parameter scoring	Supports simultaneous balancing of potency, solubility, permeability, protein binding, clearance, and toxicity	Multi-endpoint validation, uncertainty-aware ranking, assessment of property trade-off stability	Optimization improves one predicted endpoint while degrading hidden or poorly modeled properties	Use multi-objective scoring, Pareto-front analysis, and transparent endpoint-specific confidence reporting	Project teams weigh model outputs against synthesis feasibility and program priorities
Molecular generation and de novo design	Proposes novel structures under property-driven or scaffold-constrained objectives	Validity, novelty, synthesizability, patentability, biological plausibility, and experimental confirmation	Generated molecules exploit model weaknesses or violate medicinal chemistry intuition	Link generation to retrosynthesis, novelty constraints, physicochemical filters, and human review	Chemists select, modify, or reject proposed molecules before synthesis
Retrosynthesis and synthesis planning	Models bond changes, reaction centers, and multi-step route feasibility	Route plausibility, reagent availability, precedent support, success-rate estimates, and laboratory validation	Synthetic route appears plausible computationally but fails operationally in the laboratory	Integrate reaction prediction with route scoring, literature precedent, and practical chemistry constraints	Synthetic chemists assess route practicality, safety, cost, and scalability
Protein–ligand interaction and binding prediction	Combines ligand graph structure with protein sequence, pocket, or structural representations	Cross-target validation, structural validation, binding-mode consistency, and	Model learns dataset-specific shortcuts rather than transferable interaction determinants	Use structure-aware validation, target-family holdout testing, and interpretable contact-level evidence	Structural biologists and pharmacologists evaluate whether predictions fit mechanistic evidence

prospective affinity testing					
Drug repurposing and biomedical knowledge graphs	Places molecules within networks of diseases, targets, pathways, side effects, and therapeutic opportunities	Clinical plausibility checks, causal pathway support, adverse-event evidence, and real-world validation	Network proximity is mistaken for therapeutic mechanism	Combine graph prediction with mechanistic biology, literature evidence, and safety constraints	Translational teams decide whether a repurposing signal is biologically and clinically credible
Foundation models and graph-language integration	Creates reusable molecular representations linked to literature, patents, protocols, and biological evidence	Task transfer testing, hallucination controls, provenance tracking, uncertainty reporting, and domain-specific auditing	General-purpose model produces persuasive but chemically or experimentally unsupported suggestions	Use retrieval grounding, source attribution, constrained generation, and expert-in-the-loop review	Human experts verify scientific validity, novelty, regulatory relevance, and experimental feasibility
Autonomous design-make-test-analyze systems	Embeds molecular GNNs in iterative experimental loops that update from new data	Closed-loop prospective studies, assay reproducibility, decision traceability, and measurable acceleration of discovery	Automated loop amplifies biased data, narrow objectives, or poor experimental feedback	Maintain explicit stopping rules, uncertainty-triggered human review, and continuous model monitoring	Scientists define objectives, inspect failures, adjust hypotheses, and approve next experiments

The Road Ahead – Emerging Paradigms

Graph-Based Foundation Models and Generalist Molecular AI

The idea of a molecular foundation model represents the latest turn in the graph-learning story. Models such as MolE point toward graph representations that can support multiple molecular tasks through large-scale pre-training and flexible attention-based architectures [28]. This vision differs from earlier task-specific predictors because it imagines a reusable chemical intelligence layer that can be adapted across property prediction, design, screening, and prioritization. The field’s next challenge is to determine whether such generality remains chemically faithful when moved from broad pre-training corpora into specific discovery programs.

Figure 3 maps the review’s central narrative arc, showing how molecular graph neural networks evolved from message-passing representations into architecture-rich, scale-driven, and translationally constrained drug-discovery systems.




Figure 3. Historical-to-Translational Landscape of Molecular Graph Neural Networks in Drug Discovery

Integration with Large Language Models and Textual Knowledge

Molecular GNNs increasingly sit beside language models that encode literature, patents, protocols, and biological knowledge. This creates the possibility of systems that combine graph-level chemical structure with text-derived hypotheses about targets, mechanisms, adverse events, or synthetic precedent. Reviews of molecular graph learning have already suggested that the field’s future lies in linking chemistry to richer biomedical context rather than optimizing isolated structure–property tasks [15, 29]. The deepest opportunity may come from models that can reason across molecular graphs and scientific language while remaining grounded in experimental evidence. Recent work suggests that the next generation of molecular GNNs will be shaped by three converging directions: stronger pre-training objectives, explicit 3D molecular reasoning, and tighter links between molecular graphs, biological context, and scientific language. Mole-BERT showed that molecular graph pre-training requires chemistry-aware tokenization rather than direct imitation of language-model masking [30], while Uni-Mol, GraphMVP, and geometry-denoising approaches demonstrate that 3D conformational information can improve representation learning when topology alone is insufficient [31–33]. At the protein–ligand interface, physicochemical graph learning has

further shown that interpretable interaction fingerprints can be learned even when high-resolution structures are unavailable [34]. In parallel, recent explainable and multi-task molecular representation frameworks indicate that future models will need to handle incomplete annotations, cross-endpoint relationships, and imperfect real-world datasets rather than only clean benchmark tables [35]. These developments reinforce the idea that molecular GNNs are moving from isolated structure–property predictors toward broader discovery systems that combine generation, binding prediction, multimodal evidence, uncertainty awareness, and translational validation [36–38].

Active Learning and Autonomous Drug Discovery Loops

Active learning reframes GNNs as participants in an iterative discovery loop rather than static predictors trained once and deployed passively. Molecular generation and optimization studies showed how graph models could propose new compounds under desired objectives [5, 19], while virtual screening successes demonstrated how learned representations could influence experimental selection [26]. In a closed-loop design–make–test–analyze workflow, these ideas converge: the model proposes, experiments respond, and the next model update reflects newly measured chemistry. The promise is not full automation for its own sake, but a faster and more reflective dialogue between computation and experiment.

Toward Causal Graph Models

The next frontier may require molecular GNNs to move from correlation toward causal reasoning. Current models often learn associations between substructures and endpoints, but medicinal chemistry frequently asks counterfactual questions about what would happen if a substituent, conformation, protein contact, or assay condition changed. Equivariant and geometry-aware architectures have already shown that respecting physical structure improves representation learning [11, 39], but causal graph models would go further by trying to distinguish mechanistic drivers from statistical artifacts. Such models could make graph learning more useful for design decisions because they would support intervention, not merely prediction.

Strengths and Limitations of This Narrative

Strengths

The strength of this narrative review is that it connects technical developments in molecular graph learning to the practical and conceptual evolution of drug discovery. Rather than treating MPNNs, SchNet, attention, self-supervision, equivariance, and foundation models as isolated inventions, it follows how each changed the field’s sense of what molecular representation could mean [1, 4, 9, 16, 28]. It also links small-molecule property prediction to synthesis, protein–ligand modeling, repurposing, and deployment, reflecting the broadening scope described in contemporary drug discovery reviews [15, 29]. This connected view is useful because the field’s progress has been cumulative, with each architectural turn responding to limitations exposed by the previous one.

Limitations

The limitation of any narrative review is selectivity. Important adjacent areas, including reaction graph learning, materials-focused molecular GNNs, and specialized biomedical knowledge-graph methods, are represented here only insofar as they illuminate the drug discovery story [20, 27]. The synthesis is qualitative rather than quantitative, and it does not attempt to rank architectures or adjudicate benchmark performance. That choice is deliberate, because the purpose is to trace an intellectual arc rather than to reproduce the logic of a systematic or scoping review.

Conclusion

The arc from message-passing neural networks to molecular foundation models represents a remarkable decade of innovation in how researchers represent and learn from molecules. What began as a graph-based reformulation of molecular property prediction has become a broad computational language for chemistry, biology, and design.

Graph neural networks have moved from proof of concept to an indispensable component of the drug discovery toolkit. Yet their full potential remains constrained by data quality, generalization, interpretability, computational cost, and reproducibility. The future lies in tighter integration with other data modalities, prospective validation in real discovery programs, and the development of self-critiquing autonomous molecular design systems. These systems will need to combine graph learning with structural biology, synthesis awareness, uncertainty estimation, and human chemical judgment.

This narrative review has aimed to serve as both a historical record and a compass for the next phase of molecular graph learning. The most exciting future for the field will not be defined by a single architecture, but by models that help scientists ask better chemical questions and make better experimental decisions.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Gilmer J, Schoenholz SS, Riley PF, Vinyals O, Dahl GE. Neural message passing for quantum chemistry. In: International Conference on Machine Learning (ICML). PMLR; 2017. p. 1263-72.
2. Wu Z, Ramsundar B, Feinberg EN, Gomes J, Geniesse C, Pappu AS, et al. MoleculeNet: a benchmark for molecular machine learning. *Chem Sci*. 2018;9(2):513-30.
3. Jiang D, Wu Z, Hsieh CY, Chen G, Liao B, Wang Z, et al. Could graph neural networks learn better molecular representation for drug discovery? A comparison study of descriptor-based and graph-based models. *J Cheminform*. 2021;13(1):12.
4. Schütt KT, Sauceda HE, Kindermans PJ, Tkatchenko A, Müller KR. SchNet—a deep learning architecture for molecules and materials. *J Chem Phys*. 2018;148(24):241701.
5. Jin W, Barzilay R, Jaakkola T. Junction tree variational autoencoder for molecular graph generation. In: International Conference on Machine Learning (ICML). PMLR; 2018. p. 2323-32.
6. Feinberg EN, Sur D, Wu Z, Husic BE, Mai H, Li Y, et al. PotentialNet for molecular property prediction. *ACS Cent Sci*. 2018;4(11):1520-30.
7. Yang K, Swanson K, Jin W, Coley C, Eiden P, Gao H, et al. Analyzing learned molecular representations for property prediction. *J Chem Inf Model*. 2019;59(8):3370-88.
8. Atz K, Grisoni F, Schneider G. Geometric deep learning on molecular representations. *Nat Mach Intell*. 2021;3(12):1023-32.
9. Xiong Z, Wang D, Liu X, Zhong F, Wan X, Li X, et al. Pushing the boundaries of molecular representation for drug discovery with the graph attention mechanism. *J Med Chem*. 2019;63(16):8749-60.
10. Satorras VG, Hoogeboom E, Welling M. E(n) equivariant graph neural networks. In: International Conference on Machine Learning (ICML). PMLR; 2021. p. 9323-32.
11. Schütt K, Unke O, Gastegger M. Equivariant message passing for the prediction of tensorial properties and molecular spectra. In: International Conference on Machine Learning (ICML). PMLR; 2021. p. 9377-88.
12. Rong Y, Bian Y, Xu T, Xie W, Wei Y, Huang W, et al. Self-supervised graph transformer on large-scale molecular data. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2020. p. 12559-71.
13. Fang X, Liu L, Lei J, He D, Zhang S, Zhou J, et al. Geometry-enhanced molecular representation learning for property prediction. *Nat Mach Intell*. 2022;4(2):127-34.
14. Nguyen T, Le H, Quinn TP, Nguyen T, Le TD, Venkatesh S. GraphDTA: predicting drug–target binding affinity with graph neural networks. *Bioinformatics*. 2021;37(8):1140-7.
15. Gaudelet T, Day B, Jamasb AR, Soman J, Regep C, Liu G, et al. Utilizing graph machine learning within drug discovery and development. *Brief Bioinform*. 2021;22(6):bbab159.
16. Wang Y, Wang J, Cao Z, Barati Farimani A. Molecular contrastive learning of representations via graph neural networks. *Nat Mach Intell*. 2022;4(3):279-87.
17. Wieder O, Kohlbacher S, Kuenemann M, Garon A, Ducrot P, Seidel T, et al. A compact review of molecular property prediction with graph neural networks. *Drug Discov Today Technol*. 2020;37:1-2.
18. Fooladi H, Vu TN, Mathea M, Kirchmair J. Evaluating machine learning models for molecular property prediction: performance and robustness on out-of-distribution data. *J Chem Inf Model*. 2025;65(19):9871-91.
19. Zhou Z, Kearnes S, Li L, Zare RN, Riley P. Optimization of molecules via deep reinforcement learning. *Sci Rep*. 2019;9(1):10752.
20. Coley CW, Green WH, Jensen KF. Machine learning in computer-aided synthesis planning. *Acc Chem Res*. 2018;51(5):1281-9.
21. Mercado R, Rastemo T, Lindelöf E, Klambauer G, Engkvist O, Chen H, et al. Graph networks for molecular design. *Mach Learn Sci Technol*. 2021;2(2):025023.
22. Yang Z, Zhong W, Zhao L, Chen CY. MGraphDTA: deep multiscale graph neural network for explainable drug–target binding affinity prediction. *Chem Sci*. 2022;13(3):816-33.
23. Wang K, Zhou R, Tang J, Li M. GraphscoreDTA: optimized graph neural network for protein–ligand binding affinity prediction. *Bioinformatics*. 2023;39(6):btad340.
24. Zitnik M, Agrawal M, Leskovec J. Modeling polypharmacy side effects with graph convolutional networks. *Bioinformatics*. 2018;34(13):i457-66.
25. Zhu W, Zhang Y, Zhao D, Xu J, Wang L. HiGNN: A hierarchical informative graph neural network for molecular property prediction equipped with feature-wise attention. *J Chem Inf Model*. 2022;63(1):43-55.
26. Stokes JM, Yang K, Swanson K, Jin W, Cubillos-Ruiz A, Donghia NM, et al. A deep learning approach to antibiotic discovery. *Cell*. 2020;180(4):688-702.
27. Reiser P, Neubert M, Eberhard A, Torresi L, Zhou C, Shao C, et al. Graph neural networks for materials science and chemistry. *Commun Mater*. 2022;3(1):93.

28. Méndez-Lucio O, Nicolaou CA, Earnshaw B. MolE: a foundation model for molecular graphs using disentangled attention. *Nat Commun.* 2024;15(1):9431.
29. Tang M, Li B, Chen H. Application of message passing neural networks for molecular property prediction. *Curr Opin Struct Biol.* 2023;81:102616.
30. Xia J, Zhao C, Hu B, Gao Z, Tan C, Liu Y, et al. Mole-BERT: Rethinking pre-training graph neural networks for molecules. In: *International Conference on Learning Representations (ICLR).* 2023.
31. Zhou G, Gao Z, Ding Q, Zheng H, Xu H, Wei Z, et al. Uni-Mol: A universal 3D molecular representation learning framework. In: *International Conference on Learning Representations (ICLR).* 2023.
32. Liu S, Wang H, Liu W, Lasenby J, Guo H, Tang J. Pre-training molecular graph representation with 3D geometry. In: *International Conference on Learning Representations (ICLR).* 2022.
33. Liu S, Guo H, Tang J. Molecular geometry pretraining with SE(3)-invariant denoising distance matching. In: *International Conference on Learning Representations (ICLR).* 2023.
34. Koh HY, Nguyen ATN, Pan S, May LT, Webb GI. Physicochemical graph neural network for learning protein–ligand interaction fingerprints from sequence data. *Nat Mach Intell.* 2024;6:710-22.
35. Wang B, Li J, Zhou D, Li L, Li J, Hao J, et al. Unified and explainable molecular representation learning for imperfectly annotated data from the hypergraph view. *Nat Commun.* 2025;16:8717.
36. Yang N, Zeng X, Jiang Y, Chen J, Fang X, Zhang X. Molecule generation for drug design: A graph learning perspective. *Comput Struct Biotechnol J.* 2024;23:3347-66.
37. Yao R, Shen Z, Xu X, Ling G, Xiang R, Song T, et al. Knowledge mapping of graph neural networks for drug discovery: a bibliometric and visualized analysis. *Front Pharmacol.* 2024;15:1393415.
38. Wang R, Zhuang C. Graph neural networks driven acceleration in drug discovery. *Acta Pharm Sin B.* 2025;15(12):6163-77.
39. Wang Y, Wang T, Li S, He X, Li M, Wang Z, et al. Enhancing geometric representations for molecules with equivariant vector-scalar interactive message passing. *Nat Commun.* 2024;15(1):313.