


Closed-Loop Medicinal Chemistry: Active Learning with Synthesis and Resistance Constraints for Graph Neural Networks for Drug Discovery

Md. Humayan Kabir

May 3, 2026

Abstract

This article explains how graph neural networks can support a closed-loop medicinal chemistry system in which prediction, synthesis, biological testing, and resistance information continuously update one another. Because molecules can be represented as graphs, GNNs are well suited to learning structure–property relationships, but the article argues that ranking compounds by predicted activity alone is insufficient. Active learning is used to balance exploitation of promising candidates with exploration of uncertain or poorly sampled chemical regions. Synthesis feasibility, reaction yield, route accessibility, purification, and building-block availability are incorporated as practical constraints. The framework also includes resistance labels, multi-task learning, multi-fidelity data, uncertainty calibration, scaffold-based evaluation, and diversity-aware batch selection. Three-dimensional representation, stereochemistry, interpretability, experimental turnaround, and provenance are treated as important for real medicinal chemistry. Negative reactions, insolubility, assay interference, and unsuccessful compounds are retained as informative data. The article ultimately evaluates success by prospective hit rate, chemical diversity, synthesis efficiency, resistance retention, and the reduction of low-value experimental cycles.



Molecules are graphs; atoms are nodes, bonds are edges, and neighborhoods contain a large amount of chemical structure information, making graph neural networks an ideal fit for medicinal chemistry. Based on this representation, predictive models have been built that now enable property estimation, prediction of interactions and generation of molecules. However, when the graph model is coupled with an experimental loop, then a new level of usefulness is added. Closed-loop medicinal chemistry uses prediction, uncertainty, synthesis feasibility, biological testing, and resistance information—all of these to teach the model what to do next as it synthesizes a new compound.

The traditional virtual screen typically ranks thousands or millions of molecules, and picks out the highest predicted score. This presupposes the model has confidence in the entire chemical space that is not always the case. Graph neural networks work best when the structures and property regimes they are trained on are similar. Active learning tackles this issue by using the information gain criterion for selecting experiments, in addition to the predicted-winner criterion.

A seed dataset is used to start the loop. The structures are matched to the measured potency, selectivity, ADME properties, toxicity, etc. The GNN learns relationships between molecular graph features and the outcomes. Uncertainty estimation, in turn, marks out areas where predictions are not well-supported, and acquisition functions trade off high expected performance against high information value.

Synthesis constraints need to be applied to the compounds as soon as possible. A molecule

generated could be the perfect molecule in the computer, but not the one that can be synthesized via the provided route, the intermediate is unstable, the building block is not available, or it is difficult to purify. Synthetic accessibility scores are approximate filters, route-based assessment is more robust. The selection of the compounds should not be virtual decorations anymore, but it can do so with the assistance of reaction prediction, retrosynthesis tools, known chemistry and chemist review.

Active learning can be applied to optimizing the reaction. Graph-based models, based on the reaction inputs and conditions, can be used to predict yields or optimum conditions, while adaptive experimentation of temperature, catalyst, solvent, concentration and reagent combinations can be conducted efficiently. This will give a second closed loop that is nested within molecular optimization. The medicinal-chemistry platform learns to make the molecule as well as how to reliably make it.

There's an objective which is specific to resistance data. Potent compounds that are effective in antibacterial, antiviral, antifungal or oncology programs can fail due to various mechanisms such as pathway bypass, target mutation, efflux or metabolism change that restore survival. Labels for resistance are added to cell lines or evolved pathogens that are robust under adaptive pressure. GNN models can then be trained to capture structural characteristics that are also correlated with the retention of activity against resistant phenotypes as well as baseline potency.

This more holistic perspective is furthered by multi task learning. A different model can be used to predict potency for multiple targets, strains, or cell states, and another head will predict solubility, permeability, metabolism, or toxicity. In the case of varying-sized datasets, information can be transferred using shared graph representations. The problem is not being overwhelmed by a lot of large, easy tasks, in favor of a few small, but important resistance tasks or selectivity tasks.

Another way is multi-fidelity learning. The larger high-throughput screening databases have labels that are 'noisy' or less accurate, and the smaller confirmatory assays have more accurate measurements. Transfer-learning experiments have demonstrated that this hierarchy can be useful for graph models when available low-fidelity data can be used to build a representation of the chemicals and a few high-fidelity data can be then used to fine-tune the predictions. This then leads to closed-loop selection of compounds for the costly confirmatory action.

Uncertainty calibration is crucial as it is essential for active learning to know what the model does not know. Different methods of confidence estimation include deep

ensembles, Bayesian approximations, conformal prediction and distance-aware methods. Only held-out chemical series and scaffold splits should be tested for calibration, not random splits. Close analogues tend to leak in between the training and test sets, which can cause over-estimating of performance with random splitting.

There is a chemical diversity as well in acquisition. It is better to pick ten molecules from a single scaffold that are all uncertain, than to pick a set of molecules, representing several poorly understood regions, that are structurally diverse. Avoids redundant experiments using diversity-aware batch selection. The acquisition function thus trades-off predicted value, uncertainty and structural coverage, subject to a synthesis budget.

The speed of the learning rate in the loop is determined by experimental turnaround. Cycle time is shortened by automated synthesis, parallel purification, rapid analytical characterization, and miniaturized biological assays. Data pipelines must capture failed reactions, insoluble compounds, assay interference and negative biology, not just the uncaught and unwanted ones. The information of failed experiments is related to the boundary between virtual plausibility and laboratory reality.

For mechanisms in which the influence of conformation is important, one should also have stereochemical and three-dimensional awareness when drawing graphs. Two dimensional molecular graph representations generally do not convey the shape, protonation, tautomerism, metal coordination or conformational changes in a molecule. Part of that gap is being filled by three-dimensional GNNs, geometric deep learning, physics-informed descriptors and molecular simulation. The representation that should be used is based on the property to be optimized.

Interpretability helps to build chemist trust. Explanation in terms of attention maps or attributions on the atom level can provide indications on which parts of the structure are responsible for the prediction but this type of explanation should be checked with chemistry. The matched molecular pairs offer a practical test in that they can tell if the predicted contribution is in line with the measured change for all close analogues. Interpretability is best when it results in a falsifiable synthetic hypothesis.

This is because there are resistance constraints that require mechanism data. Relying on the labels of Survival/Death will not be sufficient to design around a resistance. The changes in potency are accounted for by the combination of the target engagement, the sequence of resistant clones, the activity in the pathways and the measurements of efflux. Mechanistic terminology can help the learning system differentiate between target site escape and exposure failure or downstream adaptation.

Governance becomes relevant with increasing length of the loop. Traceability is required for model versions, training datasets, preprocessing decisions and acquisition rules, compound provenance, and assay batches. Failure to keep this record means there will be no way to reproduce subsequent performance and it becomes challenging to audit apparent learning gains. Closed-loop discovery is, in part, a discipline that involves data engineering.

The final word as to whether any closed loop is viable will be given by prospective evaluation. The result of compounds that are designed and synthesized after model freezing should be predictable from a model trained on chemistry in the past. That challenge can be difficult to reproduce with a retrospective cross-validation even if the splits are used in the scaffold. Prospective cycles are used to see if indeed uncertainty estimates, synthesis filters, and acquisition rules are beneficial to experimental efficiency in real discovery.

Difficult-to-statistically encode human chemists also provide priors. The knowledge of those unstable motifs, reactive liabilities, artefacts of the assays, privileged scaffolds and route flexibility affects the decision of whether a computational suggestion should be synthesized. Recording those decisions as structured data can be another source of learning. As time goes on, the platform can differentiate between the algorithmic errors and proposals that are turned down due to practical chemical considerations.

In the case of training data bias, the bias must be actively managed. Success and familiarity of series, reaction, and project team selections are overrepresented in the historical medicinal chemistry data sets. Failing chemistry often isn't there or it is negative molecules. This is at least corrected by active learning that delves into neglected areas, when unsuccessful experiments are recorded. If not, the model keeps repeating itself and learning a sanitized version of the history of chemistry that obscures the real price of exploration.

The decision metrics should be based on the project value and should not be based on benchmark accuracy. It is not a major difference if the same compounds are used for synthesis with a small improvement in the root-mean-square error. More meaningful metrics are: hit rate per synthesized compound, reduction in the number of experimental cycles, diversity of viable leads, retention of resistance, and number of failures in reaction avoided. Closed loop is not just about numbers.

A disciplined loop also sets stopping criteria, so that the synthesis effort does not get wasted on regions of the molecule worthless to pursue.

Strongest closed loop platform isn't to automate medicinal chemistry out of existence, it's to organize decision making. Possibilities are the point of interest for graph neural networks, uncertainty is the focus of experiments, constraints of synthesis are made grounded in proposals, and resistance data are the focus of durable activity. Uncertainty decreases and the quality of candidates increases with each cycle. The large virtual library is no longer the main objective of drug discovery, but rather a selection of the fewest experiments that will be most effective in learning the chemistry to build the next one.