TRUSTWORTHY RETROSYNTHESIS: ELIMINATING HAL-LUCINATIONS WITH A DIVERSE ENSEMBLE OF REAC-TION SCORERS

Anonymous authors

000

001

002

004

006

008 009 010

011 012 013

014

016

018

019

020

021

024

025

026

027

028

029

031 032 033

034

037

038

040

041

043

044

046

047

048

050 051

052

Paper under double-blind review

ABSTRACT

Retrosynthesis is one of the domains transformed by the rise of generative models, and it is one where the problem of nonsensical or erroneous outputs (hallucinations) is particularly insidious: reliable assessment of synthetic plans is time-consuming, with automatic methods lacking. In this work, we present RetroTrim, a retrosynthesis system that successfully avoids nonsensical plans on a set of challenging drug-like targets. Compared to common baselines in the field, our system is not only the sole method that succeeds in filtering out hallucinated reactions, but it also results in the highest number of high-quality paths overall. The key insight behind RetroTrim is the combination of diverse reaction scoring strategies, based on machine learning models and existing chemical databases. We show that our scoring strategies capture different classes of hallucinations by analyzing them on a dataset of labeled retrosynthetic intermediates. To measure the performance of retrosynthesis systems, we propose a novel evaluation protocol for reactions and synthetic paths based on a structured review by expert chemists. Using this protocol, we compare systems on a set of 32 novel targets, curated to reflect recent trends in drug structures. While the insights behind our methodology are broadly applicable to retrosynthesis, our focus is on targets in the drug-like domain. By releasing our benchmark targets and the details of our evaluation protocol, we hope to inspire further research into reliable retrosynthesis.

1 Introduction

The advent of deep generative modeling has transformed a broad range of domains, and resulted in impressive applications such as photorealistic image synthesis, code generation, and automated theorem proving (Rombach et al., 2022; Li et al., 2022; Yang et al., 2023). The undisputed efficacy of generative models is nonetheless hindered by the possibility of factually wrong or outright nonsensical output, often called "hallucinations" Sahoo et al. (2024). In some applications, such as automated theorem proving, hallucinations can be curtailed with formal verification of output. In others, they can only be mitigated through partial and often imprecise verification.

Retrosynthesis, the task of constructing synthetic routes — sequences of chemical reactions that lead to a desired target molecule from simpler precursors — is another domain that has undergone significant developments with the rise of generative modeling (Coley et al., 2018). Modern systems typically generate complete synthetic routes by iteratively applying predictions from a single-step retrosynthesis (SSR) model within a graph-based search algorithm. When reactions have no known chemical precedent or display features such as unstable reagents, they jeopardize the validity of the synthetic path they are part of. Generative SSR models inherit the limitations of generative modeling: the reactions they propose are frequently hallucinated, that is, chemically unsound (Torren-Peraire et al., 2024). An example of a hallucinated reaction is shown in Figure 1

In this work, we propose RetroTrim, a retrosynthesis system designed to tackle the issue of hallucinated reactions by pairing a performant SSR generator with an ensemble of complementary reaction scorers as an in-scope filter. We show that RetroTrim is the only among common retrosynthetis solutions that successfully avoids *all* hallucinated reactions in proposed synthetic plans for a set of

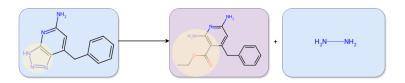


Figure 1: An example of grossly incorrect (hallucinated) reaction generated by a Single-Step Retrosynthesis model. A PhD-level chemist recognizes that the only reasonable atom mapping between the substrates and the product is one where the reaction center is an ortho-amino benzoate converting into a triazole (highlighted in yellow). It does not belong to any commonly known reaction class, and further investigation involving extensive searches of synthetic databases yields no examples that would inform what reagents and conditions could induce such a reaction. Executing this transformation would be impractical and require the development of a novel synthetic methodology, which typically entails a multi-month research program.

unpublished challenging drug-like targets. Moreover, RetroTrim does so while leading in the number of targets for which a synthetic plan without issues is found.

Our approach in developing RetroTrim is based on the principle of ensemble learning Hansen & Salamon (1990): combining diverse scorers with distinct error patterns yields a more robust and accurate assessment of reaction plausibility. The ensemble scorer (meta-scorer) functions as an external validator of the generative SSR model. It prunes all reactions below a given score threshold from the synthesis tree, curtailing expansion of hallucinated nodes. For an overview see Fig. 2. The three scorers which make up the in-scope filter are as follows:

- Reaction Prior (RP): A Transformer-based architecture Vaswani et al. (2017) whose scoring
 function is designed to mimic the considerations chemists take into account when evaluating
 reaction plausibility. Its training and scoring method lends itself to discovering a broad
 spectrum of hallucinations.
- 2. Reaction Graph Plausibility (RGP): A graph model trained to distinguish positive/feasible reactions from synthetically generated negatives. Negative reactions are generated by applying reaction templates at random in both the forward and retro- direction. The forward negatives are designed to increase fidelity in discriminating selectivity problems, while retro- negatives lead of a broader spectrum of incompatibilities between functional groups in reactants.
- 3. Reaction Retrieval Score (RRS): A mechanism that assesses the similarity of proposed reactions to known experimental precedents in reaction databases. It is designed to catch hallucinations of the more gross kind: reactions with no precedents and significant mismatches between the target and reactants.

We compare performances of RP, RGP and RRS on a dataset of retrosynthetic intermediates. We show that each scorer excels at filtering different kinds of hallucinations, and that the meta-scorer further improves performance in most areas.

To measure the performance of both retrosynthesis systems as well as reaction scorers, we employ a novel reaction validation protocol in which PhD-level chemists assess each reaction in a given synthetic plan according to a predefined labeling schema. Reactions are sorted into seven categories depending on the kind of issue they raise: *Magic*, *Selectivity*, *Functional group incompatibility*, *Reactivity*, *One pot*, *Unstable*, or *Reactants mismatch*; and according to the severity of said issue: *Worthwhile*, *Rather not*, and *Nonsense*. Reactions which present no issues are given the special labels *No problem* and *Safe bet*, respectively. Under this categorization, *Nonsense* are the reactions that correspond to hallucinations, while *Worthwhile* and *Rather not* reactions contain mild or moderate issues which makes them undesirable or unpredictable. The overall confidence in a synthetic pathway is determined by the reaction with the highest issue severity within it.

All retrosynthesis systems are evaluated on a set of thirty-two challenging drug-like targets. These targets were defined so as to match common structural features present in modern drugs while not

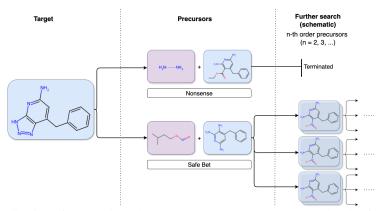


Figure 2: Visualization of RetroTrim. The generator proposes precursor molecules for a given target, which are then evaluated for reaction plausibility. Plausible reactions proceed to the next retrosynthetic step for non-commercially available precursors (blue boxes), while implausible reactions terminate the branch. The search concludes when a complete synthetic route from market-available starting materials (purple boxes) to the target molecule is found. The displayed reactions originate from retrosynthetic pathways generated by our systems and the baselines. The "Safe Bet" route is strongly supported by literature with multiple documented precedents, while the "Nonsense" route lacks any published examples.

being close analogs of any structure found in synthetic literature. This ensures no data leakage occurred in-between training and inference of compared models, while still being representative of the medicinal chemistry domain.

Our main contributions are as follows:

- We propose RetroTrim which achieves state-of-the-art accuracy, rejecting as the only method
 all hallucinated reactions for a set of challenging targets while at the same time finding
 flawless synthetic routes for the largest number of them.
- We demonstrate that the three scorers used in RetroTrim as reaction filters display different strengths across potential issues with reactions, and furthermore that they are complementary: combining them leads to the best-performing system.
- We outline a reaction labeling protocol that recognizes seven kinds of possible issues at three levels of severity, to be used by expert chemists for granular evaluation of both reactions and the synthetic routes composed of them.
- We release a set of 32 unpublished drug-like targets, designed as a challenging test set for retrosynthesis systems.

2 Related Work

To contextualize our proposed method, this section examines the workflow in automated retrosynthesis, covering both the generation of reactions by Single-Step Retrosynthesis (SSR) models and techniques used for their validation.

2.1 Multi-Step Retrosynthesis

Multi-step retrosynthesis constructs complete synthetic routes by iteratively applying single-step retrosynthesis (SSR) predictions. Classical search strategies, such as Monte Carlo tree search Segler et al. (2018) and A* search (Retro*) Chen et al. (2020), use SSR models to expand compounds during generation of synthetic pathways.

SSR models can be categorized into template-based, template-free, and semi-template-based approaches. Template-based methods (e.g., RetroSim (Coley et al., 2017), NeuralSym (Segler & Waller, 2017), GLN (Dai et al., 2019a)) use predefined reaction templates, ensuring interpretability, but

their coverage is limited and their construction involves trade-offs that can lead to invalid template application. Template-free models (e.g., sequence-to-sequence Transformers (Zhong et al., 2022), GNNs (Liu et al., 2017; Karpov et al., 2019; Sacha et al., 2021)) learn chemical rules directly from data, offering scalability and the ability to generalize to novel reactions, but with little guarantee as to the validity of output. Semi-template-based approaches (e.g., GraphRetro (Somnath et al., 2021), RetroXpert (Yan et al., 2020)) combine templates with learned representations, inheriting both the advantages and limitations of both approaches. Regardless of their category, SSR models are all prone to erroneous output Coley et al. (2018).

Evaluating single-step retrosynthesis (SSR) models remains challenging. Traditional metrics, such as top-k accuracy, measure whether the correct reactants appear among the top predicted candidates, but they provide limited insight into chemical plausibility and alternative valid reactions.

To address this, round-trip accuracy (Schwaller et al., 2020) is commonly used. This metric evaluates a predicted reaction by passing the predicted reactants through a forward reaction model and checking if the original product is recovered. Round-trip accuracy is considered a better proxy for chemical validity than top-k accuracy.

However, it remains unclear how well round-trip accuracy reflects the actual plausibility of the predicted reactions. There is currently no exhaustive data quantifying how many predictions that pass or fail the round-trip check correspond to chemically plausible reactions.

Evaluating multi-step retrosynthetic routes is typically limited to basic pathway properties, such as the number of steps, route length, branching factor, or overall synthesizability (Maziarz et al., 2025). These metrics provide little insight into the actual correctness or chemical plausibility of the complete synthetic route.

2.2 MID-SEARCH REACTION VALIDATION

The plausibility of multi-step retrosynthetic routes has drawn increasing attention from researchers, leading to the development of methods aimed at improving pathway validity. A commonly used approach is to employ a feasibility model as a filter. Such models are typically trained as classifiers to distinguish correct reactions from artificially generated negative examples, which are often created by perturbing positive reactions, for example by randomly modifying substrates or swapping the original product with a chemically similar one (Segler et al., 2018; Genheden et al., 2020a). A forward model (predicting products based on substrates) trained only on positive data was used in the same way in IBM RXN Schwaller et al. (2020).

Another strategy is to use the likelihood of a generated reaction as a scoring function during search. This can be computed from template-free models, such as the Molecular Transformer (Schwaller et al., 2019) architecture, where the model's confidence, derived from output token probabilities, has been shown to correlate with reaction correctness, or from template-based models using softmax scores, as in RetroFallback (Tripp et al., 2023). In order to improve the quality of predicted reactions, RetroGFN (Gaiński et al., 2024) was trained using feedback from a forward model and round-trip accuracy.

Finally, evidence-based validation via retrieval grounds model predictions in established chemical knowledge, mirroring a chemist's workflow of searching for literature precedents. Retrieval-augmented methods, such as the Retrieval-Augmented RetroBridge (RARB) framework, retrieve similar molecules from a database to guide the generation of reactants (Qiao et al., 2025). These approaches help ensure that generated reactions are consistent with known chemistry, further improving the plausibility of multi-step routes. In the context of these works, RetroTrim is the first to use retrieved reactions to filter predictions of a reaction generation.

To date, none of these approaches have been extensively validated through human expert evaluation.

3 METHODS

RetroTrim is built around a central, best-performing scorer called Reaction Prior, supported by two additional scorers designed to compliment it by targeting specific error types. These three scorers are aggregated into a Meta-Scorer that provides the final, robust assessment of reaction correctness.

3.1 REACTION PRIOR

The Reaction Prior (RP) is a novel method inspired by how experienced chemist reason about reactions: considering the reaction globally, assessing the reaction center, and comparing it to alternative reactions that could occur. The RP score S_{final} is thus a weighted combination of three components: $S_{final} = S_{RP}^{\alpha} \cdot S_{Regio}^{\beta} \cdot S_{RC}^{\gamma}$. Here, S_{RP} is the Reaction Prior (global) Score, S_{RC} is the Reaction Center Score, and S_{Regio} is the Regioselectivity Score, with α , β , and γ serving as weighting factors. RP is implemented as an autoregressive, encoder-decoder BART (Lewis et al., 2019) architecture, where substrates and products are processed by the decoder, trained for next-token prediction by minimizing cross-entropy loss.

Reaction Prior Score (S_{RP}) This score reflects the overall feasibility of the reaction. It is the log probability assigned by the model to the reaction sequence, normalized by the square root of the total number of tokens (T): $S_{RP} = \frac{1}{\sqrt{T}} \log P(\text{reaction})$.

Reaction Center Score (S_{RC}) This score evaluates the model's confidence in the identified reactive sites. It is the sum of log probabilities for the tokens representing atoms at the reaction's core, normalized by the number of such tokens (T_{RC}) : $S_{RC} = \frac{1}{T_{RC}} \sum_{i \in \text{reaction center}} \log P(\text{token}_i)$.

Regioselectivity Score (S_{Regio}) This component quantifies reaction site specificity. It is calculated by comparing the probability of the desired reaction (P_{desired}) to the summed probabilities of all alternative reactions at different potential sites $(P_{\text{undesired}})$: $S_{Regio} = \log\left(\frac{P_{\text{desired}}}{P_{\text{undesired}}+\epsilon}\right)$, where ϵ is a small constant to prevent division by zero.

3.2 REACTION GRAPH PLAUSIBILITY

A Graph Attention Network (GAT) (Veličković et al., 2017) is trained to differentiate chemically valid reactions from implausible ones. Training uses reaction datasets for positive examples and synthetic negative examples generated through forward and two-step backward template applications. This approach is similar to those proposed by (Segler et al., 2018; Genheden et al., 2020a), but it uses a graph neural network instead of a feedforward network with fingerprint inputs, and employs more sophisticated artificial negative reactions, generated not only by applying random templates in the forward direction but also through retro-synthetic random template applications, which increases the diversity of types of generated incorrect reactions. Details of GAT featurization are described in Appendix B.

3.3 Reference Reaction Scorer

We developed a structured retrieval pipeline that extracts chemical precedent information through a two-tiered reaction clustering procedure based on bond change patterns. First *Coarse-grained clustering* extracts connected components of the reaction center and applies atom mapping to identify the underlying transformation pattern. Reactions belong to the same cluster if their transformation patterns are identical. Then *Fine-grained clustering* extends the coarse-grained approach by incorporating chemically significant substructures — aromatic systems and conjugated double bonds - into the cluster classification.

Our Reference Reaction Retrieval Scorer (RRS) quantifies reaction plausibility through a logarithmic transformation of the unique reference reaction count within the candidate reaction's coarse-grained cluster and fine-grained cluster:

$$p(reaction) = \log(n_{ref}(reaction) + 1) \tag{1}$$

where $n_{ref}(reaction)$ represents the unique number of reference reactions in the coarse-grained and fine-grained clusters containing reaction.

3.4 Meta-Scorer Aggregation

To improve reaction filtering, our Meta-Scorer integrates scores from Reaction Prior (RP) and RGP models, and empirical precedents ($n_{\text{ref}} > 0$) retrieved via the pipeline in Sec. 3.3. This hybrid

approach mitigates the weaknesses of purely data-driven or precedent-based methods. The continuous score is described by equation $score_{META} = max(score_{RGP}, score_{RP})$ if $n_{ref} > 0$ (0 otherwise).

For binary classification tasks and search, reactions are filtered using predefined thresholds, selected through grid search to balance the recall and precision:

$$score_{META} = \begin{cases} 1 & \text{if } score_{RGP} > thr_{RGP} \text{ and } score_{RP} > thr_{RP} \text{ and } n_{ref} > 0 \\ 0 & \text{otherwise} \end{cases}$$
 (2)

By synthesizing diverse evidence types Meta-Scorer enables more reliable reaction filtering for multi-step synthesis planning. This integrated approach mitigates individual weaknesses of purely data-driven or precedent-based methods, yielding improved performance.

3.5 INTEGRATION WITH SEARCH (RETRO*)

The calibrated Meta-Scorer is used during multi-step retrosynthesis search to improve the quality of the pathways predicted by the BART generator. We integrate the scorer into the Retro* (Chen et al., 2020) search framework as a reaction filtering mechanism. Reactions are pruned from the search tree if score_{META} defined in 2 is equal to 0.

4 HUMAN EVALUATION

We curated a dataset of over 4,500 reactions generated by our SSR models. Each reaction was evaluated and labeled by PhD-level chemists into one of the expert-defined categories, creating the first comprehensive dataset of its kind. This datased provides a way to evaluate the error patterns of our reaction plausibility scorers. A subset of this dataset will be released as a benchmark for the community.

Reaction Evaluation Protocol was designed to systematically evaluate predicted reactions based on expert-defined heuristics. Reactions were rated using a four-point confidence scale: *Nonsense, Rather not, Worthwhile,* and *Safe bet. Safe bet* reactions are considered fully reliable; *Worthwhile* reactions remain plausible but carry a moderate risk of failure; *Rather not* reactions are associated with a high risk of major difficulties; and *Nonsense* reactions are effectively infeasible, i.e. hallucinated. For the system to be reliable, valid pathways should consist primarily of *Safe bets* reactions. The presence of *Nonsense* reactions effectively invalidates a pathway, while the presence of a *Rather Not* reaction may still be acceptable in target-oriented synthesis when no alternatives exist. Reactions which aren't a *Safe bet* receive an additional label specifying the cause of their incorrectness, chosen from: *Reactants mismatch, Unstable, Magic, One pot, Reactivity, Functional group incompatibility*, and *Selectivity*. These error categories correlate with confidence levels to varying degrees: for example, *Magic* errors almost always map to *Nonsense*, while *Selectivity* issues more often correspond to *Worthwhile* or *Rather Not*. Otherwise, a reaction is assigned a *No Problem* label. A detailed description of the evaluation framework is provided in Appendix C.

5 EXPERIMENTS

5.1 Dataset

All of the models are based on the proprietary Pistachio (Mayfield et al. (2017)) dataset, either used as training data (generator, RP, RGP), or as a source of reference reactions (RRS). Pistachio offers substantial advantages over the commonly used USPTO-50K Schneider et al. (2016) and USPTO-FULL Dai et al. (2019b) datasets - it features enhanced curation, resulting in higher data quality and more comprehensive coverage of chemical reaction space. For training, we preprocess the dataset through a multi-step filtering pipeline that removes duplicate reactions, reactions from unrecognized reaction classes, entries with invalid SMILES, unmapped reactions, and reactions deemed unrelated to drug-like compound synthesis (too big molecules, "separation" reaction classes), retaining approximately 4 million reactions.

5.2 PATH-LEVEL PLAUSIBILITY EVALUATION

We compared the ability of each individual scorer (RP, RGP, and RRS) as well as the Meta-Scorer to filter out implausible reactions in multi-step retrosynthesis. We applied the scorers to a Transformer model trained on root-aligned reactions (Zhong et al., 2022) as a single-step generator (RootAligned).

The thresholds for filtering implausible reactions were selected based on results from the full chemistevaluated reaction dataset described in 4, targeting 0.8 precision, which translated into approximately 90% of retrosynthetic pathways being successfully found. As baselines, we report results for RootAligned, for LocalRetro (Chen & Jung, 2021) without scoring, and for RootAligned with forward reaction scorer, enabling direct comparison with prior literature. For the search, we used the widely adopted Retro* algorithm basing on implementation from (Maziarz et al., 2025) with expansions limit set to 500. For all systems, we used the same starting material database, eMolecules. Additionally, we compared against two open-source retrosynthetic systems: IBM RXN (IBM, 2025) and AiZynthFinder (Genheden et al., 2020b). AiZynthFinder was executed in its default configuration from the official repository (template generator + filtering based on a model trained on positive and artificial negative reactions), with the only modification being an increased time limit of 15 minutes to better match the runtime of other systems. IBM RXN was accessed through its free web application (IBM, 2025). We evaluated the top-1 synthesis paths generated by all systems for 32 selected targets (listed in D). Each path was assigned a four-tier confidence score (Safe bet, Worthwhile, Rather not, Nonsense), determined by the lowest-scoring reaction in the path. This conservative scoring reflects the intuition that a single implausible step can invalidate an otherwise promising synthesis. Increasing the proportion of Safe bets, while eliminating Nonsense and reducing Rather Not paths is the goal of all retrosynthesis systems.

5.3 REACTION-LEVEL PLAUSIBILITY PREDICTION

Additionally, we compare the performance of individual scorers (RP, RGP and RRS) and the Meta-Scorer on individual reactions with ground truth labels established through expert chemist evaluations described in Section 4.

Model performance was assessed using precision-recall (PR) and receiver operating characteristic (ROC) curves, with area under the curve metrics (PR-AUC and ROC-AUC) reported for each method. Reactions with confidence rating *Safe Bet* were treated as positive examples. *Worthwhile* reactions were excluded from the test set as they represent borderline cases where chemist confidence is uncertain, making them neither clearly positive nor negative examples for evaluation purposes. All others (*Rather Not* and *Nonesense*) were labeled as negatives. We also conducted additional analysis across each failure category, reporting individual ROC-AUC and PR-AUC scores, as well as false positive counts.

To evaluate model complementarity, we analyzed the overlap in false positive predictions across individual scorers, calculated as:

$$overlap = \frac{\left| \bigcap_{scorer \in \{RGP,RP,RRS\}} FP_{scorer} \right|}{\min_{scorer \in \{RGP,RP,RRS\}} |FP_{scorer}|},$$
(3)

where FP is a set of false positives produced by a given scorer.

6 RESULTS

6.1 PATH-LEVEL PLAUSIBILITY EVALUATION

Pathway correctness comparison is presented in the figure 3. AiZynthFinder demonstrates significant limitations, failing to identify viable pathways for significant number of the target molecules while generating a substantial proportion of unreliable routes classified as *Nonsense* and *Rather Not*. IBM RXN shows improved performance by increasing the number of reliable pathways and reducing hallucinated predictions, yet fails to produce valid synthetic routes for a considerable fraction of target compounds.

Our baseline model (SSR generator without any filtering mechanism) significantly improves number of pathways found, providing solutions for all targets. However, confidence in its results is undermined



Figure 3: Comparison of our retrosynthesis systems against IBM RXN, AiZynthFinder and Local-Retro. AiZynthFinder, IBM RXN and LocalRetro fail on many targets. Our generator when used without any reaction scorer (indicated as "Our generator w/o scorer") finds pathways for all targets but includes unreliable routes. Introduction of individual scorers trades coverage for reliability, with RP eliminating all *Nonsense* pathways. The search system backed by the Meta-Scorer produces most trustworthy results.

by the significant presence of unreliable *Nonsenense* and *Rather Not* paths. Introducing individual scorers increases the fraction of targets for which no paths are found, a trade-off that can be desirable for the trustworthiness of the system — rejecting some targets is preferable to mixing reliable and unreliable pathways, as long as the remaining routes are correct. While RGP and RRS scorers reduce number of unreliable paths only modestly, our RP scorer demonstrates its value as a primary filter by eliminating all *Nonsense* reactions, though this comes at the cost of fewer *Safe Bet* and *Worthwhile* pathways. Finally, the Meta-Scorer delivers substantial improvements in reliability: significantly increasing *Safe Bet* paths, maintaining zero *Nonsense* results, and reducing *Rather Not* pathways. These results indicate that relying solely on individual scorers, especially those similar to commonly used approaches in the literature (such as RGP - feasibility classifier and forward scorers), is insufficient to achieve correct pathways. Only our RP-based solution combined with the Meta-Scorer provides routes with the highest overall plausibility while eliminating the most serious errors.

6.2 REACTION-LEVEL PLAUSIBILITY EVALUATION

Our results show that the Meta-Scorer outperforms individual scorers in both precision and recall, demonstrating effective integration of complementary signals. Figure 4 presents the ROC and precision-recall curves, with the Meta-Scorer achieving consistently higher area under the curve (AUC) values across both metrics. Similar curves broken down by reaction failure category can be found in Appendix E.

Figure 5 shows ROC-AUC values for each scorer broken down into different failure categories, illustrating that individual scorers demonstrate proficiency in filtering out reactions deemed implausible under different evaluation criteria. RGP achieves the best performance on *Selectivity* and *Reactivity* errors. RRS is most capable of detecting fundamental structural issues such as *Reactant mismatches* and *Magic*, in addition to *One pot* errors. RP shows a balanced profile, which explains its overall superior performance compared to RGP and RRS in Figure 4. By leveraging the unique strengths of each individual scorer, the Meta-Scorer maintains robust predictive performance across all failure categories.

We also analyze the overlap between false positives that each scorer fails to filter, as shown in Figure 6. The results show distinct complementarity: while RRS and RP exhibit high overlap in most categories, it is notably reduced for *One-pot*, *Magic*, and *Reactant mismatch* failure modes — the categories where Figure 5 demonstrates RRS's superior performance. RGP and RP show consistently low overlap across all failure categories, indicating that these scorers capture different aspects of reaction implausibility. Importantly, when considering all three scorers jointly, the overlap drops to

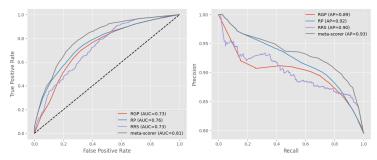
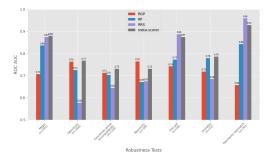


Figure 4: ROC (on the left) and precision-recall (on the right) curves comparing the performance of individual scorers versus the Meta-Scorer. The meta-scorer achieve higher AUC values for both ROC and PR curves, indicating better discrimination between plausible and implausible reactions. Among the individual scorers, RP shows the best performance.

very low levels across all categories, providing strong evidence that each scorer contributes unique discriminative value essential for building a robust Meta-Scorer.



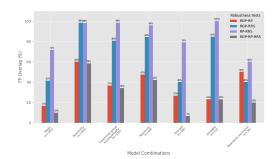


Figure 5: ROC-AUC performance of individual scorers across different failure categories, with sample sizes indicated for each category.

Figure 6: Overlap between individual pairs of scorers and triple of all scorers across different failure categories, with sample sizes indicated for each category.

7 CONCLUSIONS

In this work, we introduced RetroTrim, a retrosynthesis system designed to avoid hallucinations in the synthesis tree through a combination of three complementary reaction scoring strategies. We demonstrated its success on thirty-two unpublished drug-like targets, where no generated paths contained hallucinated reactions. Among the available methods we compared RetroTrim with, our method was the only one to completely avoid hallucinations, while at the same time finding more paths without issues than other methods. To understand the strengths and complementarity of each scoring strategy, we compared their performance across different classes of possible issues. We found evidence of synergy between the scorers, both at the level of filtering individual reaction, and in terms of the retrosynthetic paths resulting from their use.

In evaluating retrosynthesis systems and scorers, we made use of a novel labeling protocol where we leveraged chemists' expertise to produce fine-grained labels for generated reactions. To our knowledge, this is the first instance of such a granular analysis of retrosynthetis systems' output, where automated metrics and ad-hoc manual inspection were the norm. In order to facilitate further development in the field, we release the thirty-two targets used for path generation. While our evaluation process is generally applicable to retrosynthesis, RetroTrim was trained on data that biases it towards the medicinal chemistry context. Nonetheless, we hope that the insights and methodologies presented in this work lead to more reliable retrosynthesis in general.

REFERENCES

- Binghong Chen, Chengtao Li, Hanjun Dai, and Le Song. Retro*: learning retrosynthetic planning with neural guided a* search. In *Proceedings of the 37th International Conference on Machine Learning*, ICML'20. JMLR.org, 2020.
- Shuan Chen and Yousung Jung. Deep retrosynthetic reaction prediction using local reactivity and global attention. *JACS Au*, 1(10):1612–1620, 2021.
- Connor W. Coley, Luke Rogers, William H. Green, and Klavs F. Jensen. Computer-assisted retrosynthesis based on molecular similarity. *ACS Central Science*, 3(12):1237–1245, 2017. doi: 10.1021/acscentsci.7b00355. URL https://doi.org/10.1021/acscentsci.7b00355. PMID: 29296663.
- Connor W. Coley, William H. Green, and Klavs F. Jensen. Machine learning in computer-aided synthesis planning. *Accounts of chemical research*, 51 5:1281-1289, 2018. URL https://api.semanticscholar.org/CorpusID:13748494.
- Hanjun Dai, Chengtao Li, Connor Coley, Bo Dai, and Le Song. Retrosynthesis prediction with conditional graph logic network. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019a. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/0d2b2061826a5df3221116a5085a6052-Paper.pdf.
- Hanjun Dai, Chengtao Li, Connor Coley, Bo Dai, and Le Song. Retrosynthesis prediction with conditional graph logic network. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019b. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/0d2b2061826a5df3221116a5085a6052-Paper.pdf.
- Piotr Gaiński, Michał Koziarski, Krzysztof Maziarz, Marwin Segler, Jacek Tabor, and Marek Śmieja. Retrogfn: Diverse and feasible retrosynthesis using gflownets. *arXiv preprint arXiv:2406.18739*, 2024.
- Samuel Genheden, Ola Engkvist, and Esben Jannik Bjerrum. A quick policy to filter reactions based on feasibility in ai-guided retrosynthetic planning. 2020a.
- Samuel Genheden, Amol Thakkar, Veronika Chadimová, Jean-Louis Reymond, Ola Engkvist, and Esben Bjerrum. Aizynthfinder: a fast, robust and flexible open-source software for retrosynthetic planning. *Journal of cheminformatics*, 12(1):70, 2020b.
- L.K. Hansen and P. Salamon. Neural network ensembles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(10):993–1001, 1990. doi: 10.1109/34.58871.
- IBM. Rxn for chemistry. https://rxn.app.accelerate.science/, 2025. Accessed: 2025-08-26.
- Pavel Karpov, Guillaume Godin, and Igor V. Tetko. A transformer model for retrosynthesis. In Igor V. Tetko, Věra Kůrková, Pavel Karpov, and Fabian Theis (eds.), *Artificial Neural Networks and Machine Learning ICANN 2019: Workshop and Special Sessions*, pp. 817–830, Cham, 2019. Springer International Publishing. ISBN 978-3-030-30493-5.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *CoRR*, abs/1910.13461, 2019. URL http://arxiv.org/abs/1910.13461.
- Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, Thomas Hubert, Peter Choy, Cyprien de Masson d'Autume, Igor Babuschkin, Xinyun Chen, Po-Sen Huang, Johannes Welbl, Sven Gowal, Alexey Cherepanov, James Molloy, Daniel J. Mankowitz, Esme Sutherland Robson, Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu, and Oriol Vinyals. Competition-level code generation with alphacode. *Science*, 378(6624):1092–1097, 2022. doi: 10.1126/science.abq1158. URL https://www.science.org/doi/abs/10.1126/science.abq1158.

Bowen Liu, Bharath Ramsundar, Prasad Kawthekar, Jade Shi, Joseph Gomes, Quang Luu Nguyen, Stephen Ho, Jack Sloane, Paul Wender, and Vijay Pande. Retrosynthetic reaction prediction using neural sequence-to-sequence models. *ACS Central Science*, 3(10):1103–1113, 2017. doi: 10.1021/acscentsci.7b00303. URL https://doi.org/10.1021/acscentsci.7b00303. PMID: 29104927.

- John Mayfield, Daniel Lowe, and Roger Sayle. Pistachio: Search and faceting of large reaction databases. In *ABSTRACTS OF PAPERS OF THE AMERICAN CHEMICAL SOCIETY*, volume 254. AMER CHEMICAL SOC 1155 16TH ST, NW, WASHINGTON, DC 20036 USA, 2017.
- Krzysztof Maziarz, Austin Tripp, Guoqing Liu, Megan Stanley, Shufang Xie, Piotr Gaiński, Philipp Seidl, and Marwin H. S. Segler. Re-evaluating retrosynthesis algorithms with syntheseus. *Faraday Discuss.*, 256:568–586, 2025. doi: 10.1039/D4FD00093E. URL http://dx.doi.org/10.1039/D4FD00093E.
- Anjie Qiao, Zhen Wang, Jiahua Rao, Yuedong Yang, and Zhewei Wei. Advancing retrosynthesis with retrieval-augmented graph generation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(19):20004–20013, Apr. 2025. doi: 10.1609/aaai.v39i19.34203. URL https://ojs.aaai.org/index.php/AAAI/article/view/34203.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Bjorn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 10674–10685, Los Alamitos, CA, USA, June 2022. IEEE Computer Society. doi: 10.1109/CVPR52688.2022.01042. URL https://doi.ieeecomputersociety.org/10.1109/CVPR52688.2022.01042.
- Mikołaj Sacha, Mikołaj Błaż, Piotr Byrski, Paweł Dąbrowski-Tumański, Mikołaj Chromiński, Rafał Loska, Paweł Włodarczyk-Pruszyński, and Stanisław Jastrzębski. Molecule edit graph attention network: Modeling chemical reactions as sequences of graph edits. *Journal of Chemical Information and Modeling*, 61(7):3273–3284, 2021. doi: 10.1021/acs.jcim.1c00537. URL https://doi.org/10.1021/acs.jcim.1c00537. PMID: 34251814.
- Pranab Sahoo, Prabhash Meharia, Akash Ghosh, Sriparna Saha, Vinija Jain, and Aman Chadha. A comprehensive survey of hallucination in large language, image, video and audio foundation models. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), Findings of the Association for Computational Linguistics: EMNLP 2024, pp. 11709–11724, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.685. URL https://aclanthology.org/2024.findings-emnlp.685/.
- Nadine Schneider, Nikolaus Stiefl, and Gregory Landrum. What's what: The (nearly) definitive guide to reaction role assignment. *Journal of Chemical Information and Modeling*, 56, 11 2016. doi: 10.1021/acs.jcim.6b00564.
- Philippe Schwaller, Teodoro Laino, Théophile Gaudin, Peter Bolgar, Christopher A. Hunter, Costas Bekas, and Alpha A. Lee. Molecular transformer: A model for uncertainty-calibrated chemical reaction prediction. *ACS Central Science*, 5(9):1572–1583, 2019. doi: 10.1021/acscentsci.9b00576. URL https://doi.org/10.1021/acscentsci.9b00576. PMID: 31572784.
- Philippe Schwaller, Riccardo Petraglia, Valerio Zullo, Vishnu H. Nair, Rico Andreas Haeuselmann, Riccardo Pisoni, Costas Bekas, Anna Iuliano, and Teodoro Laino. Predicting retrosynthetic pathways using transformer-based models and a hyper-graph exploration strategy. *Chem. Sci.*, 11:3316–3325, 2020. doi: 10.1039/C9SC05704H. URL http://dx.doi.org/10.1039/C9SC05704H.
- Marwin H. S. Segler and Mark P. Waller. Neural-symbolic machine learning for retrosynthesis and reaction prediction. *Chemistry A European Journal*, 23(25):5966–5971, 2017. doi: https://doi. org/10.1002/chem.201605499. URL https://chemistry-europe.onlinelibrary.wiley.com/doi/abs/10.1002/chem.201605499.
- Marwin H. S. Segler, Mike Preuss, and Mark P. Waller. Planning chemical syntheses with deep neural networks and symbolic ai. *Nature*, 555(7698):604–610, Mar 2018. ISSN 1476-4687. doi: 10.1038/nature25978. URL https://doi.org/10.1038/nature25978.

Vignesh Ram Somnath, Charlotte Bunne, Connor Coley, Andreas Krause, and Regina Barzilay. Learning graph models for retrosynthesis prediction. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 9405–9415. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/4e2a6330465c8ffcaa696a5a16639176-Paper.pdf.

Paula Torren-Peraire, Alan Kai Hassen, Samuel Genheden, Jonas Verhoeven, Djork-Arné Clevert, Mike Preuss, and Igor V. Tetko. Models matter: the impact of single-step retrosynthesis on synthesis planning. *Digital Discovery*, 3:558–572, 2024. doi: 10.1039/D3DD00252G. URL http://dx.doi.org/10.1039/D3DD00252G.

Austin Tripp, Krzysztof Maziarz, Sarah Lewis, Marwin Segler, and José Miguel Hernández-Lobato. Retro-fallback: retrosynthetic planning in an uncertain world. *arXiv preprint arXiv:2310.09270*, 2023.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

Chaochao Yan, Qianggang Ding, Peilin Zhao, Shuangjia Zheng, JINYU YANG, Yang Yu, and Junzhou Huang. Retroxpert: Decompose retrosynthesis prediction like a chemist. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 11248–11258. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/819f46e52c25763a55cc642422644317-Paper.pdf.

Kaiyu Yang, Aidan Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil, Ryan J Prenger, and Animashree Anandkumar. Leandojo: Theorem proving with retrieval-augmented language models. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), Advances in Neural Information Processing Systems, volume 36, pp. 21573–21612. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/4441469427094f8873d0fecb0c4elcee-Paper-Datasets_and_Benchmarks.pdf.

Zipeng Zhong, Jie Song, Zunlei Feng, Tiantao Liu, Lingxiang Jia, Shaolun Yao, Min Wu, Tingjun Hou, and Mingli Song. Root-aligned smiles: a tight representation for chemical reaction prediction. *Chemical Science*, 13(31):9023–9034, 2022.

A EXAMPLES OF REACTION PATHWAYS

Figure 7: Example of a pathway with Safe Bet and Worthwhile reactions

Figure 8: Example of a pathway with a Rather Not reaction

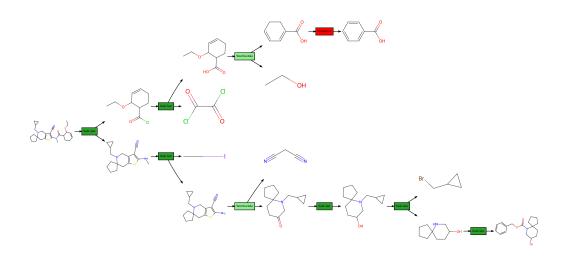


Figure 9: Example of a pathway with a Nonsense reaction

B GAT

The GAT model processes reaction graphs where individual atoms and bonds are featurized with chemically-meaningful characteristics, outputting a scalar plausibility score for each reaction. The attention mechanism is modified to ensure that attention weights between non-connected nodes approach zero, preserving the chemical connectivity structure. The key difference to the original GAT is the support of global information exchange across the entire molecular graph, ensured by an artificial supernode that connects to all other nodes in the graph.

C REACTION EVALUATION PROTOCOL

Each candidate reaction is assessed sequentially along the following dimensions. Unless specified otherwise, in each of them the reaction is scored on a four-level confidence scale: *Nonsense*, *Rather Not*, *Worthwhile*, and *Safe Bet*, indicating the plausibility of the reaction. Every reaction that is not a safe bet is assigned an additional label explaining the reason for its incorrectness.

- 1. **Reactant-Product Consistency:** Structural alignment between reactants and product is verified. Reactions in which the product contains a large fragment that is neither present in the substrate nor originates from a commonly used reagent, or in which no clear relationship between the atoms of the product and the substrates can be established, are marked as *Nonsense*, with the reason for incorrectness labeled as *Reactants mismatch*.
- 2. **Stability:** Reactions producing products or including substrates that are not isolable under the typically achievable conditions are marked as *Nonsense*, with the reason for inplausibility labeled as *Unstable*.
- 3. **Mechanistic Plausibility:** Reactions lacking a plausible mechanism are classified as *Non-sense* or *Rather Not* due to *Magic*, covering transformations requiring unknown or highly implausible reactivity. Transformations that would require more than two non-trivial steps are also placed in this category.
- 4. **Multistep Feasibility and One-Pot Potential:** Reactions not achievable in a single step are assessed for decomposability into two coherent steps. If they pass this test, feasibility in a one-pot setting is scored on a four-level scale and failing reactions are marked as *One pot*.
- 5. **Reactivity of Substrates:** Feasibility of the reaction, given the reactivity of the substrates (e.g., electron deficiency), is verified. Reactions that cannot be reasonably expected to occur are marked as implausible, with the reason for incorrectness labeled as *Reactivity*.

7. **Selectivity:** Selectivity of the reaction is verified, including competition between functional groups of the same type, regioisomeric outcomes (e.g., in electrophilic aromatic substitution), or other cases where multiple plausible products can arise. Reactions that fail this evaluation are marked as *Selectivity*.

C.1 Types of errors in pathways generated by evaluated systems



Figure 10: Detailed information on the types of errors for each reaction in the pathways from Fig. 3. Note that only reactions from the found pathways are presented here.

C.2 IMPLAUSIBILITY ANNOTATION EXAMPLES

C.2.1 REACTANTS MISMATCH

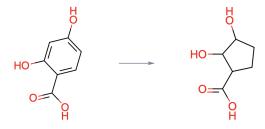


Figure 11: Nonsense: No clear relationship between atoms in the product and the substrate can be confidently proposed

Figure 12: Nonsense: The pyridyl fragment require an additional substrate, that is missing

C.2.2 Unstable

Figure 13: Nonsense: The carbon atom with amine and chlorine is not something seen in literature

Figure 14: Nonsense: The second substrate would tautomerize to phenol instantly

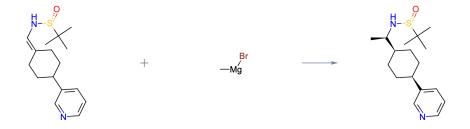


Figure 15: Nonsense: The substrate is unstable, it would tautomerize to imine

C.2.3 MAGIC

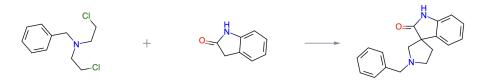


Figure 16: Nonsense: Changing length of the alkyl chain, no known precedent of such variant of carbon alkylation

Figure 17: Nonsense: An alkyl chain acting as a leaving group and bond formation by an unactviated amine carbon. No such reactivity ever demonstrated in literature

C.2.4 ONE POT

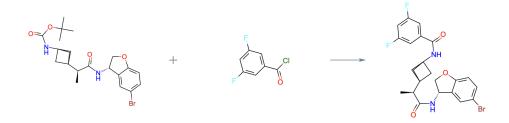


Figure 18: Rather not: 2 steps required – Boc deprotection and acylation

Figure 19: Rather not: 2 steps required - Cbz deprotection and Boc protection

C.2.5 REACTIVITY

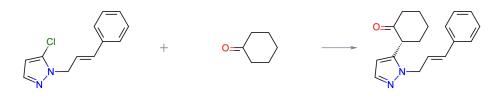


Figure 20: Rather not: Most of the references for this reaction are around electron-deficient heterocycles, only one example with pyrazole in literature

Figure 21: Rather not: High likelihood of steric hindrance

C.2.6 FUNCTIONAL GROUP INCOMPATIBILITY

$$H_2N$$
 H_2N
 H_2N
 H_2N
 H_3N
 H_3N

Figure 22: Rather not: No literature references where a bromine is located in alpha to the ester position. The alkyl bromine would most likely react more readily than the ester.

Figure 23: Nonsense: No conditions allow to cleave a methyl ether in a way that wouldn't affect the sulfonyl chloride

C.2.7 SELECTIVITY

Figure 24: Rather not: There is a considerable risk that achieving the disubstituted product at a satisfactory yield would be very difficult (especially accounting for the presence of amine in the structure).

 H_2N OH OH OH OH OH OH

Figure 25: Rather not: There are 3 equivalent hydroxyl groups, so in bromination we expect triple substitution rather than this scenario

D RETROSYNTHESIS TARGETS

D.1 SMILES

```
989
990
      Clc1ccc(-c2c(N(CC)CC)c(c(nc2C)C)CC(=0)NCC)cc1
991
992
      O(c1cc(c(N+)(-0)(0-))cc1)COC1CN(C(-0)(C00+)2C(C0)3(NC(OC3)-0)C2)C1)C1CCCC1
993
994
      FC1 (F) C (N2N=CC (=C2C) c2cc (ccc2) C#Cc2c (OC) cc (nc2) C (=O) O) C1
995
996
      O(C(C)(C)(C)C)[C@H](C(=O)Nc1nc2[C@](O)(CCc2cc1)CC)c1c(nc(cc1)C)C
997
998
      FC1(F)Oc2c(O1)cc(nc2)C(=O)NC1=NN2C(C(=O)N[C@@H]3[C@H]2CCC3)=C1
999
1000
      Clc1c(N2CCC(F)(F)CC2)c(Cl)cc(NC(=0)CC[C@]2(NC(=0)NC2=0)C2CC2)c1
1001
      S(=0) (=0) (Nc1nc2N(N(C(=0)c2cn1)CC=C)C)c1ccc([C@@H](C2=Nc3c(N2)cccc3)CCO)cc1
1002
1003
      Fc1cc(F)cc(N2[C@H](CN(CC(=0)Nc3ncnc4N(C(C)C)C=C(F)c34)CC2)C)c1
1004
1005
      Fc1c(nc2c(c(F)ccc2)c1)Nc1cc2C(OC(=0)c2cc1)(C)C
1006
1007
      O1C(=NN=C1)c1c(ncnc1)NC1C[C@H](O)[C@@H](O)C1
1008
1009
      Fc1cc2c(OB(O)[C@@H](NC(=O)C3CC3)C2)cc1
1010
1011
      S(C=1NN=NC1C(=0)NCCOCCNC(=0)C=1N=C(SC1)N1N=CC(=C1)C)c1ccccc1
1012
      O1C(Oc2c1c(ccc2C)C)([C@@H]1CC[C@@H](NC(=O)c2ncc(cc2)C#N)CC1)C
1013
1014
      S(C1=C(C(=0)NC(=C1)C)CN(c1c2c(nccc2)c(cc1)C\#N)C)C
1015
1016
      O(CC(=0) NC1CC2N(C(C1)CC2)C)CCN1c2c3N(C(=0)C1=0)CCCc3ccc2
1017
1018
      FC(F)(F)c1cc(C2=CN(C(=0)C(NC(=0)C3=NN(c4c3cccc4)C)=C2)C)ccc1
1019
1020
      Fc1cc(F)cc(C(=0)NC23CC([C@@H](C(=0)N[C@H]4c5c(OC4)ccc(-c4c(OC)ccc(c4)C)c5)C)(C2)C3)c
1021
1022
      O1c2c (cc (C3=CN4N=C (N=C4N=C3) c3cnc (C (=0) C) cc3) cc2) CCC1
1023
      S1C (N(C(=0)C2C(OCC)C=CCC2)C)=C(C2=C1CC1(N(C2)CC2CC2)CCCC1)C#N
1024
1025
      S(=0) (=0) (N[C@@H] ([C@@H] 1CC[C@H] (c2cnccc2)CC1)C)c1cc(F)cc(-c2nccc2)c1
```

```
1026
1027
      FC(F)(F)[C@@H](N1CCC2(C(=0)N(Cc3c4OC=C(c4cc(OC(C)C)c3)C)CC2)CC1)CC1[C@@H](O)[C@@H](O
1028
1029
      FC(F)(F)[C@@H]([C@H](C(=O)N[C@@H]([C@@](O)(N)CC)C)c1cc(OC)cc(OC)c1)C
1030
      FC(F)(F)clncc(-c2ncc(C(F)(F)F)c(c2)CNC(c2cc(C3=NOC(=C3CO)CC)ccc2)C2CC2)cn1
1031
1032
      O1c2c (nc (N3C (=CC=C3C) C) nc2CCC1) NC1CCC (CO) CC1
1033
1034
      O(c1ccc([N+](=0)[O-])cc1)CC[C@@](N)(CCN(C(=0)c1c2c(C(=0)c3c(C2=0)ccc3)ccc1)C)C
1035
1036
      S(=0)(c1cccc1)CCNC(=0)CN(c1ncnc([C@@H]2C[C@@H](0)C2)c1)C
1037
1038
      Fc1c(C=20C(=NN2)C=20c3c(cc4NC(0c4c3)=0)C2)cc(F)cc1
1039
1040
      O=C (N1C2C (Nc3ncc (-c4cnccc4) cn3) CC1CC2) C1C (0) C (0) CC1
1041
      S1[C@]2(C(=O)N3CC4[C@@H](NC5=NN(C=N5)CC(F)(F)F)[C@H](C3)CC4)[C@H]([C@](N=C1N)(c1cccc
1042
1043
      P(=0)(0)(0)CO[C@H]1C(C=2N(N=CC2)C/C=C/c2cccc2)CCCC1
1044
1045
      FC(F)(F)C(Nc1cncc(C(CO)C)c1)c1c(F)cc(OC2CN(C2)CCCF)cc1
1046
1047
      O=C (N1CC (N2C (=0) CNC (C2) C) C1) N [C0H]1C (=0) NC [C00H]1c1ccc (N2C [C00H] (0) CC2) cc1
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
```

D.2 VISUALIZATION

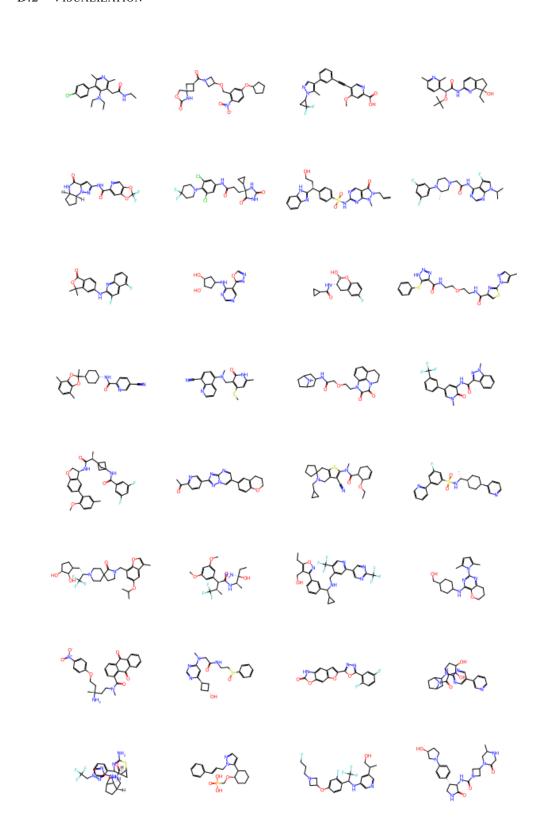


Figure 26: 32 molecules that have been used as targets for retrosynthesis.

ROC AND PRECISION-RECALL CURVES BY FAILURE CATEGORY

E.1 MAGIC

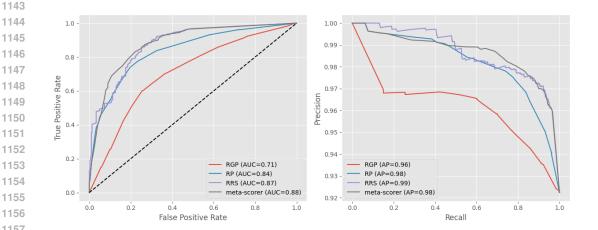


Figure 27: ROC (on the left) and precision-recall (on the right) curves comparing the performance of individual scorers versus the Meta-Scorer on Magic and No Problem reactions.

E.2 SELECTIVITY

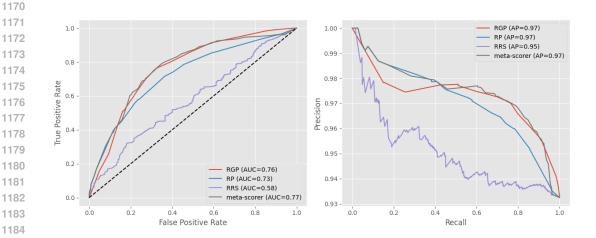


Figure 28: ROC (on the left) and precision-recall (on the right) curves comparing the performance of individual scorers versus the Meta-Scorer on Selectivity and No Problem reactions.

E.3 FUNCTIONAL GROUP INCOMPATIBILITY

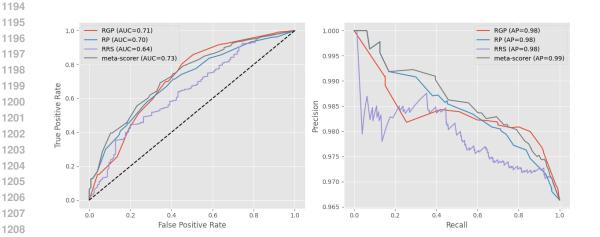


Figure 29: ROC (on the left) and precision-recall (on the right) curves comparing the performance of individual scorers versus the Meta-Scorer on *Functional group incompatibility* and *No Problem* reactions.

E.4 REACTIVITY

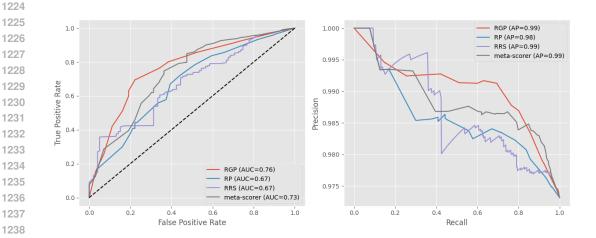


Figure 30: ROC (on the left) and precision-recall (on the right) curves comparing the performance of individual scorers versus the Meta-Scorer on *Reactivity* and *No Problem* reactions.

E.5 ONE POT

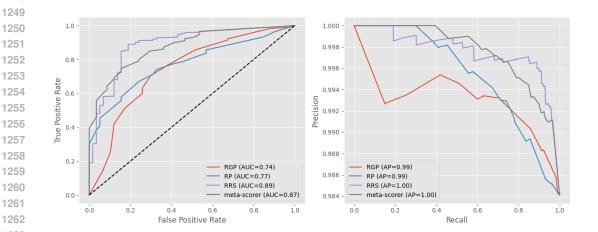


Figure 31: ROC (on the left) and precision-recall (on the right) curves comparing the performance of individual scorers versus the Meta-Scorer on *One pot* and *No Problem* reactions.

Unstable

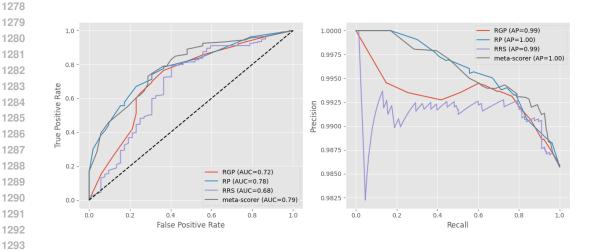


Figure 32: ROC (on the left) and precision-recall (on the right) curves comparing the performance of individual scorers versus the Meta-Scorer on *Unstable* and *No Problem* reactions.

E.7 REACTANTS MISMATCH

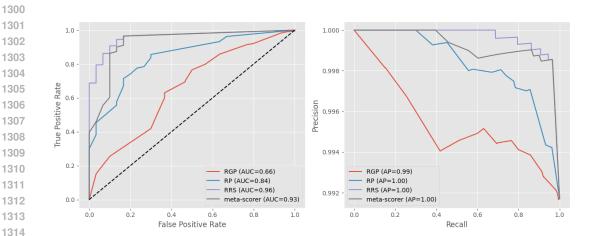


Figure 33: ROC (on the left) and precision-recall (on the right) curves comparing the performance of individual scorers versus the Meta-Scorer on *Reactants mismatch* and *No Problem* reactions.

F FALSE POSITIVES COUNTS BY FAILURE CATEGORY

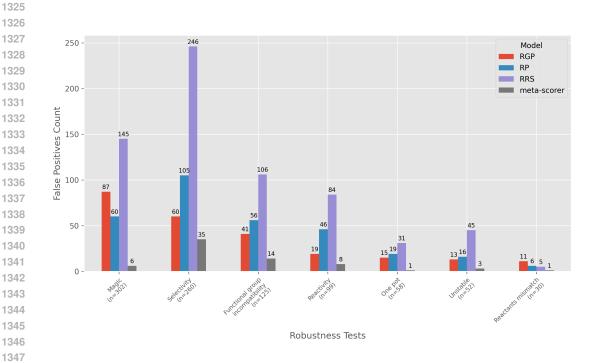


Figure 34: Counts of false positives produced by individual scorers versus the Meta-Scorer across different failure categories, with sample sizes indicated for each category.

G TRUE NEGATIVES COUNTS BY FAILURE CATEGORY

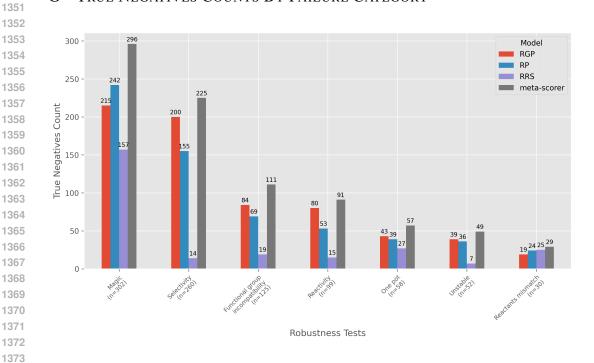


Figure 35: Counts of true negatives produced by individual scorers versus the Meta-Scorer across different failure categories, with sample sizes indicated for each category.

H LARGE LANGUAGE MODEL USAGE

We used large language models solely to polish the writing by correcting grammar and spelling errors. No part of the technical content, methodology, or results was generated or influenced by these models.