RXNCOMP: CHEMICAL REACTION COMPLETION VIA CONSTRAINED ELECTRON REDISTRIBUTION MATRIX INFERENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

Organic reaction data extracted from patents and scientific literature underpin a wide range of reaction understanding and prediction tasks. However, many recorded reactions include only major products while omitting byproducts, leaving reaction equations incomplete and often element-imbalanced. Recovering these missing species is important for constructing chemically complete reaction records, yet manual annotation is prohibitively expensive at scale. Existing completion methods largely rely on heuristic rules, while evaluation typically depends on exact matching against curated ground-truth completions, making it difficult to assess large-scale unlabeled reaction corpora or characterize candidate quality beyond reference recovery. To address these limitations, we formulate reaction completion as a conditional generation problem and introduce **RxnComp**, a constrained generative framework that completes missing byproducts through electron redistribution matrix modeling while explicitly enforcing chemical conservation during generation. We further develop a reference-free evaluation framework with complementary proxy metrics to assess the chemical validity and overall quality of generated completion candidates without requiring ground-truth annotations. Across benchmark datasets, **RxnComp** consistently produces more chemically consistent and higher-quality completion candidates than existing baselines, while maintaining competitive reference-based recovery performance. These results demonstrate the potential of constrained conditional generation for scalable and chemically grounded reaction data completion.

1 INTRODUCTION

Chemical reaction understanding is fundamental to automated organic synthesis and plays a pivotal role in accelerating both drug and materials discovery (Wang et al., 2023; Meng et al., 2023a). In recent years, a broad range of deep learning approaches have been developed for diverse reaction modeling tasks, including forward reaction prediction (Joung et al., 2025; Jin et al., 2017; Schwaller et al., 2019; Bi et al., 2021), retrosynthetic analysis (Wan et al., 2022; Yan et al., 2020; Dai et al., 2020; Segler et al., 2018), reaction yield prediction (Schwaller et al., 2021; Kwon et al., 2022; Saebi et al., 2023; Shi et al., 2025), and synthesizable molecule generation (Bradshaw et al., 2019b; Gao et al., 2025; Swanson et al., 2024; Lee et al., 2026; Luo et al., 2024), among others. These advances have been largely enabled by large-scale reaction datasets curated from the public patent literature, most notably the USPTO reaction dataset (Lowe, 2012), which has become a cornerstone benchmark for data-driven reaction modeling.

Despite their substantial scale, these datasets suffer from several inherent limitations, including biased and repetitive reaction patterns (Kovacs et al., 2021), noisy and inconsistent atom mappings (Schwaller et al., 2019; Tetko et al., 2020), and erroneous or incomplete reaction records (Lowe, 2012; Schneider et al., 2015; Schwaller et al., 2021; Saebi et al., 2023). A particularly important yet underexplored issue is the frequent omission of byproducts (Vogel et al., 2026), which results in chemically unbalanced reactions and incomplete supervision signals. This problem has received relatively little attention because most prior work focuses on top-k prediction accuracy, where correctly identifying the major product is typically sufficient, regardless of whether the full reaction is chemically balanced.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

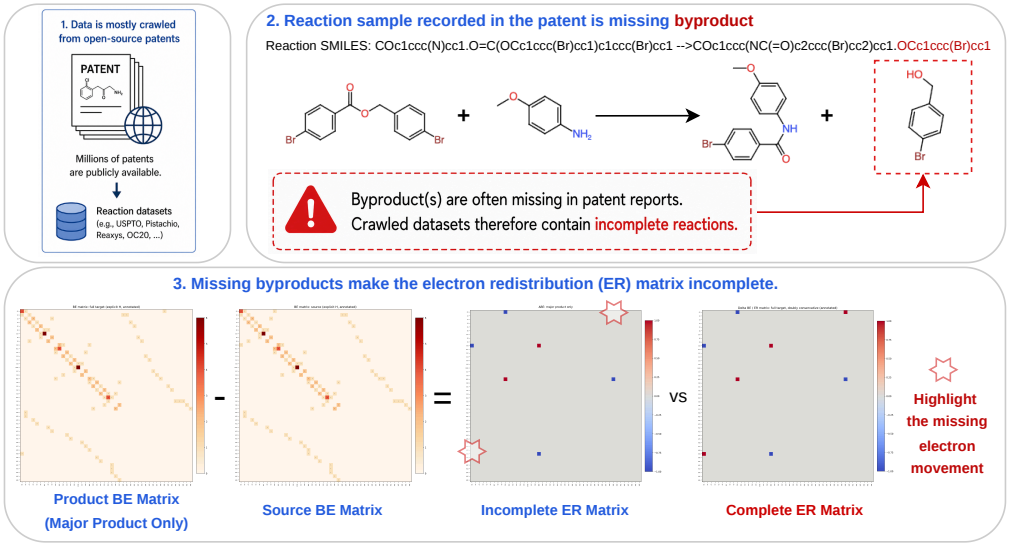

Figure 1: Illustration of the problem. (1) Public reaction datasets are largely curated from academic and patent literature. (2) Reported reactions often include only the main products, with byproducts omitted. (3) Missing byproducts lead to incomplete electron redistribution (ER) matrices, limiting reaction prediction and mechanistic understanding.

Recently, the community has increasingly recognized the importance of explicitly modeling electron redistribution for chemical reaction mechanism understanding (Joung et al., 2025; Bradshaw et al., 2019a; Bi et al., 2021; Choi, 2023; Duan et al., 2025; Bran et al., 2025). Learning such fine-grained electron-level transformations, however, typically requires complete reactions that provide full supervision over electron rearrangements. In existing reaction records, byproducts are often omitted, leaving part of the underlying electron redistribution signal unobserved, as illustrated in Figure 1. While a small number of reactions can be manually curated or completed using heuristic rules, such approaches are difficult to scale. Therefore, automatically inferring and recovering missing byproducts is important for constructing chemically complete reaction data and enabling more faithful electron-level reaction modeling, and will become increasingly critical as web-crawled reaction datasets continue to grow in scale.

Prior work (Phan et al., 2024; Toniato et al., 2021; Valenti et al., 2024) has explored rule-based and heuristic approaches for reaction dataset curation and completion, while more recent studies (Zhang et al., 2024; Vogel et al., 2026) incorporate atom-balance constraints into model decoding. However, these methods often rely on hand-crafted templates or predefined heuristics for common species, limiting their ability to generalize to novel and complex reactions and to capture diverse byproduct distributions. A key challenge is the *absence of ground-truth byproducts*: because such species are typically omitted from reaction records, they are difficult to supervise directly and to evaluate using exact-match criteria at scale. A second challenge is ensuring *electron-level physical consistency*. Chemically feasible reactions must satisfy conservation laws (Joung et al., 2025; Bi et al., 2021), including electron-change conservation (Meng et al., 2023b; Bi et al., 2021; Joung et al., 2025) together with electron-counting principles such as the duet and octet rules (Atkins & Jones, 2016; Petrucci et al., 2017). Such structured constraints are difficult to represent and enforce explicitly within SMILES-based frameworks (Schwaller et al., 2019; Lu & Zhang, 2022; Weininger, 1988).

In this work, we formulate reaction completion as a conditional generative task for inferring missing byproducts and propose **RxnComp**, an autoregressive framework that completes the missing entries of the electron redistribution (ER) matrix (Joung et al., 2025; Bi et al., 2021). By performing constrained decoding directly in ER space, **RxnComp** generates multiple chemically feasible completion hypotheses, making it well suited to large-scale reaction datasets where ground-truth byproducts are unavailable and candidate completions can be further screened by domain experts. To enable evaluation in this setting, we further introduce a reference-free framework based on complementary proxy metrics for chemical validity, conservation, and electron-redistribution plausibility. Experiments show that **RxnComp** produces substantially more chemically plausible completions than existing heuristic approaches. Our contributions are summarized as follows:

- We propose **RxnComp**, an autoregressive ER matrix completion framework that performs constrained decoding to generate diverse and chemically feasible missing-byproduct candidates.
- We systematically characterize the fundamental feasibility requirements for reaction completion, including molecular validity, atom and charge conservation, and electron-level consistency, providing a principled basis for constrained byproduct generation.
- We establish a reference-free evaluation protocol that quantifies these complementary aspects of completion quality without requiring ground-truth byproduct annotations.

2 Related Works

Organic chemical reaction modeling has been extensively studied across a variety of tasks in recent years. Representative tasks—including forward reaction prediction (Schwaller et al., 2019; Jin et al., 2017; Segler & Waller, 2017b;a; Meng et al., 2023b; Sacha et al., 2021; Bi et al., 2021; Joung et al., 2025), single-step retrosynthetic analysis (Coley et al., 2017; Dai et al., 2020; Yan et al., 2020; Wan et al., 2022; Somnath et al., 2021; Wang et al., 2021; Chen & Jung, 2021), multi-step retrosynthetic planning (Segler et al., 2018; Chen et al., 2020; Genheden et al., 2020; Schreck et al., 2019; Liu et al., 2023; Han et al., 2022; Lin et al., 2020; Genheden & Bjerrum, 2022; Hong et al., 2023; Tripp et al., 2024), reaction yield prediction (Schwaller et al., 2021; Kwon et al., 2022; Saebi et al., 2023; Shi et al., 2024; 2025), and reaction mechanism prediction (Choi, 2023; Kim et al., 2024; Duan et al., 2025; Galustian et al., 2025)—play a crucial role in advancing the roadmap toward automated chemical synthesis. Despite their diversity, these tasks are largely trained and evaluated on the same USPTO dataset (Lowe, 2012), albeit with different preprocessing strategies and experimental settings. Previous studies (Schneider et al., 2015; Kovacs et al., 2021; Phan et al., 2024; Toniato et al., 2021; Valenti et al., 2024; Zhang et al., 2024) have identified limitations of this dataset, which is curated from publicly available patent literature. In this work, we focus specifically on the omission of byproducts, which leads to incomplete supervision for reaction modeling. Rule-based completion approaches (Valenti et al., 2024; Phan et al., 2024) rely heavily on predefined templates and heuristic rules, resulting in limited coverage and generalization. In contrast, **RxnComp** is a learned generative framework that produces diverse candidates under explicit fundamental physical constraints.

Recent advances in scientific foundation models and autonomous agents have broadened the scope of machine-assisted scientific reasoning (Lála et al., 2023; Asai et al., 2026; Ren et al., 2025), with chemistry emerging as an important domain for reasoning, synthesis planning, and tool-augmented problem solving (Guo et al., 2023; Sprueill et al., 2024; Yamada et al., 2025; Polat et al., 2025; Zhang et al., 2025; Wu et al., 2025). From this perspective, reaction completion provides a potentially valuable future challenge for chemistry agents, as it requires inferring unobserved reaction components under global chemical consistency while accommodating multiple plausible hypotheses.

3 Preliminary

Problem Formulation. A chemical reaction can be described as a transformation $\mathcal{R} \rightarrow \mathcal{P}$ where a set of reactants $\mathcal{R} = \{R_1, .., R_m\}$ is converted into a set of products $\mathcal{P} = \{P_1, .., P_n\}$. Each atom in R can be uniquely mapped to an atom in P, such that $\mathcal{A}(\mathcal{R}) = \mathcal{A}(\mathcal{P})$, reflecting the law of mass conservation, which states that atoms participating in a reaction are neither created nor destroyed. However, in large-scale reaction datasets collected from public patent literature, only the main products $\mathcal{P}^{\text{main}}$ are typically recorded, while byproducts $\mathcal{P}^{\text{by}}$ are often missing, such that $\mathcal{P} = \mathcal{P}^{\text{main}} \cup \mathcal{P}^{\text{by}}$ and $\mathcal{P}^{\text{main}} \cap \mathcal{P}^{\text{by}} = \emptyset$). We therefore formulate reaction completion as a novel constrained generation task, aiming to infer the conditional distribution $p(\mathcal{P}^{\text{by}}|\mathcal{P}^{\text{main}}, \mathcal{R})$. Notably, due to mass conservation, the identities and counts of missing atoms are fully determined by the difference between the atomic compositions of $\mathcal{R}$ and $\mathcal{P}^{\text{main}}$. We denote the set of missing atoms as $\mathcal{A}^{\text{miss}} = \mathcal{A}(\mathcal{R}) \setminus \mathcal{A}(\mathcal{P}^{\text{main}})$, which induces an explicit atomic composition constraint under which the conditional distribution $p(\mathcal{P}^{\text{by}} \mid \mathcal{P}^{\text{main}}, \mathcal{R})$ must be inferred. We further define the set of observed atoms as $\mathcal{A}^{\text{obs}} = \mathcal{A}(\mathcal{R}) \cup \mathcal{A}(\mathcal{P}^{\text{main}})$.

Electron Redistribution (ER) Matrix. Beyond atom and charge conservation, a chemical reaction is fundamentally governed by the redistribution of electrons among atoms. Given a reaction $\mathcal{R} \rightarrow \mathcal{P}$ with a one-to-one atom mapping, we introduce an *electron redistribution matrix*

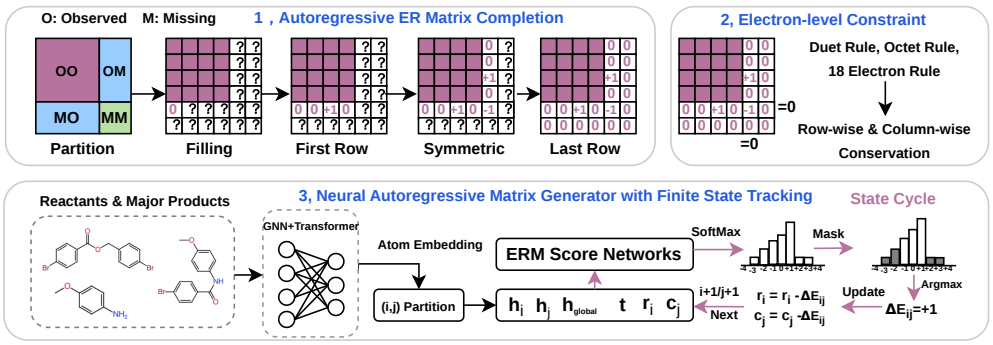

Figure 2: Overview of the **RxnComp** solution pipeline. (1) We formulate the chemical reaction completion task as an autoregressive completion problem over the electron redistribution (ER) matrix. (2) We impose a doubly conservative electron-level physical prior, ensuring that both row-wise and column-wise electron balances are strictly preserved. (3) We introduce a neural autoregressive matrix generator that predicts entry-wise electron changes while maintaining feasibility through a finite-state mechanism that tracks row and column electron budgets during decoding.

$\Delta\mathbf{E} \in \mathbb{R}^{N\times N}$ to characterize the net electron flow between atoms, where $N = |\mathcal{A}(\mathcal{R})| = |\mathcal{A}(\mathcal{P})|$ denotes the total number of atoms (including explicit hydrogens). Let $\mathbf{E}^{\mathcal{R}}$ and $\mathbf{E}^{\mathcal{P}}$ denote the electron adjacency matrices (e.g., bond-order representations) of the reactants and products, respectively. Each entry $\mathbf{E}_{ij}$ represents the number of electrons shared between atoms i and j, while the diagonal elements $\mathbf{E}_{ii}$ encode atom-wise electron properties, such as lone-pair electrons and formal charge contributions. The redistribution matrix is defined as $\mathbf{E}^{\mathcal{P}} = \mathbf{E}^{\mathcal{R}} + \Delta\mathbf{E}$, which captures the change in electron configuration induced by the reaction while preserving atom correspondence. $\Delta\mathbf{E}$ usually satisfy two fundamental electron-level physical priors. First, under the *Octet (and Duet) Rule*, atoms tend to gain, lose, or share electrons to achieve a stable electron configuration: $\sum_j \Delta\mathbf{E}_{ij} = 0, \forall i$. Second, electron transfer is symmetric between atom pairs: $\Delta\mathbf{E}_{ij} = \Delta\mathbf{E}_{ji}, \forall i, j$. Together, these constraints imply that $\Delta\mathbf{E}$ satisfies a *doubly conservative* property: $\sum_j \Delta\mathbf{E}_{ij} = 0, \sum_i \Delta\mathbf{E}_{ij} = 0, \forall i, j$, which enforces conservation of electrons at both the source and destination levels. Furthermore, chemical reactions typically involve localized electron rearrangements, implying that $\Delta\mathbf{E}$ is expected to be sparse, with only a small number of non-zero off-diagonal entries corresponding to active electron transfer pathways. We further explain these properties of the ER matrix in detail in the Appendix.

4 Methods: RxnComp

4.1 Reaction Completion as Autoregressive ER Matrix Completion

In the context of reaction completion, where only the main products $\mathcal{P}^{\text{main}}$ are observed, the ER matrix $\Delta\mathbf{E}$ provides a mechanistic constraint linking the reactants $\mathcal{R}$, the main products $\mathcal{P}^{\text{main}}$, and the missing byproducts $\mathcal{P}^{\text{by}}$. Specifically, any valid completion $\hat{\mathcal{P}}^{\text{by}}$ must induce an ER matrix $\Delta\mathbf{E}$ that satisfies both the doubly conservative constraint and the atomic valence constraint, thereby ensuring that the completed reaction $\mathcal{R} \to \mathcal{P}^{\text{main}} \cup \hat{\mathcal{P}}^{\text{by}}$ is chemically plausible at the electron level. We partition the electron redistribution matrix $\Delta\mathbf{E}$ into block components as $\Delta\mathbf{E} = \begin{pmatrix} \Delta\mathbf{E}_{oo} & \Delta\mathbf{E}_{om} \\ \Delta\mathbf{E}_{mo} & \Delta\mathbf{E}_{mm} \end{pmatrix}$, where o and m index the observed atoms $\mathcal{A}^{\text{obs}}$ and missing atoms $\mathcal{A}^{\text{miss}}$, respectively. Therefore, reaction completion can be formulated as the problem of inferring the missing entries of the ER matrix, namely $\Delta\mathbf{E}_{om}$, $\Delta\mathbf{E}_{mo}$, and $\Delta\mathbf{E}_{mm}$ conditioned on given $\Delta\mathbf{E}_{oo}$. In this work, to better incorporate electron-level physical constraints during inference, we define the matrix completion problem as an autoregressive generation process. Let the set of missing entries be indexed by an ordered sequence: $\mathcal{I} = \{(i_1, j_1), \ldots, (i_T, j_T)\}$, covering all entries in $\Delta\mathbf{E}_{miss} = \{\Delta\mathbf{E}_{om}, \Delta\mathbf{E}_{mo}, \Delta\mathbf{E}_{mm}\}$. We define the autoregressive generation process as:

$$p_\theta(\Delta\mathbf{E}_{miss}|\mathcal{R}, \mathcal{P}^{main}) = \prod_{t=1}^{T} p_\theta(\Delta\mathbf{E}_{i_t j_t}|\Delta\mathbf{E}_{<t}, \mathcal{R}, \mathcal{P}^{main}), \tag{1}$$

where $\Delta\mathbf{E}_{<t}$ denotes all previously generated entries, p_θ denotes the trainable autoregressive matrix generator networks with parameters θ. The generation order is defined according to the RDKit atom indices of the reactants, with rows processed in ascending order of index i, and columns within each row traversed in ascending order. At each step, the model predicts $\Delta\mathbf{E}_{i_t j_t}$ as a categorical variable over the discrete set $\{-4, -3, -2, -1, 0, 1, 2, 3, 4\}$, where 4 denotes the maximum allowed electron transfer magnitude. The conditional distribution for each step is $p_\theta(\Delta\mathbf{E}_{i_t j_t} | \Delta\mathbf{E}_{<t}, \mathcal{R}, \mathcal{P}^{main}) = \mathrm{Softmax}(s_\theta(i_t, j_t, v))$, $v \in \mathcal{V} = \{-4, \ldots, 4\}$. Note that, due to the symmetry constraint of ΔE, the model predicts only the upper (or lower) triangular entries, with each prediction simultaneously assigned to its symmetric counterpart, i.e., $\Delta\mathbf{E}_{i_t j_t} = \Delta\mathbf{E}_{j_t i_t}$

4.2 Budget-Constrained Autoregressive Generation as a Finite-State Process

To ensure that the generated ER matrix satisfies the doubly conservative constraint during inference, we cast the autoregressive generation process as a *structured decoding problem with a finite-state machine (FSM)* that explicitly tracks electron-level constraints.

FSM with Electron Budget Tracking. At each decoding step t, the generation state is defined by the remaining electron budgets of all atoms: $r^{(t)} = (r_1^{(t)}, \ldots, r_N^{(t)})$ and $c^{(t)} = (c_1^{(t)}, \ldots, c_N^{(t)})$, where $r_i^{(t)} = -\sum_j \Delta\mathbf{E}_{ij}^{(t)}$ and $c_j^{(t)} = -\sum_i \Delta\mathbf{E}_{ij}^{(t)}$ denote the residual row-wise and column-wise electron flows, respectively. These budgets define the FSM state, representing the remaining electron redistribution required to satisfy the doubly conservative constraint. Given the state $(r^{(t)}, c^{(t)})$, the model predicts the next entry $\Delta\mathbf{E}_{i_t j_t}$ by selecting a value $v \in \mathcal{V}$. This induces a state transition by updating the corresponding row and column budgets as $r_{i_t}^{(t+1)} = r_{i_t}^{(t)} - v$ and $c_{j_t}^{(t+1)} = c_{j_t}^{(t)} - v$ respectively. Due to the symmetry constraint $\Delta\mathbf{E}_{j_t i_t} = \Delta\mathbf{E}_{i_t j_t}$, the assignment simultaneously affects the paired entry, requiring additional updates $r_{j_t}^{(t+1)} = r_{j_t}^{(t)} - v$ and $c_{i_t}^{(t+1)} = c_{i_t}^{(t)} - v$.

Constrained Decoding Strategy. To ensure that the final state satisfies the conservation constraints, we restrict transitions to those that remain feasible. Specifically, let $R_i^{(t+1)}$ and $C_j^{(t+1)}$ denote the number of remaining undecoded entries in row i and column j after step t. Since each entry is bounded in $[-K, K]$ (with $K = 4$), a candidate value v is valid only if the updated budgets remain reachable: $r_{i_t}^{(t)} - v \in [-KR_{i_t}^{(t+1)}, KR_{i_t}^{(t+1)}]$ and $c_{j_t}^{(t)} - v \in [-KC_{j_t}^{(t+1)}, KC_{j_t}^{(t+1)}]$. This defines a conservation-aware feasible set $\mathcal{V}_t$ at each step. To unify the above constraints, we define a step-wise masking function $m_t(v)$ that encodes both budget feasibility and chemical validity:

$$m_t(v) = \prod_{a \in \{i_t, j_t\}} \mathbf{1}\Big[|r_a^{(t)} - v| \le KR_a^{(t+1)}\Big]\, \mathbf{1}\Big[|c_a^{(t)} - v| \le KC_a^{(t+1)}\Big] \cdot \mathbf{1}\big[v \in \mathcal{V}_{i_t j_t}^{\mathrm{chem}}\big]. \quad (2)$$

where $a \in \{i_t, j_t\}$ indexes the two atoms involved in the symmetric update, $r_a^{(t)}$ and $c_a^{(t)}$ denote the remaining row and column electron budgets at step t, $R_a^{(t+1)}$ and $C_a^{(t+1)}$ are the numbers of unfilled entries associated with row a and column a after the current step, and K is the maximum absolute electron-change value per entry. The first two terms ensure that, after assigning the coupled symmetric entries $\Delta E_{i_t j_t} = \Delta E_{j_t i_t} = v$, the remaining budgets of both atoms remain reachable, while $\mathcal{V}_{i_t j_t}^{\mathrm{chem}}$ denotes the set of chemically admissible values that enforce valence and bond-order constraints. Accordingly, the constrained step-wise distribution is defined as $p_\theta(\Delta E_{i_t j_t} = v \mid \cdot) = \mathrm{MaskedSoftmax}_{v \in \{-K, \ldots, K\}}\big(s_\theta(i_t, j_t, v), m_t(v)\big)$.

To further guarantee that the generated ER matrix satisfies the doubly conservative constraint, we adopt a *deterministic closure strategy* that avoids infeasible terminal states. Instead of sampling all missing entries, we reserve one *slack* entry per row and column during decoding. Specifically, for each row (and column), we leave one entry undecoded and defer its assignment to the final stage. After all other entries are generated, the remaining row and column budgets uniquely determine the values of these slack entries, which can be filled deterministically as $\Delta\mathbf{E}_{ij} = r_i^{(t)}$ or $\Delta\mathbf{E}_{ij} = c_j^{(t)}$.

Under this formulation, the entire generation process can be viewed as a sequence of state transitions in a finite-state system, where valid trajectories correspond to electron redistribution matrices that satisfy $\Delta\mathbf{E}\mathbf{1} = \mathbf{0}$ and $\Delta\mathbf{E}^\top\mathbf{1} = \mathbf{0}$. This structured decoding mechanism enforces physical con-

straints throughout generation rather than applying them post hoc, significantly reducing the search space and improving the reaction-level validity of generated byproducts.

4.3 Neural Autoregressive ER Matrix Generator

We parameterize the autoregressive distribution $p_\theta(\Delta\mathbf{E}_{\text{miss}} \mid \mathcal{R}, \mathcal{P}^{\text{main}})$ using a neural generator that predicts ER entries sequentially. We first encode the reactants $\mathcal{R}$ and main products $\mathcal{P}^{\text{main}}$ using a hybrid GNN–Transformer encoder to obtain contextualized embeddings for all observed atoms $\{\mathbf{h}_i\}_{i=1}^{N} = \text{Enc}_\theta(\mathcal{R}, \mathcal{P}^{\text{main}})$. These embeddings capture both local chemical structure and global relational dependencies. A global representation is further computed via mean pooling: $\mathbf{h}_{\text{global}} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{h}_i$. For each missing entry (i, j) at decoding step t, we construct a pairwise representation based on the embeddings of atoms i and j. To incorporate constraint-aware decoding, we additionally encode the current row and column electron budgets, as well as step information. Specifically, the score for assigning value $v \in \{-K, \ldots, K\}$ to entry (i, j) is given by

$$s_\theta(i, j, v) = \text{MLP}_\theta\Big(\big[\mathbf{h}_i,\ \mathbf{h}_j,\ \mathbf{h}_{\text{global}},\ \mathbf{r}_i^{(t)},\ \mathbf{c}_j^{(t)},\ \mathbf{t}\big]\Big),$$

where $\mathbf{h}_i, \mathbf{h}_j$ are determined based on the matrix partitioning. If atom i or j belongs to the observed set, then $\mathbf{h}_i$ or $\mathbf{h}_j$ is obtained from contextualized embeddings; otherwise, they are represented using atom-type embeddings. $\mathbf{r}_i^{(t)}$ and $\mathbf{c}_j^{(t)}$ are embeddings of the current row and column budgets, representing the remaining electron budget required to satisfy conservation constraints. $\mathbf{t}$ denotes the discrete time-step embedding, which informs the model of the relative progress of the current step within the overall generation process for each sample. The resulting logits $s_\theta(i, j, v)$ define a distribution over candidate electron-change values, which is subsequently filtered by the feasibility mask to enforce symmetry, conservation, and chemical constraints. Further details on the construction of these embeddings are provided in the Appendix.

4.4 Training & Inference

Training with Limited Supervision. We construct a modest set of complete reactions using heuristic rules, yielding over $30k$ training examples. To address the scarcity of ground-truth supervision, we further adopt a self-supervised learning strategy by randomly masking a subset of products and training the model to reconstruct them. This auxiliary reconstruction task encourages the model to learn general electron redistribution patterns beyond the limited supervised labels. The overall training objective combines the supervised autoregressive completion loss and the self-supervised reconstruction loss:

$$\mathcal{L} = \sum_{(i,j)\in\mathcal{I}_{\text{miss}}} \text{CE}(p_\theta(\cdot \mid \Delta\mathbf{E}_{<ij}), \Delta\mathbf{E}_{ij}) + \lambda_{\text{ssl}} \sum_{(i,j)\in\mathcal{I}_{\text{mask}}} \text{CE}\big(p_\theta(\cdot \mid \Delta\mathbf{E}^*), \Delta\mathbf{E}^*_{ij}\big)\,. \tag{3}$$

where $\mathcal{I}_{\text{miss}}$ denotes the set of target entries corresponding to missing ER values, with only the triangular portion considered due to symmetry, and $\Delta\mathbf{E}^*$ denotes the incomplete ER matrix induced by randomly masking byproduct-related entries. $\mathcal{I}_{\text{mask}}$ denotes the masked entries used for self-supervised reconstruction. $\lambda_{\text{ssl}} \in [0, 1]$ controls the relative weight of the self-supervised loss.

Inference. During inference, we perform beam search over the missing ER entries under FSM-guided constraints. At each step, candidate values are filtered by the feasibility mask, and the top-B partial solutions are retained according to their accumulated scores. Beam search is critical in this setting, as the strong coupling between symmetry and conservation constraints can easily lead greedy decoding into dead-end states. By maintaining multiple hypotheses, beam search significantly improves the likelihood of finding globally feasible ER matrix completions. More details about the inference process are shown in the Appendix.

5 Experimental Results

5.1 Experimental Setups

Dataset & Baselines. We evaluate **RxnComp** in a setting that reflects practical reaction completion: only a small fraction of reactions contain annotated byproducts, while large-scale reaction

Table 1: **Reference-free reaction completion performance on large-scale unlabeled USPTO-480K reactions.** Models are evaluated using coverage and proxy metrics of chemical validity and consistency. "–" denotes settings that are not applicable because some methods cannot generate multiple byproduct candidates per reaction. The best and second-best results are highlighted in **bold** and underline, respectively. **RxnComp** is highlighted in light green, and the key COV-JVR metric in light red. Complete Top-1/5 results are provided in Table 12.

Model	COV-JVR↑	COV↑	SVR↑		ACR↑		CCR↑		ECR↑		ESR↑		JVR↑	
			@1	@5	@1	@5	@1	@5	@1	@5	@1	@5	@1	@5
Rule-based Completion	0.131	0.132	**1.000**	-	**1.000**	-	**0.999**	-	**1.000**	-	**1.000**	-	**0.999**	-
SynRBL	0.579	0.708	**1.000**	-	0.818	-	0.997	-	0.818	-	0.818	-	0.818	-
Differential-AR	0.002	0.023	**1.000**	1.000	0.118	0.120	0.631	0.974	0.118	0.120	0.118	0.120	0.113	0.113
Molecular Transformer	0.019	0.983	**1.000**	1.000	0.020	0.036	0.972	0.972	0.020	0.036	0.020	0.036	0.020	0.036
T5Chem	0.001	**1.000**	0.073	0.138	0.019	0.019	0.919	0.923	0.001	0.001	0.001	0.001	0.001	0.001
RxnComp *(lightweight)*	0.673	0.760	0.952	0.973	0.971	0.971	0.947	0.947	**1.000**	1.000	**0.997**	0.696	0.885	0.906
RxnComp	**0.738**	0.841	0.947	0.957	0.971	0.971	0.940	0.940	**1.000**	1.000	**0.996**	0.583	0.877	0.886

Table 2: **Reference-based reaction completion performance on RMechDB**. COV is computed over the full test set, while EM@k and JVR@k are evaluated on reactions for which each method produces at least one candidate. EM@k measures recovery of the ground-truth byproduct, whereas JVR@k measures the validity of the completed reaction under the joint feasibility criteria. (CRB) denotes the variant of molecular transformer with constrained beam search.

Method	**COV ↑**	**EM@1 ↑**	**EM@3 ↑**	**EM@5 ↑**	**JVR@1 ↑**	**JVR@3 ↑**	**JVR@5 ↑**
SynRBL	37.35	1.56	6.32	10.97	7.73	15.14	18.52
Molecular Transformer	**100.00**	0.31	0.43	0.50	1.21	1.74	1.92
Molecular Transformer (CRB)	93.57	1.10	3.78	5.63	2.91	6.18	9.29
RxnComp	91.16	**23.35**	**81.94**	**82.38**	**98.24**	**98.24**	**98.24**

corpora predominantly record major products alone. We use reactions with ground-truth byproducts from USPTO-480K for supervised training, and treat the remaining reactions without byproduct annotations as a large-scale inference set for reference-free evaluation. We further evaluate direct byproduct recovery on RMechDB (Tavakoli et al., 2023), where the lower-molecular-weight product is treated as the ground-truth byproduct and withheld during inference. This setup assesses learning from limited labeled data, large-scale completion of unlabeled reactions, and recovery of known byproducts when references are available. Our experiments consider three categories of baseline models. (1) **Rule-based heuristic methods**: SynRBL (Phan et al., 2024) and Rule-based Completion (Valenti et al., 2024), which use predefined heuristic rules and template matching. (2) **SMILES-based sequence models**: Molecular Transformer (Schwaller et al., 2019) and T5Chem (Lu & Zhang, 2022), which operate directly on SMILES representations. Differential-AR is an internally developed baseline that applies difference prediction to reaction completion; further details are provided in the Appendix.

Evaluation Metrics. We use two complementary evaluation settings depending on whether ground-truth byproducts are available. On RMechDB, where a reference byproduct can be identified, we report **EM@k** to measure whether the ground-truth byproduct is recovered among the top-k predictions. For large-scale evaluation on USPTO, however, most reactions do not contain annotated byproducts, making exact-match evaluation infeasible. We therefore insert each predicted byproduct back into the reaction and assess the resulting completion using a set of reference-free proxy metrics. **SVR@k** measures whether the predicted byproducts form valid molecular structures, while **ACR@k** and **CCR@k** evaluate atom and charge conservation, respectively. To assess *electron-level consistency*, **ECR@k** measures whether the inferred ER matrix satisfies the required conservation structure, and **ESR@k** measures the sparsity of electron changes, reflecting the prior that reactions typically involve localized transformations. We further report **JVR@k**, which measures the fraction of predictions that simultaneously satisfy structural validity, atom conservation, charge conservation, and electron-level consistency, providing an overall measure of reaction feasibility. All top-k metrics are evaluated for $k \in \{1, 3, 5\}$; the main tables report $k \in \{1, 5\}$ for compactness, with complete results provided in the appendix. Finally, **COV** measures the fraction of reactions for which a method produces at least one completion candidate, while **COV-JVR** jointly captures coverage and validity. Together, these metrics enable direct reference-based evaluation on RMechDB and scalable reference-free evaluation over large unlabeled USPTO reaction corpora.

Table 3: **Ablation results for different variants of RxnComp**. We evaluate both **RxnComp** and its lightweight version to assess the contribution of each component. Top-1 and Top-5 proxy metrics are reported; complete Top-1/3/5 results are provided in Table 13.

Model	Variant	COV-JVR↑	COV↑	SVR↑		ACR↑		CCR↑		ECR↑		ESR↑		JVR↑	
				@1	@5	@1	@5	@1	@5	@1	@5	@1	@5	@1	@5
RxnComp *(lightweight)*	Full model	**0.673**	0.760	0.952	0.973	0.971	0.971	0.947	0.947	1.000	1.000	0.997	0.696	0.885	0.906
	w/o step embedding	0.642	0.732	0.953	0.972	0.969	0.969	0.939	0.939	1.000	1.000	0.996	0.665	0.878	0.889
	w/o budget embedding	0.640	0.722	0.946	0.969	0.966	0.966	0.955	0.955	1.000	1.000	0.996	0.679	0.887	0.907
	w/o atom-pair embedding	0.579	0.650	0.947	0.972	0.969	0.969	0.956	0.956	1.000	1.000	0.996	0.688	0.891	0.906
	w/o SSL loss	0.543	0.600	0.959	0.963	0.963	0.963	0.963	0.963	1.000	1.000	0.996	0.654	0.905	0.908
	w/o electron budget constraint	0.018	1.000	0.707	0.739	0.976	0.976	0.915	0.915	0.051	0.051	0.997	0.993	0.018	0.018
RxnComp	Full model	**0.738**	0.841	0.947	0.957	0.971	0.971	0.940	0.940	1.000	1.000	0.996	0.582	0.877	0.886
	w/o step embedding	0.705	0.797	0.959	0.962	0.969	0.969	0.939	0.939	1.000	1.000	0.996	0.782	0.885	0.888
	w/o GNN	0.672	0.768	0.942	0.960	0.968	0.968	0.947	0.947	1.000	1.000	0.996	0.611	0.876	0.892
	w/o SSL loss	0.402	0.445	0.954	0.968	0.963	0.963	0.968	0.968	1.000	1.000	0.996	0.612	0.904	0.918
	w/o electron budget constraint	0.018	1.000	0.707	0.715	0.976	0.976	0.915	0.915	0.051	0.053	0.997	0.993	0.018	0.018

5.2 Main Empirical Results.

Table 1 first evaluates reaction completion at scale on unlabeled USPTO reactions using reference-free metrics. **RxnComp** achieves the highest **COV-JVR** score of 0.738, demonstrating stronger overall candidate-inference quality by jointly maintaining broad coverage and reaction feasibility. In contrast, competing approaches generalize poorly beyond the regimes in which they are effective. Rule-based methods produce highly valid completions when applicable but suffer from limited coverage; for example, Rule-based Completion attains a JVR@1 of 0.999 while covering only 13.2% of the test reactions. Sequence-based models show the opposite behavior: Molecular Transformer and T5Chem achieve near-complete coverage but extremely low JVR, indicating that their candidates often violate fundamental feasibility requirements. **RxnComp** instead maintains both broad coverage (0.841) and high joint validity (0.877/0.886 at Top-1/5), indicating stronger generalization across diverse reactions. The lightweight variant achieves slightly higher validity but lower coverage, whereas the full model obtains the best COV-JVR, suggesting that increased capacity mainly improves the discovery of feasible completions over a broader reaction space. This advantage is further supported by the reference-based evaluation on RMechDB in Table 2. Although Molecular Transformer achieves complete coverage, its reference recovery and joint validity remain extremely low, with only 0.31% EM@1 and 1.21% JVR@1. Introducing constrained beam search substantially improves both metrics, increasing EM@5 to 5.63% and JVR@5 to 9.29%, but the resulting performance remains far below that of **RxnComp**. SynRBL exhibits a different trade-off, achieving somewhat higher EM and JVR among the reactions it can complete, but covering only 37.35% of the test set. In contrast, **RxnComp** covers 91.16% of the reactions while achieving 23.35% EM@1 and 98.24% JVR@1. With multiple hypotheses, reference recovery rises sharply to 81.94% at Top-3 and 82.38% at Top-5, while JVR remains at 98.24%, showing that the additional candidates predominantly expand ground-truth recovery without sacrificing feasibility. The large improvement from Top-1 to Top-3 further highlights the benefit of hypothesis generation when multiple chemically plausible completions may exist. Together, these results show that **RxnComp** achieves a substantially stronger balance among coverage, physical validity, and reference recovery across both large-scale unlabeled reactions and an independent mechanistic benchmark.

5.3 Ablation Studies

We conduct ablation studies on variants of **RxnComp**, including its lightweight version, with results reported in Table 3. We draw four conclusions. **(1) Closure constraints are critical for validity.** Removing the electron budget constraint yields perfect coverage (COV = 1.000), since the model can always produce a candidate without restriction. However, JVR@1 collapses (e.g., to 0.018), showing that these candidates rarely satisfy physical validity. Thus, high coverage alone is trivial; the key challenge is enforcing electron-level constraints during generation rather than relying on post-hoc filtering. **(2) Self-supervised loss improves coverage.** Removing the auxiliary reconstruction loss substantially reduces coverage (e.g., $0.841 \rightarrow 0.445$), while JVR remains stable or slightly improves. This suggests that the self-supervised objective exposes the model to broader electron-redistribution patterns and expands the feasible solution space. Without it, the model focuses on a narrower subset of reactions. **(3) Structural components enhance searchability.** Removing step embeddings, budget embeddings, atom-pair embeddings, or GNN layers consistently reduces coverage while leaving JVR largely unchanged. Since validity is enforced by the constraint mechanism,

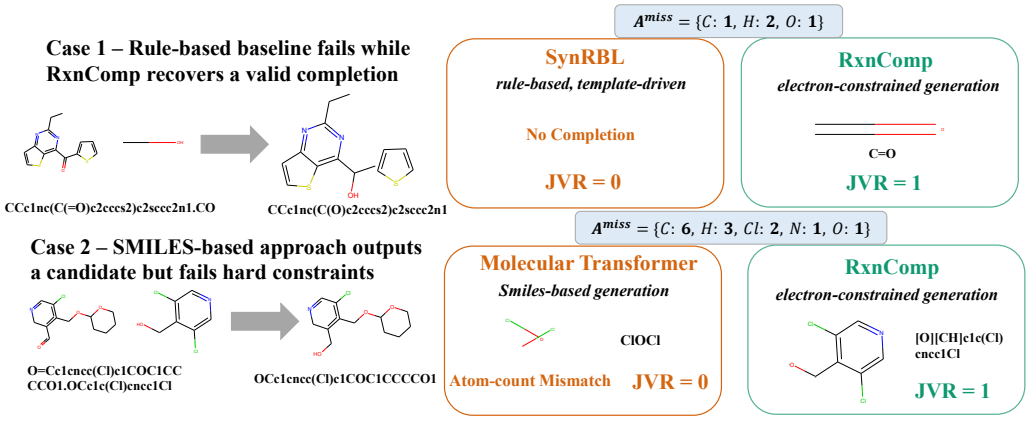

Figure 3: Representative case studies comparing **RxnComp** with SynRBL and sequence-based baselines, highlighting their characteristic strengths and failure modes.

these components mainly improve the discovery of feasible solutions and make constrained search more effective. **(4) Coverage and validity are largely decoupled.** Overall, the ablations reveal a separation of roles: closure constraints primarily determine validity (JVR), whereas model capacity and supervision mainly determine coverage (COV). This enables **RxnComp** to preserve high validity while improving coverage through stronger representations and training signals.

5.4 Case Study

Rather than showing only successful examples, we present two representative cases highlighting the characteristic limitations of existing approaches and the corresponding advantages of **RxnComp**: (i) a failure of rule-based methods and (ii) a failure of sequence-based methods. The corresponding structures are shown in Figure 3. **Case 1: Rule-based limitation (SynRBL).** This case illustrates the limited coverage of rule-based methods. The reaction leaves a residual atom budget $\mathcal{A}^{\text{miss}} = \{\text{C:1}, \text{H:2}, \text{O:1}\}$, corresponding to a CH_2O fragment. SynRBL fails because no template in its rule library matches the transformation and therefore returns an empty completion. In contrast, **RxnComp** correctly predicts `C=O` (formaldehyde) as the top-1 byproduct, exactly matching the residual budget. The resulting ER matrix is sparse and satisfies all validity constraints, showing that **RxnComp** can generalize beyond fixed rule sets by searching directly in electron-redistribution space. **Case 2: Sequence-based model failure under weak constraints.** This case highlights a different failure mode. The reaction leaves a residual atom budget $\mathcal{A}^{\text{miss}} = \{\text{C:6}, \text{H:3}, \text{Cl:2}, \text{N:1}, \text{O:1}\}$. The sequence-based baseline predicts `ClOCl`, which is syntactically valid but contains only two chlorines and one oxygen, violating atom conservation and leaving the remaining atoms unaccounted for. In contrast, **RxnComp** predicts `[O][CH]c1c(Cl)cncc1Cl`, which exactly matches the required composition and satisfies all conservation constraints. This illustrates a key limitation of unconstrained sequence generation: valid molecular strings need not correspond to globally consistent reactions, whereas **RxnComp** enforces conservation throughout decoding.

6 Conclusion and Future Work.

In this work, we formulate reaction completion as a conditional generation problem and introduce **RxnComp**, an autoregressive framework that infers missing byproducts through constrained completion of the electron redistribution (ER) matrix. By explicitly modeling electron-level consistency during generation, **RxnComp** produces high-quality completion candidates while maintaining broad coverage across diverse reactions. We further establish a reference-free evaluation protocol for large-scale unlabeled reaction data, together with reference-based evaluation when ground-truth byproducts are available. Empirical results demonstrate that **RxnComp** consistently outperforms existing reaction-completion baselines in both candidate quality and generalization, establishing reaction completion as a practical and scalable setting for improving the completeness of reaction datasets. In future work, we plan to explore more expressive generative architectures, improved reference-free metrics, and post-training strategies to further enhance coverage and chemical plausibility.

AI USE STATEMENT

In this work, we employed generative AI tools to support writing and language refinement. Generative AI was not used to retrieve references, and no automated research tools were employed in this study. We also used generative AI to support method implementation. All AI-assisted outputs, including implementation assistance and writing suggestions, were thoroughly reviewed and verified by the authors. The authors assume full responsibility for the final content of this work, including all text, claims, implementations, and other artifacts created with the assistance of generative AI.

ETHICS STATEMENT

This study focuses on computational modeling for chemical reaction completion and does not involve human participants, personally identifiable information, or sensitive personal data. The proposed framework is intended to improve the completeness and chemical consistency of reaction datasets by inferring plausible missing byproducts under explicit conservation constraints. Generated completions should be regarded as computational hypotheses rather than experimentally verified reaction outcomes and may require further validation before being used in downstream chemical analysis or experimental decision-making. We are not aware of any significant ethical risks arising directly from the datasets, methodology, or experimental procedures considered in this work.

REPRODUCIBILITY STATEMENT

We evaluate **RxnComp** on the USPTO-480K reaction dataset using the data splits and evaluation protocol described in the manuscript. The paper specifies the main components required to reproduce our framework, including the construction and representation of the electron redistribution matrix, constrained autoregressive generation, validity constraints, decoding procedures, and evaluation metrics. To support reproducibility, we plan to release the source code, preprocessing pipeline, model configurations and hyperparameters, trained checkpoints, inference scripts, and evaluation code. We will also provide the processed dataset splits and detailed procedures used to construct reaction-completion instances and compute the reported metrics.

REFERENCES

Akari Asai, Jacqueline He, Rulin Shao, Weijia Shi, Amanpreet Singh, Joseph Chee Chang, Kyle Lo, Luca Soldaini, Sergey Feldman, Mike D'Arcy, David Wadden, Matt Latzke, Jenna Sparks, Jena D. Hwang, Varsha Kishore, Minyang Tian, Pan Ji, Shengyan Liu, Hao Tong, Bohao Wu, Yanyu Xiong, Luke Zettlemoyer, Graham Neubig, Daniel S. Weld, Doug Downey, Wen tau Yih, Pang Wei Koh, and Hannaneh Hajishirzi. Synthesizing scientific literature with retrieval-augmented language models. *Nature*, 650(8103):857–863, 2026. doi: 10.1038/s41586-025-10072-4.

Peter Atkins and Loretta Jones. *Chemical Principles: The Quest for Insight*. W.H. Freeman, 7th edition, 2016. ISBN 978-1464183959.

Hangrui Bi, Hengyi Wang, Chence Shi, Connor Coley, Jian Tang, and Hongyu Guo. Non-autoregressive electron redistribution modeling for reaction prediction. In *International Conference on Machine Learning*, pp. 904–913. PMLR, 2021.

John Bradshaw, Matt J. Kusner, Brooks Paige, Marwin H. S. Segler, and José Miguel Hernández-Lobato. A generative model for electron paths. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019a.

John Bradshaw, Matt J. Kusner, Brooks Paige, Marwin H. S. Segler, and José Miguel Hernández-Lobato. Generating molecules via chemical reactions, 2019b. URL https://openreview.net/forum?id=BJlQEILY_N.

Andres M. Bran, Théo A. Neukomm, Daniel P. Armstrong, Zlatko Joncev, and Philippe Schwaller. Chemical reasoning in llms unlocks steerable synthesis planning and reaction mechanism elucidation. *CoRR*, abs/2503.08537, 2025. doi: 10.48550/ARXIV.2503.08537. URL https://doi.org/10.48550/arXiv.2503.08537.

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 2020, 13-18 July 2020, Virtual Event*, Proceedings of Machine Learning Research, pp. 1608–1616. PMLR, 2020. URL http://proceedings.mlr.press/v119/chen20k.html.

Shuan Chen and Yousung Jung. Deep retrosynthetic reaction prediction using local reactivity and global attention. *JACS Au*, 1(10):1612–1620, 2021.

Sunghwan Choi. Prediction of transition state structures of gas-phase chemical reactions via machine learning. *Nature Communications*, 14(1):1168, 2023. doi: 10.1038/s41467-023-36823-3. URL https://doi.org/10.1038/s41467-023-36823-3.

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.

Hanjun Dai, Chengtao Li, Connor W. Coley, Bo Dai, and Le Song. Retrosynthesis prediction with conditional graph logic network. *CoRR*, abs/2001.01408, 2020.

Chenru Duan, Guan-Horng Liu, Yuanqi Du, Tianrong Chen, Qiyuan Zhao, Haojun Jia, Carla P. Gomes, Evangelos A. Theodorou, and Heather J. Kulik. Optimal transport for generating transition states in chemical reactions. *Nature Machine Intelligence*, 7:615–626, 2025. doi: 10.1038/s42256-025-01010-0. URL https://doi.org/10.1038/s42256-025-01010-0.

Leonard Galustian et al. Goflow: efficient transition state geometry prediction with flow matching and e (3)-equivariant neural networks. *Digital Discovery*, 2025. doi: 10.1039/d5dd00283d. URL https://pubs.rsc.org/en/content/articlelanding/2025/dd/d5dd00283d.

Wenhao Gao, Shitong Luo, and Connor W Coley. Generative ai for navigating synthesizable chemical space. *Proceedings of the National Academy of Sciences*, 122(41):e2415665122, 2025.

Samuel Genheden and Esben Jannik Bjerrum. Paroutes: towards a framework for benchmarking retrosynthesis route predictions. *Digital Discovery*, 1:410–420, 2022. doi: 10.1039/D2DD00015F. URL https://pubs.rsc.org/en/content/articlelanding/2022/dd/d2dd00015f.

Samuel Genheden, Amol Thakkar, Veronika Chadimová, Jean-Louis Reymond, Ola Engkvist, and Esben Jannik Bjerrum. Aizynthfinder: a fast, robust and flexible open-source software for retrosynthetic planning. *J. Cheminformatics*, 12(1):70, 2020. doi: 10.1186/S13321-020-00472-1. URL https://doi.org/10.1186/s13321-020-00472-1.

Taicheng Guo, Kehan Guo, Bozhao Nan, Zhenwen Liang, Zhichun Guo, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. What can large language models do in chemistry? A comprehensive benchmark on eight tasks. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023.

Peng Han, Peilin Zhao, Chan Lu, Junzhou Huang, Jiaxiang Wu, Shuo Shang, Bin Yao, and Xiangliang Zhang. Gnn-retro: Retrosynthetic planning with graph neural networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(4), 2022.

Siqi Hong, Hankz Hankui Zhuo, Kebing Jin, Guang Shao, and Zhanwen Zhou. Retrosynthetic planning with experience-guided monte carlo tree search. *Communications Chemistry*, 6(1):120, 2023.

Wengong Jin, Connor Coley, Regina Barzilay, and Tommi Jaakkola. Predicting organic reaction outcomes with weisfeiler-lehman network. *Advances in neural information processing systems*, 30, 2017.

Joonyoung F Joung, Mun Hong Fong, Nicholas Casetti, Jordan P Liles, Ne S Dassanayake, and Connor W Coley. Electron flow matching for generative reaction mechanism prediction. *Nature*, 645(8079):115–123, 2025.

Seonghwan Kim, Jeheon Woo, and Woo Youn Kim. Diffusion-based generative ai for exploring transition states from 2d molecular graphs. *Nature Communications*, 15(1):341, 2024. doi: 10.1038/s41467-023-44629-6. URL https://doi.org/10.1038/s41467-023-44629-6.

David Kovacs, William McCorkindale, and Alpha Lee. Quantitative interpretation explains machine learning models for chemical reaction prediction and uncovers bias. *Nature Communications*, 12, 03 2021. doi: 10.1038/s41467-021-21895-w.

Youngchun Kwon, Dongseon Lee, Youn-Suk Choi, and Seokho Kang. Uncertainty-aware prediction of chemical reaction yields with graph neural networks. *Journal of Cheminformatics*, 14(1):2, 2022.

Jakub Lála, Odhran O'Donoghue, Aleksandar Shtedritski, Sam Cox, Samuel G. Rodriques, and Andrew D. White. Paperqa: Retrieval-augmented generative agent for scientific research. *CoRR*, abs/2312.07559, 2023. doi: 10.48550/ARXIV.2312.07559. URL https://doi.org/10.48550/arXiv.2312.07559.

Seul Lee, Karsten Kreis, Srimukh Veccham, Meng Liu, Danny Reidenbach, Saee Paliwal, Weili Nie, and Arash Vahdat. Exploring synthesizable chemical space with iterative pathway refinements. In *International conference on learning representations*, 2026.

Kangjie Lin, Youjun Xu, Jianfeng Pei, and Luhua Lai. Automatic retrosynthetic route planning using template-free models. *Chemical Science*, 11, 03 2020. doi: 10.1039/C9SC03666K.

Songtao Liu, Zhengkai Tu, Minkai Xu, Zuobai Zhang, Lu Lin, Rex Ying, Jian Tang, Peilin Zhao, and Dinghao Wu. Fusionretro: molecule representation fusion via in-context learning for retrosynthetic planning. In *International Conference on Machine Learning*, pp. 22028–22041. PMLR, 2023.

D. M. Lowe. Extraction of chemical structures and reactions from the literature. 2012.

Jieyu Lu and Yingkai Zhang. Unified deep learning model for multitask reaction predictions with explanation. *J. Chem. Inf. Model.*, 62(6):1376–1387, 2022.

Shitong Luo, Wenhao Gao, Zuofan Wu, Jian Peng, Connor W. Coley, and Jianzhu Ma. Projecting molecules into synthesizable chemical spaces. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, Proceedings of Machine Learning Research, pp. 33289–33304. PMLR / OpenReview.net, 2024. URL https://proceedings.mlr.press/v235/luo24a.html.

Ziqiao Meng, Peilin Zhao, Yang Yu, and Irwin King. A unified view of deep learning for reaction and retrosynthesis prediction: Current status and future challenges. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, pp. 6775–6783, 8 2023a. doi: 10.24963/ijcai.2023/753. URL https://www.ijcai.org/proceedings/2023/0753.pdf.

Ziqiao Meng, Peilin Zhao, Yang Yu, and Irwin King. Doubly stochastic graph-based non-autoregressive reaction prediction. In *IJCAI 2023*. ijcai.org, 2023b.

Ralph H Petrucci, F Geoffrey Herring, Jeffry D Madura, and Carey Bissonnette. *General Chemistry: Principles and Modern Applications*. Pearson, Toronto, Canada, 11th edition, 2017. ISBN 978-0132931281.

Tieu-Long Phan, Klaus Weinbauer, Thomas Gärtner, Daniel Merkle, Jakob L Andersen, Rolf Fagerberg, and Peter F Stadler. Reaction rebalancing: a novel approach to curating reaction databases. *Journal of Cheminformatics*, 16(1):82, 2024.

Can Polat, Mehmet Tuncel, Mustafa Kurban, Erchin Serpedin, and Hasan Kurban. xchemagents: Agentic ai for explainable quantum chemistry. *arXiv preprint arXiv:2505.20574*, 2025.

Shuo Ren, Can Xie, Pu Jian, Zhenjiang Ren, Chunlin Leng, and Jiajun Zhang. Towards scientific intelligence: A survey of llm-based scientific agents. *arXiv preprint arXiv:2503.24047*, 2025.

Mikołaj Sacha, Mikołaj Błaz, Piotr Byrski, Paweł Dabrowski-Tumanski, Mikołaj Chrominski, Rafał Loska, Paweł Włodarczyk-Pruszynski, and Stanisław Jastrzebski. Molecule edit graph attention network: modeling chemical reactions as sequences of graph edits. *Journal of Chemical Information and Modeling*, 61(7):3273–3284, 2021.

Mandana Saebi, Bozhao Nan, John E Herr, Jessica Wahlers, Zhichun Guo, Andrzej M Zurański, Thierry Kogej, Per-Ola Norrby, Abigail G Doyle, Nitesh V Chawla, et al. On the use of real-world datasets for reaction yield prediction. *Chemical science*, 14(19):4997–5005, 2023.

Nadine Schneider, Daniel M. Lowe, Roger A. Sayle, and Gregory A. Landrum. Development of a novel fingerprint for chemical reactions and its application to large-scale reaction classification and similarity. *Journal of Chemical Information and Modeling*, 55(1):39–53, 2015.

John S. Schreck, Connor W. Coley, and Kyle J. M. Bishop. Learning retrosynthetic planning through simulated experience. *ACS Central Science*, 5(6):970–981, 2019.

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.

Philippe Schwaller, Alain C Vaucher, Teodoro Laino, and Jean-Louis Reymond. Prediction of chemical reaction yields using deep learning. *Machine learning: science and technology*, 2(1):015016, 2021.

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, 2017a.

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, 2018. doi: 10.1038/nature25978. URL https://doi.org/10.1038/nature25978.

Marwin HS Segler and Mark P Waller. Modelling chemical reasoning to predict and invent reactions. *Chemistry–A European Journal*, 23(25):6118–6128, 2017b.

Runhan Shi, Gufeng Yu, Xiaohong Huo, and Yang Yang. Prediction of chemical reaction yields with large-scale multi-view pre-training. *Journal of Cheminformatics*, 16(1):22, 2024.

Runhan Shi, Gufeng Yu, Letian Chen, and Yang Yang. Yieldfcp: enhancing reaction yield prediction via fine-grained cross-modal pre-training. *Artificial Intelligence Chemistry*, 3(1):100085, 2025.

Vignesh Ram Somnath, Charlotte Bunne, Connor W. Coley, Andreas Krause, and Regina Barzilay. Learning graph models for retrosynthesis prediction. In *NeurIPS 2021*, 2021.

Henry W. Sprueill, Carl Edwards, Khushbu Agarwal, Mariefel V. Olarte, Udishnu Sanyal, Conrad Johnston, Hongbin Liu, Heng Ji, and Sutanay Choudhury. CHEMREASONER: heuristic search over a large language model's knowledge space using quantum-chemical feedback. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, Proceedings of Machine Learning Research, pp. 46351–46374. PMLR / OpenReview.net, 2024. URL https://proceedings.mlr.press/v235/sprueill24a.html.

Kyle Swanson, Gary Liu, Denise B Catacutan, Autumn Arnold, James Zou, and Jonathan M Stokes. Generative ai for designing and validating easily synthesizable and structurally novel antibiotics. *Nature machine intelligence*, 6(3):338–353, 2024.

Mohammadamin Tavakoli, Yin Ting T Chiu, Pierre Baldi, Ann Marie Carlton, and David Van Vranken. Rmechdb: A public database of elementary radical reaction steps. *Journal of chemical information and modeling*, 63(4):1114, 2023.

Igor Tetko, Pavel Karpov, Ruud Deursen, and Guillaume Godin. State-of-the-art augmented nlp transformer models for direct and single-step retrosynthesis. *Nature Communications*, 11, 11 2020. doi: 10.1038/s41467-020-19266-y.

Alessandra Toniato, Philippe Schwaller, Antonio Cardinale, Joppe Geluykens, and Teodoro Laino. Unassisted noise reduction of chemical reaction datasets. *Nature Machine Intelligence*, 3(6): 485–494, 2021.

Austin Tripp, Krzysztof Maziarz, Sarah Lewis, Marwin H. S. Segler, and José Miguel Hernández-Lobato. Retro-fallback: retrosynthetic planning in an uncertain world. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL https://openreview.net/forum?id=dl0u4ODCuW.

Andrea Valenti, Davide Bacciu, and Antonio Vergari. Chemalgebra: Algebraic reasoning on chemical reactions. In *International Joint Conference on Neural Networks, IJCNN 2024, Yokohama, Japan, June 30 - July 5, 2024*, pp. 1–8. IEEE, 2024.

Gabriel Vogel, Minouk Noordsij, Evgeny Pidko, and Jana M Weber. Completerxn: Toward completing open chemical reaction databases. *arXiv preprint arXiv:2605.00222*, 2026.

Yue Wan, Chang-Yu Hsieh, Ben Liao, and Shengyu Zhang. Retroformer: Pushing the limits of end-to-end retrosynthesis transformer. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pp. 22475–22490. PMLR, 2022.

Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, Anima Anandkumar, Karianne Bergen, Carla P. Gomes, Shirley Ho, Pushmeet Kohli, Joan Lasenby, Jure Leskovec, Tie-Yan Liu, Arjun K. Manrai, Debora S. Marks, Bharath Ramsundar, Le Song, Jimeng Sun, Jian Tang, Petar Velickovic, Max Welling, Linfeng Zhang, Connor W. Coley, Yoshua Bengio, and Marinka Zitnik. Scientific discovery in the age of artificial intelligence. *Nat.*, 620(7972):47–60, 2023. doi: 10.1038/S41586-023-06221-2. URL https://doi.org/10.1038/s41586-023-06221-2.

Xiaorui Wang, Yuquan Li, Jiezhong Qiu, Guangyong Chen, Huanxiang Liu, Benben Liao, Chang-Yu Hsieh, and Xiaojun Yao. Retroprime: A diverse, plausible and transformer-based method for single-step retrosynthesis predictions. *Chemical Engineering Journal*, 420:129845, 2021.

David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *J. Chem. Inf. Comput. Sci.*, 28(1):31–36, 1988.

Mengsong Wu, YaFei Wang, Yidong Ming, Yuqi An, Yuwei Wan, Wenliang Chen, Binbin Lin, Yuqiang Li, Tong Xie, and Dongzhan Zhou. Chematagent: Enhancing llms for chemistry and materials science through tree-search based tool learning. *CoRR*, abs/2506.07551, 2025. doi: 10.48550/ARXIV.2506.07551. URL https://doi.org/10.48550/arXiv.2506.07551.

Yutaro Yamada, Robert Tjarko Lange, Cong Lu, Shengran Hu, Chris Lu, Jakob N. Foerster, Jeff Clune, and David Ha. The AI scientist-v2: Workshop-level automated scientific discovery via agentic tree search. *CoRR*, abs/2504.08066, 2025. doi: 10.48550/ARXIV.2504.08066. URL https://doi.org/10.48550/arXiv.2504.08066.

Chaochao Yan, Qianggang Ding, Peilin Zhao, Shuangjia Zheng, Jinyu Yang, Yang Yu, and Junzhou Huang. Retroxpert: Decompose retrosynthesis prediction like A chemist. In *NeurIPS 2020*, 2020.

Chonghuan Zhang, Adarsh Arun, and Alexei A Lapkin. Completing and balancing database excerpted chemical reactions with a hybrid mechanistic-machine learning approach. *ACS Omega*, 9 (16):18385–18399, 2024. doi: 10.1021/acsomega.4c00262.

Yu Zhang, Yang Han, Shuai Chen, Ruijie Yu, Xin Zhao, Xianbin Liu, Kaipeng Zeng, Mengdi Yu, Jidong Tian, Feng Zhu, et al. Large language models to accelerate organic chemistry synthesis. *Nature Machine Intelligence*, 7(7):1010–1022, 2025.

A APPENDIX

A.1 ELECTRON REDISTRIBUTION MATRIX REPRESENTATION

The electron redistribution matrix representation $\Delta\mathbf{E}$ provides a unified way to characterize electron flow in chemical reactions. To justify the use of the doubly conservative property as a practical physical constraint, we first revisit classical electron-counting rules that govern chemical stability.

- **Duet Rule.** For light elements such as hydrogen and helium, stability is achieved when the valence shell contains two electrons. For example, hydrogen typically forms one covalent bond to satisfy this requirement.
- **Octet Rule.** For most main-group elements, atoms tend to achieve a stable configuration with eight valence electrons by gaining, losing, or sharing electrons. For instance, carbon commonly forms four bonds (e.g., CH_4) and oxygen forms two bonds (e.g., H_2O). While widely applicable, the octet rule admits known exceptions such as incomplete octets, expanded octets, and radical species.
- **18-Electron Rule.** For transition metal complexes, stability is often achieved when the total valence electron count at the metal center reaches 18, corresponding to filled s, p, and d orbitals. This count includes contributions from both the metal and its coordinated ligands.

These rules can be viewed as electron-counting constraints that guide chemically valid configurations:

$$\text{Valence electrons} = \begin{cases} 2, & \text{(duet rule)} \\ 8, & \text{(octet rule)} \\ 18, & \text{(18-electron rule)} \end{cases} \tag{4}$$

Electron-Counting Perspective. Electron-counting rules reflect a fundamental principle: atoms in stable molecules tend to maintain preferred valence electron configurations. As a result, chemical reactions primarily redistribute electrons among atoms while preserving these local stability conditions. From this perspective, the doubly conservative property of $\Delta\mathbf{E}$—i.e., zero row and column sums—naturally encodes electron conservation at the atomic level, while the pairwise symmetry of electron exchange further constrains how electrons can be redistributed between atoms. For example, a carbon atom in organic molecules typically maintains an octet configuration (four bonds); during a reaction, although bonding partners may change, the total electron count around carbon remains approximately constant. Meanwhile, any electron transferred from atom i to atom j must be reflected consistently in the interaction between the two atoms, inducing a structured, balanced redistribution pattern. Together, these electron-counting and symmetry constraints explain why $\Delta\mathbf{E}$ should satisfy a doubly conservative structure, making it a practical and physically grounded constraint for modeling chemically valid reaction transformations.

Practical Considerations. We note that not all chemical reactions strictly satisfy the doubly conservative constraint, particularly in cases involving radical species, charged intermediates, or complex electron delocalization effects. In such scenarios, transient violations of local electron-counting rules or asymmetric electron redistribution may occur. Nevertheless, the doubly conservative assumption remains highly effective in practice, as the majority of stable reactions—especially in organic chemistry—closely adhere to these electron-counting and symmetry principles. As a result, enforcing this constraint significantly reduces the search space and filters out chemically implausible solutions, while still retaining a large portion of valid reaction outcomes. Therefore, despite its idealized nature, the doubly conservative property provides a powerful and practical inductive bias for modeling and recovering a wide range of chemically meaningful reactions.

A.2 USPTO-480K DATASET ANALYSIS

Table 4 reports the sample count at each stage of data preparation. The mixed splits are obtained by union: $33{,}946 = 33{,}122 + 824$ for training and $2{,}417 = 2{,}361 + 56$ for validation. The inference-only test split contains $39{,}687$ reactions obtained from the cleaned USPTO test split by removing

Table 4: Sample counts at each stage of the dataset pipeline. The mixed training and validation splits union ground-truth and heuristically completed reactions; the test split removes all ground-truth reactions so that the proxy metrics in Table 1 measure byproduct *inference* rather than recall against an annotated answer.

Subset	Train	Valid	Test
Original cleaned USPTO-480K split	409,035	30,000	40,000
Ground-truth byproducts (all rows)	3,123	208	313
Ground-truth byproducts (non-empty)	824	56	86
Heuristic completion (script-completable)	33,122	2,361	3,133
Mixed training set	**33,946**	**2,417**	—
Inference-only test set (no ground-truth)	—	—	**39,687**

every reaction with a ground-truth byproduct annotation; the proxy metrics reported in Table 1 are therefore evaluated only on reactions for which no answer is available in the source literature.

Heuristic completion templates. The heuristic-completion subset is produced by exact atom-balance matching against a small library of byproduct templates: we accept the first template whose atomic composition equals the residual atom budget $\mathcal{A}^{\text{miss}}$ of the reaction. The templates fall into five chemical families, each chosen to cover a high-frequency byproduct pattern in patent reactions:

- *Small inorganic byproducts*: H_2O, H_2, O_2, N_2, CO_2, NH_3.
- *Hydrogen halides and free halogens*: HF, HCl, HBr, HI, F_2, Cl_2, Br_2, I_2.
- *Common small organic fragments*: HCN, H_2CO, acetic acid, methanol, ethanol, acetaldehyde.
- *Sulfur and phosphorus species*: H_2S, SO_2, sulfuric acid, sulfurous acid, phosphoric acid.
- *Ionic salts and acid–base pairs*: alkali halides (NaCl, KCl, NaBr, KBr, NaF, KF, NaI, KI), alkali hydroxides (NaOH, KOH), and ammonium halides (NH_4Cl, NH_4Br).

This library is intentionally small. Its purpose is to provide weak electron-level supervision at scale, and a larger or more elaborate template set would risk biasing the model toward complex completions that do not generalise to the open-ended distribution of byproducts encountered at test time.

A.3 Algorithm Details

Embedding construction. The **RxnComp** scorer uses learned embeddings for both atom-level and entry-level discrete attributes. For atom i, we represent its input attributes by the tuple

$$(\text{atomic_num}_i,\ \text{formal_charge}_i,\ \text{degree}_i,\ \mathbb{I}[i \in A^{\text{obs}}]),$$

where A^{obs} denotes the set of observed atoms. The initial atom representation is constructed as

$$\mathbf{h}_i^{(0)} = \text{MLP}\Big(\text{Emb}_{\text{atom}}(\text{atomic_num}_i) \parallel \text{Emb}_{\text{charge}}(\text{formal_charge}_i) \parallel \text{Emb}_{\text{degree}}(\text{degree}_i) \parallel \text{Emb}_{\text{obs}}(\mathbb{I}[i \in A^{\text{obs}}])\Big).$$

where the four embedding dimensions are 64, 16, 16, and 8, respectively, giving a 104-dimensional concatenated vector before projection to the hidden size.

For each candidate entry (i, j), the scorer further concatenates entry-level features, including the pair-type embedding, row- and column-budget embeddings, reactant bond-order embeddings and partial product-bond embeddings derived only from the observed main product, step embedding, normalized scalar budgets. The residual budget IDs are obtained by integer rounding and clipping to $[-B_{\max}, B_{\max}]$, followed by a shift of $+B_{\max}$, with $B_{\max} = 256$ in the current implementation. This clipping range is used only for embedding residual budgets and is independent of the ER entry-value bound $K = 4$.

Candidate values $v \in [-K, K]$ are represented by a dedicated value embedding $\text{Emb}_{\text{value}}(v + K)$, and infeasible values are removed by the FSM feasibility mask before applying the log-softmax.

Algorithm 1 Sample construction.

Require: Mapped reaction $r \gg p_{\text{main}}$
Ensure: Sample with observed atoms $\mathcal{A}^{\text{obs}}$, residual atoms $\mathcal{A}^{\text{miss}}$, observed ER block $\Delta\mathbf{E}_{oo}$, and ordered decoding plan $\mathcal{I}$

Parse r and p_{main} with explicit hydrogens added
Compute the atom mapping π from reactants to main product using map numbers, falling back to MCS if necessary
$\mathcal{A}^{\text{obs}} \leftarrow \text{dom}(\pi)$; $\mathcal{A}^{\text{miss}} \leftarrow \mathcal{A}(r) \setminus \mathcal{A}^{\text{obs}}$
Reorder atoms as $[\mathcal{A}^{\text{obs}}, \mathcal{A}^{\text{miss}}]$
Build $\mathbf{B}^{\mathcal{R}}$ on the reordered indexing
Build $\mathbf{B}^{\mathcal{P}}_{oo}$ from the mapped main product
$\Delta\mathbf{E}_{oo} \leftarrow \mathbf{B}^{\mathcal{P}}_{oo} - \mathbf{B}^{\mathcal{R}}_{oo}$, and apply closure
Plan $\mathcal{I}$ by enumerating upper-triangular missing entries in row-major order; reserve one slack entry per row and per column for deterministic closure
Initialise per-atom electron budgets $r_a, c_a \leftarrow -\sum_j \Delta E_{aj}$ for every $a \in \mathcal{A}^{\text{obs}}$
return Sample $(\mathcal{A}^{\text{obs}}, \mathcal{A}^{\text{miss}}, \Delta\mathbf{E}_{oo}, \mathcal{I}, \mathbf{r}, \mathbf{c})$

Algorithm 2 FSM-guided beam search.

Require: Sample, beam width B, value range $\mathcal{V} = \{-K, \ldots, K\}$, scorer s_θ
Ensure: Up to B completed ER matrices $\Delta\mathbf{E}$

Initialise beam $\mathcal{H}_0 \leftarrow \{(\Delta\mathbf{E}_{oo}, \mathbf{r}, \mathbf{c}, 0)\}$
for $t = 1$ to $|\mathcal{I}_{\text{free}}|$ **do**
    Let (i_t, j_t) be the next entry in $\mathcal{I}$
    $\mathcal{H}_t \leftarrow \emptyset$
    for each state $(\Delta\mathbf{E}^{(<t)}, \mathbf{r}^{(t-1)}, \mathbf{c}^{(t-1)}, \ell) \in \mathcal{H}_{t-1}$ **do**
        for each $v \in \mathcal{V}$ **do**
            Evaluate $m_t(v)$ from Eq. (2): value-domain ∩ non-negative bond order ∩ budget reachability ∩ chemical valence
            if $m_t(v) = 1$ **then**
                Tentatively set $\Delta E_{i_t j_t} = \Delta E_{j_t i_t} = v$
                Update $r_{i_t}, r_{j_t}, c_{i_t}, c_{j_t}$ by symmetric coupling
                Push extended state with score $\ell + \log p_\theta(v \mid \cdot)$ into $\mathcal{H}_t$
            end if
        end for
    end for
    $\mathcal{H}_t \leftarrow \text{top-B}(\mathcal{H}_t)$ by score
    if $\mathcal{H}_t = \emptyset$ **then**
        return $\emptyset$
    end if
end for
for each surviving state **do**
    Fill the reserved row- and column-slack entries from the residual budgets and close the corner
    Project to a non-negative integer bond-order matrix and discard states that fail the chemical valence check
end for
return surviving states sorted by score

Architecture variants and motivation. We report two architecture variants in the main experiments: **RxnComp (lightweight)** and **RxnComp**. The motivation for reporting a lightweight variant is to separate the effect of the proposed electron-constrained generation framework from the effect of model capacity. If strong results were observed only with a larger encoder, the improvement could be partly attributed to representation scaling rather than to the constrained ER-matrix completion formulation itself. The lightweight variant therefore serves as a compact instantiation of the same framework, allowing us to verify whether the finite-state electron-budget decoder, feasibility-masked scoring, and self-supervised ER reconstruction remain effective under a simpler atom encoder.

Algorithm 3 One epoch of training.

Require: Supervised set $\mathcal{D}$, masked counterpart $\mathcal{D}_{\text{mask}}$ derived from $\mathcal{D}$, weight λ_{ssl}

1: **for** mini-batch $\mathcal{B} \subset \mathcal{D}$ with paired $\widetilde{\mathcal{B}} \subset \mathcal{D}_{\text{mask}}$ **do**
2: Build the supervised teacher-forcing trajectory along $\mathcal{I}$
3: Build the self-supervised trajectory by masking a random subset of product entries to obtain $\Delta\mathbf{E}^{\star}$
4: Compute the joint loss as in Eq. (3):
5: $$\mathcal{L} = \sum_{(i,j)\in\mathcal{I}_{\text{miss}}} \text{CE}(p_\theta(\cdot \mid \Delta\mathbf{E}_{<ij}), \Delta\mathbf{E}_{ij}) + \lambda_{\text{ssl}} \sum_{(i,j)\in\mathcal{I}_{\text{mask}}} \text{CE}\big(p_\theta(\cdot \mid \Delta\mathbf{E}^{\star}), \Delta\mathbf{E}^{\star}_{ij}\big)$$
6: Backpropagate $\mathcal{L}$, clip the gradient norm, and take one AdamW step
7: **end for**
8: Validate by beam decoding; retain the checkpoint with the highest validation COV-JVR

Table 5: Architecture comparison between the two reported **RxnComp** variants.

Component	**RxnComp** (lightweight)	**RxnComp**
Atom encoder backbone	Transformer encoder	Bond-message GNN + Transformer encoder
Temporal conditioning	Scalar decoding-progress conditioning	Learned step-bucket conditioning
Entry scoring	Both variants score discrete candidate values $v \in [-K, K]$ using pair/context features and value embeddings, followed by masked log-softmax under FSM feasibility constraints.	
Training objective	Both variants use the same supervised completion loss and self-supervised reconstruction loss.	
Inference	Both variants use the same FSM-guided beam-search decoding procedure.	

Both variants share the same reaction-completion formulation, finite-state constrained decoding mechanism, candidate value space, feasibility-masked scoring rule, self supervised training objective, and beam-search inference procedure. Their main architectural difference lies in the atom encoder and the temporal conditioning used by the entry scorer. **RxnComp** (lightweight) uses a simpler Transformer-based atom encoder, while the full **RxnComp** model adopts a bond-message GNN combined with a Transformer encoder to obtain richer contextual atom representations.

Overall, **RxnComp** (lightweight) demonstrates that the proposed electron-constrained generation framework remains effective even with a compact atom encoder, while the full **RxnComp** model shows the additional benefit of stronger molecular representation learning. This comparison helps distinguish the contribution of the constrained decoding framework from that of increased encoder capacity.

Evaluation metrics. Let

$$\mathcal{D} = \{(\mathcal{R}^{(t)}, P^{(t)}_{\text{main}})\}_{t=1}^{T}$$

denote the evaluation set, where $\mathcal{R}^{(t)}$ denotes the reactant multiset and $P^{(t)}_{\text{main}}$ is the observed major product. For reaction t, a method returns an ordered set of coproduct candidates

$$\widehat{\mathcal{P}}^{(t)} = \{\widehat{P}^{(t,1)}_{\text{by}}, \ldots, \widehat{P}^{(t,n_t)}_{\text{by}}\},$$

where n_t denotes the number of candidates retained by the common output protocol before metric-specific validity filtering. We define the covered subset as

$$\mathcal{C} = \{t \mid n_t > 0\},$$

and, for Top-k evaluation,

$$K_t = \min(k, n_t),$$

such that only actually generated candidates are considered, without padding missing candidates or altering their original ranking.

(i) **Coverage (COV)** measures the fraction of test reactions for which at least one candidate is returned:

$$\mathrm{COV} = \frac{|\mathcal{C}|}{T}.$$

(ii) **SMILES Validity Rate (SVR)** evaluates whether at least one available Top-k candidate is chemically parseable and sanitizable by RDKit. Let $s_{t,i} \in \{0, 1\}$ denote this indicator. Then

$$\mathrm{SVR@}k = \frac{1}{|\mathcal{C}|} \sum_{t \in \mathcal{C}} \mathbb{I}\left[\exists\, i \le K_t : s_{t,i} = 1\right].$$

(iii) **Atom Conservation Rate (ACR)** evaluates whether at least one Top-k completion preserves the element-wise atom inventory. Let $\mathbf{A}(\cdot)$ denote the element-count vector and define

$$a_{t,i} = \mathbb{I}\left[\mathbf{A}(\mathcal{R}^{(t)}) = \mathbf{A}\left(P^{(t)}_{\mathrm{main}} \uplus \widehat{P}^{(t,i)}_{\mathrm{by}}\right)\right],$$

where $\uplus$ denotes multiset union. For candidates whose molecular structure cannot be parsed, we set $a_{t,i} = 0$. We then compute

$$\mathrm{ACR@}k = \frac{1}{|\mathcal{C}|} \sum_{t \in \mathcal{C}} \mathbb{I}\left[\exists\, i \le K_t : a_{t,i} = 1\right].$$

(iv) **Charge Conservation Rate (CCR)** evaluates conservation of total formal charge. Let $Q(M)$ denote the formal charge of molecule M, and define

$$c_{t,i} = \mathbb{I}\left[\sum_{M \in \mathcal{R}^{(t)}} Q(M) = \sum_{M \in P^{(t)}_{\mathrm{main}} \uplus \widehat{P}^{(t,i)}_{\mathrm{by}}} Q(M)\right].$$

For candidates for which the charge cannot be evaluated, we set $c_{t,i} = 0$. Accordingly,

$$\mathrm{CCR@}k = \frac{1}{|\mathcal{C}|} \sum_{t \in \mathcal{C}} \mathbb{I}\left[\exists\, i \le K_t : c_{t,i} = 1\right].$$

(v) **Electron Consistency Rate (ECR)** evaluates whether the electron-redistribution matrix reconstructed from a completed reaction satisfies the doubly conservative constraint. Let $\Delta\mathbf{E}^{(t,i)}$ denote the reconstructed ER matrix and define

$$e_{t,i} = \mathbb{I}\left[\left(\forall a,\ \sum_b \Delta E^{(t,i)}_{ab} = 0\right) \wedge \left(\forall b,\ \sum_a \Delta E^{(t,i)}_{ab} = 0\right)\right].$$

If an ER matrix cannot be reconstructed for a candidate, we set $e_{t,i} = 0$. ERR@k is then

$$\mathrm{ECR@}k = \frac{1}{|\mathcal{C}|} \sum_{t \in \mathcal{C}} \mathbb{I}\left[\exists\, i \le K_t : e_{t,i} = 1\right].$$

ESR measures conservation consistency rather than ground-truth correctness.

(vi) **Electron Sparsity Rate (ESR)** measures how localized the reconstructed electron redistribution is. For candidate i of reaction t, we define

$$\mathrm{ESR}_{t,i} = 1 - \frac{\sum_{a \neq b} \mathbb{I}\left[\Delta E^{(t,i)}_{ab} \neq 0\right]}{N_t(N_t - 1)},$$

where N_t is the number of atoms in the corresponding ER matrix. Since ERS is a continuous candidate-level score rather than a binary criterion, we first average over the actually available Top-k candidates for each reaction and subsequently across covered reactions:

$$\mathrm{ESR@}k = \frac{1}{|\mathcal{C}|} \sum_{t \in \mathcal{C}} \frac{1}{K_t} \sum_{i=1}^{K_t} \mathrm{ESR}_{t,i}.$$

Table 6: Training and inference hyperparameters.

Item	Value
Optimiser	AdamW
Weight decay	1×10^{-4}
Learning rate	1×10^{-3}
Epochs	4
Batch size	16
Gradient-norm clipping	1.0
Maximum entry magnitude K	4
Self-supervision weight λ_{ssl}	0.5
Beam width B	10

(vii) **Joint Validity Rate (JVR)** requires the *same candidate* to simultaneously satisfy SMILES validity, atom conservation, charge conservation, and electron-redistribution rationality. Defining

$$v_{t,i} = s_{t,i} a_{t,i} c_{t,i} e_{t,i},$$

we compute

$$\text{JVR@}k = \frac{1}{|\mathcal{C}|} \sum_{t \in \mathcal{C}} \mathbb{I}\left[\exists\, i \le K_t : v_{t,i} = 1\right].$$

Finally, we define

$$\text{COV-JVR@}k = \text{COV} \times \text{JVR@}k.$$

Since $\text{COV} = |\mathcal{C}|/T$, this is equivalently

$$\text{COV-JVR@}k = \frac{1}{T} \sum_{t \in \mathcal{C}} \mathbb{I}\left[\exists\, i \le K_t : v_{t,i} = 1\right].$$

All binary Top-k metrics are reaction-level existential metrics: a reaction is counted as successful if at least one of its actually available Top-k candidates satisfies the corresponding criterion. Missing candidates are never padded, and candidate rankings are never modified according to ground-truth labels or proxy validity scores.

Differential-AR baseline. To provide a controlled autoregressive baseline, we additionally construct **Differential-AR**, an in-house model that predicts reaction-completion changes autoregressively without the explicit electron-budget constraints used by **RxnComp**. Given the reactants and the observed major product, we first identify the residual atoms that are not accounted for by the major product and represent the missing reaction information as a sequence of discrete structural changes. Differential-AR then predicts these changes sequentially using a learned autoregressive scorer and generates multiple completion hypotheses through beam search.

Importantly, Differential-AR does not employ the finite-state electron-budget tracking or the doubly conservative feasibility mask of **RxnComp** during decoding. Its predictions are therefore selected according to the learned autoregressive scores without explicitly enforcing row- and column-wise electron conservation. After generation, the predicted structural changes are converted into coproduct candidates and evaluated using exactly the same method-independent evaluation pipeline as the other baselines. We introduce Differential-AR as a controlled baseline to distinguish the benefit of autoregressive structured generation from that of explicitly enforcing electron-redistribution constraints during decoding.

Hyperparameters. Table 6 lists the hyperparameters used to produce the results in Table 1. The two **RxnComp** variants reported in the main paper, the lightweight and the full model, share the same training schedule and self-supervision weight, and differ only in encoder capacity.

Table 7: Direct reaction-completion evaluation on the original 86 reactions with non-empty recorded coproducts. COV denotes the fraction of reactions for which a completion is produced, and EM@1 evaluates exact recovery of the recorded coproduct.

Method	**COV** ↑	**EM@1** ↑
Molecular Transformer	1.00	0.00
SynRBL	1.00	0.35
Rule-based Completion	0.59	0.59
RxnComp	0.95	0.62
RxnComp (lightweight)	0.95	**0.95**

Table 8: Agreement between SynRBL and **RxnComp** on reactions where SynRBL exactly recovers the recorded coproduct in the original GT86 subset.

Evaluation subset	**SynRBL exact**	**RxnComp Top-1 agreement**
SynRBL-correct reactions	30/86	30/30 (1.00)

A.4 Additional Evaluation and Robustness Analysis

This subsection provides additional analyses that complement the large-scale proxy evaluation in the main paper. We focus on direct reference-based evaluation, expert assessment, generalization beyond heuristic supervision, and robustness to alternative molecular and decoding serializations.

Direct evaluation on reactions with recorded coproducts. The large-scale USPTO evaluation necessarily relies on proxy metrics because recorded coproduct annotations are unavailable for most patent reactions. We therefore additionally evaluate reaction completion on the subset for which non-empty coproduct annotations are available.

The original annotated test subset contains 86 reactions with non-empty recorded coproducts. Table 7 reports the evaluation performed on this original subset. The full **RxnComp** model achieves an EM@1 of 0.62 on the original annotated subset, while the lightweight variant reaches 0.95. The latter result should be interpreted cautiously because this reference set is small and contains relatively simple completion patterns.

Agreement with SynRBL on the recorded-reference subset. We further examine whether **RxnComp** merely disagrees with the rule-based completion system on reactions where the latter is already correct. Among the 86 originally annotated reactions, SynRBL exactly recovers the recorded coproduct for 30 reactions. For all 30 of these cases, the top-1 prediction produced by **RxnComp** agrees with the SynRBL completion.

This agreement shows that the learned model preserves the common completion patterns captured by SynRBL when those rules match the recorded coproduct, while the broader coverage analyses below evaluate whether the model can generalize beyond that rule-supported regime.

Blinded human-expert evaluation. Automatic metrics do not fully capture whether a generated coproduct is chemically reasonable when several feasible completions exist. We therefore perform an additional blinded manual evaluation on 50 reactions randomly sampled from the inference test set without filtering by prediction validity.

For each reaction, a chemist evaluates the anonymized top-1 prediction from each method using the reactants and recorded major product as context. Invalid, unparsable, or chemically implausible predictions are counted as unsuccessful. We report the Expert Acceptance Rate, defined as the fraction of all sampled reactions for which the predicted coproduct set is judged chemically plausible and acceptable without further correction.

The expert assessment provides complementary evidence of end-to-end chemical plausibility rather than merely structural validity. Because the study contains only 50 reactions and relies on a single expert, we regard it as preliminary supporting evidence rather than a comprehensive human-evaluation benchmark.

Table 9: Blinded expert assessment on 50 randomly sampled reactions. The evaluation includes all sampled reactions rather than only successfully decoded or structurally valid cases.

Method	**Expert Acceptance Rate** ↑
Molecular Transformer	0%
SynRBL	42%
Rule-based Completion	32%
RxnComp	**65%**
RxnComp (lightweight)	61%

Table 10: Generalization analysis on a fixed 500-reaction subset. COV/JVR@1 is reported for the rule-based baseline and **RxnComp** under different measures of proximity to the heuristic training distribution.

Stratification	**Subset**	N	**Rule-based COV/JVR@1**	**RxnComp COV/JVR@1**
Overall	All reactions	500	0.128 / 0.128	**0.828 / 0.726**
Template coverage	Rule-covered	32	1.000 / 1.000	1.000 / 1.000
	Rule-uncovered	468	0.068 / 0.068	**0.816 / 0.707**
Residual formula	Seen	34	0.941 / 0.941	**1.000 / 1.000**
	Unseen	466	0.069 / 0.069	**0.815 / 0.706**
Reaction similarity	$[0.00, 0.25)$	231	0.022 / 0.022	**0.766 / 0.645**
	$[0.25, 0.50)$	182	0.088 / 0.088	**0.852 / 0.753**
	$[0.50, 0.75)$	58	0.345 / 0.345	**0.931 / 0.879**
	$[0.75, 1.00]$	29	0.793 / 0.793	**0.966 / 0.897**
Residual heavy atoms	0–2	52	0.654 / 0.654	**0.827 / 0.788**
	3–5	45	0.378 / 0.378	**0.867 / 0.822**
	6–10	135	0.096 / 0.096	**0.822 / 0.704**
	> 10	268	0.000 / 0.000	**0.825 / 0.709**

Generalization beyond heuristic supervision. Most weakly supervised training examples are constructed from heuristic completion rules, raising the possibility that the learned model merely distills the heuristic baseline. We examine this hypothesis on a fixed 500-reaction subset, stratified by proximity to the heuristic training distribution.

The largest differences occur outside the support of the heuristic baseline. For example, on 468 rule-uncovered reactions, the rule-based method achieves 0.068/0.068 COV/JVR@1, whereas **RxnComp** achieves 0.816/0.707. Similarly, on 466 reactions with unseen residual formulas, **RxnComp** maintains 0.815/0.706 COV/JVR@1. The performance gap persists for low-similarity reactions and reactions with large residual fragments. These results indicate that **RxnComp** does not simply reproduce the discrete template inventory used to construct the weak supervision, although heuristic completion remains an important source of inductive bias.

Robustness to molecular and decoding order. **RxnComp** uses an autoregressive factorization over missing ER-matrix entries, which raises the question of whether predictions are overly sensitive to arbitrary molecular or serialization orderings. We therefore evaluate several perturbations that preserve the underlying reaction while altering its representation or decoding order.

The results show only small variations under the tested perturbations, indicating limited empirical sensitivity to molecular serialization and to the alternative internal decoding orders considered here.

Several aspects of the representation help explain this robustness. First, each generated ER value is associated with an explicit atom pair (i, j) rather than with an anonymous sequence position. Its score therefore depends on the identities and reaction contexts of the corresponding atoms in addition to its autoregressive position.

Second, ER entries are structurally coupled through shared rows and columns. Entries involving the same atom share chemical context and residual electron budgets, which provides a strong structural signal beyond their serialization positions.

Table 11: Robustness to alternative molecular serializations and decoding orders. Randomized-SMILES and internal-order results are averaged over three seeds.

Input / decoding setting	COV	JVR@1	COV-JVR
Original	0.841	0.877	0.738
Fragment reorder	0.845	0.873	0.738
Randomized SMILES (3 seeds)	0.843	0.880	0.742
Alternative ER-entry order (3 seeds)	0.841	0.873	0.738

Table 12: Complete reaction completion performance on USPTO-480K. "–" denotes results that are not applicable, as some models cannot generate multiple byproduct candidates for each test sample. The best and second-best results are highlighted in **bold** and underline, respectively. Our **RxnComp** is highlighted in light green, while the key metric COV-JVR is highlighted in light red.

Model	COV-JVR↑	COV↑	SVR↑			ACR↑			CCR↑			ECR↑			ESR↑			JVR↑		
			@1	@3	@5	@1	@3	@5	@1	@3	@5	@1	@3	@5	@1	@3	@5	@1	@3	@5
Rule-based Completion	0.131	0.132	**1.000**	-	-	**1.000**	-	-	**0.999**	-	-	**1.000**	-	-	**1.000**	-	-	**0.999**	-	-
SynRBL	0.579	0.708	**1.000**	-	-	0.818	-	-	0.997	-	-	0.818	-	-	0.818	-	-	0.818	-	-
Differential-AR	0.002	0.023	**1.000**	1.000	1.000	0.118	0.120	0.120	0.631	0.970	0.974	0.118	0.120	0.120	0.118	0.120	0.120	0.113	0.113	0.113
Molecular Transformer	0.019	0.983	**1.000**	1.000	1.000	0.020	0.031	0.036	0.972	0.972	0.972	0.020	0.031	0.036	0.020	0.031	0.036	0.020	0.031	0.036
T5Chem	0.001	**1.000**	0.073	0.125	0.138	0.019	0.019	0.019	0.919	0.922	0.923	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
RxnComp *(lightweight)*	0.673	0.760	0.952	0.973	0.973	0.971	0.971	0.971	0.947	0.947	0.947	**1.000**	1.000	1.000	**0.997**	0.816	0.696	0.885	0.906	0.906
RxnComp	**0.738**	0.841	0.947	0.957	0.957	0.971	0.971	0.971	0.940	0.940	0.940	**1.000**	1.000	1.000	**0.996**	0.782	0.583	0.877	0.886	0.886

Third, the FSM restricts the admissible values of each entry using the remaining electron budgets, non-negative bond-order constraints, and chemical valence constraints. This substantially reduces the effective branching space at each generation step.

Finally, only one triangular portion of the symmetric ER matrix is generated, with its counterpart reconstructed deterministically, removing redundant generation decisions.

Sensitivity to atom correspondence. **RxnComp** does not use the original USPTO atom-map identifiers as direct neural features. For recorded coproduct labels, the major product is removed from the complete product side and the remaining fragments are exported with atom mapping removed. For heuristic labels, the residual elemental composition is matched against the candidate coproduct inventory.

During sample construction, atom correspondence is reconstructed using available positive atom-map identifiers when present, with maximum-common-substructure matching and same-element matching used as fallbacks. The resulting correspondence is then used to align the reactant and product graphs and construct the ER target.

We therefore expect simple relabeling of map identifiers to have little effect once the same atom correspondence is reconstructed. Nevertheless, ambiguous correspondences, particularly for symmetric structures, remain a possible source of preprocessing noise, and we do not claim complete robustness to atom-mapping errors.

A.5 Additional Experimental Analysis and Visualisations

This subsection provides additional quantitative and qualitative analyses complementing Section 5. For readability, the main paper reports Top-1 and Top-5 results for the primary comparison and ablation studies. Here, we provide the complete Top-1/3/5 results, followed by additional analyses of the diversity and sparsity of ER-matrix completions, hyperparameter sensitivity, training dynamics, and computational cost.

Complete reaction-completion results. Table 12 provides the complete Top-1/3/5 results corresponding to Table 1 in the main paper. All methods are evaluated under the same protocol and with the same proxy metrics. For methods that return only a single completion, Top-3 and Top-5 results are not applicable and are therefore denoted by "–". The complete results are consistent with the observations in the main text: coverage alone does not guarantee a valid reaction completion, whereas **RxnComp** achieves the strongest overall COV–JVR trade-off. The Top-k results further show how allowing multiple completion hypotheses affects molecular validity and joint validity across different methods.

Table 13: Complete ablation results for different variants of **RxnComp**. We evaluate both **RxnComp** and its lightweight version to assess the contribution of each component. All Top-1/3/5 proxy metrics are reported.

Model	Variant	COV-JVR↑	COV↑	SVR↑			ACR↑			CCR↑			ECR↑			ESR↑			JVR↑		
				@1	@3	@5	@1	@3	@5	@1	@3	@5	@1	@3	@5	@1	@3	@5	@1	@3	@5
RxnComp *(lightweight)*	Full model	**0.673**	0.760	0.952	0.973	0.973	0.971	0.971	0.971	0.947	0.947	0.947	1.000	1.000	1.000	0.997	0.816	0.696	0.885	0.906	0.906
	w/o step embedding	0.642	0.732	0.953	0.972	0.972	0.969	0.969	0.969	0.939	0.939	0.939	1.000	1.000	1.000	0.996	0.796	0.665	0.878	0.889	0.889
	w/o budget embedding	0.640	0.722	0.946	0.969	0.969	0.966	0.966	0.966	0.955	0.955	0.955	1.000	1.000	1.000	0.996	0.777	0.679	0.887	0.907	0.907
	w/o atom-pair embedding	0.579	0.650	0.947	0.972	0.972	0.969	0.969	0.969	0.956	0.956	0.956	1.000	1.000	1.000	0.996	0.798	0.688	0.891	0.906	0.906
	w/o SSL loss	0.543	0.600	0.959	0.963	0.963	0.963	0.963	0.963	0.963	0.963	0.963	1.000	1.000	1.000	0.996	0.780	0.654	0.905	0.908	0.908
	w/o electron budget constraint	0.018	1.000	0.707	0.728	0.739	0.976	0.976	0.976	0.915	0.915	0.915	0.051	0.051	0.051	0.997	0.994	0.993	0.018	0.018	0.018
RxnComp	Full model	**0.738**	0.841	0.947	0.957	0.957	0.971	0.971	0.971	0.940	0.940	0.940	1.000	1.000	1.000	0.996	0.783	0.582	0.877	0.886	0.886
	w/o step embedding	0.705	0.797	0.959	0.962	0.962	0.969	0.969	0.969	0.939	0.939	0.939	1.000	1.000	1.000	0.996	0.898	0.782	0.885	0.888	0.888
	w/o GNN	0.672	0.768	0.942	0.955	0.960	0.968	0.968	0.968	0.947	0.947	0.947	1.000	1.000	1.000	0.996	0.768	0.611	0.876	0.886	0.892
	w/o SSL loss	0.402	0.445	0.954	0.968	0.968	0.963	0.963	0.963	0.968	0.968	0.968	1.000	1.000	1.000	0.996	0.767	0.612	0.904	0.918	0.918
	w/o electron budget constraint	0.018	1.000	0.707	0.715	0.715	0.976	0.976	0.976	0.915	0.915	0.915	0.051	0.053	0.053	0.997	0.994	0.993	0.018	0.018	0.018

Complete ablation results. Table 13 expands the ablation study in Table 3 by reporting all Top 1/3/5 proxy metrics for both **RxnComp** and **RxnComp** *(lightweight)*. The detailed results reveal that different components primarily affect different aspects of the completion process. Removing the representation and learning components, including the step embedding, budget embedding, atom-pair embedding, GNN encoder, or self-supervised objective, mainly reduces coverage and consequently degrades the overall COV–JVR score. The self-supervised objective has a particularly pronounced effect on coverage, especially for the full **RxnComp** model.

In contrast, removing the electron-budget constraint produces a qualitatively different failure mode: coverage increases to 1.000, while ECR and JVR collapse to approximately 0.05 and 0.02, respectively. This result confirms that unconstrained decoding can readily produce completion candidates, but most of them fail to satisfy the required electron-redistribution and joint-validity constraints. The electron-budget mechanism therefore acts primarily as a feasibility constraint rather than merely improving output coverage.

Top-k diversity. A central design motivation of **RxnComp** is that beam search produces *diverse* valid completions rather than a single deterministic point estimate. In our test reactions, the top-5 candidates may contain more than one distinct canonical-SMILES byproduct set. This diversity is especially pronounced for reactions with multi-fragment residual budgets, where multiple atom partitions can be physically admissible and the FSM constraints alone cannot fully disambiguate among them. This empirical pattern is consistent with the role of beam search in our pipeline: it widens the reachable set of conservation-feasible completions, while the joint validity criterion supplies the final ranking signal.

ER-matrix sparsity. The completed ER matrices produced by **RxnComp** are typically sparse, with non-zero entries concentrated on a small number of atom pairs in the missing block. This observation agrees with the localised electron-rearrangement assumption discussed in Section 3: most reactions only require a limited number of bond-order changes to explain the missing byproduct structure. Figure 4 provides a qualitative beam-search example. The displayed candidates share the same observed ER block but differ slightly in the predicted missing-block entries, showing how beam search preserves multiple conservation-feasible hypotheses in the native electron-redistribution space.

Hyperparameter Sensitivity Analysis. Metrics. Following the main experiments, we report coverage (COV), joint validity (JVR@k), and their product COV-JVR@1:

$$\text{COV-JVR@1} = \text{COV} \times \text{JVR@1}.$$

Here, COV measures the fraction of reactions for which the decoder returns at least one completion candidate, while JVR@k measures the fraction of generated candidates that jointly satisfy molecular validity, atom conservation, charge conservation, and ER consistency among the top-k predictions. All sensitivity results are evaluated on the full inference-only test split used in the main experiments.

Beam-width sweep. We first vary the beam width B while fixing the maximum ER-entry magnitude to $K = 4$. The row with $B = 10$ corresponds to the default setting used for the main **RxnComp** results.

Self-supervision weight λ_{ssl}. We ablate the self-supervision weight λ_{ssl} introduced in Eq. (3), which controls the relative contribution of the self-supervised reconstruction loss to the overall train-

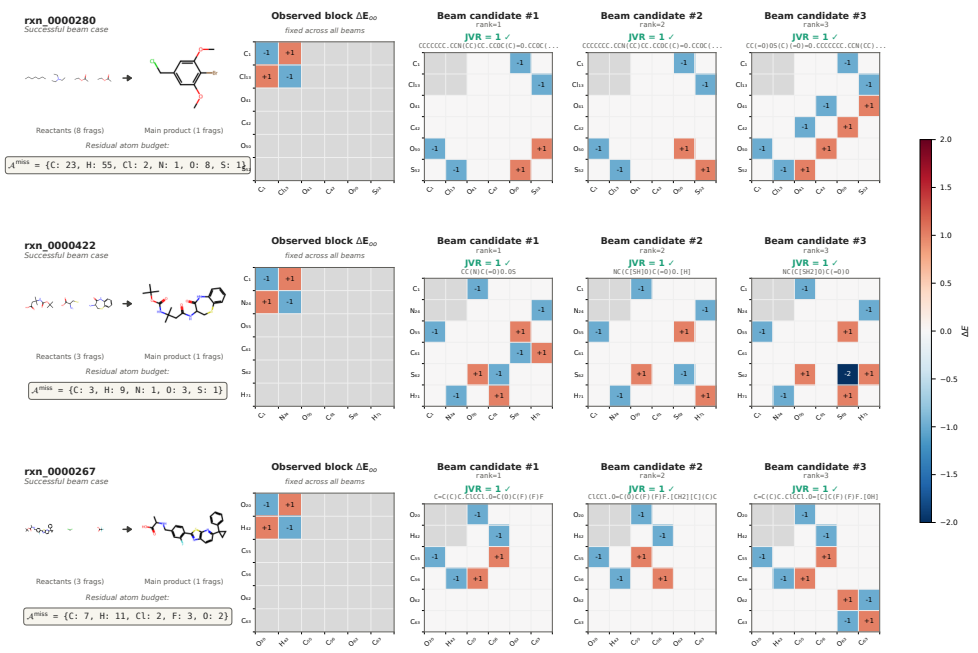

Figure 4: Qualitative visualisation of beam-search ER-matrix outputs produced by **RxnComp** for a representative test reaction. The candidates share the same observed ER block but differ in a small number of missing-block entries, corresponding to alternative conservation-feasible byproduct completions.

Table 14: Beam-width sensitivity on the full inference-only test split, with the maximum ER-entry magnitude fixed to $K = 4$.

Beam width B	COV	JVR@1	JVR@3	JVR@5	COV-JVR@1
1	0.624	0.889	0.889	0.889	0.555
3	0.722	0.879	0.890	0.890	0.634
5	0.754	0.876	0.887	0.887	0.661
10	0.841	0.877	0.886	0.886	0.738
20	0.846	0.885	0.894	0.894	0.748

ing objective. Table 15 reports COV, JVR@1, and COV-JVR@1 on the full inference-only test split with all other hyperparameters fixed to their default values (beam width B=10, K=4).

Table 15: Sensitivity to the self-supervision weight λ_{ssl} on the full inference-only test split. Beam width and maximum entry magnitude are fixed to B=10 and K=4, respectively. The default setting used in the main experiments is highlighted in grey.

λ_{ssl}	COV↑	JVR@1↑	COV-JVR@1↑
0.1	0.717	0.875	0.628
0.3	0.620	0.872	0.541
0.5	0.841	0.877	0.738
0.7	0.822	0.856	0.704
1.0	0.618	0.835	0.516

As shown in Table 15, coverage (COV) is sensitive to λ_{ssl}: too small a weight (e.g., λ_{ssl}=0.1) provides insufficient self-supervised signal for the model to generalise across the diverse distribution of reactions encountered at test time, resulting in reduced coverage. Conversely, an excessively large weight (e.g., λ_{ssl}=1.0) over-emphasises reconstruction at the expense of the supervised completion

Table 16: Sensitivity to the maximum ER-entry magnitude K on the full inference-only test split, with beam width fixed to $B = 10$.

Maximum magnitude K	COV	JVR@1	JVR@3	JVR@5	COV-JVR@1
2	0.608	0.900	0.900	0.900	0.547
4	0.841	0.877	0.886	0.886	0.738
6	0.720	0.862	0.873	0.873	0.620

objective, which also degrades coverage. Joint validity (JVR@1) remains relatively stable across all settings, consistent with our earlier finding that validity is primarily governed by the electron-budget constraint rather than by training supervision. We therefore adopt λ_{ssl}=0.5 as the default, which achieves the best COV-JVR@1 trade-off.

Maximum-magnitude sweep. We then vary the maximum ER-entry magnitude K while fixing the beam width to $B = 10$. The row with $K = 4$ corresponds to the default setting used for the main **RxnComp** results.

Training curve. Figure 5 shows the supervised completion loss (`train_loss`) and the self-supervised reconstruction loss (`mask_loss`) over four training epochs. Both losses decrease rapidly in the first two epochs and continue to converge steadily thereafter, confirming stable optimisation under the joint training objective in Eq. (3).

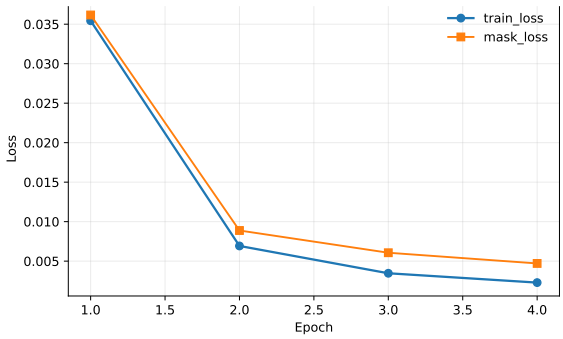

Figure 5: Training and mask loss curves over four epochs on the mixed training set. `train_loss` denotes the supervised ER-completion loss; `mask_loss` denotes the self-supervised reconstruction loss.

Compute resources. All experiments were conducted on a single NVIDIA RTX 5090 GPU (32 GB VRAM) with a 25-core Intel Xeon Platinum 8470Q CPU and 90 GB system memory, running PyTorch 2.7.0 with CUDA 12.8. Training **RxnComp** for four epochs on the mixed training set of 33,946 samples takes approximately **2** hours on this hardware configuration.

Discussion. Increasing the beam width mainly improves coverage while keeping JVR relatively stable, indicating that a wider beam helps the FSM decoder discover more feasible completions without substantially weakening validity. The main experiments use $B = 10$ as a practical trade-off between decoding quality and computational cost. The magnitude sweep shows that $K = 4$ provides the best coverage–validity trade-off among the tested value ranges: $K = 2$ is overly restrictive and reduces coverage, whereas $K = 6$ enlarges the search space and weakens both coverage and validity.

A.6 Limitations

The main limitation of this work lies in the evaluation protocol. Although we design a set of comprehensive proxy metrics, expert assessment remains the most reliable standard. Nevertheless, proxy metrics are essential, as the model must first be able to generate chemically feasible byproducts at scale before expert validation becomes practical. In addition, further improving coverage remains

challenging for **RxnComp**. The constrained generation process can be viewed as a combinatorial search over the ER matrix space, where feasible solutions may be difficult to identify for complex reactions. Addressing these challenges represents an important direction for future work.

A.7 BROADER IMPACT

This work proposes an electron-constrained generative framework for chemical reaction completion, aiming to improve the completeness and physical consistency of reaction data. It has the potential to benefit chemical discovery, data curation, and downstream tasks such as reaction prediction and synthesis planning by providing more reliable and interpretable reaction representations. However, as with other generative models in chemistry, there are potential risks of misuse, such as assisting in the design or completion of harmful or hazardous reactions. In addition, reliance on heuristic data construction may introduce biases that affect model reliability. We emphasize that this work is intended for research purposes. Future directions include incorporating safety-aware constraints and improving data quality to mitigate potential risks and ensure responsible use.