

TRACE: Transformation-Aware Graph Refinement for Reaction Condition Prediction

Yujie Chen^{1,2}, Tengfei Ma^{1,2}, Yuansheng Liu^{1,2}, Leyi Wei³, Shu Wu⁴, Dongsheng Cao⁵,
Yiping Liu^{1,2*}, Xiangxiang Zeng^{1,2}

¹State Key Lab. of Chemo & Biosensing, Coll. of Comp. Sci. & Electron. Eng., Hunan University, China

²MOE Key Lab. of Fusion Comput. of Supercomput. & Artif. Intell., Hunan University, China

³Centre for Artificial Intelligence driven Drug Discovery, Faculty of Applied Science, Macao Polytechnic University, China

⁴NLPR, MAIS, Institute of Automation, Chinese Academy of Sciences, China

⁵Xiangya School of Pharmaceutical Sciences, Central South University, China

yjchen@hnu.edu.cn, yipingliu@gmail.com

Abstract

Identifying suitable reaction conditions is critical for chemical synthesis, as they directly affect yield, selectivity, and transformation feasibility. While recent methods have shown promising results, most approaches either encode reactants and products independently or rely on rule-based reaction graphs, both of which constrain the ability of the model to capture condition-relevant structural transformations. In this work, we propose TRACE, a transformation-aware graph refinement framework for reaction condition prediction. TRACE constructs atom-level joint graphs that integrate both reactant and product structures to represent condition-relevant transformations. A structure-aware encoder enriches atom features with local chemical context, followed by a dynamic interaction refinement module that adaptively infers task-specific edges. To further guide the model toward condition-relevant patterns, a mechanism-regularized graph encoder incorporates reaction center information, enabling more accurate modeling of transformation mechanisms. Experiments on benchmark datasets show that TRACE achieves state-of-the-art performance across multiple condition types. The integration of transformation-aware refinement leads to improvements in prediction accuracy and generalization, while maintaining robust performance in challenging and realistic synthesis planning scenarios.

Code — <https://github.com/chenyujie1127/TRACE>

Introduction

Selecting appropriate reaction conditions, such as catalysts, solvents, and reagents, is critical for determining the feasibility (Coley et al. 2019), yield (Schwaller et al. 2021b), and selectivity (Ahneman et al. 2018) of reactions. With the rise of automated synthesis platforms and computer-assisted synthesis planning (CASP) systems (Segler, Preuss, and Waller 2018), condition prediction has become a central task for autonomous experimentation and accelerated reaction optimization (Szymanski et al. 2023).

Traditional methods for predicting reaction conditions primarily relied on expert intuition (Walker et al. 2019),

*Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

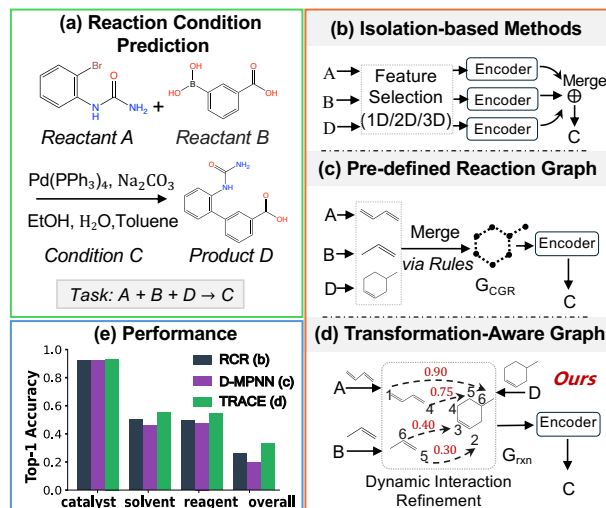


Figure 1: (a) Reaction condition prediction task. (b–d) Overview of modeling paradigms. (e) TRACE outperforms baselines in (b) and (c) across all condition types.

handcrafted heuristics (Marcou et al. 2015; Angello et al. 2022), and quantum chemical calculations (Struebing et al. 2013). Although chemically insightful, these methods suffered from subjectivity, high computational cost, and limited scalability across diverse reaction types. Recent advances in molecular representation learning (Gilmer et al. 2017; Chithrananda, Grand, and Ramsundar 2020; Rong et al. 2020; Thölke and De Fabritiis 2022) and access to large reaction corpora (Lowe 2012) have enabled data-driven approaches to become the dominant paradigm.

Early data-driven models for reaction condition prediction represent reactants and products independently using fingerprints (Gao et al. 2018; Chen and Li 2024), SMILES (Simplified Molecular Input Line Entry System) string (Nair, Schwaller, and Laino 2019; Wang et al. 2023), or molecular graphs (Ryou et al. 2020; Maser et al. 2021; Kwon et al. 2022; Zhang et al. 2022). As shown in Figure 1b, this separate encoding of reactants and products limits the ability of the model to capture cross-molecular interactions

and structural transformations essential to chemical reactivity. To capture molecular structural transformations, predefined reaction graphs such as the Condensed Graph of Reaction (CGR) (Heid and Green 2021; Wang et al. 2025) encode bond-level changes into unified molecular structures. These structures (Figure 1c) rely on atom mapping (Schwaller et al. 2021a) and fixed handcrafted rules, which prevent adaptive learning of condition-relevant interactions and obscure the distinct structural contexts of reactants and products. This motivates us design a transformation-aware framework that captures condition-driven structural changes for structure-informed condition prediction (Figure 1d).

To this end, we propose **TRACE**, a transformation-aware graph refinement framework specifically designed for reaction condition prediction. TRACE explicitly models chemical transformations by jointly encoding reactants and products into a unified interaction graph enriched with condition-relevant structural information. To better capture atom-level structural transformations, we augment the *Structure-Aware Encoder* with chemically informative features that reflect local environments relevant to reactivity (e.g., lone pairs, hybridization). Subsequently, a *Dynamic Interaction Refinement* module employs an information bottleneck to retain condition-relevant inter-molecular interactions while suppressing task-irrelevant ones. These induced graphs are further encoded by a *Mechanism-Regularized Graph Encoder*, which leverages reaction-center supervision to emphasize chemically reactive regions. The resulting representations are finally used by a *Joint Condition Predictor* to simultaneously predict suitable catalysts, solvents, and reagents. Extensive experiments validate the effectiveness (Figure 1e) and robustness of our approach across diverse scenarios. Our main contributions are as follows:

- We propose a transformation-aware graph refinement framework for reaction condition prediction, which jointly encodes reactants and products into interaction graphs and adaptively refines inter-molecular structures to capture condition-relevant transformations.
- We introduce an edge selection mechanism that adaptively infers task-relevant cross-molecular interactions, guided by the reactive center to better align with chemically meaningful regions.
- TRACE achieves state-of-the-art performance on benchmark datasets. Its transformation-aware graphs encode condition-relevant patterns that support generalization to temporal, low-resource, and real-world scenarios.

Related Work

Reaction Condition Prediction. Predicting reaction conditions has become a central challenge in computer-aided synthesis, due to the complex dependencies between molecular structure and reactivity. Early approaches relied on rule-based systems or quantum simulations, but were limited to specific reaction families, such as Menschutkin, Michael addition, or cross-coupling (Struebing et al. 2013; Marcou et al. 2015; Afonina et al. 2021). The availability of large-scale datasets like USPTO (Lowe 2012) enabled a shift toward data-driven approaches. Existing approaches

are commonly categorized by their reaction representations. **Descriptor-based methods.** These approaches represent reactions using fixed-size vectors derived from circular fingerprints (Gao et al. 2018; Walker et al. 2019; Chen and Li 2024) or CGR-based fragment counts (Afonina et al. 2021). While efficient, they rely on handcrafted features and ignore molecular topology, limiting their ability to capture structural changes and reaction-specific patterns. **Sequence-based methods.** SMILES representations are modeled as token sequences using language models. Recent work spans autoregressive generation (Jaume-Santero et al. 2023; Andronov et al. 2023; Wang et al. 2023), retrieval-augmented modeling (Qian et al. 2023), and multi-modal fusion from patent data (Zhang et al. 2024). However, these token-based models lack explicit structural encoding, making it difficult to capture structure-sensitive features critical to condition selection. **Graph-based methods.** Reactions are represented as molecular graphs to model structure. AR-GCN (Maser et al. 2021) processes each molecule independently and aggregates node features for multilabel classification. CIMG (Zhang et al. 2022) enriches molecular graphs with physicochemical priors (e.g., NMR shifts) to improve condition prediction. D-MPNN (Heid and Green 2021) fuses atom-mapped reactants and products into a superposition graph that captures changed bonds. Reacon (Wang et al. 2025) clusters reactions by expert-defined templates and trains separate D-MPNN models within each cluster for condition ranking. Atom-mapped models use fixed graph structures and are sensitive to mapping errors, while independent GNNs miss cross-reactant transformations critical for condition prediction. These limitations motivate the development of dynamic, transformation-aware graph refinement methods that can adaptively infer task-relevant reactivity patterns.

Preliminaries

Reaction Graphs. We represent a chemical reaction as a transformation from a set of reactant molecules $\mathcal{R} = \{R_1, R_2, \dots\}$ to a set of product molecules $\mathcal{P} = \{P_1, P_2, \dots\}$. Each molecule in chemical reaction ($\mathcal{R} \cup \mathcal{P}$) is represented as a graph $\mathcal{G}_i = (V_i, E_i)$, where V_i and E_i denotes the atoms and covalent bonds of the i -th molecule, respectively. We extract atom features (e.g., type, hybridization) for each $v \in V_i$, and bond features (e.g., type, aromaticity) for each $(u, v) \in E_i$. The full reaction is thus formulated as a pair of molecular graphs $(\mathcal{R}, \mathcal{P})$, which serves as the input to our model. Full details are in Appendix C.1.

Problem Definition. The reaction conditions are represented as a tuple $\mathbf{c} = \{c_t\}_{t=1}^5$, corresponding to five essential types: catalyst, two solvents, and two reagents. Given the molecular graphs of reactants and products $(\mathcal{R}, \mathcal{P})$, the goal is to predict the condition tuple \mathbf{c} based on reaction-specific structural features, as illustrated in Figure 1a. These condition types are often chemically interdependent (e.g., solvent polarity affects reagent solubility), so we model this as a joint multi-label classification task to ensure compatibility across condition types. The main challenge lies in capturing how structural transformations between reactants and products influence the selection of each condition component.

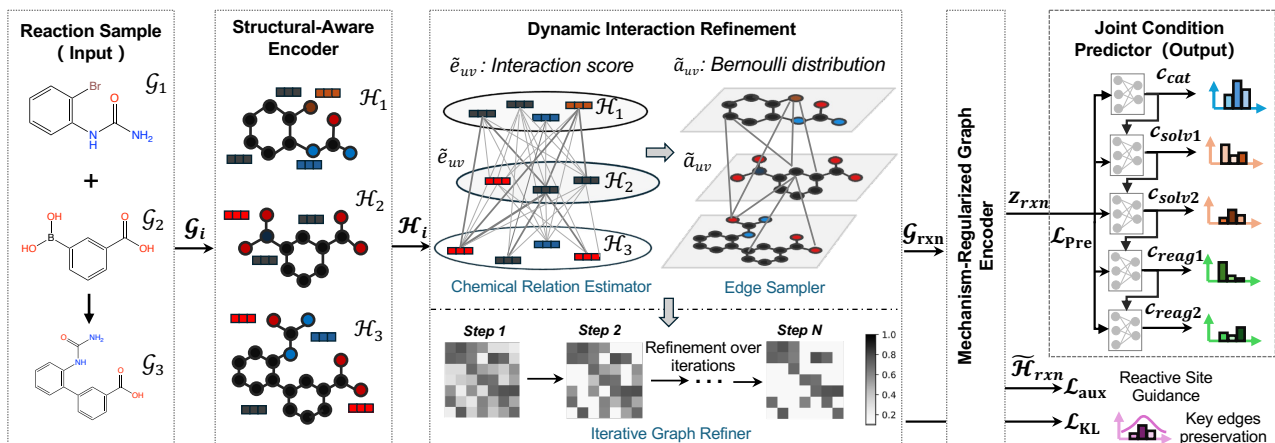


Figure 2: Overview of the TRACE framework. Given a reaction sample, atom-level representations \mathcal{H}_i are extracted via a structure-aware encoder and fed into the Dynamic Interaction Refinement module, which estimates interaction scores, samples a sparse graph \mathcal{G}_{rxn} , and iteratively refines it to retain transformation-relevant edges. The refined graph is encoded into \mathbf{z}_{rxn} via mechanism-regularized learning, guided by auxiliary losses (\mathcal{L}_{aux} , \mathcal{L}_{KL}), and used to jointly predict suitable reaction conditions.

The Proposed TRACE

This section introduces **TRACE**, a transformation-aware graph refinement framework for reaction condition prediction. As shown in Figure 2, TRACE first encodes reactants and products using a **Structure-Aware Encoder** to extract chemically relevant representations. A **Dynamic Interaction Refinement** module then constructs a sparse, transformation-relevant reaction graph through three components: a **Chemical Relation Estimator** to compute cross-molecular interaction scores, an **Edge Sampler** to sample candidate edges, and an **Iterative Graph Refiner** to enhance graph sparsity and chemical relevance. The resulting graph is processed by a **Mechanism-Regularized Encoder** guided by reactive centers. A **Joint Condition Predictor** finally performs cascaded multi-label prediction over catalysts, solvents, and reagents. Each component is described below.

Structure-Aware Encoder

To capture atom-level chemical environments and bonding patterns, we employ a Message Passing Neural Network (MPNN) (Gilmer et al. 2017) over each molecular graph \mathcal{G}_i , producing atom embeddings:

$$\mathcal{H}_i = \text{MPNN}(\mathcal{G}_i), \quad \forall \mathcal{G}_i \in \mathcal{R} \cup \mathcal{P} \quad (1)$$

where \mathcal{H}_i denotes the set of atom embeddings for molecule \mathcal{G}_i . These representations are passed to the interaction refinement module. See Appendix B.1 for implementation details.

Dynamic Interaction Refinement

To model transformation-specific dependencies between reactants and products, we introduce the **Dynamic Interaction Refinement (DIR)** module to construct a sparse, transformation-aware graph capturing cross-molecular reactivity. It first estimates atom-level interaction scores and samples transformation-relevant edges, then iteratively refines the graph by pruning irrelevant connections to enhance chemical relevance and better support condition prediction.

Chemical Relation Estimator. Given the atom embeddings $\{\mathcal{H}_i\}_{i=1}^{|\mathcal{R} \cup \mathcal{P}|}$ from the structure encoder, we assign two identifiers to each atom u : a reaction role $r_u \in \{\text{reactant}, \text{product}\}$, and a molecular index $\mathcal{I}(u) \in \{1, \dots, |\mathcal{R} \cup \mathcal{P}|\}$, where $|\mathcal{R} \cup \mathcal{P}|$ denotes the total number of molecular graphs involved in the reaction. To capture fine-grained reactivity across molecules, we introduce a multi-perspective attention mechanism to estimate pairwise interaction scores. For each perspective $p \in \{1, \dots, P\}$, the interaction score between atoms u and v is computed as:

$$e_{uv}^{(p)} = \mathbf{a}^\top \text{LeakyReLU} \left(\mathbf{W}^{(p)} \mathbf{h}_u \parallel \mathbf{W}^{(p)} \mathbf{h}_v \right), \quad (2)$$

where $\mathbf{h}_u, \mathbf{h}_v \in \mathbb{R}^d$ are atom embeddings from $\{\mathcal{H}_i\}_{i=1}^{|\mathcal{R} \cup \mathcal{P}|}$, $\mathbf{W}^{(p)} \in \mathbb{R}^{d' \times d}$ is a perspective-specific projection, $\mathbf{a} \in \mathbb{R}^{2d'}$ is a learnable scoring vector, and \parallel denotes concatenation. The final score is averaged over all perspectives: $e_{uv} = \left(\sum_{p=1}^P e_{uv}^{(p)} \right) / P$. To restrict attention to chemically meaningful pairs, we apply a *mechanism-aware role masking*:

$$m_{uv} = \begin{cases} 1, & \text{if } (r_u = \text{reactant} \wedge r_v = \text{product}) \\ & \vee (r_u = r_v = \text{reactant} \wedge \mathcal{I}(u) \neq \mathcal{I}(v)) \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

This mask focuses attention on cross-molecular interactions that are likely to reflect underlying reaction mechanisms. The masked interaction score is then computed as $\tilde{e}_{uv} = e_{uv} m_{uv}$, ensuring that only chemically meaningful atom pairs contribute to the induced interaction graph. We denote the masked interaction matrix as $\tilde{\mathbf{E}} = [\tilde{e}_{uv}] \in \mathbb{R}^{|V| \times |V|}$, where each entry reflects the transformation-specific interaction strength between atom pairs, and $V = \bigcup_i V_i$ represents the set of all atoms across reactants and products.

Edge Sampler. To construct a sparse, transformation-aware reaction graph \mathcal{G}_{rxn} , we perform probabilistic edge selection based on the masked interaction scores $\tilde{e}_{uv} \in [0, 1]$, which

reflect the transformation-specific relevance of atom pairs. To enable gradient-based optimization over discrete graph structures, we adopt the Gumbel-Sigmoid relaxation (Jang, Gu, and Poole 2017), where each score \tilde{e}_{uv} is used to parameterize the inclusion probability of edge (u, v) :

$$\hat{a}_{uv} = \sigma \left(\frac{1}{\tau} \left(\log \left(\frac{\tilde{e}_{uv}}{1 - \tilde{e}_{uv}} \right) + \log \left(\frac{\epsilon}{1 - \epsilon} \right) \right) \right), \quad (4)$$

where $\tau \in \mathbb{R}^+$ controls the relaxation sharpness, $\sigma(\cdot)$ is the sigmoid function, and $\epsilon \sim \mathcal{U}(0, 1)$ is sampled from a uniform distribution. As $\tau \geq 0$, the function is smoothly relaxed with a well-defined gradient $\frac{\partial \hat{a}_{uv}}{\partial \tilde{e}_{uv}}$, enabling end-to-end optimization of learnable reaction graph structure. This yields a relaxed adjacency matrix $\hat{A} = [\hat{a}_{uv}] \in \mathbb{R}^{|V| \times |V|}$, which serves as a soft, differentiable reaction graph for downstream refinement and condition prediction.

Iterative Graph Refiner. To promote sparsity and task-relevant structure in the learned graph, we impose an information bottleneck on the edge distribution. Inspired by the Generalized Information Bottleneck (GIB) framework (Wu et al. 2020), we treat the relaxed adjacency \hat{A} as a stochastic latent structure and apply Kullback–Leibler (KL) regularization (Hershey and Olsen 2007) to constrain its complexity:

$$\mathcal{L}_{\text{KL}} = \sum_{u,v} \text{D}_{\text{KL}}(\hat{a}_{uv} \parallel \text{Bernoulli}(\pi)), \quad (5)$$

where $\pi \in (0, 1)$ controls the expected edge density. Unlike prior approaches using symmetric priors (e.g., $\pi = 0.5$) for general compression, we adopt a sparse prior ($\pi = 0.1$) to encourage selective, transformation-aware connectivity. This design aligns with chemical intuition that only a limited subset of atom pairs participate in structural transformation, guiding the model toward reactivity-relevant topologies. We define the refined reaction graph $\mathcal{G}_{\text{rxn}} = (V, E_{\text{rxn}})$, with edge set $E_{\text{rxn}} = E_{\text{bond}} \cup E_{\text{learn}}$, where $E_{\text{bond}} = \bigcup_i E_i$ and $E_{\text{learn}} = \{(u, v) \mid \hat{a}_{uv} > 0\}$, with \hat{a}_{uv} defined in Eq. (4).

Mechanism-Regularized Graph Encoder

To promote mechanistic consistency, we regularize the graph encoder with reaction-center supervision, encouraging the model to focus on chemically reactive regions.

Reaction Graph Representation. Each node u from \mathcal{G}_{rxn} is initialized with $\tilde{\mathbf{h}}_u^{(0)} \in \{\mathcal{H}_i\}_{i=1}^{|\mathcal{R} \cup \mathcal{P}|}$, and updated through L layers of edge-weighted message passing:

$$\tilde{\mathbf{h}}_u^{(l+1)} = \sigma \left(\sum_{v \in \mathcal{N}(u)} \frac{\hat{a}_{uv}}{c_{uv}} \tilde{\mathbf{h}}_v^{(l)} \mathbf{W}^{(l)} \right), \quad l \in \{0, \dots, L-1\} \quad (6)$$

where $\mathcal{N}(u)$ denotes the neighbors of node u , $c_{uv} = \sqrt{|\mathcal{N}(u)| |\mathcal{N}(v)|}$ is a degree-based normalization term, and $\mathbf{W}^{(l)} \in \mathbb{R}^{d_l \times d_{l+1}}$ is a learnable projection matrix. The final atom representations, denoted as $\tilde{\mathbf{H}}_{\text{rxn}} = \{\tilde{\mathbf{h}}_u^{(L)} \mid u \in V\}$, are used for subsequent reaction center supervision. A permutation-invariant readout is applied over intermediate layers to derive the reaction-level embedding \mathbf{z}_{rxn} .

Reaction-center Guided Regularization. To guide the encoder toward chemically reactive regions, we incorporate an auxiliary loss supervised by annotated reaction centers. Each atom u and bond (u, v) is associated with a binary label $y_u, y_{uv} \in \{0, 1\}$, indicating its ground-truth reactivity. Two binary classifiers predict atom- and bond-level reactivity from the refined representations, with the loss defined as:

$$\mathcal{L}_{\text{aux}} = \lambda_{\text{bond}} [-y_{uv} \log \hat{y}_{uv} - (1 - y_{uv}) \log(1 - \hat{y}_{uv})] + \lambda_{\text{atom}} [-y_u \log \hat{y}_u - (1 - y_u) \log(1 - \hat{y}_u)], \quad (7)$$

where λ_{atom} and λ_{bond} are weighting coefficients. The predictions are obtained as $\hat{y}_u = \sigma(\mathbf{W}_{\text{atom}} \tilde{\mathbf{h}}_u^{(L)})$ for atoms and $\hat{y}_{uv} = \sigma(\mathbf{W}_{\text{bond}} [\tilde{\mathbf{h}}_u^{(L)} \parallel \tilde{\mathbf{h}}_v^{(L)}])$ for bonds, where \mathbf{W}_{atom} and \mathbf{W}_{bond} are learnable projection applied to the final refined atom representations $\tilde{\mathbf{H}}_{\text{rxn}}$.

Joint Condition Predictor

Accurately predicting reaction conditions requires modeling interdependencies among components (Gao et al. 2018; Wang et al. 2023). Inspired by these findings, we introduce a cascaded predictor that captures the joint distribution of condition types via sequential conditioning. At each step $t \in \{1, 2, 3, 4, 5\}$, corresponding to *catalyst*, *solvent*₁, *solvent*₂, *reagent*₁, and *reagent*₂, the prediction is conditioned on the reaction-level embedding \mathbf{z}_{rxn} and the embeddings of previously predicted labels $\hat{c}_{<t}$. The logits are computed as:

$$\hat{c}_t = \text{MLP}_t([\text{MLP}_{\text{proj}}(\mathbf{z}_{\text{rxn}}) \parallel \text{Embed}(\hat{c}_{<t})]) + \text{MLP}_t^{\text{aux}}(\text{MLP}_{\text{proj}}(\mathbf{z}_{\text{rxn}})) \quad (8)$$

where $\text{Embed}(\hat{c}_{<t})$ denotes the concatenated one-hot vectors of preceding labels, MLP_{proj} is a shared projection network, and MLP_t , $\text{MLP}_t^{\text{aux}}$ are the main and auxiliary classifiers for step t . The auxiliary head provides a fallback prediction independent of prior labels to enhance robustness.

Training and Decoding Strategy. We train the cascaded predictor with teacher forcing, where each condition is predicted based on the ground-truth of preceding ones. The objective is a weighted multi-task cross-entropy loss:

$$\mathcal{L}_{\text{pre}} = - \sum_{t=1}^5 \lambda_t \sum_{i=1}^{|\mathcal{V}_t|} \mathbb{1}_{[c_t=i]} \log \hat{c}_{t,i}, \quad (9)$$

where $\hat{c}_{t,i}$ is the predicted probability assigned to label i for type t , $\mathbb{1}_{[c_t=i]}$ indicates the ground-truth, and $|\mathcal{V}_t|$ denotes the number of candidate labels for type t . At inference, predictions are generated sequentially with beam search.

Overall Objective. The overall training objective integrates losses from three modules: (i) *Joint Condition Predictor* (\mathcal{L}_{pre}), (ii) *Reaction-center Guided Regularization* (\mathcal{L}_{aux}), and (iii) *Iterative Graph Refiner* (\mathcal{L}_{KL}). Formally,

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{pre}} + \lambda_{\text{aux}} \mathcal{L}_{\text{aux}} + \lambda_{\text{KL}} \mathcal{L}_{\text{KL}}, \quad (10)$$

where λ_{aux} and λ_{KL} are balancing coefficients for auxiliary supervision and structural regularization.

Theoretical Analysis

We formalize the structural regularization in Eq. (5) under the Information Bottleneck (IB) principle (Tishby, Pereira, and Bialek 2000). Specifically, we treat the sampled reaction graph \hat{A} as a stochastic latent variable inferred from the masked interaction matrix $\tilde{\mathbf{E}}$, aiming to preserve predictive signals while suppressing irrelevant structure. Suppose $\tilde{\mathbf{E}} = \tilde{\mathbf{E}}_{\text{rel}} + \tilde{\mathbf{E}}_{\text{noise}}$, where $\tilde{\mathbf{E}}_{\text{rel}}$ contains task-relevant information. Under standard assumptions (Achille and Soatto 2018), the dependency on noise satisfies:

$$I(\hat{A}; \tilde{\mathbf{E}}_{\text{noise}}) \leq I(\hat{A}; \tilde{\mathbf{E}}) - I(\hat{A}; \mathbf{c}). \quad (11)$$

Minimizing the IB objective thus reduces the influence of task-irrelevant structure, promoting sparse, informative interactions for condition prediction.

Definition 1 (Information Bottleneck for Reaction Graphs). Under the IB framework, we define the objective:

$$\min_{q(\hat{A}|\tilde{\mathbf{E}})} -I(\hat{A}; \mathbf{c}) + \beta I(\hat{A}; \tilde{\mathbf{E}}), \quad (12)$$

where $\beta > 0$ balances prediction and compression.

Variational Bound. Following (Alemi et al. 2017), we upper bound the second term via KL divergence:

$$I(\hat{A}; \tilde{\mathbf{E}}) \leq \mathbb{E}_{\tilde{\mathbf{E}}} \left[D_{\text{KL}} \left(q(\hat{A} | \tilde{\mathbf{E}}) \| p(\hat{A}) \right) \right], \quad (13)$$

where $p(\hat{A}) = \prod_{(u,v)} \text{Bernoulli}(\pi)$ is a sparsity-inducing prior that assumes edge-wise independence. A small value of π encourages compact graphs by reflecting the intuition that only a limited subset of atom pairs participates in the reaction. We adopt Gumbel-Sigmoid relaxation to enable gradient-based sampling, i.e., $\hat{a}_{uv} \sim q(\hat{a}_{uv} | \tilde{e}_{uv})$. The detailed proof refers to Appendix B.2.

Experiments

We explore the following research questions: (1) **RQ1:** Can TRACE accurately predict reaction conditions? (2) **RQ2:** Does TRACE generalize to diverse and challenging reaction scenarios? (3) **RQ3:** Are the learned interactions chemically meaningful and practically useful?

Experimental Settings

Dataset. We evaluate our method on the publicly available USPTO-Condition benchmark (Wang et al. 2023), which comprises 680,741 patent-derived reactions annotated with five condition components: catalyst, two solvents, and two reagents. Each reaction is converted from SMILES to molecular graphs using RDKit (Landrum 2013), which are then used as inputs to our model. Reaction centers for supervision are heuristically extracted from atom-mapped reactions, with atom mappings obtained via RXNMapper (Schwaller et al. 2021a) and post-processed following GLN (Dai et al. 2019). To assess both in-distribution and temporally shifted generalization, we adopt two data splits: (1) a random split with 80% for training, 10% for validation, and 10% for testing, and (2) a temporal split with reactions from 1976–2014 for training, 2015 for validation, and 2016 for testing. The latter simulates temporal shift, aligning with real-world divergence between future and past reactions.

Method	Top-1 (↑)	Top-3 (↑)	Top-5 (↑)	Top-10 (↑)
RCR	25.96	37.71	42.06	46.12
Parrot-LM-E	26.91	40.35	<u>45.10</u>	49.14
AR-GCN	14.60	23.74	27.33	31.21
CIMG-Condition	18.39	27.14	30.26	33.91
D-MPNN	19.85	29.91	33.29	36.89
Reacon	<u>27.52</u>	<u>43.90</u>	–	58.63
TRACE (Ours)	33.13 (0.01)	45.59 (0.02)	49.61 (0.02)	<u>53.15</u> (0.02)

Table 1: Top- k accuracy (%) for full-condition prediction on USPTO-Condition under *random split*. Best results in bold, second-best underlined. “–” denotes unreported values. Standard deviations in parentheses.

Baselines. We compare TRACE with three categories of baselines: (1) descriptor- and sequence-based models: RCR (Gao et al. 2018), Parrot-LM-E (Wang et al. 2023); (2) isolation-based GNNs without reaction-level modeling: AR-GCN (Maser et al. 2021), CIMG-Condition (Zhang et al. 2022); (3) predefined reaction graph methods using atom-mapped CGRs: D-MPNN (Heid and Green 2021), Reacon (Wang et al. 2025).

Implementation Details. TRACE is trained with Adam optimizer (learning rate $1e-3$, weight decay $5e-5$, batch size 64) for up to 200 epochs with early stopping on validation loss. The learning rate is adjusted using a ReduceLROnPlateau scheduler (decay factor 0.7, patience 2, minimum learning rate $1e-6$). The KL regularization coefficient λ_{KL} is set to 0.1, and the auxiliary weight λ_{aux} to 0.5. All task-specific weights, including λ_t (condition types), λ_{atom} , and λ_{bond} , are set to 1.0. For sparsity prior, the Bernoulli parameter is set to $\pi = 0.1$, encouraging the model to learn compact and task-relevant interaction graphs. At inference, beam search uses widths $\{1, 3, 1, 5, 1\}$ for each condition type. All experiments run on a single NVIDIA RTX 3090 GPU. Further details are provided in Appendix C.1 and C.2.

Comparison with Baselines (RQ1)

We evaluate the performance of TRACE against a range of representative baselines across Top- k settings, considering both overall and component-wise condition prediction (Tables 1 and 2). TRACE achieves consistently strong performance, demonstrating the effectiveness of transformation-aware graph refinement for condition prediction. We highlight three key observations: (1) **TRACE** achieves the best Top-1 accuracy (33.13%) on full-condition prediction, outperforming the strongest baseline (Reacon) by 5.61%. The gains are particularly notable in solvent and reagent prediction, where condition spaces are large and structurally diverse, highlighting the ability of TRACE to model fine-grained, context-sensitive interactions. (2) Models like D-MPNN and Reacon outperform AR-GCN and CIMG, likely due to their use of atom-mapped bond-change graphs that explicitly capture structural transformations between reactants and products. (3) Our dynamic interaction modeling offers improved adaptability and precision over rule-based approaches (e.g. D-MPNN, Reacon) by learning condition-relevant structures without predefined templates. Notably,

Type	Method	Top-1 (↑)	Top-3 (↑)	Top-5 (↑)	Top-10 (↑)
Catalyst	RCR	92.19	92.19	92.19	92.19
	Parrot-LM-E	<u>92.50</u>	<u>92.50</u>	<u>92.50</u>	<u>92.50</u>
	AR-GCN	90.24	90.24	90.24	90.24
	CIMG-Condition	91.46	91.46	91.46	91.46
	D-MPNN	90.18	90.18	90.18	90.18
	Reacon	92.44	–	–	–
	TRACE (Ours)	93.27 _(0.01)	93.27 _(0.01)	93.27 _(0.01)	93.27 _(0.01)
Solvent1	RCR	50.15	66.40	70.55	73.40
	Parrot-LM-E	50.18	68.58	73.11	<u>75.36</u>
	AR-GCN	41.14	57.87	62.95	66.35
	CIMG-Condition	42.18	61.39	65.42	67.80
	D-MPNN	49.33	66.50	71.18	74.46
	Reacon	<u>50.39</u>	–	–	–
	TRACE (Ours)	55.16 _(0.04)	71.62 _(0.02)	75.40 _(0.01)	77.36 _(0.01)
Solvent2	RCR	81.30	83.69	84.61	85.25
	Parrot-LM-E	80.96	<u>84.26</u>	<u>85.21</u>	<u>85.85</u>
	AR-GCN	80.93	80.93	80.93	80.93
	CIMG-Condition	81.10	81.10	81.10	81.10
	D-MPNN	81.31	81.31	81.31	81.31
	Reacon	<u>81.58</u>	–	–	–
	TRACE (Ours)	81.82 _(0.02)	85.76 _(0.01)	86.65 _(0.01)	87.15 _(0.01)
Reagent1	RCR	49.72	65.97	74.02	81.84
	Parrot-LM-E	<u>50.39</u>	<u>68.20</u>	<u>76.29</u>	<u>84.36</u>
	AR-GCN	42.00	57.40	66.67	75.15
	CIMG-Condition	43.51	56.85	66.65	74.62
	D-MPNN	48.72	64.03	72.16	79.00
	Reacon	50.02	–	–	–
	TRACE (Ours)	54.60 _(0.05)	71.31 _(0.01)	78.54 _(0.02)	85.60 _(0.02)
Reagent2	RCR	76.22	84.08	86.64	88.75
	Parrot-LM-E	76.48	<u>84.86</u>	<u>87.74</u>	<u>89.98</u>
	AR-GCN	74.86	74.86	74.86	74.86
	CIMG-Condition	75.74	75.74	75.74	75.74
	D-MPNN	76.21	76.21	76.21	76.21
	Reacon	<u>77.95</u>	–	–	–
	TRACE (Ours)	77.96 _(0.03)	86.58 _(0.02)	89.12 _(0.01)	91.13 _(0.0)

Table 2: Top- k accuracy (%) for component-wise condition prediction on the USPTO-Condition dataset under the *random split* setting. “–” denotes unreported values. Standard deviations in parentheses.

the high Top-10 accuracy (58.63%) of Reacon reflects effective template ranking within covered reactions, but its reliance on human-crafted priors may hinder generalization beyond covered reactions. These results demonstrate that modeling cross-molecular transformations through dynamic graph refinement yields consistent improvements in condition prediction accuracy, particularly for challenging tasks.

Ablation Study (RQ1)

To assess the role of each component in **TRACE**, we perform ablations by removing the Dynamic Interaction Refinement module (**w/o DIR**), the Reaction-center guided regularization (**w/o RC**), and the Iterative Graph Refiner (**w/o IGR**). Table 3 summarizes the Top- k accuracy across the full-condition prediction task. Removing the DIR leads to the most substantial degradation, with Top-1 accuracy dropping from 33.13% to 27.19%, highlighting its critical role in capturing transformation-aware, inter-molecular interactions. Without the IGR, Top-1 accuracy drops by 3.11%, showing the role of iterative refinement in suppressing irrelevant edges. Excluding RC primarily affects Top-3 and Top-5 metrics, underscoring its utility in guiding the model toward chemically plausible regions during candidate ranking. These findings validate the effectiveness of each design

Top- k	w/o DIR	w/o IGR	w/o RC	TRACE
Top-1	27.19	30.02	30.33	33.13
Top-3	40.44	42.85	43.24	45.59
Top-5	45.16	46.83	47.58	49.61
Top-10	49.23	50.42	51.37	53.15

Table 3: Top- k accuracy (%) of different variants of TRACE on overall condition prediction (random split).

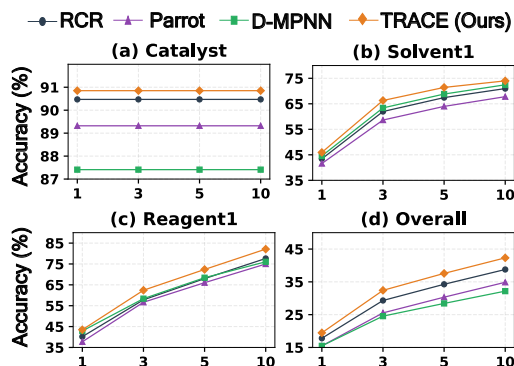


Figure 3: Top- k accuracy for condition types (time-split).

component in modeling task-specific interaction patterns.

Generalization Ability of TRACE (RQ2)

Time-Split Evaluation. We design a time-split evaluation protocol to assess the ability of the model to generalize to future, previously unseen reactions. As shown in Figure 3, TRACE consistently outperforms all baselines across three key condition components and overall prediction. It achieves 19.46% Top-1 accuracy in full-condition prediction, compared to 17.79% (RCR), 15.48% (Parrot), and 15.48% (D-MPNN). Performance gains are particularly evident for *solvent*₁ and *reagent*₁, which exhibit greater structural diversity. TRACE reaches 45.94% and 43.48% Top-1 accuracy on these components, outperforming D-MPNN (44.56%, 42.91%), RCR (43.42%, 40.17%) and Parrot (41.58%, 37.71%). Parrot shows reduced robustness in this setting, possibly due to its reliance on token-level statistics that generalize poorly under temporal shifts. By contrast, TRACE benefits from transformation-aware graph refinement, providing stronger inductive bias for modeling condition-relevant chemical transformations.

Few-shot Evaluation. We design a few-shot evaluation protocol that simulates settings with limited training examples per reaction class. To avoid data leakage from ambiguous or undefined reaction types, we exclude reactions annotated with unknown class labels (0.0). For classes 7–11, we keep 10% of reactions for training and evaluate on the corresponding test sets to simulate low-resource scenarios. Figure 5 presents Top-1 and Top-10 accuracies across rare reaction classes. TRACE consistently outperforms all baselines, achieving the highest overall Top-1 and Top-10 performance at 19.43% and 40.37%, respectively. On class




Figure 4: Prediction of reaction conditions for a real synthesis route (MK-8189). Ground truth conditions are shown above arrows (black). Model predictions are marked as correct (green), reasonable alternative (blue), or incorrect (red).




Figure 5: Few-shot results (%) on classes 7–11 and overall.




Figure 6: Graph disruption analysis with Top-1 and Top-10 accuracy under four perturbation settings.

8, it attains 30.49% Top-1 accuracy, surpassing D-MPNN (22.45%), Parrot (15.69%), and RCR (2.25%). A similar trend is observed for Top-10 accuracy. These results validate the generalization of transformation-aware modeling in low-resource settings.

Robustness of Learned Reaction Graph (RQ3)

Graph Disruption Analysis. To assess the functional relevance of the learned interaction graphs, we conduct a perturbation study under four inference-time settings: (1) **Full** (original TRACE), (2) **Top-k** (removing the top $k\%$ high-confidence edges), (3) **Random** (removing $k\%$ edges at random), and (4) **No Learned Edges** (removing all learned cross-molecular edges). As shown in Figure 6, removing the top 20% high-confidence edges leads to a substantial performance drop, with Top-1 accuracy on *reagent*₁ and *solvent*₁ reduced by 18.13% and 21.85%, respectively. In contrast, randomly masking the same proportion of edges causes minimal degradation. Discarding all learned edges results in the lowest performance across components, highlighting the

utility of the refined graph structure. We observe consistent trends across masking ratios from 10% to 50%, indicating that TRACE depends on a compact and informative set of edges to support condition prediction. This demonstrates its ability to construct compact interaction graphs that preserve key predictive signals. These findings are further supported by attention visualizations in Appendix D.3, which show that TRACE consistently attends to atom-mapped regions and newly formed bonds, such as C–O, C–N, and S–N.

Evaluation on Real-World Synthesis Routes. To assess TRACE in realistic settings, we adopt the evaluation setup from Reacon (Wang et al. 2025), using 100 reactions extracted from the published synthesis of 12 drug molecules in the Journal of Medicinal Chemistry. This selection presents practical challenges, including rare reagents, lab-specific practices, and partial vocabulary coverage (78%). TRACE achieves a Top-10 accuracy of 41%, frequently recovering ground-truth or functionally equivalent conditions. Figure 4 illustrates a representative example of MK-8189 (Layton et al. 2023), which involves diverse transformation types. TRACE correctly predicts steps such as *toluene*, and $\{NaH, THF\}$, and identifies a reasonable solvent pair $\{TEA, n-BuOH\}$ ranked within the top-4. One step deviates from the reported $Pd(dppf)_2Cl_2$ system, potentially due to limited condition coverage. This evaluation highlights the robustness of TRACE in handling multi-step synthesis scenarios with real-world complexity. Additional analysis and visualizations are provided in Appendix D.4.

Conclusion

We present **TRACE**, a transformation-aware graph refinement framework for reaction condition prediction. By jointly encoding reactants and products, TRACE constructs condition-specific interaction graphs via adaptive edge selection, while reactivity supervision guides the model toward chemically meaningful transformations. This design enables learning reactivity-relevant structures beyond static topologies or pre-defined rules. Experiments across multiple benchmarks demonstrate state-of-the-art performance, with strong generalization under temporal shift, few-shot, and real-synthesis settings. These findings highlight the scalability of TRACE and open opportunities for applying transformation-aware modeling to broader synthesis tasks.

Acknowledgments

The work was supported by the National Science and Technology Major Project (Grant No. 2023ZD0120902), the National Natural Science Foundation of China (Grant Nos. 62522110, 62472152, 62425204, U22A2037, 62450002, 62432011, and 62372159), Hunan Provincial Natural Science Foundation of China (Grant No. 2024JJ4015), and the Science and Technology Development Fund of MACAU (no. 0133/2024/RIB2).

References

- Achille, A.; and Soatto, S. 2018. Information dropout: Learning optimal representations through noisy computation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12): 2897–2905.
- Afonina, V. A.; Mazitov, D. A.; Nurmukhametova, A.; Shevelev, M. D.; Khasanova, D. A.; Nugmanov, R. I.; Buriolov, V. A.; Madzhidov, T. I.; and Varnek, A. 2021. Prediction of optimal conditions of hydrogenation reaction using the likelihood ranking approach. *International Journal of Molecular Sciences*, 23(1): 248.
- Ahneiman, D. T.; Estrada, J. G.; Lin, S.; Dreher, S. D.; and Doyle, A. G. 2018. Predicting reaction performance in C–N cross-coupling using machine learning. *Science*, 360(6385): 186–190.
- Alemi, A. A.; Fischer, I.; Dillon, J. V.; and Murphy, K. 2017. Deep Variational Information Bottleneck. In *International Conference on Learning Representations*.
- Andronov, M.; Voinarovska, V.; Andronova, N.; Wand, M.; Clevert, D.-A.; and Schmidhuber, J. 2023. Reagent prediction with a molecular transformer improves reaction data quality. *Chemical Science*, 14(12): 3235–3246.
- Angello, N. H.; Rathore, V.; Beker, W.; Wołos, A.; Jira, E. R.; Roszak, R.; Wu, T. C.; Schroeder, C. M.; Aspuru-Guzik, A.; Grzybowski, B. A.; et al. 2022. Closed-loop optimization of general reaction conditions for heteroaryl Suzuki-Miyaura coupling. *Science*, 378(6618): 399–405.
- Chen, L.-Y.; and Li, Y.-P. 2024. Enhancing chemical synthesis: a two-stage deep neural network for predicting feasible reaction conditions. *Journal of Cheminformatics*, 16(1): 11.
- Chithrananda, S.; Grand, G.; and Ramsundar, B. 2020. ChemBERTa: Large-Scale Self-Supervised Pretraining for Molecular Property Prediction. *CoRR*.
- Coley, C. W.; Thomas III, D. A.; Lummiss, J. A.; Jaworski, J. N.; Breen, C. P.; Schultz, V.; Hart, T.; Fishman, J. S.; Rogers, L.; Gao, H.; et al. 2019. A robotic platform for flow synthesis of organic compounds informed by AI planning. *Science*, 365(6453): eaax1566.
- Dai, H.; Li, C.; Coley, C.; Dai, B.; and Song, L. 2019. Retrosynthesis prediction with conditional graph logic network. *Advances in Neural Information Processing Systems*, 32.
- Gao, H.; Struble, T. J.; Coley, C. W.; Wang, Y.; Green, W. H.; and Jensen, K. F. 2018. Using machine learning to predict suitable conditions for organic reactions. *ACS Central Science*, 4(11): 1465–1476.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, 1263–1272. PMLR.
- Heid, E.; and Green, W. H. 2021. Machine learning of reaction properties via learned representations of the condensed graph of reaction. *Journal of Chemical Information and Modeling*, 62(9): 2101–2110.
- Hershey, J. R.; and Olsen, P. A. 2007. Approximating the Kullback Leibler divergence between Gaussian mixture models. In *IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP'07*, volume 4, IV–317. IEEE.
- Jang, E.; Gu, S.; and Poole, B. 2017. Categorical Reparameterization with Gumbel-Softmax. In *International Conference on Learning Representations*.
- Jaume-Santero, F.; Bornet, A.; Valery, A.; Naderi, N.; Vicente Alvarez, D.; Proios, D.; Yazdani, A.; Bournez, C.; Fessard, T.; and Teodoro, D. 2023. Transformer performance for chemical reactions: analysis of different predictive and evaluation scenarios. *Journal of Chemical Information and Modeling*, 63(7): 1914–1924.
- Kwon, Y.; Kim, S.; Choi, Y.-S.; and Kang, S. 2022. Generative modeling to predict multiple suitable conditions for chemical reactions. *Journal of Chemical Information and Modeling*, 62(23): 5952–5960.
- Landrum, G. 2013. RDKit: Open-source cheminformatics software. *Online*.
- Layton, M. E.; Kern, J. C.; Hartingh, T. J.; Shipe, W. D.; Raheem, I.; Kandebo, M.; Hayes, R. P.; Huszar, S.; Eddins, D.; Ma, B.; et al. 2023. Discovery of MK-8189, a highly potent and selective PDE10A inhibitor for the treatment of schizophrenia. *Journal of Medicinal Chemistry*, 66(2): 1157–1171.
- Lowe, D. M. 2012. *Extraction of chemical structures and reactions from the literature*. Ph.D. thesis, University of Cambridge.
- Marcou, G.; Aires de Sousa, J.; Latino, D. A.; de Luca, A.; Horvath, D.; Rietsch, V.; and Varnek, A. 2015. Expert system for predicting reaction conditions: the Michael reaction case. *Journal of Chemical Information and Modeling*, 55(2): 239–250.
- Maser, M. R.; Cui, A. Y.; Ryou, S.; DeLano, T. J.; Yue, Y.; and Reisman, S. E. 2021. Multilabel classification models for the prediction of cross-coupling reaction conditions. *Journal of Chemical Information and Modeling*, 61(1): 156–166.
- Nair, V. H.; Schwaller, P.; and Laino, T. 2019. Data-driven chemical reaction prediction and retrosynthesis. *Chimia*, 73(12): 997–997.
- Qian, Y.; Li, Z.; Tu, Z.; Coley, C. W.; and Barzilay, R. 2023. Predictive Chemistry Augmented with Text Retrieval. In *Conference on Empirical Methods in Natural Language Processing*.
- Rong, Y.; Bian, Y.; Xu, T.; Xie, W.; Wei, Y.; Huang, W.; and Huang, J. 2020. Self-supervised graph transformer on

large-scale molecular data. *Advances in Neural Information Processing Systems*, 33: 12559–12571.

Ryou, S.; Maser, M. R.; Cui, A. Y.; DeLano, T. J.; Yue, Y.; and Reisman, S. E. 2020. Graph neural networks for the prediction of substrate-specific organic reaction conditions. *arXiv preprint arXiv:2007.04275*.

Schwaller, P.; Hoover, B.; Reymond, J.-L.; Strobel, H.; and Laino, T. 2021a. Extraction of organic chemistry grammar from unsupervised learning of chemical reactions. *Science Advances*, 7(15): eabe4166.

Schwaller, P.; Probst, D.; Vaucher, A. C.; Nair, V. H.; Kreutter, D.; Laino, T.; and Reymond, J.-L. 2021b. Mapping the space of chemical reactions using attention-based neural networks. *Nature Machine Intelligence*, 3(2): 144–152.

Segler, M. H.; Preuss, M.; and Waller, M. P. 2018. Planning chemical syntheses with deep neural networks and symbolic AI. *Nature*, 555(7698): 604–610.

Struebing, H.; Ganase, Z.; Karamertzanis, P. G.; Siougkrou, E.; Haycock, P.; Piccione, P. M.; Armstrong, A.; Galindo, A.; and Adjiman, C. S. 2013. Computer-aided molecular design of solvents for accelerated reaction kinetics. *Nature Chemistry*, 5(11): 952–957.

Szymanski, N. J.; Rendy, B.; Fei, Y.; Kumar, R. E.; He, T.; Milsted, D.; McDermott, M. J.; Gallant, M.; Cubuk, E. D.; Merchant, A.; et al. 2023. An autonomous laboratory for the accelerated synthesis of novel materials. *Nature*, 624(7990): 86–91.

Thölke, P.; and De Fabritiis, G. 2022. Equivariant Transformers for Neural Network based Molecular Potentials. In *International Conference on Learning Representations*.

Tishby, N.; Pereira, F. C.; and Bialek, W. 2000. The information bottleneck method. *arXiv preprint physics/0004057*.

Walker, E.; Kammeraad, J.; Goetz, J.; Robo, M. T.; Tewari, A.; and Zimmerman, P. M. 2019. Learning to predict reaction conditions: relationships between solvent, molecular structure, and catalyst. *Journal of Chemical Information and Modeling*, 59(9): 3645–3654.

Wang, X.; Hsieh, C.-Y.; Yin, X.; Wang, J.; Li, Y.; Deng, Y.; Jiang, D.; Wu, Z.; Du, H.; Chen, H.; et al. 2023. Generic interpretable reaction condition predictions with open reaction condition datasets and unsupervised learning of reaction center. *Research*, 6: 0231.

Wang, Z.; Lin, K.; Pei, J.; and Lai, L. 2025. Reacon: a template-and cluster-based framework for reaction condition prediction. *Chemical Science*, 16(2): 854–866.

Wu, T.; Ren, H.; Li, P.; and Leskovec, J. 2020. Graph information bottleneck. *Advances in Neural Information Processing Systems*, 33: 20437–20448.

Zhang, B.; Zhang, X.; Du, W.; Song, Z.; Zhang, G.; Zhang, G.; Wang, Y.; Chen, X.; Jiang, J.; and Luo, Y. 2022. Chemistry-informed molecular graph as reaction descriptor for machine-learned retrosynthesis planning. *Proceedings of the National Academy of Sciences*, 119(41): e2212711119.

Zhang, Y.; Yu, R.; Zeng, K.; Li, D.; Zhu, F.; Yang, X.; Jin, Y.; and Xu, Y. 2024. Text-augmented multimodal llms for chemical reaction condition recommendation. *arXiv preprint arXiv:2407.15141*.