Novelty Assessment Report

Paper: DrugTrail: Explainable Drug Discovery via Structured Reasoning and Druggability-Tailored Preference Optimization
PDF URL: https://openreview.net/pdf?id=1pAW0y8WLH
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2025-12-29

Abstract

Machine learning promises to revolutionize drug discovery, but its "black-box" nature and narrow focus limit adoption by experts. While Large Language Models (LLMs) offer a path forward with their broad knowledge and interactivity, existing methods remain data-intensive and lack transparent reasoning. To address these issues, we present DrugTrail, an LLM-based framework for explainable drug discovery that integrates structured reasoning trajectories with a Druggability-Tailored Preference Optimization (DTPO) strategy. It not only introduces structured reasoning traces to articulate the "how" and "why" behind its conclusions but also serve to guide task-specific reasoning pathways within the LLM's vast knowledge space, thereby enhancing its interpretability and reliability of its final outputs. Furthermore, based on the fact that optimizing for binding affinity alone does not equate to optimizing for druggability, DTPO explicitly moves beyond single-metric optimization and opens up a broader search space that balances affinity with other essential factors. Extensive experiments demonstrate the effectiveness of our approach and its generalizability to a wider range of biomolecular optimization domains, bridging the gap between LLM reasoning capabilities and trustworthy AI-assisted drug discovery.

Core Task Landscape

This paper addresses: **Explainable Drug Discovery via Structured Reasoning and Preference Optimization**
A total of **10 papers** were analyzed and organized into a taxonomy with **7 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Preference Learning and Optimization for Molecular Design**
- **Structured Reasoning and Explainability in Drug Discovery**
- **Reinforcement Learning and Causal Inference for Treatment Optimization**
- **Cross-Domain AI Methodologies and Reviews**

Complete Taxonomy Tree

- Explainable Drug Discovery via Structured Reasoning and Preference Optimization Survey Taxonomy
- Preference Learning and Optimization for Molecular Design
  - Human Chemist Preference Modeling (2 papers)
  - [2] Extracting medicinal chemistry intuition via preference machine learning (Oh-Hyeon Choung, 2023) View paper
  - [4] Towards Interpretable Models of Chemist Preferences for Human-in-the-Loop Assisted Drug Discovery (Yasmine Nahal, 2024) View paper
  - Druggability-Tailored Preference Optimization ★ (1 papers)
  - [0] DrugTrail: Explainable Drug Discovery via Structured Reasoning and Druggability-Tailored Preference Optimization (Anon et al., 2026) View paper
- Structured Reasoning and Explainability in Drug Discovery
  - Knowledge Graph-Based Reasoning (2 papers)
  - [5] Zero-shot drug repurposing with geometric deep learning and clinician centered design (K Huang, 2023) View paper
  - [6] Explainable Reasoning Path Inference of Anti-cancer Drug Sensitivity on Genomic Knowledge Graph via Macro-Micro Agent Collaborative Reinforcement Learning (Minhua Feng, 2025) View paper
  - LLM-Based Chemical Reasoning (2 papers)
  - [7] Chem-R: Learning to Reason as a Chemist (Wang, 2025) View paper
  - [8] Reasoning-Enhanced Large Language Models for Molecular Property Prediction (Jiaxi Zhuang, 2025) View paper
- Reinforcement Learning and Causal Inference for Treatment Optimization
  - Clinical Reasoning Stabilization via Preference Optimization (1 papers)
  - [3] Stabilizing Reasoning in Medical LLMs with Continued Pretraining and Reasoning Preference Optimization (Kawakami Wataru, 2025) View paper
  - Causal Treatment Policy Learning (1 papers)
  - [10] Causal inference and interpretable machine learning for personalised medicine (Parbhoo, 2019) View paper
- Cross-Domain AI Methodologies and Reviews (2 papers)
  - [1] Artificial intelligence in drug development (Kang Zhang, 2025) View paper
  - [9] The Role of Reinforcement Learning in Large Language Models: A Survey (Liu, 2025) View paper

Narrative

Core task: explainable drug discovery via structured reasoning and preference optimization. The field structure suggested by the taxonomy reflects a convergence of machine learning techniques tailored to molecular design and therapeutic decision-making. The top-level branches organize work into Preference Learning and Optimization for Molecular Design, which focuses on aligning generative models with chemist preferences and druggability criteria; Structured Reasoning and Explainability in Drug Discovery, which emphasizes

interpretable pathways and mechanistic insights; Reinforcement Learning and Causal Inference for Treatment Optimization, addressing personalized medicine and dynamic treatment regimes; and Cross-Domain AI Methodologies and Reviews, capturing broader methodological advances and survey perspectives. Representative works such as Chemist Preferences[4] and Preference Machine Learning[2] illustrate how preference-based frameworks guide molecular generation, while Medical LLM Reasoning[3] and Chem-R[7] exemplify efforts to inject structured reasoning into chemical and clinical contexts.

Particularly active lines of work explore the tension between generative flexibility and interpretability: some studies prioritize end-to-end optimization for druggability, while others emphasize transparent reasoning chains that domain experts can audit. Within this landscape, DrugTrail[0] sits squarely in the Druggability-Tailored Preference Optimization cluster, combining preference learning with structured explanations to guide molecule design. Its emphasis on both optimization and explainability distinguishes it from purely generative approaches like Chemist Preferences[4], which focus on preference alignment without explicit reasoning traces, and from reasoning-centric methods like Chem-R[7], which prioritize interpretability but may not directly optimize for druggability metrics. This positioning reflects an emerging consensus that effective drug discovery systems must balance predictive performance with the transparency required for regulatory and scientific validation.

Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

Taxonomy-Level Summary

Both subtopics address preference optimization in drug discovery, moving beyond simple binding affinity metrics to incorporate broader considerations. The key distinction lies in the source of preferences: Druggability-Tailored Preference Optimization focuses on automated, computational strategies that balance multiple druggability criteria (ADMET, selectivity, synthesizability), while Human Chemist Preference Modeling explicitly captures and learns from medicinal chemist expertise and intuition through human feedback mechanisms.

Similarities: - Both aim to optimize molecular candidates beyond single-objective metrics like binding affinity - Both recognize that successful drug discovery requires balancing multiple competing criteria - Both employ machine learning approaches to guide molecular optimization - Both address the multi-objective nature of lead optimization in drug discovery

Differences: - Druggability-Tailored uses automated computational criteria (ADMET, physicochemical properties) while Human Chemist Preference relies on expert human judgment - Human Chemist Preference Modeling requires human-in-the-loop feedback mechanisms, while Druggability-Tailored operates autonomously - Druggability-Tailored focuses on quantifiable druggability metrics, while Human Chemist Preference captures tacit knowledge and intuition - The former emphasizes structured, predefined criteria; the latter learns implicit preferences from chemist decisions

Suggested Search Directions: - Hybrid approaches combining automated druggability filters with human chemist feedback loops - Methods for translating tacit chemist knowledge into explicit computational druggability criteria - Comparative studies on alignment between computational druggability predictions and expert chemist preferences

Sibling Subtopics

- **Human Chemist Preference Modeling** (leaves: 1, papers: 2)
- Scope: Techniques capturing medicinal chemist intuition and preferences through machine learning for lead optimization.
- Exclude: Excludes automated preference optimization without human feedback; see Druggability-Tailored Preference Optimization.

Contributions Analysis

Overall novelty summary. The paper introduces DrugTrail, a framework combining structured reasoning trajectories with Druggability-Tailored Preference Optimization (DTPO) for explainable drug discovery. According to the taxonomy, it occupies the 'Druggability-Tailored Preference Optimization' leaf under 'Preference Learning and Optimization for Molecular Design'. Notably, this leaf contains no sibling papers—the original paper is the sole occupant. This suggests the specific combination of preference optimization explicitly balancing affinity with broader druggability criteria, rather than single-metric optimization, represents a relatively sparse research direction within the surveyed literature.

The taxonomy reveals neighboring work in 'Human Chemist Preference Modeling' (two papers capturing medicinal chemist intuition) and 'LLM-Based Chemical Reasoning' (two papers training language models to emulate chemist reasoning). The exclude notes clarify boundaries: the original paper's leaf excludes human-centered preference learning, while the reasoning subtopic excludes preference-based optimization without reasoning traces. DrugTrail appears to bridge these directions by integrating structured reasoning with preference optimization, positioning itself at the intersection of interpretability and multi-objective molecular design rather than purely within either neighboring cluster.

Among 21 candidates examined across three contributions, none were identified as clearly refuting the work. The DRUGTRAIL framework examined 10 candidates with zero refutable matches; the Clinical Chemistry-Informed Reasoning module similarly examined 10 with none refuting; DTPO examined only 1 candidate with no overlap. These statistics reflect a limited search scope—top-K semantic matches plus citation expansion—rather than exhaustive coverage. The absence of refutable prior work across all contributions suggests that, within this bounded search, the specific integration of structured reasoning with druggability-tailored preference optimization has not been directly addressed by the examined literature.

Given the limited search scope (21 candidates, not hundreds), the analysis indicates the work occupies a relatively unexplored niche combining preference optimization and structured reasoning for druggability. The taxonomy structure shows active neighboring areas but no direct siblings in the same leaf. While this suggests potential novelty, the small candidate pool and sparse taxonomy leaf mean the assessment is provisional—broader literature searches or domain expert review could reveal closer prior work not captured by semantic similarity or citation links in this analysis.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: DRUGTRAIL framework for interpretable drug discovery

Description: The authors introduce DRUGTRAIL, a novel framework that combines structured reasoning trajectories with a specialized optimization strategy to enable transparent and interpretable drug discovery using large language models. The framework addresses the black-box nature of existing methods by making the reasoning process explicit.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Effective and Explainable Molecular Property Prediction by Chain-of-Thought Enabled Large Language Models and Multi-Modal Molecular Information Fusion

URL: View paper

Brief Assessment

Chain-of-Thought Molecular[22] focuses on molecular property prediction using multimodal fusion (SMILES, graphs, text) with CoT for interpretability, not on structured reasoning trajectories for drug discovery with druggability-tailored optimization as in DRUGTRAIL.

2. Beyond Chemical QA: Evaluating LLM's Chemical Reasoning with Modular Chemical Operations

URL: View paper

Brief Assessment

Chemical Operations Evaluation[31] focuses on evaluating LLM reasoning through modular chemical operations (addition, deletion, substitution) for molecular optimization and reaction prediction, not on developing an interpretable drug discovery framework with structured reasoning trajectories and druggability-tailored optimization.

3. PharmAgents: Building a Virtual Pharma with Large Language Model Agents

URL: View paper

Brief Assessment

PharmAgents[29] focuses on multi-agent collaboration for simulating the full drug discovery workflow, while the original paper presents a structured reasoning trajectory framework with druggability-tailored optimization. These represent different technical approaches to interpretability in drug discovery.

4. Molreasoner: Toward effective and interpretable reasoning for molecular llms

URL: View paper

Brief Assessment

MolReasoner[28] focuses on molecule-text translation tasks using chain-of-thought reasoning for molecular language models, not on drug discovery with structured reasoning trajectories for binding pocket analysis and druggability optimization as in DRUGTRAIL.

5. Mol-LLaMA: Towards General Understanding of Molecules in Large Molecular Language Model

URL: View paper

Brief Assessment

Mol-LLaMA[27] focuses on general molecular understanding and analysis across chemistry and biology, not on structured reasoning trajectories for drug discovery or druggability-tailored optimization strategies.

6. Reasoning-Driven Retrosynthesis Prediction with Large Language Models via Reinforcement Learning

URL: View paper

Brief Assessment

Retrosynthesis Reinforcement Learning[23] focuses on retrosynthesis prediction (predicting reactants from products) rather than drug discovery with structured reasoning trajectories for pocket-ligand interactions. The tasks and domains are fundamentally different.

7. Llm agent swarm for hypothesis-driven drug discovery

URL: View paper

Brief Assessment

LLM Agent Swarm[24] focuses on multi-agent orchestration for hypothesis generation across diverse data sources (literature, omics, market intelligence), whereas DRUGTRAIL emphasizes structured reasoning trajectories within a single LLM framework combined with druggability-tailored preference optimization for molecular generation.

8. Concept Bottleneck Language Models For protein design

URL: View paper

Brief Assessment

Concept Bottleneck Proteins[25] focuses on interpretable protein language models for protein design with concept-based control, not on drug discovery frameworks with structured reasoning trajectories for small molecule generation.

9. K-Paths: Reasoning over Graph Paths for Drug Repurposing and Drug Interaction Prediction

URL: View paper

Brief Assessment

K-Paths[30] focuses on biomedical knowledge graph path extraction for drug-drug and drug-disease interaction prediction, not on LLM-based structured reasoning trajectories with druggability-tailored optimization for molecular generation.

10. DDI-GPT: Explainable Prediction of Drug-Drug Interactions using Large Language Models enhanced with Knowledge Graphs

URL: View paper

Brief Assessment

DDI-GPT[26] focuses on predicting drug-drug interactions using knowledge graphs and LLMs for safety assessment, not on general interpretable drug discovery with structured reasoning trajectories for molecular design.

Contribution 2: Clinical Chemistry-Informed Reasoning (CCIR) module

Description: The authors design a module that generates structured reasoning trajectories following five clinical chemistry dimensions: physicochemical profiling, structural integrity, prior knowledge guidance, conservation analysis, and multi-attribute optimization. This module enables the model to articulate the how and why behind its molecular design decisions.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Analysis of the uncharted, druglike property space by self-organizing maps

URL: View paper

Brief Assessment

Druglike Property Space[17] focuses on analyzing physicochemical properties of compound libraries using self-organizing maps for chemical space navigation, not on structured reasoning for molecular design using clinical chemistry dimensions.

2. Quantitative structureâ□□activity relationship (QSAR) studies as strategic approach in drug discovery

URL: View paper

Brief Assessment

QSAR Drug Discovery[20] focuses on quantitative structure-activity relationship methodologies for drug discovery using molecular descriptors and statistical correlations. It does not address structured reasoning trajectories following clinical chemistry dimensions

(physicochemical profiling, structural integrity, prior knowledge guidance, conservation analysis, multi-attribute optimization) as proposed in the original paper's CCIR module.

---

3. Revisiting methotrexate and phototrexate Zinc15 library-based derivatives using deep learning in-silico drug design approach

URL: View paper

Brief Assessment

Methotrexate Deep Learning[13] focuses on QSAR-based virtual screening and ADMET prediction for anticancer drug discovery, not on structured reasoning frameworks for molecular design using clinical chemistry dimensions.

---

4. CNS drug design: balancing physicochemical properties for optimal brain exposure

URL: View paper

Brief Assessment

CNS Drug Design[14] focuses on physicochemical properties for brain-penetrating drugs, not on structured reasoning trajectories for molecular design using clinical chemistry dimensions as a computational framework.

---

5. Prediction of oral bioavailability in rats: Transferring insights from in vitro correlations to (deep) machine learning models using in silico model outputs and chemical â¦

URL: View paper

Brief Assessment

Oral Bioavailability Prediction[16] focuses on predicting oral bioavailability in rats using physicochemical parameters and in vitro correlations, not on structured reasoning for molecular design following clinical chemistry dimensions.

---

6. Basic concepts of artificial neural network (ANN) modeling and its application in pharmaceutical research

URL: View paper

Brief Assessment

ANN Pharmaceutical Research[18] focuses on basic artificial neural network modeling concepts and applications in pharmaceutical research, without addressing structured reasoning trajectories or clinical chemistry dimensions for molecular design. The candidate paper discusses neural networks and physicochemical descriptors as analytical tools, not as components of an interpretable reasoning framework.

---

7. Chemical predictive modelling to improve compound quality

URL: View paper

Brief Assessment

Chemical Predictive Modelling[15] focuses on QSAR methods for predicting compound quality and ADMET properties, not on structured reasoning trajectories following clinical chemistry dimensions for molecular design decisions.

---

8. Training a Scientific Reasoning Model for Chemistry

URL: View paper

Brief Assessment

Scientific Reasoning Chemistry[12] focuses on training reasoning models for chemistry tasks using reinforcement learning on verifiable problems (molecular design, synthesis, property prediction). The candidate does not describe structured reasoning trajectories following clinical chemistry dimensions (physicochemical profiling, structural integrity, prior knowledge guidance, conservation analysis, multi-attribute optimization) as specified in the original paper's CCIR module. The candidate's approach centers on chain-of-thought reasoning for chemistry problems rather than the specific five-dimensional clinical chemistry framework claimed by the original authors.

---

9. Abstract A016: A computational chemistry and AI-driven framework for structure-based drug design informed by underlying factors of mutation-induced drug resistance: A study of KRAS

URL: View paper

Brief Assessment

KRAS Drug Resistance[19] focuses on computational chemistry and ML for analyzing protein dynamics and pharmacophore extraction in drug-resistant KRAS mutations, not on structured reasoning trajectories following clinical chemistry dimensions for molecular design decisions.

---

10. Computational Profiling of Monoterpenoid Phytochemicals: Insights for Medicinal Chemistry and Drug Design Strategies

URL: View paper

Brief Assessment

Monoterpenoid Profiling[11] focuses on computational screening of natural products using physicochemical parameters and ADMET profiles, not on structured reasoning trajectories for molecular design guided by clinical chemistry dimensions.

Contribution 3: Druggability-Tailored Preference Optimization (DTPO) strategy

Description: The authors develop DTPO, a reinforcement learning optimization strategy that moves beyond single-metric binding affinity optimization by incorporating a hybrid reward function. This reward combines ligand-based similarity to bioactive compounds with rule-based druggability indicators, enabling efficient online computation while maintaining strong connections to drug-likeness.

This contribution was assessed against **1 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. CRISPR-tica.ai: A function-informed generative modeling pipeline for prioritizing drug discovery in AML

URL: View paper

Brief Assessment

CRISPR-tica AML[21] uses direct preference optimization to fine-tune a diffusion model for binding profiles and drug-likeness, but focuses on structure-based drug design for AML rather than the ORIGINAL paper's LLM-based reasoning framework with hybrid reward functions combining ligand-based similarity and rule-based druggability indicators for general drug discovery.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] DrugTrail: Explainable Drug Discovery via Structured Reasoning and Druggability-Tailored Preference Optimization View paper
- [1] Artificial intelligence in drug development View paper
- [2] Extracting medicinal chemistry intuition via preference machine learning View paper
- [3] Stabilizing Reasoning in Medical LLMs with Continued Pretraining and Reasoning Preference Optimization View paper
- [4] Towards Interpretable Models of Chemist Preferences for Human-in-the-Loop Assisted Drug Discovery View paper
- [5] Zero-shot drug repurposing with geometric deep learning and clinician centered design View paper
- [6] Explainable Reasoning Path Inference of Anti-cancer Drug Sensitivity on Genomic Knowledge Graph via Macro-Micro Agent Collaborative Reinforcement Learning View paper
- [7] Chem-R: Learning to Reason as a Chemist View paper
- [8] Reasoning-Enhanced Large Language Models for Molecular Property Prediction View paper
- [9] The Role of Reinforcement Learning in Large Language Models: A Survey View paper
- [10] Causal inference and interpretable machine learning for personalised medicine View paper
- [11] Computational Profiling of Monoterpenoid Phytochemicals: Insights for Medicinal Chemistry and Drug Design Strategies View paper
- [12] Training a Scientific Reasoning Model for Chemistry View paper
- [13] Revisiting methotrexate and phototrexate Zinc15 library-based derivatives using deep learning in-silico drug design approach View paper
- [14] CNS drug design: balancing physicochemical properties for optimal brain exposure View paper
- [15] Chemical predictive modelling to improve compound quality View paper
- [16] Prediction of oral bioavailability in rats: Transferring insights from in vitro correlations to (deep) machine learning models using in silico model outputs and chemical â¦ View paper
- [17] Analysis of the uncharted, druglike property space by self-organizing maps View paper
- [18] Basic concepts of artificial neural network (ANN) modeling and its application in pharmaceutical research View paper
- [19] Abstract A016: A computational chemistry and AI-driven framework for structure-based drug design informed by underlying factors of mutation-induced drug resistance: A study of KRAS View paper
- [20] Quantitative structureâactivity relationship (QSAR) studies as strategic approach in drug discovery View paper
- [21] CRISPR-tica.ai: A function-informed generative modeling pipeline for prioritizing drug discovery in AML View paper
- [22] Effective and Explainable Molecular Property Prediction by Chain-of-Thought Enabled Large Language Models and Multi-Modal Molecular Information Fusion View paper
- [23] Reasoning-Driven Retrosynthesis Prediction with Large Language Models via Reinforcement Learning View paper
- [24] Llm agent swarm for hypothesis-driven drug discovery View paper
- [25] Concept Bottleneck Language Models For protein design View paper
- [26] DDI-GPT: Explainable Prediction of Drug-Drug Interactions using Large Language Models enhanced with Knowledge Graphs View paper
- [27] Mol-LLaMA: Towards General Understanding of Molecules in Large Molecular Language Model View paper
- [28] Molreasoner: Toward effective and interpretable reasoning for molecular llms View paper
- [29] PharmAgents: Building a Virtual Pharma with Large Language Model Agents View paper
- [30] K-Paths: Reasoning over Graph Paths for Drug Repurposing and Drug Interaction Prediction View paper
- [31] Beyond Chemical QA: Evaluating LLM's Chemical Reasoning with Modular Chemical Operations View paper