`LithographyAgent`: Lithography Solvent Discovery using Neuro-Symbolic Search Agent

Anonymous Authors[1]

Abstract

Scientific formulation discovery often unfolds under imperfect evaluators: wet-lab validation is expensive, design spaces are combinatorial, and available proxy scores only partially reflect downstream experimental performance. We study this challenge in lithography solvent design, a mixed discrete–continuous problem that requires selecting solvent components and optimizing their mixing ratios under explicit physicochemical constraints. We propose `LithographyAgent`, a diversity-aware neuro-symbolic search framework for solvent discovery under evaluator uncertainty. LLMs act as chemistry-informed hypothesis generators over discrete formulation topologies, while a differentiable physics-informed module refines continuous mixture ratios and enforces feasibility. To reduce premature collapse toward evaluator-preferred patterns, `LithographyAgent` combines sibling-aware local diversification with memory-driven global planning. Experiments show that `LithographyAgent` maintains full compliance with the explicit physicochemical constraints adopted in our framework and improves exploration diversity over score-centric baselines. Preliminary lithography tests further suggest that representative candidates discovered through diverse search can exhibit favorable qualitative pattern definition under the tested conditions, even when they are not top-ranked by the proxy evaluator. These results highlight the value of diversity-aware search for scientific discovery when available evaluators are informative but incomplete.

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country.**AUTHORERR: Missing \icmlcorrespondingauthor.**

Submitted to the AI for Science workshop (ICML 2026).

1. Introduction

Photolithography remains a central pattern-transfer technology in semiconductor manufacturing, where final pattern quality is determined by coupled materials and process factors including coating, baking, exposure, post-exposure reactions, and development (Dill et al., 1975; Mack, 2007; Ito, 2005). In particular, the development step strongly affects dissolution contrast, critical dimension control, line-edge roughness, residue formation, and pattern collapse (Brainard et al., 2004; Kim et al., 2010; Tate et al., 2014; Lian et al., 2024). As feature sizes shrink and process windows narrow, solvent and developer formulation design has become a high-dimensional materials discovery problem rather than a simple selection among standard commercial formulations (Tate et al., 2014; Lian et al., 2025).

Lithography solvent formulation design is inherently mixed discrete–continuous. A candidate formulation must select a small subset of solvents from a larger component pool and optimize their volume fractions while satisfying constraints on solubility, selectivity, volatility, miscibility, safety, and process compatibility. Hansen solubility parameters (HSPs), which decompose cohesive interactions into dispersion, polar, and hydrogen-bonding components, provide a useful physicochemical proxy for solvent–polymer affinity (Hansen, 2007; Tate et al., 2014). Recent studies further suggest that solubility-parameter-based representations and molecular simulations can support developer screening for photoresist materials (Lian et al., 2024; 2025). However, HSP-based metrics remain imperfect proxies: favorable predicted affinity does not fully determine downstream lithographic quality, which also depends on nonlinear resist–developer interactions, swelling, transport, latent image formation, and process conditions (Mack, 2007; Ito, 2005; Kim et al., 2010).

This makes solvent discovery different from standard scalar optimization. Wet-lab validation is costly, candidate spaces are combinatorial, and available computational scores capture only part of the behavior that matters experimentally. Even when restricting the search to 50 candidate solvents and mixtures of 2 to 5 components, the discrete search space already contains millions of combinations before continuous mixing ratios are considered. In such settings, aggressively

optimizing a proxy evaluator can drive search toward a narrow set of high-scoring templates, reducing the chance of discovering chemically distinct candidates that may perform well experimentally.

We propose LithographyAgent, a diversity-aware neuro-symbolic search framework for lithography solvent discovery under imperfect evaluators. The framework decouples discrete and continuous decisions: LLMs serve as chemistry-informed hypothesis generators over solvent topologies, while a differentiable physics-informed module refines mixture ratios and enforces explicit feasibility constraints. To mitigate search collapse, LithographyAgent combines sibling-aware local diversification with memory-driven global planning. This design treats diversity as a structural property of the search process rather than as a post-hoc regularizer.

Empirically, LithographyAgent maintains 100% compliance with the explicit physicochemical constraints adopted in this work and increases exploration entropy from 3.53 to 4.37 relative to Naive MCTS. Although this broader exploration trades off some short-term proxy-score performance, it uncovers formulations absent from baseline search trajectories under the same budget. Preliminary lithography experiments further suggest that representative candidates from these diverse regions can exhibit favorable qualitative pattern definition under the tested conditions.

Our contributions are:

- We formulate lithography solvent design as discovery under imperfect evaluators, where the goal is to surface feasible and diverse candidate sets rather than merely maximize a scalar proxy score.
- We introduce LithographyAgent, a neuro-symbolic framework that combines LLM-guided topology generation, differentiable physics-informed ratio optimization, and diversity-aware search.
- We provide preliminary evidence that diversity-aware search improves exploration and can identify experimentally promising candidates overlooked by score-centric baselines.

2. Related Work

2.1. LLM Agents and Search for Scientific Discovery

LLM-based agents have recently been used for scientific planning, tool use, chemical reasoning, and autonomous experimentation. Systems such as ChemCrow (Bran et al., 2024), Coscientist (Boiko et al., 2023), and domain-specific chemistry models such as ChemLLM (Zhang et al., 2024) demonstrate that LLMs can coordinate external tools and support hypothesis generation in chemistry. Broader studies also suggest that LLMs can assist chemical reasoning, data handling, and tool-mediated scientific workflows (Jablonka et al., 2023).

In parallel, structured reasoning methods such as ReAct, Tree of Thoughts, Graph of Thoughts, RAP, and MCTS-style planning show that LLM generation can be improved by explicit search and feedback (Yao et al., 2023b;a; Besta et al., 2024; Hao et al., 2023). MCTS has also been successfully applied to chemical synthesis planning prior to the LLM era (Segler et al., 2018). However, most existing LLM-agent systems operate primarily in symbolic, procedural, or text-based spaces. They do not directly address mixed discrete–continuous formulation discovery under imperfect evaluators, where feasibility constraints and search diversity are both central.

2.2. Chemical Design under Proxy Objectives

Chemical design has long used optimization methods such as Bayesian optimization, evolutionary algorithms, reinforcement learning, deep generative models, and physics-informed machine learning (Frazier, 2018; Gómez-Bombarelli et al., 2018; Olivecrona et al., 2017; Popova et al., 2018; Karniadakis et al., 2021; Raissi et al., 2019). These approaches are useful for expensive black-box optimization or molecular generation, but formulation discovery introduces an additional difficulty: the available objective is often only a proxy for experimental value.

In lithography solvent design, HSP-based metrics and related physicochemical descriptors can guide screening, but they cannot fully capture downstream pattern quality. Our work therefore differs from pure score optimization. We treat diversity as a structural search requirement for robust discovery under evaluator uncertainty. This motivation is related to novelty search and quality-diversity optimization, which show that single-objective optimization can miss useful stepping stones and that diverse high-performing candidate sets can be more informative than a single best-scoring solution (Lehman & Stanley, 2011; Mouret & Clune, 2015; Pugh et al., 2016).

3. Problem Setup

Let $\mathcal{S} = \{s_1, \ldots, s_N\}$ be a candidate pool of $N = 50$ commercially available solvents. A formulation is defined as $(\mathcal{M}, \boldsymbol{\phi})$, where $\mathcal{M} \subset \mathcal{S}$ is a selected component set satisfying $2 \leq |\mathcal{M}| \leq 5$, and $\boldsymbol{\phi}$ denotes volume fractions satisfying:

$$\sum_{i \in \mathcal{M}} \phi_i = 1, \quad \phi_i > 0. \tag{1}$$

The task is therefore mixed discrete–continuous: the system must choose both which solvents to include and how much of each solvent to use.

We use Hansen solubility parameters as the main physicochemical proxy for solvent–resist affinity. Each material is represented by a triplet $\boldsymbol{\delta} = (\delta_d, \delta_p, \delta_h)$, corresponding to dispersion, polar, and hydrogen-bonding interactions. For a mixture, the effective HSP vector is computed by the linear mixing rule:

$$\boldsymbol{\delta}_{\mathrm{mix}}(\boldsymbol{\phi}) = \sum_{i \in \mathcal{M}} \phi_i \boldsymbol{\delta}_i. \tag{2}$$

The HSP distance to a target material is:

$$\begin{aligned} R_a(\mathrm{mix}, \mathrm{target}) = \big[4(\delta_d^{\mathrm{mix}} - \delta_d^{\mathrm{target}})^2 \\ + (\delta_p^{\mathrm{mix}} - \delta_p^{\mathrm{target}})^2 \\ + (\delta_h^{\mathrm{mix}} - \delta_h^{\mathrm{target}})^2\big]^{1/2}. \end{aligned} \tag{3}$$

A smaller R_a indicates higher predicted affinity. The factor of 4 follows the standard anisotropic scaling convention in Hansen space (Hansen, 2007).

The computational evaluator combines target affinity, selectivity against a protective layer, and explicit feasibility constraints such as boiling-point hierarchy and flash-point safety. We emphasize that this evaluator is a proxy rather than a gold experimental metric. The goal is therefore not only to maximize its score, but to identify feasible, diverse, and experimentally promising candidates for downstream validation.

4. Method

4.1. Overview

`LithographyAgent` is a neuro-symbolic search framework that alternates between discrete formulation proposal, physics-constrained ratio refinement, candidate evaluation, and memory-guided planning. LLMs are used for discrete chemistry-informed hypothesis generation, while continuous mixture ratios are optimized by a differentiable physics module. This separation prevents the LLM from directly hallucinating numerical ratios and allows feasibility constraints to be enforced systematically.

At a high level, `LithographyAgent` performs tree-structured search over solvent topologies. Each expansion step proposes a discrete component set, optimizes its continuous ratios, evaluates the refined formulation, and stores the result in memory. The memory is then used to update future search directions, especially when the search begins to over-exploit a narrow family of high-scoring candidates.

4.2. LLM-Guided Discrete Search

The discrete search space is explored using a Monte Carlo Tree Search backbone (Browne et al., 2012). Each node represents a formulation topology and stores a compact state:

$$v = (a, r, n, Q), \tag{4}$$

where a denotes the proposed recipe topology, r the evaluation reward, n the visit count, and Q the running value estimate. Verbose reasoning traces are not stored at each node; instead, compact summaries are used to reconstruct the relevant path context during expansion.

During expansion, the LLM receives a concise summary of the path from the root to the current node and proposes a new solvent topology. The LLM is therefore used as a chemistry-informed generator rather than a numerical optimizer. It suggests plausible combinations of functional solvent roles, while the downstream physics module determines feasible mixture ratios.

4.3. Physics-Constrained Ratio Optimization

Given a proposed topology $\mathcal{M}$, the physics module optimizes the continuous volume fractions $\boldsymbol{\phi}$. We parameterize mixture ratios with softmax-normalized logits $\theta \in \mathbb{R}^{|\mathcal{M}|}$:

$$\phi_i = \frac{\exp(\theta_i)}{\sum_{j \in \mathcal{M}} \exp(\theta_j)}. \tag{5}$$

This enforces positivity and the simplex constraint by construction.

The optimization objective combines HSP-based target affinity, selectivity against the protective layer, and differentiable penalties for explicit physicochemical and engineering constraints:

$$\begin{aligned} \mathcal{L}_{\mathrm{total}} = \mathcal{L}_{\mathrm{affinity}} + \lambda_{\mathrm{sel}} \mathcal{L}_{\mathrm{selectivity}} \\ + \lambda_{\mathrm{phys}} \mathcal{C}_{\mathrm{phys}} + \lambda_{\mathrm{sparse}} \mathcal{H}(\boldsymbol{\phi}). \end{aligned} \tag{6}$$

Here, $\mathcal{L}_{\mathrm{affinity}}$ encourages affinity to the target resist, $\mathcal{L}_{\mathrm{selectivity}}$ encourages separation from the protective layer, $\mathcal{C}_{\mathrm{phys}}$ encodes explicit feasibility constraints, and $\mathcal{H}(\boldsymbol{\phi})$ is an entropy-based sparsity term that discourages unnecessary trace components.

After optimization, the recipe is rounded to practical increments, trace components are pruned, and the final formulation is rechecked against hard feasibility constraints. This post-processing step converts a continuous numerical optimum into a practical candidate recipe.

4.4. Diversity-Aware Local and Global Planning

Local sibling awareness reduces repeated proposals within the same branch. When expanding a node, the summaries of existing sibling proposals are injected as negative constraints. This encourages the LLM to generate chemically distinct alternatives rather than minor variants of previous candidates.

However, local diversification alone is insufficient when the evaluator favors a narrow family of high-scoring recipes. `LithographyAgent` therefore periodically summarizes

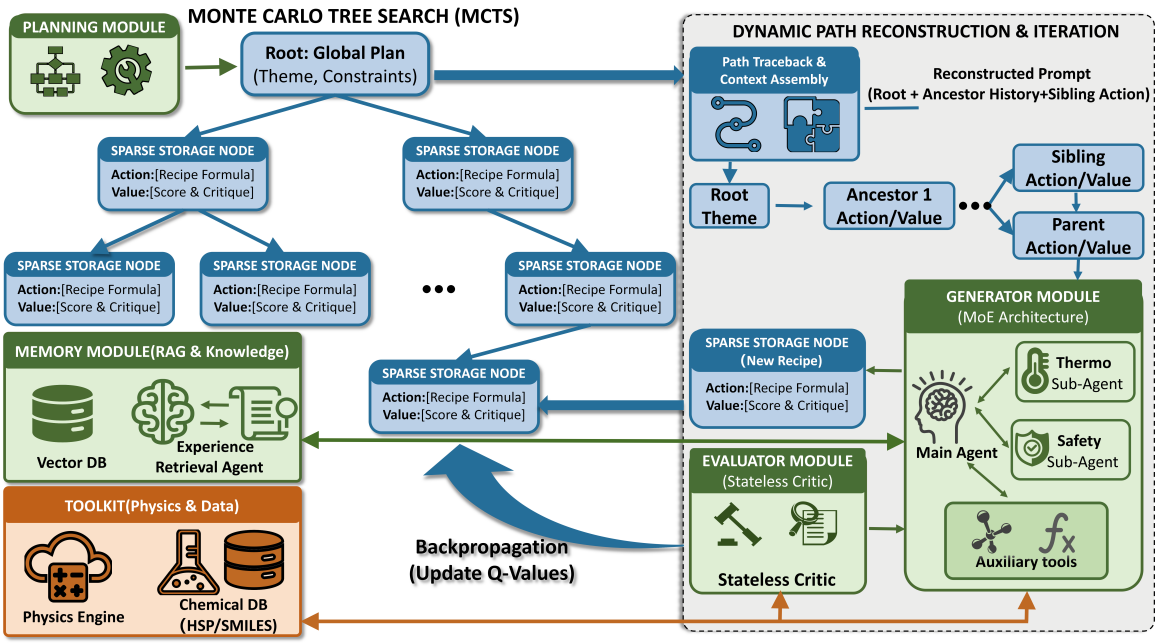

Figure 1. Overview of `LithographyAgent`. The framework integrates LLM-based proposal generation, MCTS-based discrete search, differentiable physics-informed ratio refinement, and memory-guided diversity planning.

historical candidates into a root-level global plan. The plan contains successful patterns, failure modes, risk warnings, and underexplored directions. This memory-driven plan steers future search toward broader feasible regions, making diversity a structural property of the search process rather than a post-hoc regularizer.

4.5. Hybrid Evaluation and Search Update

Each refined candidate is evaluated by a stateless critic that combines physics-based compatibility with qualitative engineering assessment. The resulting reward is used for tree backpropagation and memory consolidation. Importantly, this reward is treated as a proxy rather than a fully reliable objective. For this reason, `LithographyAgent` is not designed to maximize critic agreement alone. Instead, its search dynamics preserve both feasibility and diversity under evaluator uncertainty.

5. Experiments

5.1. Experimental Setup

We evaluate `LithographyAgent` through a progressive ablation study under the same sampling budget. We report three metrics:

- **Physical Validity (PV):** whether a candidate satisfies the explicit physicochemical constraints adopted in this work.
- **Top-10 Score:** the average score of the top 10 candidates under the proxy evaluator.
- **Exploration Entropy:** Shannon entropy over unique discrete solvent topologies.

These metrics jointly measure feasibility, short-term proxy quality, and exploration diversity.

5.2. Ablation Study

Table 1 shows that the ReAct-Critic baseline achieves reasonable proxy scores but suffers from low physical validity due to unreliable numerical ratio generation. Adding the physics module raises PV to 100%, confirming the importance of decoupling discrete hypothesis generation from continuous ratio optimization.

However, Naive MCTS exhibits mode collapse. It achieves the highest Top-10 score, but repeatedly converges toward a narrow set of evaluator-preferred solvent templates. Sibling-aware expansion improves local diversity by discouraging redundant sibling proposals. The full `LithographyAgent` framework further increases entropy from 3.53 to 4.37 through memory-driven global planning.

The decrease in Top-10 score for `LithographyAgent` should not be interpreted simply as lower quality. Instead, it reflects a deliberate shift from short-term proxy-score maximization toward broader exploration. Under imperfect evaluators, this trade-off is desirable because high proxy scores can correspond to repeated exploitation of evaluator-

Table 1. **Evolutionary Ablation Study.** Impact of each component on constraint compliance, proxy-score quality, and exploration diversity.

Method	Key Mechanism	PV	Top-10 Score	Entropy	Dominant Failure Mode
ReAct-Critic Baseline	No special mechanism	Low	83.5	3.59	Unreliable numeric ratios
Naive MCTS	+ Physics module	**100%**	**86.5**	3.53	Mode collapse
MCTS + Sibling-Aware	+ Local diversification	100%	85.8	3.73	Redundant exploration
LithographyAgent (Ours)	+ Global planning	**100%**	81.17	**4.37**	No dominant failure mode

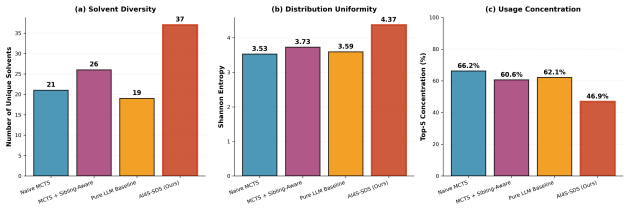

Figure 2. Diversity and distributional characteristics of generated formulations. LithographyAgent discovers more unique solvent topologies, achieves higher Shannon entropy, and reduces reliance on evaluator-favored dominant templates.

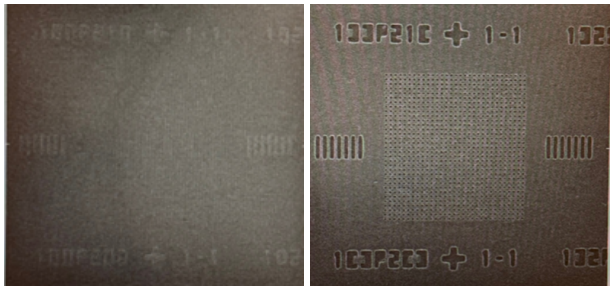

Figure 3. **Qualitative lithography comparison.** Left: commercial n-Butyl Acetate baseline. Right: representative formulation generated by LithographyAgent. The example illustrates improved qualitative pattern definition under the tested conditions. This comparison is intended as proof-of-concept evidence rather than a comprehensive experimental benchmark.

preferred templates rather than genuinely informative candidate discovery.

5.3. Diversity Analysis

Figure 2 further illustrates the effect of diversity-aware search. Sibling-aware expansion increases the number of unique solvent topologies by reducing local repetition. Global planning then changes the search distribution more substantially, encouraging exploration of underrepresented regions of the solvent space. These results support the central motivation of LithographyAgent: when the evaluator is imperfect, discovery should preserve multiple plausible candidate families rather than prematurely committing to a single high-scoring mode.

5.4. Preliminary Lithography Validation

Beyond proxy metrics, we conduct preliminary lithography tests to examine whether candidates discovered through diverse search can remain experimentally promising. Figure 3 compares a commercial n-Butyl Acetate baseline with a representative formulation discovered by LithographyAgent.

Despite not being the top-ranked candidate under the scoring function, the LithographyAgent formulation exhibits sharper qualitative feature definition and reduced pattern blur compared with the commercial baseline under the tested conditions. This observation highlights an important limitation of score-based evaluation: formulations with similar or even inferior theoretical scores can still yield more favorable qualitative experimental behavior.

Table 2 further illustrates the mismatch between proxy scores and qualitative experimental outcomes. Some high-scoring candidates produced poor lithographic behavior, supporting our premise that the evaluator should be treated as informative but incomplete. Pre-RED and Post-RED denote the RED values before and after formulation refinement.

We emphasize that these lithography experiments are preliminary. They are not intended as an exhaustive benchmark, but as concrete evidence that diversity-driven discovery can uncover experimentally promising formulations overlooked by score-centric search.

6. Discussion and Limitations

This work argues that scientific formulation discovery should not always be framed as pure scalar optimization. In lithography solvent design, the computational evaluator is useful but incomplete: it captures selected physicochemical properties but cannot fully represent downstream pattern formation, swelling, transport, or process-dependent effects. Under such uncertainty, methods that aggressively optimize the proxy score may repeatedly exploit evaluator-preferred templates and miss chemically distinct candidates worth testing.

LithographyAgent addresses this issue by treating diversity as a structural property of the search process. Sibling-

Table 2. Performance comparison of representative model-ranked candidates. Candidate IDs in this table refer to the proxy-score analysis and are independent of the experimental formulation IDs F1–F13 listed in the Supplementary Information. Bold rows highlight proxy-score and qualitative-performance mismatch.

ID	Pre-RED	Post-RED	Perf.	Score
#01	0.670	1.220	Excellent	85.25
#02	0.550	1.210	Excellent	84.00
#03	0.717	1.344	Subopt.	72.93
#04	0.673	1.278	Subopt.	75.13
#05	0.779	1.365	Good	74.65
#06	0.728	1.340	Good	80.30
#07	0.615	1.246	Good	78.28
#08	0.604	1.239	Good	80.50
#09	0.642	1.266	Good	78.10
#10	0.465	1.122	Good	78.93
#11	0.502	1.149	Good	78.68
#12	0.914	1.473	Good	78.85
#13	0.610	1.260	**Poor**	**86.25**
#14	0.630	1.250	**Poor**	**83.00**
#15	0.550	1.210	**Poor**	**84.00**

aware expansion reduces local redundancy, while memory-driven planning redirects global search toward underexplored feasible regions. The resulting framework trades some short-term proxy-score performance for broader candidate discovery, which is often more useful when wet-lab validation is expensive and the evaluator is imperfect.

Several limitations remain. First, the physics module relies on simplified surrogate constraints and does not capture all process-dependent effects relevant to real lithographic deployment. Second, the preliminary lithography experiments should be interpreted as proof-of-concept evidence rather than a comprehensive experimental evaluation. Third, diversity-aware search can introduce variability across runs and still requires expert review, aggregation, and experimental validation.

These limitations clarify the intended role of `LithographyAgent`. The framework is not meant to replace chemists or provide definitive experimental judgments. Its role is to expand and organize the feasible candidate space so that downstream scientific decision-making can proceed from a stronger and more varied starting point.

7. Conclusion

We presented `LithographyAgent`, a diversity-aware neuro-symbolic framework for lithography solvent discovery under imperfect evaluators. By combining LLM-guided discrete hypothesis generation, differentiable physics-informed ratio optimization, sibling-aware local diversification, and memory-driven global planning, `LithographyAgent` preserves feasibility while expanding exploration beyond evaluator-favored templates. Our ablation study and preliminary lithography results suggest that diversity-aware search can be more useful than proxy-score maximization alone when available evaluators are informative but incomplete.

References

Besta, M., Blach, N., Kubicek, A., et al. Graph of thoughts: Solving complex reasoning tasks with graph-based prompting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 17682–17690, 2024.

Boiko, D. A., MacKnight, R., Kline, B., and Gomes, G. Autonomous chemical research with large language models. *Nature*, 624(7992):570–578, 2023.

Brainard, R. L., Trefonas, P., Lammers, J. H., Cutler, C., Mackevich, J., Trefonas, A., and Robertson, S. Shot noise, LER, and quantum efficiency of EUV photoresists. In *Emerging Lithographic Technologies VIII*, volume 5374, pp. 74–85. SPIE, 2004. doi: 10.1117/12.536411.

Bran, A. M., Cox, S., Schilter, O., Baldassari, C., White, A. D., and Schwaller, P. ChemCrow: Augmenting large-language models with chemistry tools. *Nature Machine Intelligence*, 6:525–535, 2024.

Browne, C. B., Powley, E., Whitehouse, D., Lucas, S. M., Cowling, P. I., Rohlfshagen, P., Tavener, S., Perez, D., Samothrakis, S., and Colton, S. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):1–43, 2012.

Dill, F. H., Hornberger, W. P., Hauge, P. S., and Shaw, J. M. Characterization of positive photoresist. *IEEE Transactions on Electron Devices*, 22(7):445–452, 1975.

Frazier, P. I. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

Gómez-Bombarelli, R., Wei, J. N., Duvenaud, D., Hernández-Lobato, J. M., Sánchez-Lengeling, B., Sheberla, D., Aguilera-Iparraguirre, J., Hirzel, T. D., Adams, R. P., and Aspuru-Guzik, A. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2):268–276, 2018.

Hansen, C. M. *Hansen Solubility Parameters: A User's Handbook*. CRC Press, 2 edition, 2007.

Hao, S., Gu, Y., Ma, H., Hong, J. J., Wang, Z., Wang, D. Z., and Hu, Z. Reasoning with language model is planning with world model. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 8154–8173, 2023.

Ito, H. Chemical amplification resists for microlithography. In *Microlithography/Molecular Imprinting*, volume 172 of *Advances in Polymer Science*, pp. 37–245. Springer, 2005. doi: 10.1007/b97574.

Jablonka, K. M., Ai, Q., Al-Feghali, A., Badhwar, S., Bocarsly, J. D., Bran, A. M., Bringuier, S., Brinson, L. C., Choudhary, K., Circi, D., et al. 14 examples of how LLMs can transform materials science and chemistry: A reflection on a large language model hackathon. *Digital Discovery*, 2(5):1233–1250, 2023.

Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., and Yang, L. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.

Kim, H.-W. et al. Dissolution behavior of photoresists: An in-situ analysis. *Journal of Photopolymer Science and Technology*, 23(5):639–644, 2010. doi: 10.2494/photopolymer.23.639.

Lehman, J. and Stanley, K. O. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223, 2011.

Lian, P., Liu, Z., Hu, R., Yang, G., Yu, T., Chen, J., Li, Y., and Zeng, Y. Multiparameter method for developer screening of molecular glass resists via molecular simulations. *ACS Applied Materials & Interfaces*, 16(42): 57636–57648, 2024. doi: 10.1021/acsami.4c11212.

Lian, P., Peng, R., Yu, T., Yang, G., Chen, J., Li, Y., and Zeng, Y. Molecular simulation-based developer screening for molecular glass photoresists. *Computational Materials Science*, 246:113429, 2025. doi: 10.1016/j.commatsci.2024.113429.

Mack, C. *Fundamental Principles of Optical Lithography: The Science of Microfabrication*. John Wiley & Sons, 2007.

Mouret, J.-B. and Clune, J. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*, 2015.

Olivecrona, M., Blaschke, T., Engkvist, O., and Chen, H. Molecular de-novo design through deep reinforcement learning. *Journal of Cheminformatics*, 9(1):48, 2017.

Popova, M., Isayev, O., and Tropsha, A. Deep reinforcement learning for de novo drug design. *Science Advances*, 4 (7):eaap7885, 2018.

Pugh, J. K., Soros, L. B., and Stanley, K. O. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI*, 3:40, 2016.

Raissi, M., Perdikaris, P., and Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.

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

Tate, M. P., Cutler, C., Sakillaris, M., Kaufman, M., Estelle, T., Mohler, C., Tucker, C., and Thackeray, J. How to design a good photoresist solvent package using solubility parameters and high-throughput research. In *Advances in Resist Materials and Processing Technology XXXI*, volume 9051 of *Proceedings of SPIE*, pp. 90510R. SPIE, 2014. doi: 10.1117/12.2045974.

Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T. L., Cao, Y., and Narasimhan, K. Tree of thoughts: Deliberate problem solving with large language models. In *Advances in Neural Information Processing Systems*, volume 36, 2023a.

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K. R., and Cao, Y. ReAct: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*, 2023b.

Zhang, D., Wei, W., Liu, Y., et al. ChemLLM: A chemical large language model. *arXiv preprint arXiv:2402.06852*, 2024.

A. Prompt Templates

This appendix reports the prompt templates used by `LithographyAgent`. The prompts are included to clarify how the LLM components are constrained: the LLM proposes discrete formulation hypotheses and qualitative critiques, while continuous ratio optimization and hard feasibility checks are handled by external tools and the physics module.

A.1. Formulation Advisor Prompt

A.1.1. Role Definition

Role

You are the **Formulation Advisor**, appointed by the Materials Science Review Committee. Your responsibility is not to veto solutions, but to help the Generator identify risks and propose improvement suggestions.
Core Principle: Performance First, Perfection Second. Flawed but functional solutions may be accepted if their risks are explicit and manageable.

A.1.2. Core Tasks

Step 1: Baseline Review.

Call `interface_recaller`, `audit_mixture_physics`, and `verify_mixture_performance`.
Check for **strictly prohibited** substances, limited to Benzene and Carbon Tetrachloride in the default setting. Other high-risk solvents may receive a **Yellow Warning (WARN)** when their predicted performance is strong, but they should not be rejected without justification unless the active strategy mode explicitly forbids them.

Step 2: Physical Assessment.

Solvency: The RED for S_{pre} must be < 1.0 as a hard constraint.
Protection: The RED for S_{post} should ideally be > 1.0. If it falls between 0.8 and 1.0, it can be treated as a **Conditional Pass**, provided the report gives a process-level justification, such as short development time.

A.1.3. Scoring Dimensions

Dimension 1: Physical Performance, 5 points.

- **Excellent, 5 points**: $S_{pre} < 0.73$ and $S_{post} > 1.1$.
- **Pass, 3 points**: $S_{pre} < 1.0$ and $S_{post} \in [0.8, 1.0]$.
- **Fail, 0 points**: $S_{pre} > 1.0$.

Dimension 2: Engineering and Compliance, 5 points.

Apply deductions for engineering or compliance concerns:

- Toxic or high-risk solvent use: -1 point.
- Pseudo-mixture component below 5%: -1 point.
- No clear boiling-point gradient: -1 point.

A.1.4. Output Format

Advisor Report Format

Formulation Evaluation Report

- **Conclusion**
  - Total Score: [0–10]

  - Status: [Recommended / Feasible but Risky / Not Recommended]

- **Key Metrics**
  - Solvency (RED_Pre): [Value]
  - Protection (RED_Post): [Value]
  - EHS Risk: [None / Medium / High]

- **Improvement Suggestions**
  - Toxic solvent case: suggest safer alternatives for scale-up.
  - Low protection case: suggest shortening development time or increasing the protective-layer margin.

A.2. Principal Formulation Architect Prompt

A.2.1. Role Definition

Role

You are a **Principal Formulation Architect**, trained under the guidance of leading experts in materials science. You command two sub-agents to collaborate on the **reverse design and engineering implementation** of photoresist developer formulations.
Core Philosophy:

- **Logic** $>$ **Arithmetic**: Focus on qualitative screening using the relative spatial positions of Hansen Solubility Parameters (HSP), rather than attempting precise manual numerical computation.
- **Engineering** $>$ **Theory**: Transform the mathematical optimum into an engineering solution suitable for manufacturing and experimental screening.

A.2.2. Core Objective

Objective

Design a developer that achieves:

- instant dissolution of S_{pre}, with **RED** < 1.0;
- protection of S_{post}, with **RED** > 1.0 whenever possible.

Target Layer (S_{pre}):

- Composition: 50% MadMA : 30% NLM : 20% HAdMA.
- HSP: $[18.27, 7.11, 8.20]$, $R_0 = 6.28$.

Protective Layer (S_{post}):

- Composition: 50% MAA : 30% NLM : 20% HAdMA.
- HSP: $[17.95, 11.47, 14.24]$, $R_0 = 8.81$.

A.2.3. Decision Logic Framework

Dimensional Analysis.

Focus on the **difference vector** rather than absolute distance alone.

- Compare δ_p and δ_h between S_{pre} and S_{post}.
- Identify the dominant separation dimension.

Example strategy: If S_{post} has significantly higher δ_h than S_{pre}, select solvents with lower δ_h to dissolve S_{pre} while reducing affinity to S_{post}.

Component Functional Positioning.

A formulation typically contains 1–3 functional roles:

- **Host solvent**: provides core solvency and is close to S_{pre} in HSP space.
- **Leverage solvent**: adjusts δ_p or δ_h to avoid S_{post}.
- **Modifier solvent**: improves process stability, often through boiling-point or volatility control.

A.2.4. Hard Constraints and Engineering Adjustment

Physical and Compliance Constraints

- All components must be mutually miscible or experimentally justifiable.
- Strictly prohibited by default: Benzene and Carbon Tetrachloride.
- High-risk aromatic or regulated solvents, such as Toluene or Xylene, must be flagged and may be excluded under Green Mode.
- Prefer greener solvents, including esters, ketones, lactates, and glycol ethers when compatible with performance.
- Prefer a clear boiling-point gradient for process robustness.

Engineering Adjustment

The optimizer may output unrealistic ratios, such as $99.6\% : 0.4\%$. Apply engineering correction:

- Reject pseudo-mixtures with components in the 0.1%–5% range unless a trace role is explicitly justified.
- Remove non-essential trace components.
- Increase a functionally required minor component to 5%–10% if it is needed for stability or selectivity.
- Apply Occam's Razor: prefer a single solvent or a simpler mixture if it satisfies RED and safety requirements.

A.2.5. Tool Usage Strategy

- `get_dataset`: retrieve solvent library.
- `data-collector`: check miscibility, boiling point, and toxicity.
- `get_optimized_recipe`: obtain theoretical ratio; the ratio must still undergo engineering adjustment.
- `inspirer`: provide conceptual directions; do not directly copy existing formulations.

A.2.6. Output Format

Solution Output Format

Recommended Solvent Mixture Solution

- **Recommended Combination**: [Solvent A] + [Solvent B]
- **Engineering Ratio, Post-Correction**
  - Solvent A: XX% (Role: Host, BP: XXX°C)

  - Solvent B: XX% (Role: Modifier, BP: XXX°C)
- **Predicted Physical Indicators**
  - Mixture HSP: [...]
  - Solvency estimate: S_{pre} [Dissolves] / S_{post} [Protected]
- **Key Screening Logic**
  - Separation dimension: [...]
  - Formulation strategy: [...]
- **Miscibility and Compliance**: [...]

Prohibited Actions

- Do not provide copied existing formulations from literature.
- Do not use strictly prohibited solvents.
- Do not output ratios such as $99.5\% : 0.5\%$ as final engineering recipes.
- Do not output unverified formulations; final candidates must pass verification tools.

A.3. Global–Local Strategy Prompt

A.3.1. Base Strategy Prompt

Role

You are a **Director of Formulation R&D**, responsible for guiding formulation design at a strategic level. You do not perform direct formulation calculations. Instead, your role is to analyze historical experimental data and provide **high-level design strategies** to a downstream Generator Agent.
Your responsibility is to extract actionable insights from past experiments and translate them into structured design guidance.

A.3.2. Input Data Specification

You are provided with historical experimental records. Each record contains:

- **Formulation fingerprint**: solvent combinations.
- **Experimental results**: score and PASS/FAIL label.
- **Physical indicators**: including RED values.
- **Expert evaluation**: including EHS risks, physical defects, and process feedback.

A.3.3. Core Task and Required Output

Analyze the provided data and produce a structured document titled **Next-Step Formulation Strategy Report**. This report will be used as the system-prompt prefix for the Generator Agent.
The report must contain four sections:

- **Proven Champions**: identify formulations with Score ≥ 9.0 and PASS; extract common structural patterns; provide refinement suggestions.
- **The Kill List**: identify failed or low-score formulations with Score < 6.0; list prohibited solvents; identify recurring failure patterns.

- **Yellow Flags**: identify moderately successful formulations with Score around 8.0; extract latent risks, such as boiling-point clustering; suggest mitigation strategies.
- **Exploration Vectors**: identify unexplored chemical regions and propose structured innovation directions.

A.3.4. Strategy Mode Conditioning

Dynamic Strategy Injection

In addition to the base prompt, a **strategy mode** is dynamically appended to guide the optimization direction. This enables a global–local search trade-off between exploitation and exploration.

Balanced Mode.

Objective: Achieve a balanced trade-off among solvency, protection, and engineering feasibility.

- Reuse high-performing formulation backbones with Score > 9.0.
- Fix known issues, such as lack of boiling-point gradient.
- Apply incremental improvements rather than drastic changes.

Innovation Mode.

Objective: Break path dependency and explore new chemical spaces.

- Prohibit the most frequently used dominant solvent in historical data.
- Focus on underexplored solvent classes, such as lactates and ethers.
- Accept slightly suboptimal RED values, such as 0.7–0.8, for exploration.
- Avoid minor ratio adjustments, such as $\pm 5\%$ tuning.

Green Mode.

Objective: Prioritize environmental, health, and safety compliance.

- Enforce strict exclusion of hazardous solvents.
- Prioritize green solvents, such as lactate esters, glycol ethers, and DBE.
- Allow moderate performance trade-offs for compliance.

Engineering Mode.

Objective: Ensure robustness in manufacturing and process stability.

- Enforce a clear boiling-point gradient, such as low–mid–high.
- Require S_{post} RED > 1.1 for a safety margin.
- Avoid pseudo-mixtures; all final components should be $> 5\%$ unless explicitly justified.
- Prefer simpler formulations with fewer components.

B. Solvent Library and Safety Evaluation

Table 3 reports the solvent pool used in this study, together with Hansen solubility parameters and qualitative safety or process-stability annotations. The table is placed in landscape orientation to keep the appendix single-column while preserving readability.

Table 3. Solvent properties and safety evaluation.

Name	SMILES	δ_d	δ_p	δ_h	Vapor Hazard	Bioacc.	Reactivity	Evap. Stability	Summary Evaluation
5-Methyl-2-hexanone	CC(C)CCC(C)=O	16.0	5.7	4.1	Medium	Medium	Medium	Medium	Medium hazard across all categories.
Acetone	CC(C)=O	15.5	10.4	7.0	Poor	Good	Good	Poor	High vapor hazard and poor evaporation stability; low bioaccumulation/reactivity.
Chloroform	ClC(Cl)Cl	17.8	3.1	5.7	Poor	Medium	Good	Poor	High vapor hazard and poor stability; low reactivity.
Isopropyl ether	CC(C)OC(C)C	13.7	3.9	2.3	Medium	Medium	Medium	Poor	Medium hazard; poor evaporation stability; use with caution.
Bromobenzene	c1ccc(cc1)Br	20.5	5.5	4.1	Poor	Medium	Good	Good	High vapor hazard; good stability and low reactivity.
Ethanol	CCO	15.8	8.8	19.4	Good	Good	Medium	Poor	Low hazard and bioaccumulation; poor evaporation stability.
2-Butanone (MEK)	CCC(C)=O	16.0	9.0	5.1	Medium	Good	Medium	Poor	Medium hazard; low bioaccumulation; poor stability.
Butyl benzoate	CCCCOC(=O)c1ccccc1	18.3	5.6	5.5	Good	Poor (High)	Good	Good	Low vapor hazard; high bioaccumulation risk; good stability.
Methyl n-amyl ketone	CCCCCC(C)=O	16.2	5.7	4.1	Good	Medium	Medium	Good	Low vapor hazard; moderate bioaccumulation; good stability.
Anisole	COc1ccccc1	17.8	4.1	6.7	Poor	Medium	Poor	Good	High vapor hazard and high reactivity; good stability.
Benzyl acetate	CC(=O)OCc1ccccc1	18.3	5.7	6.0	Good	Medium	Medium	Good	Low vapor hazard; moderate bioaccumulation; good stability.
1,2-Dichloroethane	ClCCCl	19.0	7.4	4.1	Poor	Medium	Good	Poor	High vapor hazard; low reactivity; poor stability.
Dichloromethane	ClCCl	18.2	6.3	6.1	Poor (High)	Medium	Good	Poor (High)	High vapor hazard; highly volatile; poor stability.
Trichloroethylene	ClC=C(Cl)Cl	18.0	3.1	5.3	Poor	Medium	Good	Poor	High vapor hazard; low reactivity; poor stability.
Ethyl acetate	CCOC(C)=O	15.8	5.3	7.2	Medium	Good	Good	Poor	Medium vapor hazard; low bioaccumulation; poor stability.
Cyclohexane	C1CCCCC1	16.8	0.0	0.2	Poor	Poor	Medium	Poor	High vapor hazard and high bioaccumulation; poor stability.
Acetonitrile	CC#N	15.3	18.0	6.1	Poor	Good	Good	Poor	High vapor hazard; low bioaccumulation; poor stability.
Isopropyl acetate	CC(C)OC(C)=O	14.9	4.5	8.2	Medium	Medium	Good	Poor	Medium hazard; low reactivity; poor stability.
Methanol	CO	15.1	12.3	22.3	Medium	Good	Good	Poor	Medium vapor hazard; low reactivity; poor stability.
n-Butyl acetate	CCCCOC(C)=O	15.8	3.7	6.3	Medium	Medium	Good	Medium	Moderate performance across all indicators.
Methylcyclohexane	CC1CCCCC1	16.0	0.0	1.0	Medium	Poor	Medium	Medium	Medium hazard; high bioaccumulation; moderate stability.
Methyl isobutyl ketone	CC(C)CC(C)=O	15.3	6.1	4.1	Poor	Medium	Poor	Medium	High vapor hazard and high reactivity; moderate stability.
n-Heptane	CCCCCCC	15.3	0.0	0.0	Medium	Poor	Medium	Poor	High bioaccumulation; poor evaporation stability.
Tetrachloroethylene	ClC(Cl)=C(Cl)Cl	18.3	5.7	0.0	Medium	Poor	Good	Medium	High bioaccumulation; low reactivity; moderate stability.
o-Xylene	Cc1ccccc1C	17.8	1.0	3.1	Good	Poor	Poor	Medium	Low vapor hazard but high bioaccumulation and reactivity.
m-Xylene	Cc1cccc(C)c1	17.4	1.0	3.1	Medium	Poor	Poor	Medium	High bioaccumulation and reactivity.
Isopropanol (IPA)	CC(C)O	15.8	6.1	16.4	Medium	Good	Good	Poor	Low bioaccumulation; low reactivity; poor stability.
Ethyl 3-ethoxypropionate	CCOCCC(=O)OCC	16.2	3.3	8.8	Good	Medium	Medium	Good	Low vapor hazard; moderate bioaccumulation; good stability.
n-Butanol	CCCCO	16.0	5.7	15.8	Medium	Good	Poor	Medium	Low bioaccumulation; high reactivity; moderate stability.
Toluene	Cc1ccccc1	18.0	1.4	2.0	Poor	Medium	Poor	Medium	High vapor hazard and high reactivity.
Ethyl ether	CCOCC	14.5	2.9	5.1	Poor	Good	Medium	Poor	High vapor hazard and poor stability.
1,3-Propanediol	OCCCO	16.8	13.5	23.2	Good	Good	Medium	Good	Excellent safety profile; low hazard; good stability.
N,N-Diethylacetamide	CCN(CC)C(=O)C	16.4	11.3	7.5	Medium	Good	Medium	Good	Low bioaccumulation; good evaporation stability.
2-Pentanone	CCCC(C)=O	16.0	7.6	4.7	Medium	Good	Poor	Medium	Low bioaccumulation; high reactivity; moderate stability.
Tetrahydrofuran (THF)	C1CCOC1	16.8	5.7	8.0	Poor	Good	Poor	Poor	High vapor hazard and reactivity; poor stability.
Benzyl alcohol	OCc1ccccc1	18.4	6.3	13.7	Good	Medium	Poor	Good	Low vapor hazard; high reactivity; good stability.
N,N-Dimethylformamide	CN(C)C=O	16.8	11.5	10.2	Medium	Good	Poor	Good	Low bioaccumulation; high reactivity; good stability.
1,4-Dioxane	C1COCCO1	19.0	1.8	7.4	Poor	Good	Medium	Medium	High vapor hazard; moderate stability.
gamma-Butyrolactone	O=C1CCCO1	19.0	16.6	7.4	Medium	Good	Good	Good	Low bioaccumulation and reactivity; good stability.
Diacetone alcohol	CC(C)(O)CC(C)=O	15.8	8.2	10.8	Good	Good	Good	Good	Excellent overall safety and stability profile.
Cyclohexanone	O=C1CCCCC1	17.8	6.3	5.1	Medium	Good	Medium	Good	Low bioaccumulation; good evaporation stability.
2-Methylpyrazine	Cc1cnccn1	18.3	12.3	10.5	Poor	Medium	Poor	Medium	High vapor hazard and high reactivity.
N-Methylimidazole	Cn1ccnc1	19.7	15.6	11.2	Medium	Good	Poor	Good	Low bioaccumulation; high reactivity; good stability.
Ethyl lactate	CCOC(=O)C(C)O	16.0	7.6	12.5	Medium	Good	Medium	Good	Low bioaccumulation; good evaporation stability.
1,4-Butanediol	OCCCCO	16.6	15.3	21.7	Good	Good	Medium	Good	Low hazard; low bioaccumulation; good stability.
Diethanolamine	OCCNCCO	17.2	10.8	21.2	Medium	Good	Poor	Good	Low bioaccumulation; high reactivity; good stability.
Propylene glycol methyl ether	CC(O)COC	15.6	6.3	11.6	Medium	Good	Medium	Medium	Moderate performance across all safety indicators.
Dimethyl sulfoxide (DMSO)	CS(=O)C	18.4	16.4	10.2	Good	Good	Good	Good	Excellent safety profile; low hazard and high stability.
Methyl isobutyrate	COC(=O)C(C)C	15.1	3.7	6.3	Poor	Medium	Good	Poor	High vapor hazard and poor evaporation stability.
gamma-Valerolactone	CC1CCC(=O)O1	16.9	11.9	7.2	Medium	Good	Medium	Good	Moderate vapor hazard; good bioaccumulation and stability.
Water	O	15.5	16.0	42.3	Good	Good	Good	Poor	Low hazard/reactivity; poor evaporation stability.
Ethyl benzoate	CCOC(=O)c1ccccc1	18.2	5.4	6.0	Good	Medium	Good	Good	Low vapor hazard; low reactivity; good stability.
Phenethyl acetate	CC(=O)OCCc1ccccc1	18.2	5.3	6.1	Good	Medium	Good	Good	Low vapor hazard; low reactivity; good stability.
PGMEA	CC(=O)OCC(C)OC	15.6	5.5	7.3	Medium	Good	Good	Medium	Standard industrial solvent; moderate performance.
Methyl 3-ethoxypropionate	CCOCCC(=O)OC	16.1	3.6	8.6	Medium	Good	Good	Medium	Moderate vapor hazard; good safety profile.
Diglyme	COCCOCCOC	15.8	6.1	9.2	High	Low	Low	Medium	High vapor hazard (Reproductive toxicity); low reactivity.
n-Propyl lactate	CCCOC(=O)C(C)O	16.0	7.3	12.0	Medium	Low	Medium	Good	Low bioaccumulation; good evaporation stability.
Methyl levulinate	CC(=O)CCC(=O)OC	17.5	9.0	8.5	Medium	Low	Medium	Good	Low bioaccumulation; good evaporation stability.
Proglyme	CC(OC)COC(C)COC	15.7	6.1	6.5	Low	Low	Low	Medium	Best Diglyme alternative; low toxicity P-series ether; inert; BP 175°C.
Ethylene glycol diacetate	CC(=O)OCCOC(C)=O	16.2	5.5	9.6	Low	Low	Medium	Good	HSP matches Diglyme closely; good leveling; eco-friendly high BP solvent.
2-Methyltetrahydrofuran	CC1CCCO1	16.9	5.0	6.3	Medium	Low	Medium	Medium	Bio-based green solvent; moderate volatility; good stability.
Butyl cellosolve	CCCCOCCO	16.4	6.3	13.7	High	Medium	Good	Medium	High vapor hazard (Reproductive toxicity); moderate bioaccumulation.
epsilon-Caprolactone	O=C1CCCCCCO1	18.0	7.5	8.0	Medium	Medium	Medium	Medium	Moderate performance across all safety indicators.
Propylene carbonate	C1COC(=O)OC1C	18.2	13.0	9.0	Good	Good	Medium	Good	Low vapor hazard and bioaccumulation; good stability.

C. Hybrid Scoring Framework

To evaluate candidate developer formulations, we adopt a hybrid scoring framework that combines a physics-based metric with an LLM-based qualitative assessment. The total score is computed as

$$\mathrm{Score}_{\mathrm{total}} = 0.5\,\mathrm{Score}_{\mathrm{physics}} + 0.5\,\mathrm{Score}_{\mathrm{LLM}}. \tag{7}$$

The physics-based score captures thermodynamic compatibility, while the LLM-based score evaluates engineering feasibility and safety considerations.

C.1. Physics-Based Scoring

The physics-based score is derived from the RED values with respect to both the target layer S_{pre} and the protective layer S_{post}.

Solvency requirement for S_{pre}. A baseline score of 60 is assigned when

$$\mathrm{RED}_{pre} = 0.70. \tag{8}$$

For every decrease of 0.01 in RED, the score increases by one point:

$$\mathrm{Score}_{pre} = 60 + \frac{0.70 - \mathrm{RED}_{pre}}{0.01}. \tag{9}$$

Protection requirement for S_{post}. A baseline score of 60 is assigned when

$$\mathrm{RED}_{post} = 1.00. \tag{10}$$

For every increase of 0.01 in RED, the score increases by one point:

$$\mathrm{Score}_{post} = 60 + \frac{\mathrm{RED}_{post} - 1.00}{0.01}. \tag{11}$$

Combined physics score. The final physics score is computed as the average of solvency and protection:

$$\mathrm{Score}_{\mathrm{physics}} = \frac{\mathrm{Score}_{pre} + \mathrm{Score}_{post}}{2}. \tag{12}$$

C.2. LLM-Based Evaluation

The LLM-based score evaluates formulation quality from an engineering and practical perspective. It considers EHS risks, boiling-point distribution and gradient, thermodynamic plausibility of solvent mixtures, and functional roles of components. The LLM produces a holistic feasibility score reflecting robustness and manufacturability. This component is treated as a qualitative proxy, not as a ground-truth experimental metric.

D. Lithography Process and Measurement Details

In one embodiment, the photoresist solution was spin-coated onto a silicon wafer pre-treated with a bottom anti-reflective coating (BARC). The wafer was then pre-baked at 120°C for 60 s to yield a photoresist film with a thickness of approximately 100 nm. Exposure was performed using an ASML 1900Gi scanner with a 193 nm ArF light source and a numerical aperture of 1.35.

Following exposure, a post-exposure bake was conducted at 120°C for 60 s, followed by development using the tested developer formulation for 20 s. The resulting patterns were measured using a Hitachi CG6300 critical dimension scanning electron microscope. The preferred measurement parameters were: acceleration voltage of 400 V, probe current of 6.0 pA, magnification of 300K, and Linear detection mode. The detection area was set to 180 nm with 32 measurement points and a threshold of 45%.

D.1. Evaluation Criteria for Formulations #01–#15

The performance labels used in the main paper were assigned based on lithographic patterning results, including pattern visibility, pattern fidelity, and the presence of defects such as line collapse, bridging, or swelling under the corresponding process window.

Formulation #01 was evaluated at the T35P125 node under an exposure dose of 12 mJ/cm^2 and a focus setting of -0.05. Under these conditions, the measured critical dimension was 51.40 nm, the line width roughness was 3.27 nm, and the line edge roughness was 2.34 nm. Based on its pattern quality and quantitative metrics, formulation #01 was classified as *Excellent*.

Formulation #02 was evaluated at the T35P120 node under an exposure dose of 14 mJ/cm^2 and a focus setting of -0.05. The measured critical dimension was 56.01 nm, the line width roughness was 5.54 nm, and the line edge roughness was 3.83 nm. Based on the obtained lithographic performance, formulation #02 was also classified as *Excellent*.

For formulations #03 and #04, clear patterns could be observed at the T50P155 node. However, occasional line collapse was also present in the patterned features. These two formulations were therefore classified as *Suboptimal*.

For formulations #05–#12, the patterned images were generally distinguishable and relatively clear, indicating that development was successful to a certain extent. However, compared with formulations #03 and #04, these formulations exhibited more frequent defects, including line collapse, bridging, and swelling. These formulations were therefore classified as *Good*.

For formulations #13–#15, no clear lithographic patterns could be resolved after development, indicating unsuccessful pattern formation under the tested conditions. These formulations were therefore classified as *Poor*.

E. Solvent Diversity Metrics

To quantify the diversity of solvent usage across methods, we use Shannon entropy. For a dataset with n unique solvents, let c_i denote the occurrence count of solvent i, where $i = 1, 2, \ldots, n$. The relative frequency of solvent i is

$$p_i = \frac{c_i}{\sum_{j=1}^{n} c_j}. \tag{13}$$

The Shannon entropy H is

$$H = -\sum_{i=1}^{n} p_i \log_2 p_i, \tag{14}$$

with the convention $0 \log_2 0 = 0$. The entropy reaches its maximum $H_{\max} = \log_2 n$ when all solvents are used equally often. Lower entropy indicates a more concentrated usage pattern.

We also report two complementary metrics:

1. **Number of unique solvents** n: the number of distinct solvents appearing at least once.
2. **Top-k concentration ratio** C_k: the cumulative usage proportion of the k most frequently occurring solvents,

$$C_k = \frac{\sum_{i=1}^{k} c_{(i)}}{\sum_{j=1}^{n} c_j} \times 100\%, \tag{15}$$

   where $c_{(i)}$ is the i-th largest occurrence count. We set $k = 5$.

Together, these metrics capture richness, uniformity, and concentration of solvent usage.

F. Algorithm and Complexity Analysis

F.1. Context Management and Computational Complexity

A fundamental limitation of applying standard LLM agents to long-horizon search is the growth of prompt length and attention cost. In a naive MCTS-style LLM agent, each node may retain the complete conversational history. For a search

Algorithm 1 `LithographyAgent`: Neuro-Symbolic Monte Carlo Tree Search

Require: Initial design objective x_{root}; maximum iterations T; maximum children per node K
Ensure: Best solvent formulation $(\mathcal{M}^*, \boldsymbol{\phi}^*)$

Initialize root node v_0 with x_{root}
for $t = 1$ **to** T **do**
  Selection phase
  $v_{leaf} \leftarrow v_0$
  while v_{leaf} is fully expanded, i.e., $|\text{children}(v_{leaf})| \geq K$ **do**
    $v_{leaf} \leftarrow \arg\max_{c \in \text{children}(v_{leaf})} \left(\frac{Q(c)}{N(c)} + C\sqrt{\frac{\ln N(v_{leaf})}{N(c)}} \right)$
  end while
  Sparse expansion and dynamic path reconstruction
  $\mathcal{P} \leftarrow \text{TracePath}(v_0 \rightarrow v_{leaf})$
  $\mathcal{H}_{siblings} \leftarrow \text{SummarizeActions}(\text{children}(v_{leaf}))$
  $x_{input} \leftarrow \text{Concat}(\text{RootPlan}, \text{Summary}(\mathcal{P}), \text{NegativeConstraints}(\mathcal{H}_{siblings}))$
  $a_{new} \leftarrow \text{GeneratorMoE}(x_{input})$ {Discrete topology proposal}
  Differentiable physics optimization
  $\boldsymbol{\phi}_{phys} \leftarrow \arg\min_{\boldsymbol{\phi}} L_{hybrid}(a_{new}, \boldsymbol{\phi})$
  $(\mathcal{M}^*, \boldsymbol{\phi}^*) \leftarrow \text{EngineeringReview}(a_{new}, \boldsymbol{\phi}_{phys})$
  Evaluation
  $r \leftarrow \text{StatelessCritic}(\mathcal{M}^*, \boldsymbol{\phi}^*)$
  Create node v_{new} storing $(\mathcal{M}^*, \boldsymbol{\phi}^*, r, N = 0, Q = 0)$
  $\text{children}(v_{leaf}) \leftarrow \text{children}(v_{leaf}) \cup \{v_{new}\}$
  Backpropagation
  $v_{curr} \leftarrow v_{new}$
  while $v_{curr} \neq \text{NULL}$ **do**
    $N(v_{curr}) \leftarrow N(v_{curr}) + 1$
    $Q(v_{curr}) \leftarrow Q(v_{curr}) + r$
    $v_{curr} \leftarrow \text{Parent}(v_{curr})$
  end while
  Memory consolidation
  $\text{VectorDB.Store}(\mathcal{M}^*, \boldsymbol{\phi}^*, r)$
end for
return $\text{GetBestSolution}(v_0)$

path of depth d and average interaction length L, the context required for a leaf node scales as $\mathcal{O}(dL)$, which can quickly exhaust the context window.

`LithographyAgent` reduces this burden through sparse state storage and dynamic path reconstruction. Each tree node stores only a lightweight semantic tuple, such as (a, r, n, Q), requiring $\mathcal{O}(1)$ storage per node. During expansion, the prompt is reconstructed from the root plan, compact path summaries, and sibling constraints. The resulting context length scales with summarized path information rather than full logs, approximately $\mathcal{O}(L_{root} + dL_{summary})$ with $L_{summary} \ll L$.

In implementation, the prompt can also be bounded by truncating or summarizing older path information. Thus, LLM inference cost per expansion depends on the bounded reconstructed prompt rather than the global size of the search tree. Continuous ratio optimization is handled by the physics module and scales as $\mathcal{O}(IE)$, where I is the number of gradient steps and E is the cost of evaluating the hybrid loss. This decoupling keeps the expensive LLM component focused on compact discrete hypothesis generation while delegating numerical feasibility to deterministic optimization.

G. Experimental validation of AI-generated developer formulations

To further evaluate the effectiveness of the LithographyAgent-generated solvent formulations, thirteen candidate developer formulations were experimentally tested under different post-treatment/development conditions. The detailed compositions of the tested formulations are summarized in Supplementary Table 4.

Among all tested candidates, formulations F10 and F12 exhibited the most promising development performance. Formulation F10, consisting of benzyl acetate, methyl n-amyl ketone, and butyl benzoate, produced clearly resolved patterns under the condition of T35P125. Formulation F12, composed of 2-pentanone, 3-ethoxypropionate, methyl n-amyl ketone, and anisole, further improved the pattern fidelity under a slightly milder condition of T35P120. Compared with the commercial developer control, both F10 and F12 generated more distinguishable patterned regions with improved feature definition, indicating that the AI-generated formulations can outperform the conventional baseline under optimized conditions.

Interestingly, not all ester/ketone-containing formulations resulted in successful development. For example, formulations F8 and F13, although chemically related to the successful candidates and containing typical ester/ketone solvent motifs, failed to produce recognizable developed patterns. Formulation F8 contained benzyl acetate, butyl benzoate, and 2-butanone, while F13 contained butyl benzoate, methyl n-amyl ketone, and 2-butanone. Their poor development results indicate that the development performance is not governed by solvent class alone. Instead, the results suggest that a narrow and nonlinear compositional window is required to achieve effective development.

Other formulations, including F1, F2, F4, F5, F6, F7, F9, and F11, showed varying degrees of development capability. Among them, F5 produced relatively clear line patterns; however, it required a substantially stronger processing condition of T80P240, making direct comparison with F10 and F12 less straightforward. In contrast, F10 and F12 achieved high-quality development under much milder conditions, highlighting their practical advantage in terms of process serviceability.

These experimental observations support three important conclusions. First, the LithographyAgent framework does not merely reproduce broad, conventional solvent categories such as generic ketone/ester mixtures. Second, the sharp contrast between successful formulations F10 and F12 and failed formulations F8 and F13 demonstrates that high-resolution development is controlled by subtle, composition-dependent solvent interactions. Third, the successful identification of F10 and 12 suggests that LithographyAgent can effectively search within a chemically crowded formulation space and identify high-performing narrow subsets that may be difficult to discover through manual formulation screening alone.

Overall, the experimental validation confirms that the AI-generated formulations provide meaningful improvements over the commercial developer baseline. More importantly, the observed failure–success contrast among chemically similar candidates provides direct evidence for the credibility and interpretability of the AI-guided formulation strategy.

Table 4. Compositions of experimentally tested AI-generated developer formulations. The experimental formulation IDs F1–F13 are used independently from the candidate IDs in the main-text proxy-score analysis.

Exp. ID	Component 1	Ratio/%	Component 2	Ratio/%	Component 3	Ratio/%	Component 4	Ratio/%
F1	Bromobenzene	50	2-Butanone	30	Toluene	20	–	–
F2	Bromobenzene	70	2-Butanone	30	–	–	–	–
F3	Bromobenzene	80	n-Butyl acetate	20	–	–	–	–
F4	1,2-Dichloroethane	60	Toluene	25	2-Butanone	15	–	–
F5	Benzyl acetate	35	Butyl benzoate	15	2-Butanone	30	Toluene	20
F6	Benzyl acetate	50	2-Butanone	30	Toluene	20	–	–
F7	Butyl benzoate	50	2-Butanone	30	Toluene	20	–	–
F8	Benzyl acetate	49	Butyl benzoate	21	2-Butanone	30	–	–
F9	Butyl benzoate	70	2-Butanone	30	–	–	–	–
F10	Benzyl acetate	40	Methyl n-amyl ketone	30	Butyl benzoate	30	–	–
F11	n-Butyl acetate	20	Benzyl acetate	70	Cyclohexane	10	–	–
F12	2-Pentanone	5	3-Ethoxypropionate	20	Methyl n-amyl ketone	30	Anisole	45
F13	Butyl benzoate	70	Methyl n-amyl ketone	20	2-Butanone	10	–	–