Adaptive Human-in-the-Loop Optimization Using Language-Guided Priors for Chemical Synthesis

Amirreza Mottafegh [1] **Gwang-Noh Ahn** [1]

[1]*Digital Chemical Research Center, Korea Research Institute of Chemical Technology, Daejeon, Republic of Korea.*
Correspondence to: Amirreza Mottafegh armtfgh@krict.re.kr, Gwang-Noh Ahn gnahn@krict.re.kr.

1. Introduction

The acceleration of discovery has been transformed by Self-Driving Laboratories (SDLs) and automated High-Throughput Experimentation (HTE) [1]. By integrating robotic execution with Bayesian Optimization (BO), these platforms can navigate complex chemical spaces efficiently. However, a critical inefficiency remains: the cold start problem.

Standard autonomous workflows typically treat every new campaign as a *tabula rasa*, ignoring the vast reservoir of institutional knowledge and heuristic intuition possessed by domain experts. This exclusion stems from the translation gap (converting natural language intuition into mathematical priors) and the trust gap (fear that human bias will lock the algorithm into suboptimal regions).

In this work, we introduce the Human-in-the-Loop Optimization (HILO) framework. Our approach utilizes Large Language Models (LLMs) to translate unstructured expert text into structured priors for Gaussian Processes (GPs). Crucially, to address the trust gap, we introduce the Adaptive Weight / Credibility Detection (AWCD) mechanism. Functioning as a probabilistic safety trigger. AWCD continuously evaluates the validity of expert constraints against observed data, effectively assessing the divide between expert intuition and autonomous discovery.

2. Methodology

The HILO framework operates through two primary stages: Language-Driven Prior Specification and Ensemble Bayesian Optimization with AWCD.

2.1 Language-Driven Priors

Experts often express knowledge in natural language (e.g., "yields peak around 130°C"). We employ an LLM to parse these inputs into a structured JSON schema, which constructs a prior mean function $m_0(\mathbf{x})$. This function is an additive composition of three primitives:

1. **Monotone Trends:** Sigmoid functions for parameters believed to increase/decrease the objective.
2. **Unimodal Peaks:** Gaussian radial basis functions (RBF) for specific sweet spots.
3. **Constraints:** Soft penalty functions representing negative knowledge (where *not* to look).

2.2 The AWCD Safety Feature

To prevent prior negative effect, we employ an ensemble strategy maintaining two hypotheses: a Skeptic (purely data-driven zero-mean GP) and a Believer (prior-augmented GP). The credibility weight $w \in [0, 1]$ is dynamically updated based on the Marginal Log-Likelihood (MLL) of the observed data.

The AWCD mechanism monitors "constraint pressure"—instances where the data-driven model predicts high potential in regions forbidden by the expert. If the data contradicts the prior (e.g., high yield discovered in a forbidden zone), the algorithm trips the safety feature ($w \to 0$), disabling the prior and recovering the unbiased performance of standard BO.

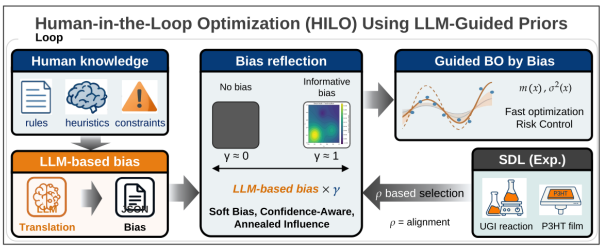

Fig. 1: **Overview of the HILO Framework.** (Left) LLM translates text to structured priors. (Middle) Construction of the mean function $m(x)$. (Right) The AWCD safety mechanism acts as a circuit breaker if data contradicts the prior.

3. Results and Discussion

We validated the framework against two distinct optimization modalities: the continuous parameter search of the Ugi multicomponent reaction (using a dataset of ~15,000 points) [2] and the finite library screening of **P3HT-CNT composites** (233 experimental points) [3].

3.1 Acceleration and Efficiency

In both continuous and discrete landscapes, the inclusion of language-guided priors significantly resolved the cold start problem. The HILO framework demonstrated a ~3x reduction in experimental iterations required to reach 90% of the global optimum compared to baseline BO. Notably, in high-knowledge regimes (expert providing priors based on 90% of historical data), the algorithm with a single initialization point ($n_{init} = 1$) matched the performance of a standard BO initialized with 40 random experiments.

3.2 Robustness to "Bad" Knowledge

A key contribution of this work is safety against deceptive priors. We simulated a Bad Knowledge scenario where the expert explicitly forbade the region containing the global optimum.

- **Unguarded BO:** Permanently failed to locate the optimum, locked out by the false constraint.
- **HILO with AWCD:** Successfully detected the conflict between data and prior. The circuit breaker was triggered at approximately iteration 10, disabling the false constraint and allowing the system to converge to the true global optimum.

4. Conclusion

We have presented a robust framework for integrating natural language expert intuition into autonomous chemical synthesis. By combining LLM-based translation with the novel AWCD safety mechanism, HILO resolves the cold start problem without sacrificing the theoretical convergence guarantees of Bayesian Optimization. This approach transforms expert knowledge from a passive resource into an active, safe, and efficient accelerator for materials discovery.

Acknowledgments

This research was supported by a major research project from the Korea Research Institute of Chemical Technology (KK2551-10 and BSK25-141) and the Nano Material Technology Development Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Science and ICT (RS-2024-00446683).

References

[1] Gary Tom, Stefan P. Schmid, Sterling G. Baird, Yang Cao, Kourosh Darvish, Han Hao, Stanley Lo, Sergio Pablo-García, Ella M. Rajaonson, Marta Skreta, Naruki Yoshikawa, Samantha Corapi, Gun Deniz Akkoc, Felix Strieth-Kalthoff, Martin Seifrid, and Alán Aspuru-Guzik. Self-driving laboratories for chemistry and materials science. *Chemical Reviews*, 124(16):9633–9732, 2024. PMID: 39137296.

[2] Yankai Jia, Rafał Frydrych, Yaroslav I Sobolev, Wai-Shing Wong, Bibek Prajapati, Daniel Matuszczyk, Yasemin Bilgi, Louis Gadina, Juan Carlos Ahumada, Galymzhan Moldagulov, et al. Robot-assisted mapping of chemical reaction hyperspaces and networks. *Nature*, 645(8082):922–931, 2025.

[3] Daniil Bash, Yongqiang Cai, Vijila Chellappan, Swee Liang Wong, Xu Yang, Pawan Kumar, Jin Da Tan, Anas Abutaha, Jayce JW Cheng, Yee-Fun Lim, Siyu Isaac Parker Tian, Zekun Ren, Flore Mekki-Berrada, Wai Kuan Wong, Jiaxun Xie, Jatin Kumar, Saif A. Khan, Qianxiao Li, Tonio Buonassisi, and Kedar Hippalgaonkar. Multi-fidelity high-throughput optimization of electrical conductivity in p3ht-cnt composites. *Advanced Functional Materials*, 31(36):2102606, 2021.