

# La Agente Optima – orchestrated Bayesian optimization and active learning for accelerated *in silico* compound discovery

Marcel Müller<sup>1,2</sup> Jiaru Bai<sup>2,3</sup> Willi Gottstein<sup>4</sup> Jan Gerit Brandenburg<sup>5\*</sup> Sissi Feng<sup>4\*</sup>  
 Martin Fitzner<sup>5\*</sup> Sergio Pablo-García<sup>4\*</sup> Jesús Valdés Hernández<sup>1\*</sup> Ilya Yakavets<sup>4\*</sup> Yimu Zhao<sup>4</sup>  
 Varinia Bernales<sup>3,4</sup> Alán Aspuru-Guzik<sup>1,2,3,4,6,7,8,9</sup>

\*Equal contribution <sup>1</sup>Department of Chemistry, University of Toronto, 80 St. George St., Toronto, ON M5S 3H6, Canada <sup>2</sup>Vector Institute for Artificial Intelligence, W1140-108 College St., Schwartz Reisman Innovation Campus, Toronto, ON M5G 0C6, Canada <sup>3</sup>Department of Computer Science, University of Toronto, 40 St George St., Toronto, ON M5S 2E4, Canada <sup>4</sup>Acceleration Consortium, 700 University Ave., Toronto, ON M7A 2S4, Canada <sup>5</sup>Merck KGaA, Frankfurter Str. 250, 64293 Darmstadt, Germany <sup>6</sup>Department of Materials Science & Engineering, University of Toronto, 184 College St., Toronto, ON M5S 3E4, Canada <sup>7</sup>Department of Chemical Engineering & Applied Chemistry, University of Toronto, 200 College St., Toronto, ON M5S 3E5, Canada <sup>8</sup>Canadian Institute for Advanced Research (CIFAR), 661 University Ave., Toronto, ON M5G 1M1, Canada <sup>9</sup>NVIDIA, 431 King St. W #6th, Toronto, ON M5V 1K4, Canada. Correspondence to: Varinia Bernales [varinia@bernales.org](mailto:varinia@bernales.org), Alán Aspuru-Guzik [alan@aspuru.com](mailto:alan@aspuru.com).

## 1. Abstract

Bayesian optimization (BO) plays a key role in data-efficient exploration and optimization for chemistry, materials science, and other scientific fields, particularly in settings where objective function evaluations are expensive [2, 3, 4]. Such evaluations are usually time-intensive and/or costly experiments, which can be noisy, asynchronous, and heterogeneous in fidelity, as encountered in physical experiments or high-accuracy computational simulations [5, 6]. BO formalizes closed-loop discovery as a problem of sequential decision-making under uncertainty, in which probabilistic beliefs over objective functions are updated via Bayesian inference and experimental actions are selected to optimize expected outcomes. In practice, BO encompasses a broad design space of models, priors, acquisition strategies, and optimization choices, whose effective coordination is critical in realistic closed-loop discovery settings [7].

At present, such workflows, where optimization algorithms are directly coupled to experimental evaluations, are typically realized through complex, manually assembled pipelines. This requires substantial technical expertise across multiple layers of the stack, spanning electronic-structure software, experimental platforms, and the coordination of optimization algorithms with orchestration engines [8, 9].

In parallel, recent advances in artificial intelligence (AI) agentic systems, driven by large language models (LLMs), have enabled the autonomous orchestration of complex scientific workflows by combining reasoning, tool use, and domain-specific execution modules, such as electronic-structure codes or laboratory automation interfaces [10]. As demonstrated in recent works from our lab with EL AGENTE Q and more recently EL AGENTE GRÁFICO, complex quantum chemistry (QC) workflows can be executed fully autonomously end-to-end through natural-

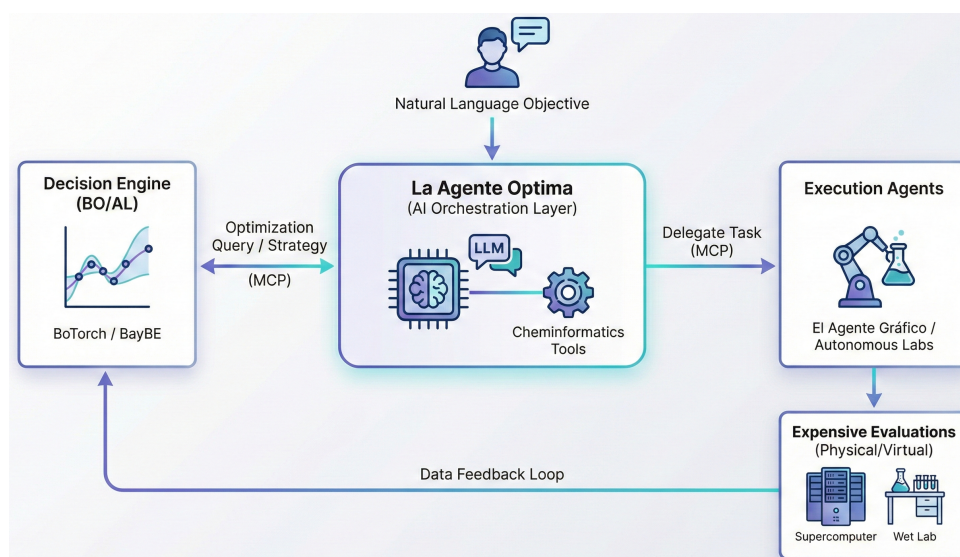


Fig. 1: Schematic of the LA AGENTE OPTIMA architecture [1], bridging Bayesian optimization (BO) and active learning (AL) frameworks through a central agentic artificial intelligence (AI) orchestration layer. The system translates natural-language objectives into closed-loop discovery workflows, delegating tasks to heterogeneous execution agents via model context protocol (MCP).

language interaction with a multi-agent system [11, 12]. This concept has since been extended to additional domains, including quantum computing [13], drug discovery [14], solid-state chemistry, and others.

Despite sharing a common abstraction of sequential decision-making under uncertainty, BO frameworks and agentic AI systems have largely evolved independently and have not yet been systematically integrated within unified discovery frameworks.

Here, we outline an early-stage concept to bridge these two paradigms (see Fig. 1). BO frameworks such as BoTorch [15] and BayBE [3], together with active learning (AL) algorithms [16], are integrated via a central agentic AI orchestration layer: LA AGENTE OPTIMA translates high-level scientific objectives into optimization and exploration queries handled by the BO/AL engines, while expensive function evaluations are delegated to execution agents such as EL AGENTE GRÁFICO for electronic structure calculations. For the BO backend, we employ and compare BoTorch and BayBE, both exposed through in-house hosted Model Context Protocols (MCPs) servers. Results from expensive physical or virtual evaluations are asynchronously ingested by the decision engine via a dedicated data pipeline, allowing BO or AL models to update independently of the central agent, which only intermittently polls or supervises the feedback loop. By equipping LA AGENTE OPTIMA with cheminformatics and electronic structure tools, the system can reason about appropriate (macro-)molecular representations and optimization strategies.

We demonstrate the capabilities of OPTIMA through four diverse case studies. The first one is a fully autonomous, natural-language-orchestrated optimization of an emitter material, inspired by Ref. [8]. We further apply BO to conformer searches in small molecules, using a search space defined over internal coordinates. We then tackle a challenging catalyst optimization task that extends beyond standard optimization and routine electronic-structure calculations by requiring secondary, LLM-driven post-processing steps. In collaboration with a Human Organ Mimicry Self-Driving Laboratory at the Acceleration Consortium, we also demonstrate fully autonomous, multi-objective optimization of a biofabrication workflow using purely natural-language orchestration [17]. Beyond these case studies, OPTIMA naturally extends to applications such as force-field and semiempirical method parameter calibration, as well as multi-objective materials screening [e.g., for metal-organic frameworks *aka* MOFs].

## 2. Conclusion

LA AGENTE OPTIMA enables closed-loop scientific discovery workflows by coupling agentic reasoning with Bayesian decision-making over domain-specific evaluation functions. While demonstrated here primarily in computational chemistry, we also present early-stage experimental orchestration, and the architecture is designed to generalize beyond a

single scientific domain. This work marks an initial design and collaboration phase, laying the groundwork for broader validation, benchmarking, and joint development across multiple research groups, with the longer-term goal of enabling interconnected, AI-driven scientific workflows across institutions. As part of this effort, the EL AGENTE ecosystem [11, 13, 14] will be made available through a shared platform, with community access and registration provided via [elagente.ca](http://elagente.ca).

## 3. Related work

LLMs have also been leveraged to warm-start BO workflows, for example, by improving candidate initialization strategies or by imposing structured priors on the search space [18, 19]. Han et al. showed accelerated BO using LLM-based multi-agent systems [20]. BORA integrates LLMs to enhance BO with domain knowledge and reasoning capabilities of LLMs [21]. As discussed in the main text, a growing body of work employs multi-agent or agentic AI systems to orchestrate scientific tools and workflows. Additional representative examples include Refs. [22, 23, 24].

## Acknowledgments

We gratefully acknowledge the longstanding contributions of the Matter Lab’s current and past group members ([matter.toronto.edu](http://matter.toronto.edu)), in particular the El Agente subgroup. This research is part of the University of Toronto’s Acceleration Consortium, which receives funding from the CFREF-2022-00042 Canada First Research Excellence Fund. This work was supported by the AI2050 program of Schmidt Sciences. This research was enabled in part by support provided by the SciNet HPC Consortium ([scinethpc.ca](http://scinethpc.ca)) and the Digital Research Alliance of Canada ([alliancecan.ca](http://alliancecan.ca)).

## References

- [1] Google. Gemini Nano Banana Pro (Gemini 3 Thinking). Large language model, 2026.
- [2] Sterling Baird, Mehrad Ansari, Zartashia Afzal, Qianxiang Ai, Alexander Al-Feghali, Mathieu Alain, Matias Altamirano, Thomas Andrews, Andy Sode Anker, Rija Ansari, Samuel Ampofo Appiah, Raul Astudillo, Ruhana Azam, Mohammed Azzouzi, Suneel Kumar Bvs, Ben Blaiszik, Anna Borisova, Andres Bran, Pengfei Cai, Ting-Yeh Chen, Curtis Chong, Samantha Corapi, Mark Croxall, Gbetondji Dovonon, Jose Manuel Napoles Duarte, Andrew Falkowski, Giuseppe Fisicaro, Martin Fitzner, Quinn Gallagher, Sabah Gaznaghi, Jerome Genzling, Christoph Griebel, Ryan-Rhys Griffiths, Taicheng Guo, Kehan Guo, Nipun Gupta, Ankur Gupta, Mohammad Haddadnia, Yuyang Han, Joscha Hoche, Alexander V. Hopp, Marko Huang, Ayodeji Ijishakin, Ramsey Issa, Yeonghun Kang, Jungtaek Kim, Akshay Kudva, Ruben Laplaza,

- Magdalena Lederbauer, Shi Xuan Leong, Paul W. Leu, Viola Muning Li, Mingxuan Li, Tao Liu, Stanley Lo, Jakub Lala, Osman Mamun, Owen Melville, Michail Mitsakis, Cameron Movassaghi, Madhav Reddy Muthyala, Marcel Muller, Bozhao Nan, Duc Nguyen, Daniele Ongari, Anthony Onwuli, Can Ozkan, Sergio Pablo-Garcia, Elton Pan, Ratish Panda, Sean Park, Jaehee Park, Dieter Plessers, Tobias Plotz, Ella M. Rajaonson, Bojana Rankovic, Jarett Ren, Rim Rihana, Jurgis Ruza, Akhil S. Nair, Carter Salbego, Erick Lopez Saldivar, Arifn San, Christina Schenk, Stefan P. Schmid, Dylan Schubert, Philippe Schwaller, Cher-Tian Ser, Maitreyee Sharma Priyadarshini, Yuxin Shen, Kevin Shen, Jiale Shi, Farshud Sorourifar, Adrian Susic, Taylor Sparks, Jan Christopher Spies, Felix Strieth-Kalthoff, Suraj Sudhakar, Aditya Sundar, Alessio Tamburro, Clara Tamura, Yifeng Tang, Dandan Tang, Nikhil Thota, Mohammad Erfan Toloue Sadegh Azadi, Gary Tom, Sang Truong, Ricardo Valencia Albornoz, Luis Walter, Lawrence Wang, Fanjin Wang, Andrew Wang, Yiran Wang, Jeffrey Watchorn, Benjamin Weiser, Geemi Wellawatte, Alexander Wiczorek, Tim Wurger, Ilya Yakavets, Jakob Zeitler, Sylvester Zhang, Yimu Zhao, Yanqiao Zhu, Ruijie Zhu, and Yunheng Zou. Bayesian Optimization Hackathon for Chemistry and Materials, June 2025.
- [3] Martin Fitzner, Adrian Šošić, Alexander V. Hopp, Marcel Müller, Rim Rihana, Karin Hrovatin, Fabian Liebig, Mathias Winkel, Wolfgang Halter, and Jan Gerit Brandenburg. BayBE: A Bayesian Back End for experimental planning in the low-to-no-data regime. *Digital Discovery*, 4(8):1991–2000, August 2025.
- [4] Johannes P. Dürholt, Thomas S. Asche, Johanna Kleinekorte, Gabriel Mancino-Ball, Benjamin Schiller, Simon Sung, Julian Keupp, Aaron Osburg, Toby Boyne, Ruth Misener, Rosona Eldred, Wagner Steuer Costa, Chrysoula Kappatou, Robert M. Lee, Dominik Linzner, David Walz, Niklas Wulkow, and Behrang Shafei. BoFire: Bayesian Optimization Framework Intended for Real Experiments, August 2024.
- [5] Tom Savage, Nausheen Basha, Jonathan McDonough, James Krassowski, Omar Matar, and Ehecatl Antonio del Rio Chanona. Machine learning-assisted discovery of flow reactor designs. *Nature Chemical Engineering*, 1(8):522–531, August 2024.
- [6] Maituo Yu, Shuyang Yang, Chunzhi Wu, and Noa Marom. Machine learning the Hubbard U parameter in DFT+U using Bayesian optimization. *npj Computational Materials*, 6(1):180, November 2020.
- [7] Roman Garnett. *Bayesian optimization*. Cambridge University Press, 2023.
- [8] Felix Strieth-Kalthoff, Han Hao, Vandana Rathore, Joshua Derasp, Théophile Gaudin, Nicholas H. Angello, Martin Seifrid, Ekaterina Trushina, Mason Guy, Junliang Liu, Xun Tang, Masashi Mamada, Wesley Wang, Tuul Tsagaantsooj, Cyrille Lavigne, Robert Pollice, Tony C. Wu, Kazuhiro Hotta, Leticia Bodo, Shangyu Li, Mohammad Haddadnia, Agnieszka Wołos, Rafał Roszak, Cher Tian Ser, Carlota Bozal-Ginesta, Riley J. Hickman, Jenya Vestfrid, Andrés Aguilar-Granda, Elena L. Klimareva, Ralph C. Sigerson, Wenduan Hou, Daniel Gahler, Slawomir Lach, Adrian Warzybok, Oleg Borodin, Simon Rohrbach, Benjamin Sanchez-Lengeling, Chihaya Adachi, Bartosz A. Grzybowski, Leroy Cronin, Jason E. Hein, Martin D. Burke, and Alán Aspuru-Guzik. Delocalized, asynchronous, closed-loop discovery of organic laser emitters. *Science*, 384(6697):eadk9227, May 2024.
- [9] Jiaru Bai, Sebastian Mosbach, Connor J. Taylor, Dogancan Karan, Kok Foong Lee, Simon D. Rihm, Jethro Akroyd, Alexei A. Lapkin, and Markus Kraft. A Dynamic Knowledge Graph Approach to Distributed Self-Driving Laboratories. *Nature Communications*, 15:462, 2024.
- [10] Hongliang Xin, John R. Kitchin, and Heather J. Kulik. Towards agentic science for advancing scientific discovery. *Nature Machine Intelligence*, pages 1–3, September 2025.
- [11] Yunheng Zou, Austin H. Cheng, Abdulrahman Aldossary, Jiaru Bai, Shi Xuan Leong, Jorge Arturo Campos-Gonzalez-Angulo, Changhyeok Choi, Cher Tian Ser, Gary Tom, Andrew Wang, Zijian Zhang, Ilya Yakavets, Han Hao, Chris Crebolder, Varinia Bernales, and Alán Aspuru-Guzik. El Agente: An autonomous agent for quantum chemistry. *Matter*, 8(7):102263, July 2025.
- [12] Jiaru Bai, Abdulrahman Aldossary, Thomas Swanick, Zijian Zhang, Jin Won Lee, Marcel Müller, Tsz Wai Ko, Yeonghun Kang, Varinia Bernales, and Alán Aspuru-Guzik. El agente gráfico: Type-safe abstractions for agentic scientific execution. Manuscript in preparation, 2026.
- [13] Ignacio Gustin, Luis Mantilla Calderón, Juan B. Pérez-Sánchez, Jérôme F. Gonthier, Yuma Nakamura, Karthik Panicker, Manav Ramprasad, Zijian Zhang, Yunheng Zou, Varinia Bernales, and Alán Aspuru-Guzik. El Agente Cuántico: Automating quantum simulations, January 2026.
- [14] Marcel Müller, Mohammad Ghazi Vakili, Ella Miray Rajaonson, Ella Ashoori, Yunheng Zou, Varinia Bernales, and Alán Aspuru-Guzik. El Agente Fármaco, 2026. Manuscript in preparation.
- [15] Maximilian Balandat, Brian Karrer, Daniel R. Jiang, Samuel Daulton, Benjamin Letham,

Andrew Gordon Wilson, and Eytan Bakshy. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization, December 2020.

- [16] Tivadar Danka and Peter Horvath. modAL: A modular active learning framework for Python, December 2018.
- [17] Kevin Angers, Kouros Darvish, Naruki Yoshikawa, Sargol Okhovatian, Dawn Bannerman, Ilya Yakavets, Florian Shkurti, Alán Aspuru-Guzik, and Milica Radisic. RoboCulture: A Robotics Platform for Automated Biological Experimentation, May 2025.
- [18] Tension Liu, Nicolás Astorga, Nabeel Seedat, and Mihaela van der Schaar. Large Language Models to Enhance Bayesian Optimization, March 2024.
- [19] Andrej Schwanke, Lyubomir Ivanov, David Salinas, Fabio Ferreira, Aaron Klein, Frank Hutter, and Arber Zela. Improving LLM-based Global Optimization with Search Space Partitioning, May 2025.
- [20] Dong Han, Zhehong Ai, Pengxiang Cai, Shanya Lu, Jianpeng Chen, Zihao Ye, Shuzhou Sun, Ben Gao, Lingli Ge, Weida Wang, Xiangxin Zhou, Xihui Liu, Mao Su, Wanli Ouyang, Lei Bai, Dongzhan Zhou, Tao Xu, Yuqiang Li, and Shufei Zhang. ChemBOMAS: Accelerated BO in Chemistry with LLM-Enhanced Multi-Agent System, November 2025.
- [21] Abdoulatif Cissé, Max E Cooper, Mengjia Zhu, Xenophon Evangelopoulos, and Andrew I Cooper. Can We Automate Scientific Reasoning in Closed-Loop Experiments using Large Language Models?, January 2026.
- [22] Behnam Yousefi, Nora Constanze Laubach, Sven Heins, Lucia Testa, Søren W. Gersting, and Stefan Bonn. BioChemAgent: An AI-driven Protein Modeling and Docking Framework for Structure-Based Drug Discovery, December 2025.
- [23] Aikaterini Vriza, Uma Kornu, Aditya Koneru, Henry Chan, and Subramanian K. R. S. Sankaranarayanan. Multi-agentic AI framework for end-to-end atomistic simulations. *Digital Discovery*, December 2025.
- [24] Jaehyung Lee, Justin Ely, Kent Zhang, Akshaya Ajith, Charles Rhys Campbell, and Kamal Choudhary. AGAPI-Agents: An Open-Access Agentic AI Platform for Accelerated Materials Design on AtomGPT.org, December 2025.