

# LLMs for Optimization: From Automated Modeling to Algorithmic Discovery

Connor Lawless<sup>1</sup>, Fei Liu<sup>2,3</sup>, Hanzhang Qin<sup>4</sup> and Ellen Vitercik<sup>1</sup>

<sup>1</sup>Stanford University

<sup>2</sup>University of Zurich, <sup>3</sup>ETH Zurich

<sup>4</sup>National University of Singapore

lawlessc@stanford.edu, fei.liu2@uzh.ch, hzqin@nus.edu.sg, vitercik@stanford.edu

## 1 Description

2 This tutorial surveys how large language models (LLMs)  
3 are reshaping mathematical optimization, from translating  
4 natural-language problem descriptions into formal optimiza-  
5 tion models to assisting with solver configuration and al-  
6 gorithm design. It introduces the core ideas, representative  
7 systems, evaluation challenges, and open research questions  
8 needed for a broad AI audience to understand this rapidly  
9 emerging interface between LLMs and optimization.

## 10 Overview

11 Mathematical optimization is a foundational pillar of mod-  
12 ern artificial intelligence, underpinning a wide range of real-  
13 world decision-making systems. It plays a central role in ap-  
14 plications such as training machine learning models, supply  
15 chain and logistics planning, energy systems, finance, and  
16 scheduling, where complex objectives and constraints must  
17 be translated into actionable solutions. Despite its impor-  
18 tance, building and deploying optimization models remains  
19 a challenging, expert-driven process that requires significant  
20 domain knowledge and technical expertise. Recent advances  
21 in LLMs offer a transformative opportunity to lower these  
22 barriers by enabling more intuitive, automated, and scalable  
23 interaction with optimization systems.

24 This tutorial provides a comprehensive overview of the  
25 emerging intersection between LLMs and optimization,  
26 structured around two synergistic themes. First, we exam-  
27 ine how LLMs can act as powerful copilots across the opti-  
28 mization pipeline, assisting with problem formulation, model  
29 construction, solver configuration, and validation. Second,  
30 we explore the growing role of LLMs in algorithmic discov-  
31 ery, where they are used to generate, refine, and discover new  
32 optimization algorithms and heuristics. Designed for a broad  
33 AI audience, the tutorial bridges machine learning and opti-  
34 mization, covering foundational concepts, surveying state-  
35 of-the-art methods and systems, and highlighting key chal-  
36 lenges such as correctness, robustness, and handling ambigu-  
37 ous problem specifications. The tutorial concludes with a  
38 forward-looking discussion of open research directions and  
39 opportunities for impactful real-world applications.

## History

A preliminary version of this tutorial was successfully deliv-  
ered at **AAAI 2026** attracting approximately 200 attendees  
and fostering substantial engagement from both academic re-  
searchers and industry practitioners. The original tutorial ma-  
terials and slides are publicly available<sup>1</sup>.

Building on this foundation, the content has been signifi-  
cantly expanded through recent tutorials and invited talks on  
*automated algorithm design with LLMs*, presented at **CEC  
2025**, **IJCCI 2025**, and the **IEEE Web Seminar 2025**. These  
updated materials, including slides and video recordings, are  
also accessible online<sup>2</sup>. The strong positive reception and  
subsequent scholarly publications underscore the commu-  
nity's growing interest in this rapidly evolving intersection  
of LLMs and optimization.

## Proposed Length

Half day (3 hours 30 minutes total, consisting of two 1.5-hour  
slots with a 30-minute break in between).

## Estimated number of participants

We estimate that there would be 50-120 audience members.

## Target Audience and Prerequisite Knowledge

The tutorial is intended for three overlapping audiences: AI  
researchers and graduate students with little or no formal  
background in optimization; optimization and operations re-  
search researchers who want to understand how LLMs can  
assist modeling and algorithm design; and practitioners build-  
ing AI-enabled decision-support systems in domains such as  
logistics, planning, scheduling, and resource allocation.

We assume only basic mathematical maturity, familiarity  
with standard undergraduate-level concepts such as variables,  
constraints, and functions, and general exposure to modern  
AI or machine learning. No prior expertise in mathemati-  
cal optimization, integer programming, solver engineering,  
or LLM training is required. The tutorial will introduce the  
necessary optimization and LLM concepts before covering  
recent research results.

<sup>1</sup>[https://conlaw.github.io/llm\\_opt\\_tutorial/tutorial.html#tutorial](https://conlaw.github.io/llm_opt_tutorial/tutorial.html#tutorial)

<sup>2</sup>[https://cis.taskforce.ieee.org/esco/webinar-series/  
esco-webinar-26/](https://cis.taskforce.ieee.org/esco/webinar-series/esco-webinar-26/)

76	<b>Detailed Outline</b>	
77	<b>Total Duration: 3.5 hours (including 30-minute break)</b>	
78	<b>Part 1: Introduction (15 mins)</b>	The tutorial will start with
79		a broad overview of mathematical optimization, focusing on
80		mixed integer linear programming (MILP)—a core technique
81		in AI and operations research widely applied to planning,
82		scheduling, and resource allocation [Conforti <i>et al.</i> , 2014].
83		We will walk through the key stages of the optimization
84		pipeline, including (1) understanding the problem context,
85		(2) translating it into a MILP model, (3) tuning the model
86		and solver configuration for efficiency, and (4) validating the
87		model, highlighting where current bottlenecks lie and setting
88		the stage for how LLMs can help.
89	<b>Part 2: Optimization Model Formulation (45 mins)</b>	A central bottleneck in applying optimization to real-world prob-
90		lems is model formulation: translating informal problem
91		descriptions into mathematically precise optimization mod-
92		els. This process is especially challenging for domain ex-
93		perts without formal training in optimization. LLMs offer a
94		promising interface for bridging this gap, but they often pro-
95		duce incomplete or inconsistent formulations, omit key con-
96		straints, or misinterpret problem semantics.
97		
98		• The tutorial will start with a broad overview of mathe-
99		matical optimization.
100		• We survey recent progress in LLM-based auto-
101		formulation, including agentic frameworks such as Opti-
102		MUS [AhmadiTeshnizi <i>et al.</i> , 2024] and LEAN-LLM-
103		OPT [Liang <i>et al.</i> , 2026], Chain-of-Experts [Xiao <i>et al.</i> ,
104		2023], MCTS-based approaches [Astorga <i>et al.</i> , 2025],
105		as well as fine-tuned models specialized for optimiza-
106		tion such as ORLM [Huang <i>et al.</i> , 2025], and LLM-OPT
107		[Jiang <i>et al.</i> , 2025].
108		• We also review emerging datasets for benchmarking
109		(e.g., textbook-style, synthetic, and real-world col-
110		lections such as IndustryOR) and discuss evaluation
111		methods for assessing formulation correctness, includ-
112		ing execution-based metrics, graph-based equivalence
113		checking, and recent frameworks such as EquivaMap
114		[Zhai <i>et al.</i> , 2025].
115	<b>Part 3: Optimization Model Evaluation (30 mins)</b>	
116		One of the primary challenges in auto formulation is
117		evaluating whether a given formulation is correct (i.e.,
118		whether it solves the correct optimization problem). We
119		will present recent frameworks for evaluating whether a
120		generated formulation matches the intended problem, in-
121		cluding graph isomorphism-based approaches [Xing <i>et al.</i> ,
122		2024], execution accuracy [AhmadiTeshnizi <i>et al.</i> , 2024], and
123		EquivaMap [Zhai <i>et al.</i> , 2025].
124	<b>Coffee Break (30 mins)</b>	
125	<b>Part 4: Optimization Model Solving (45 mins)</b>	
126		Modern optimization solvers such as Gurobi and CPLEX
127		expose a wide range of configuration parameters that can
128		dramatically affect performance. Identifying effective con-
129		figurations is challenging and time-consuming, even for ex-
130		perienced practitioners. This module explores how machine
		learning and LLMs can assist in solver configuration and im-
		prove optimization performance [Lawless <i>et al.</i> , 2025].
		• We discuss how LLMs can leverage knowledge from
		documentation, code, and prior research to recommend
		high-quality configurations with little or no additional
		training.
		• We contrast this approach with traditional data-
		driven configuration methods that require large training
		datasets, and we highlight the strengths, limitations, and
		open challenges of LLM-based solver configuration in
		practice.
	<b>Part 5: Optimization Algorithmic Discovery with LLMs (45 mins)</b>	
		Recent advances in LLMs have opened new opportunities
		for automating and accelerating algorithm design. Designing
		effective algorithms has traditionally required significant do-
		main expertise, creativity, and iterative experimentation, of-
		ten making it a slow and resource-intensive process. LLMs
		offer a promising new paradigm by lowering the barrier to
		entry, enabling rapid prototyping, and exploring vast design
		spaces that are difficult for humans to navigate manually [Liu
		<i>et al.</i> , 2026].
		• We provide a structured overview of how LLMs are be-
		ing used for algorithm design, generation, and improve-
		ment.
		• We introduce a taxonomy of existing approaches, high-
		lighting four main paradigms for integrating LLMs into
		algorithmic discovery: LLM as optimizers [Yang <i>et al.</i> ,
		], LLM as extractors [Wu <i>et al.</i> , 2024], LLM as predic-
		tors [Pan <i>et al.</i> , 2025], and LLM as designers [Liu <i>et al.</i> ,
		2024].
		• We then review state-of-the-art methods such as
		EoH [Liu <i>et al.</i> , 2024], ReEvo [Ye <i>et al.</i> , 2024],
		Llamea [Van Stein and Bäck, 2024], HEEvo [Dat <i>et</i>
		<i>al.</i> , 2025], MEoH [Yao <i>et al.</i> , 2025], MLES [Hu <i>et</i>
		<i>al.</i> , 2026]. and applications across domains, includ-
		ing combinatorial optimization, black-box optimization,
		machine learning and science discovery.
		• We conclude by discussing open challenges and future
		directions, including domain LLMs, benchmarking, and
		human–AI collaboration in algorithm design.
	<b>Relevance to the AI Community and Objectives Served</b>	
		This tutorial addresses a timely topic at the intersection of
		large language models and mathematical optimization. As
		LLMs become increasingly capable in reasoning, coding, and
		tool use, they are emerging as promising assistants for opti-
		mization modeling, solver interaction, and algorithm design.
		These developments are of broad interest to the AI commu-
		nity because optimization underlies many important AI sys-
		tems and applications.
		The tutorial serves several IJCAI objectives: it introduces
		expert non-specialists in LLMs and optimization to each
		other’s areas, explains an AI topic of emerging importance,

185 surveys a rapidly developing research direction, and presents  
186 a synthesis connecting LLMs, optimization, and algorithm  
187 design.

## 188 Ethical Concerns

189 This tutorial addresses several ethical and societal concerns  
190 that arise when using LLMs in optimization workflows. First,  
191 incorrect or hallucinated formulations may lead to unsafe or  
192 harmful decisions in high-stakes applications such as logis-  
193 tics, energy, finance, or healthcare. Second, LLMs may in-  
194 herit biases from training data or from problem descriptions,  
195 which can propagate into optimization objectives, constraints,  
196 or recommendations.

197 The tutorial will explicitly discuss these risks and em-  
198 phasize that LLM-generated optimization artifacts must be  
199 treated as proposals requiring verification rather than author-  
200 itative solutions. We will cover the importance of evaluation,  
201 validation, transparency, reproducibility, and human over-  
202 sight, and we will highlight responsible-use considerations  
203 for deploying these methods in real-world decision-making  
204 systems.

## 205 References

206 [AhmadiTeshnizi *et al.*, 2024] Ali AhmadiTeshnizi, Wenzhi  
207 Gao, and Madeleine Udell. Optimus: Scalable optimiza-  
208 tion modeling with (mi) lp solvers and large language  
209 models. In *Forty-first International Conference on Ma-  
210 chine Learning*, 2024.

211 [Astorga *et al.*, 2025] Nicolás Astorga, Tennison Liu,  
212 Yuanzhang Xiao, and Mihaela van der Schaar. Autoform-  
213 ulation of mathematical optimization models using llms.  
214 In *Forty-second International Conference on Machine  
215 Learning*, 2025.

216 [Conforti *et al.*, 2014] Michele Conforti, Gérard Cornuéjols,  
217 Giacomo Zambelli, et al. *Integer Programming*, volume  
218 271. Springer, 2014.

219 [Dat *et al.*, 2025] Pham Vu Tuan Dat, Long Doan, and  
220 Huynh Thi Thanh Binh. Hsevo: Elevating automatic  
221 heuristic design with diversity-driven harmony search and  
222 genetic algorithm using llms. In *Proceedings of the AAAI  
223 Conference on Artificial Intelligence*, volume 39, pages  
224 26931–26938, 2025.

225 [Hu *et al.*, 2026] Qinglong Hu, Tong Xialiang, Mingxuan  
226 Yuan, Fei Liu, Zhichao Lu, and Qingfu Zhang. Mul-  
227 timodal LLM-assisted evolutionary search for program-  
228 matic control policies. In *The Fourteenth International  
229 Conference on Learning Representations*, 2026.

230 [Huang *et al.*, 2025] Chenyu Huang, Zhengyang Tang, Shixi  
231 Hu, Ruoqing Jiang, Xin Zheng, Dongdong Ge, Benyou  
232 Wang, and Zizhuo Wang. Orlm: A customizable frame-  
233 work in training large models for automated optimization  
234 modeling. *Operations Research*, 2025.

235 [Jiang *et al.*, 2025] Caigao Jiang, Xiang Shu, Hong Qian,  
236 Xingyu Lu, Jun Zhou, Aimin Zhou, and Yang Yu. Ll-  
237 mopt: Learning to define and solve general optimization

problems from scratch. *Proceedings of the International  
238 Conference on Learning Representations (ICLR)*, 2025. 239

[Lawless *et al.*, 2025] Connor Lawless, Yingxi Li, Anders  
240 Wikum, Madeleine Udell, and Ellen Vitercik. Llms for  
241 cold-start cutting plane separator configuration. In *In-  
242 ternational Conference on the Integration of Constraint  
243 Programming, Artificial Intelligence, and Operations Re-  
244 search*, pages 51–69. Springer, 2025. 245

[Liang *et al.*, 2026] Kuo Liang, Yuhang Lu, Jianming Mao,  
246 Shuyi Sun, Chunwei Yang, Congcong Zeng, Xiao Jin,  
247 Hanzhang Qin, Ruihao Zhu, and Chung-Piaw Teo. Large-  
248 scale optimization model auto-formulation: Harnessing  
249 llm flexibility via structured workflow. *arXiv preprint  
250 arXiv:2601.09635*, 2026. 251

[Liu *et al.*, 2024] Fei Liu, Tong Xialiang, Mingxuan Yuan,  
252 Xi Lin, Fu Luo, Zhenkun Wang, Zhichao Lu, and Qingfu  
253 Zhang. Evolution of heuristics: Towards efficient auto-  
254 matic algorithm design using large language model. In  
255 *Forty-first International Conference on Machine Learning*,  
256 2024. 257

[Liu *et al.*, 2026] Fei Liu, Yiming Yao, Ping Guo, Zhiyuan  
258 Yang, Xi Lin, Zhe Zhao, Xialiang Tong, Kun Mao,  
259 Zhichao Lu, Zhenkun Wang, et al. A systematic survey on  
260 large language models for algorithm design. *ACM Com-  
261 puting Surveys*, 58(8):1–32, 2026. 262

[Pan *et al.*, 2025] Qihong Pan, Haofei Tan, Guojiang Shen,  
263 Xiangjie Kong, Mengmeng Wang, and Chenyang Xu.  
264 Llm-tpf: Multiscale temporal periodicity-semantic fusion  
265 llms for time series forecasting. In *Proceedings of the  
266 Thirty-Fourth International Joint Conference on Artificial  
267 Intelligence*, pages 6030–6038, 2025. 268

[Van Stein and Bäck, 2024] Niki Van Stein and Thomas  
269 Bäck. Llamea: A large language model evolutionary algo-  
270 rithm for automatically generating metaheuristics. *IEEE  
271 Transactions on Evolutionary Computation*, 29(2):331–  
272 345, 2024. 273

[Wu *et al.*, 2024] Xingyu Wu, Yan Zhong, Jibin Wu, Bing-  
274 bing Jiang, and Kay Chen Tan. Large language model-  
275 enhanced algorithm selection: towards comprehensive al-  
276 gorithm representation. In *Proceedings of the Thirty-Third  
277 International Joint Conference on Artificial Intelligence*,  
278 pages 5235–5244, 2024. 279

[Xiao *et al.*, 2023] Ziyang Xiao, Dongxiang Zhang, Yangjun  
280 Wu, Lilin Xu, Yuan Jessica Wang, Xiongwei Han, Xiaojin  
281 Fu, Tao Zhong, Jia Zeng, Mingli Song, et al. Chain-of-  
282 Experts: When LLMs meet complex operations research  
283 problems. In *Proceedings of the International Conference  
284 on Learning Representations (ICLR)*, 2023. 285

[Xing *et al.*, 2024] Linzi Xing, Xinglu Wang, Yuxi Feng,  
286 Zhenan Fan, Jing Xiong, Zhijiang Guo, Xiaojin Fu,  
287 Rindra Ramamonjison, Mahdi Mostajabdaveh, Xiongwei  
288 Han, et al. Towards human-aligned evaluation for lin-  
289 ear programming word problems. In *Proceedings of the  
290 2024 Joint International Conference on Computational  
291 Linguistics, Language Resources and Evaluation (LREC-  
292 COLING 2024)*, pages 16550–16556, 2024. 293

- 294 [Yang *et al.*, ] Chengrun Yang, Xuezhi Wang, Yifeng Lu,  
295 Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen.  
296 Large language models as optimizers. In *The Twelfth In-*  
297 *ternational Conference on Learning Representations*.
- 298 [Yao *et al.*, 2025] Shunyu Yao, Fei Liu, Xi Lin, Zhichao Lu,  
299 Zhenkun Wang, and Qingfu Zhang. Multi-objective evo-  
300 lution of heuristic using large language model. In *Pro-*  
301 *ceedings of the AAAI Conference on Artificial Intelligence*,  
302 volume 39, pages 27144–27152, 2025.
- 303 [Ye *et al.*, 2024] Haoran Ye, Jiarui Wang, Zhiguang Cao,  
304 Federico Berto, Chuanbo Hua, Haeyeon Kim, Jinkyoo  
305 Park, and Guojie Song. Reevo: Large language models  
306 as hyper-heuristics with reflective evolution. *Advances in*  
307 *neural information processing systems*, 37:43571–43608,  
308 2024.
- 309 [Zhai *et al.*, 2025] Haotian Zhai, Connor Lawless, Ellen  
310 Vitercik, and Liu Leqi. Equivamap: Leveraging llms for  
311 automatic equivalence checking of optimization formula-  
312 tions. In *International Conference on Machine Learning*  
313 *(ICML)*, 2025.