

promptolution: A Unified, Modular Framework for Prompt Optimization

Tom Zehle^{1,2,*}, Timo Heiß^{3,4,*}, Moritz Schlager^{5,*},
Matthias Aßenmacher^{3,4}, Matthias Feurer^{6,7}

*Equal contribution ¹ELLIS Institute, Tübingen, Germany ² University of Freiburg, Germany

³LMU Munich, Germany ⁴Munich Center for Machine Learning (MCML), Germany

⁵Technical University of Munich, Germany ⁶TU Dortmund University, Germany

⁷Lamarr Institute for Machine Learning and Artificial Intelligence, Dortmund, Germany

Correspondence: tom.zehle@tue.ellis.eu, moritz.schlager@tum.de, timo.heiss@stat.uni-muenchen.de

Abstract

Prompt optimization has become crucial for enhancing the performance of large language models (LLMs) across a broad range of tasks. Although many research papers show its effectiveness, practical adoption is hindered as existing implementations are often tied to unmaintained and isolated research codebases. To address this, we introduce promptolution, a unified and modular open-source framework that provides all components required for prompt optimization within a single extensible system for both practitioners and researchers. It integrates multiple contemporary discrete prompt optimizers while remaining agnostic to the underlying LLM implementation.

 [AutoML/Promptolution](#)
 [System Demonstration](#)

1 Introduction

Modern large language models (LLMs) exhibit impressive general-purpose capabilities, being able to solve a wide variety of tasks (Radford et al., 2019; Ouyang et al., 2022; Bai et al., 2023; Touvron et al., 2023). They can be adapted to solve specific tasks

through in-context learning, i.e., simply by a textual instruction and optionally few-shot examples provided to the LLM as input (Brown et al., 2020). Since this input (referred to as *prompt*) steers the output of the LLM (Karmaker Santu and Feng, 2023; White et al., 2023), the LLM’s performance on a given task highly depends on it – in terms of quality, formulation, and the choice and order of examples (Zhao et al., 2021; Lu et al., 2022; Zhou et al., 2023). Table 1 illustrates this with two similar prompts for the GMS8K math dataset (Cobbe et al., 2021). Despite their strong semantic similarity (matching parts in the same color), their performances differ substantially. This sensitivity highlights the potential of optimizing prompts for specific tasks, much like how hyperparameter tuning boosts performance in classical machine learning (Kohavi and John, 1995). Just as manual hyperparameter search, crafting and refining good prompts by hand (manual prompt engineering) is tedious and unreliable (Jiang et al., 2020; Liu et al., 2023). Similar to how AutoML automates the search over large hyperparameter spaces (Feurer and Hutter, 2019), automatic prompt optimization recently emerged to systematically search the




Figure 1: Overview of the promptolution framework. promptolution takes a dataset (dev set + few-shot examples), token budget constraints, a description of the task, and optionally initial prompts as input. In an iterative process, a user-selected prompt optimizer refines the prompt(s) by evaluating the LLM’s prediction performance on the task’s development set and adapting the prompts accordingly (e.g., through another LLM). Detailed logging and callbacks enable tracking the entire process. The optimized prompts are returned and can be evaluated on test data.

Prompt	Accuracy
Tackle this elementary math problem by breaking it into logical steps. When you reach the solution, enclose the final answer with <code><final_answer></code> and <code></final_answer></code> markers for clarity.	37.6%
Assist with solving the elementary or grade school level math problem that requires multiple steps and provide the solution within <code><final_answer></code> <code></final_answer></code> tags for easy identification.	53.8%

Table 1: Example prompts and their test set accuracy on GSM8K with Llama-3.3-70B (colors = similar phrases).

combinatorial space of prompts (Li et al., 2025). Especially nowadays, when multi-agent systems have become central in both industry and research, LLMs specialized for individual tasks are increasingly important (He et al., 2025; Li et al., 2024). While fine-tuning entire models is expensive and data-intensive, automatic prompt optimization is a lightweight alternative and potential addition, often compatible with black-box LLMs (Cheng et al., 2024). A diverse collection of prompt optimizers has thus been introduced (see Cui et al., 2025; Li et al., 2025). However, several hurdles arise when applying these optimizers in practice. Each optimizer lives in its own research repository. Trying and comparing several optimizers requires juggling with multiple code bases and conflicting requirements. Additionally, these repositories are typically not actively maintained, and the research code often lacks proper software tests, documentation, and robustness. Moreover, their setups are inflexible, and deviations from their use cases (often limited to simple classification tasks) and LLM deployments require considerable effort. While existing libraries and tools for prompt optimization partly address these issues, they are either commercial and closed-source (Amazon Bedrock, 2025; Anthropic, 2025; Jina AI, 2025; Lee and Nardini, 2024), only implement a single optimizer (Adalflow, 2025; Agarwal et al., 2024; Hinthorn and Nishimi, 2025; Yuksekgonul et al., 2024), or have a high abstraction level designed mainly for end-to-end AI application development (Khatab et al., 2024; Kulin et al., 2025).

Contribution. We introduce **promptolution**, a modular, lightweight, and extensible open-source framework for automatic prompt optimization in Python, providing thoroughly tested and stable implementations. Our library implements multi-

ple LLM interfaces, NLP tasks, and contemporary discrete prompt optimization methods, which can be used interchangeably or replaced by custom implementations of the components. While promptolution facilitates single-prompt optimization for practitioners, a low abstraction level and helpers for systematic, reproducible experiments also make it geared toward researchers. Figure 1 shows an overview of how the framework operates.

Outline. We position our work within the landscape of automatic prompt optimization tools and libraries (§2), describe the design of our framework (§3), evaluate its performance compared to unoptimized prompts and other libraries to demonstrate its utility and competitiveness (§4), and outline future directions of the framework (§5).

2 Background & Related Works

Prompt Optimization. Automatic prompt optimization refers to systematically exploring prompt spaces using various automated optimization strategies, which may optimize both the instruction and the few-shot example components of a prompt (Cui et al., 2025; Li et al., 2025; Wan et al., 2024). Prompt optimization is commonly categorized by the nature of the prompt space into continuous (Li and Liang, 2021; Lester et al., 2021; Qin and Eisner, 2021) and discrete approaches (Guo et al., 2024; Yang et al., 2024; Zhou et al., 2023). For overviews of the field, we refer to Li et al. (2025) and Cui et al. (2025). promptolution focuses on the LLM-agnostic and interpretable discrete prompt optimization, which iteratively refines textual prompts directly, often using another “meta”-LLM¹ to generate improved candidates. The following optimizers are implemented in promptolution:

1. **OPRO** (Yang et al., 2024) uses LLMs as optimizers by providing a task description, examples, and previously scored candidates to the meta-LLM, which then proposes refined instructions.
2. **EvoPrompt** (Guo et al., 2024) optimizes instructions using evolutionary algorithms, based on (a) a genetic algorithm (GA) and (b) on differential evolution (DE). In both cases, the meta-LLM performs crossover and mutation.
3. **CAPO** (Zehle et al., 2025) is a recent GA-based alternative that leverages AutoML techniques to improve cost-efficiency and jointly optimizes both instructions and few-shot examples, outperforming the discrete optimizers above.

¹Can be the same as the one we optimize prompts for.

Framework	Extensible	Low Abstraction	Multiple Optimizers
promptolution	✓	✓	✓
DSPy	(✓)	✗	✓
CoolPrompt	(✓)	✗	✓
promptomatix	✗	✗	✓
prompt-ops	✗	✓	✓
TextGrad	✗	✓	✗
adalflow	✗	✓	✗
PromptWizard	✗	✓	✗
PromptIM	✗	✓	✗

Table 2: Comparison of open-source frameworks.

Other promising prompt optimizers not yet implemented in promptolution include *GEPA* (Agrawal et al., 2025), *TextGrad* (Yuksekgonul et al., 2025), *MIPRO* (Opsahl-Ong et al., 2024), and *PromptWizard* (Agarwal et al., 2024).

Existing Libraries, Frameworks & Tools. Most optimizers above are implemented in siloed research repositories with hard-coded experimental setups. Oftentimes not actively maintained, lacking software tests and proper documentation, they are inherently difficult to use for both scientific benchmark experiments with other optimizers and practical use cases with specific requirements. In the prompt optimization landscape, many actively maintained libraries and tools have emerged, including both open- and closed-source solutions:

Open-Source. Several libraries only implement a single prompt optimizer, including TextGrad (Yuksekgonul et al., 2024) using the method from Yuksekgonul et al. (2025), AdalFlow (Adalflow, 2025) with *LLM-AutoDiff* (Yin and Wang, 2025), Microsoft’s PromptWizard (Agarwal et al., 2024) based on the eponymous optimization algorithm, and PromptIM (Hinthorn and Nishimi, 2025) with an own iterative optimization strategy. prompt-ops (Meta-Llama, 2025) implements two optimizers, but focuses on prompt optimization for Llama models. In contrast, promptolution is a framework with multiple optimizers that can be used interchangeably for arbitrary LLMs.

DSPy is arguably the most popular existing framework in prompt optimization (Khattab et al., 2024) for building modular, declarative LLM pipelines, and includes an embedded prompt optimization component. It encompasses multiple optimizers like *MIPROv2* (building on Opsahl-Ong et al., 2024) or *GEPA* (Agrawal et al., 2025), and can also combine LLM training with prompt optimization (Soylu et al., 2024). While DSPy in-

tegrates prompt optimization as only one component within a high-level program compiler, promptolution solely focuses on prompt optimization with relatively low-level abstraction, extensible design, and support for systematic scientific benchmark studies and is thus more geared toward researchers and advanced practitioners than end-to-end AI application development (see also Appendix A.1). promptomatix (Murthy et al., 2025) builds on DSPy through its structured prompt compilation backend while also offering a lighter meta-prompt optimizer.

CoolPrompt (Kulin et al., 2025) is a very recent LLM-agnostic prompt optimization framework that supports multiple optimizers and emphasizes “zero-configuration” in contrast to promptolution, where researchers maintain close control over the setup. Table 2 summarizes how promptolution differs from all these open-source tools.

Other related tools with a slightly different focus include PromptBench (Zhu et al., 2024), which targets LLM evaluation supporting only simple optimization techniques, and OpenPrompt (Ding et al., 2022), designed for prompt learning for language models predating modern LLMs.

Closed-Source. Proprietary tools range from commercial web-platforms like PromptPerfect (Jina AI, 2025), cloud integrations such as Google Cloud’s *Vertex AI Prompt Optimizer* (Lee and Nardini, 2024) and the *AWS Bedrock Prompt Engineering Playground* (Amazon Bedrock, 2025), to vendor-specific solutions like *Anthropic Claude Prompt Tools* (Anthropic, 2025).

Positioning of promptolution. Our library is an open-source, LLM-agnostic, and highly modular framework. It focuses exclusively on prompt optimization rather than constructing full LLM pipelines. It already includes relevant LLM interfaces and NLP tasks, along with evaluation metrics, and provides contemporary discrete prompt optimizers in a single, unified system. The framework is highly customizable, offering fine-grained control over optimizers, tasks, logging, evaluation, and experiment configuration. As a result, it is suitable for both practitioners performing single-prompt optimization and researchers conducting systematic and reproducible large-scale benchmark studies.

3 System Design

promptolution is designed as a modular and extensible framework consisting of four key com-




Figure 2: Core components of promptlution. The upper part of each box lists the component implementations, the lower part important functions. Arrows between components indicate conceptual connections.

ponents (see Figure 2). All components follow a unified interface defined through corresponding Base-classes, ensuring that implementations can be used interchangeably, remain fully compatible with the framework, and automatically inherit shared functionality. While each component can be configured individually (see §3.1–3.4), a separate ExperimentConfig together with associated helper functions enables convenient parameterization of all components in a single object (see §3.5).

3.1 LLM

The LLM component provides an interface for obtaining responses from any LLM implementation. Its base class enables parallelization and monitors token usage. Three classes are implemented:

1. **APILLM** enables calls to LLMs hosted via an API, covering common vendors such as OpenAI and Anthropic. Listing 3 (L.1–5) illustrates the setup through the DeepInfra API.
2. **LocalLLM** allows using a local model using transformers (Wolf et al., 2020).
3. **VLLM** integrates the vllm library (Kwon et al., 2023) for efficient high-throughput inference and serving, and deployment on GPU clusters.

3.2 Predictor

The Predictor component defines how predictions are extracted from the LLM’s output. For classification tasks, we extract only the predicted class, whereas for math reasoning tasks, we extract the final solution. The FirstOccurrencePredictor searches for the first occurrence of any possible class label in the response and extracts it. The more robust MarkerBasedPredictor (Listing 3, L.6) extracts predictions between predefined HTML-like markers (see Table 3).

To solve this problem, we need to calculate the total number of fish in the fishbowls at all the tables. First, [...] Then, we add the 3 fish from the table that has 3 fish: 62 fish + 3 fish = 65 fish. <final_answer> 65 </final_answer>

Table 3: Example of the MarkerBasedPredictor extraction from a LLM response for a GSM8K sample.

3.3 Task

The Task component holds the dataset and other task-related information, including a textual description of the task, and defines how prompts are evaluated. It controls how subsampling is performed, which is crucial for efficiency, as not every prompt needs evaluation on the full data for a reasonable performance estimate.² We implement the following tasks relevant to NLP:

1. **ClassificationTask** targets discrete class labels and evaluates predictions using standard classification metrics. In our example in Listing 3 (L. 7–13), the task is specified using accuracy as metric. The task is easily created from a Pandas DataFrame by specifying the input and label columns, along with a task description.
2. **JudgeTask** is based on LLM-as-a-judge (Zheng et al., 2023), where another LLM scores the quality of the output according to the task description. This enables optimizing for subjective, creative tasks, both supervised and unsupervised, as ground-truth labels are optional.
3. **RewardTask** allows optimization w.r.t. custom reward functions. The reward can represent any measurable objective, such as code execution time or business metrics. Users define their own reward functions that compute a score (“reward”) directly from a predictor’s output.

3.4 Optimizer

The Optimizer component combines all other components by using a Predictor and a Task to determine the optimal prompt for the specified setup. Depending on the chosen optimization method, it may also rely on a (meta-) LLM for prompt alteration. As the core of the framework, it iteratively evaluates LLM predictions for a given task and refines the prompt(s) according to the respective optimization strategy.

promptlution currently implements four es-

²Depending on the subsampling strategy, evaluation is either performed on randomly drawn samples, a slice of the dataset (block), or the full dataset.

tablished prompt optimizers: OPRO (Yang et al., 2024), both EvoPromptDE and EvoPromptGA (Guo et al., 2024), and the current SOTA discrete prompt optimizer CAPO (Zehle et al., 2025). promptolution caches previously evaluated prompts out of the box, and constrains prompt evaluation during optimization to a subset of the available data, making algorithms faster and more efficient. In Listing 3, we set up CAPO (L. 14–22) and let it optimize prompts for 12 steps (L. 23).

Furthermore, new prompt optimizers can be added with minimal effort by inheriting from the base optimizer class and implementing a custom `_step()` method that defines the iterative optimization scheme. This makes our framework particularly useful for researchers developing and benchmarking new optimization algorithms.

```
llm = APILLM(
    api_url="api.deepinfra.com/v1/...",
    model_id="google/gemma-3-27b-it",
    api_key="...",
)
predictor = MarkerBasedPredictor(llm=llm)
task = ClassificationTask(
    df,
    task_description="The_task_is...",
    x_column="text",
    y_column="label_text",
    metric=accuracy_score,
)
optim = CAPO(
    predictor,
    task,
    meta_llm=llm,
    init_prompts=[
        "Classify_the_text_based_on...",
        # ...
    ],
)
prompts = optim.optimize(n_steps=12)
```

Figure 3: Setup of promptolution’s core components.

3.5 Experiment Configuration & Helper

promptolution not only supports optimizing prompts for a single specific setup, but also provides an `ExperimentConfig` framework that enables a convenient and structured configuration of larger benchmark experiments. Additional helper functions enable the running and evaluation of such experiments with just a few lines of code.

Listing 4 (L. 1–7) illustrates how the previous example (Listing 3) can be expressed with a single configuration class. Arguments that are

not specified resort to carefully chosen defaults. The associated helper functions allow users to run the optimization process according to the configuration (`run_optimization`) and to evaluate the prompts on unseen test data (`run_evaluation`). The `run_experiment()` function (Listing 4, L. 8) combines both steps, aiming to support researchers who want to perform extensive benchmark studies across multiple datasets and optimizers, such as for newly proposed prompt optimizers.

```
config = ExperimentConfig(
    optimizer="capo",
    task_description="The_task_is...",
    n_steps=12,
    api_url="api.deepinfra.com/v1/...",
    model_id="google/gemma-3-27b-it",
)
prompts = run_experiment(df, config)
```

Figure 4: Running an experiment via the `ExperimentConfig` abstraction.

3.6 Supporting Modules and Utilities

Exemplar Selection: Some prompt optimization algorithms (e.g., OPRO or EvoPrompt) do not consider few-shot examples. However, they can substantially improve LLM performance (Brown et al., 2020), even with simple selection strategies (Wan et al., 2024). promptolution offers post-hoc exemplar selection in an additional module, implementing random selection and random search³ to add few-shot examples to a fixed instruction.

Initial Prompt Creation: Many prompt optimizers require an initial pool of prompts to start with. We offer functions to automatically create prompts from a task description, a base prompt (following Zhou et al., 2023), or samples from a dataset.

Callbacks: The base class for the optimizers supports callbacks, allowing for easy tracking of the optimization progress. Callbacks can access the state of the optimizer at every optimization step, and optionally terminate the process. Important implementations include the `TokenCountCallback`, which tracks the accumulated token budget and terminates optimization if a specified threshold is exceeded, and the `FileOutputCallback`, which writes prompts and their scores to a file, enabling easy post-hoc analysis of the process.

³Random selection selects exemplars at random, whereas random search generates multiple sets of random examples, evaluates them, and selects the best performing set.

4 Evaluation

Setup. To evaluate promptolusion and contextualize its performance relative to other prompt optimization tools, we perform prompt optimization on the popular *GSM8K* (grade school math word problems; Cobbe et al., 2021) and *SST-5* dataset (sentiment classification; Socher et al., 2013). We use gemma-3-27B instruction tuned (Kamath et al., 2025) as the downstream LLM, and for optimizers that require one also as the meta-LLM. Further details on datasets and implementation choices are provided in Appendix A.2.

For our comparison, we employ the optimizers *CAPO*, *EvoPromptGA*, and *OPRO* from the promptolusion library, and additionally evaluate the leading optimizers from two other frameworks AdalFlow (*LLM-AutoDiff*) and DSPy (*GEPA*). To assess the impact of prompt optimization itself, we also evaluate three manually created prompts per dataset and report their average performance. We use the default parameterization for each optimizer to ensure practical comparability. We further restrict the token budget to at most one million in- and output tokens combined, which corresponds to a cost below \$0.15 for this LLM. All frameworks are initialized from a single task description without providing any initial prompts to test the full automation workflow. Both the unoptimized prompts and best prompts per optimizer (based on validation performance) are evaluated on a test set.

Results. A summary of the results is presented in Table 4. With the exception of *GEPA* on *SST-5* and *OPRO* on *GSM8K*, all optimizers substantially outperform the unoptimized baseline with improvements of up to 15%p in accuracy (CAPO on *GSM8K*). This demonstrates the utility of prompt optimization in general and of our framework in particular. The best-performing optimizer on both datasets, CAPO, as well as each runner-up, is imple-

mented in promptolusion, underscoring the competitiveness of our framework compared to other prompt optimization tools. Although not every optimizer included in promptolusion performs optimally on every task (e.g., *OPRO* on *GSM8K*), the library consistently provides at least one strong optimizer per task. Combined with its modular design, this allows users to switch easily to an alternative optimizer if the current one yields unsatisfying performance. We further emphasize that comparing optimizers within promptolusion required minimal manual effort due to its dedicated support for systematic benchmark experiments.

5 Conclusion & Future Directions

In this work, we introduced promptolusion, a unified and modular open-source Python framework for automatic prompt optimization, designed for both practitioners performing single-task prompt optimization and for researchers conducting large-scale benchmark studies. We highlighted the unique position of our framework within the landscape of prompt optimization tools, emphasizing its extensible, modular design and its low abstraction level, focused specifically on optimizing prompts. Furthermore, we demonstrated the role of each component in the framework and how they interact to support effective prompt optimization. Finally, through a comparative evaluation, we verified the utility and competitiveness of promptolusion.

Looking ahead, we plan to develop interoperability interfaces with higher-level frameworks such as DSPy, enabling a combination of strengths from both ecosystems. To further simplify the management of complex, large-scale experimental setups, we intend to introduce an interface to configuration management frameworks such as hydra (Yadan, 2019). We also aim to improve accessibility by implementing a lightweight graphical interface for real-time experiment tracking and visual analysis of the optimization process. In light of the recent rise of multi-agent systems, we plan to support the optimization of system prompts and interaction protocols across multiple agents. Additionally, inspired by AutoML, we intend to explore ensembling strategies for prompt optimizers, akin to ELPO (Zhang et al., 2025). More broadly, the extensibility of our framework ensures that we can continue to incorporate new state-of-the-art optimization methods as the field evolves.

Framework	Optimizer	GSM8K	SST5
Baseline	unoptimized	78.1	44.6
AdalFlow	AutoDiff	88.7	55.7
DSPy	GEPA	84.7	42.0
promptolusion	OPRO	69.7	<u>56.0</u>
	EvoPrompt	<u>91.0</u>	53.3
	CAPO	93.7	56.3

Table 4: Test set accuracy of optimized prompts using Gemma3-27B-it. Bold values indicate the best, underlined values the second-best performance per dataset.

Broader Impact

By unifying multiple prompt optimization methods that were previously scattered across separate research repositories, promptolution makes the benefits of prompt optimization for improving LLM performance broadly accessible, allowing practitioners to leverage these capabilities in real-world industry applications. At the same time, its modular and extensible design, combined with experiment-friendly implementations, makes the library a powerful tool for researchers benchmarking new prompt optimization algorithms. Since reimplementing competing methods and setting up rigorous benchmark experiments is typically time-consuming, promptolution offers the potential to accelerate research progress in prompt optimization and to enhance methodological comparability across the field.

References

- Adalflow. 2025. [Build and Optimize LM Workflows](#). Last accessed: 11/30/2025.
- Eshaan Agarwal, Joykirat Singh, Vivek Dani, Raghav Magazine, Tanuja Ganu, and Akshay Nambi. 2024. PromptWizard: Task-aware prompt optimization framework. *arXiv:2405.18369 [cs.CL]*.
- Lakshya A. Agrawal, Shangyin Tan, Dilara Soylu, Noah Ziemis, Rishi Khare, Krista Opsahl-Ong, Arnav Singhvi, Herumb Shandilya, Michael J. Ryan, Meng Jiang, Christopher Potts, Koushik Sen, Alexandros G. Dimakis, Ion Stoica, Dan Klein, Matei Zaharia, and Omar Khattab. 2025. GEPA: Reflective prompt evolution can outperform reinforcement learning. *arXiv:2507.19457 [cs.CL]*.
- Amazon Bedrock. 2025. [User Guide: Optimize a prompt](#). Last accessed: 11/25/2025.
- Anthropic. 2025. [Prompt engineering: Use our prompt improver to optimize your prompts](#). Last accessed: 11/25/2025.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, and 29 others. 2023. Qwen technical report. *arXiv:2309.16609 [cs.CL]*.
- T. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, and 12 others. 2020. Language models are few-shot learners. In *Proceedings of the 33rd International Conference on Advances in Neural Information Processing Systems (NeurIPS'20)*, pages 1877–1901. Curran Associates.
- Jiale Cheng, Xiao Liu, Kehan Zheng, Pei Ke, Hongning Wang, Yuxiao Dong, Jie Tang, and Minlie Huang. 2024. Black-box prompt optimization: Aligning large language models without model training. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3201–3219, Bangkok, Thailand. Association for Computational Linguistics.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv:2110.14168 [cs.LG]*.
- Wendi Cui, Jiaxin Zhang, Zhuohang Li, Hao Sun, Damien Lopez, Kamalika Das, Bradley A. Malin, and Sricharan Kumar. 2025. Heuristic-based search algorithm in automatic instruction-focused prompt optimization: A survey. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 22093–22111, Vienna, Austria. Association for Computational Linguistics.
- Ning Ding, Shengding Hu, Weilin Zhao, Yulin Chen, Zhiyuan Liu, Haitao Zheng, and Maosong Sun. 2022. OpenPrompt: An open-source framework for prompt-learning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 105–113, Dublin, Ireland. Association for Computational Linguistics.
- Matthias Feurer and Frank Hutter. 2019. Hyperparameter optimization. In Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren, editors, *Automated Machine Learning: Methods, Systems, Challenges*, pages 3–33. Springer International Publishing.
- Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yujiu Yang. 2024. Connecting large language models with evolutionary algorithms yields powerful prompt optimizers. In *The Twelfth International Conference on Learning Representations (ICLR'24)*. ICLR. Published online: iclr.cc.
- Junda He, Christoph Treude, and David Lo. 2025. LLM-based multi-agent systems for software engineering: Literature review, vision, and the road ahead. *ACM Transactions on Software Engineering and Methodology*, 34(5):1–30.
- William F. Hinthorn and Masahiro Nishimi. 2025. [Promptim](#). Last accessed: 11/25/2025.
- Zhengbao Jiang, Frank Xu, Jun Araki, and Graham Neubig. 2020. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 8:423–438.

- Jina AI. 2025. [PromptPerfect: AI Prompt Optimizer](#). Last accessed: 11/25/2025.
- Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre RamAl, Morgane RiviAlre, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean-bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, GaÁnl Liu, and 196 others. 2025. Gemma 3 technical report. *arXiv:2503.19786 [cs.CL]*.
- Shubhra Karmaker Santu and Dongji Feng. 2023. TELeR: A general taxonomy of LLM prompts for benchmarking complex tasks. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14197–14203. Association for Computational Linguistics.
- Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan, Saiful Haq, Ashutosh Sharma, Thomas Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. 2024. DSPy: Compiling declarative language model calls into state-of-the-art pipelines. In *The Twelfth International Conference on Learning Representations (ICLR’24)*. ICLR. Published online: iclr.cc.
- Ron Kohavi and George H. John. 1995. Automatic parameter selection by minimizing estimated error. In *Proceedings of the Twelfth International Conference on Machine Learning (ICML’95)*. Morgan Kaufmann Publishers.
- Nikita Kulin, Viktor Zhuravlev, Artur Khairullin, Alena Sitkina, and Sergey Muravyov. 2025. CoolPrompt: Automatic prompt optimization framework for Large Language Models. In *Proceedings of the 38th Conference of Open Innovations Association FRUCT, Issue 1 (Full papers)*, pages 158–166, Helsinki, Finland. FRUCT Oy.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles (SOSP ’23)*, pages 611–626. Association for Computing Machinery.
- George Lee and Ivan Nardini. 2024. [Announcing Public Preview of Vertex AI Prompt Optimizer](#). Last accessed: 11/25/2025.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3045–3059. Association for Computational Linguistics.
- Wenwu Li, Xiangfeng Wang, Wenhao Li, and Bo Jin. 2025. A survey of automatic prompt engineering: An optimization perspective. *arXiv:2502.11560 [cs.AI]*.
- Xiang Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597. Association for Computational Linguistics.
- Xinyi Li, Sai Wang, Siqi Zeng, Yu Wu, and Yi Yang. 2024. A survey on LLM-based multi-agent systems: workflow, infrastructure, and challenges. *Vicinity*, 1(1):9.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9):195:1–195:35.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2022. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098. Association for Computational Linguistics.
- Meta-Llama. 2025. [prompt-ops: An open-source tool for LLM prompt optimization](#). Last accessed: 11/30/2025.
- Rithesh Murthy, Ming Zhu, Liangwei Yang, Jieli Qiu, Juntao Tan, Shelby Heinecke, Caiming Xiong, Silvio Savarese, and Huan Wang. 2025. Promptomatrix: An automatic prompt optimization framework for Large Language Models. *arXiv:2507.14241 [cs.CL]*.
- Krista Opsahl-Ong, Michael J Ryan, Josh Purtell, David Broman, Christopher Potts, Matei Zaharia, and Omar Khattab. 2024. Optimizing instructions and demonstrations for multi-stage language model programs. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9340–9366. Association for Computational Linguistics.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *Proceedings of the 35th International Conference on Advances in Neural Information Processing Systems (NeurIPS’22)*. Curran Associates.
- Guanghui Qin and Jason Eisner. 2021. Learning how to ask: Querying LMs with mixtures of soft prompts. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5203–5212. Association for Computational Linguistics.

- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642. Association for Computational Linguistics.
- Dilara Soylu, Christopher Potts, and Omar Khattab. 2024. Fine-tuning and prompt optimization: Two great steps that work better together. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 10696–10710. Association for Computational Linguistics.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. LLaMA: Open and efficient foundation language models. *arXiv:2302.13971 [cs.CL]*.
- Xingchen Wan, Ruoxi Sun, Hootan Nakhost, and Serkan Arik. 2024. Teach better or show smarter? on instructions and exemplars in automatic prompt optimization. In *Proceedings of the 37th International Conference on Advances in Neural Information Processing Systems (NeurIPS’24)*, pages 58174–58244. Curran Associates.
- Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas Schmidt. 2023. A prompt pattern catalog to enhance prompt engineering with ChatGPT. In *Proceedings of the 30th Conference on Pattern Languages of Programs, PLoP’23*, pages 1–31, USA. The Hillside Group.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. Transformers: State-of-the-art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Omry Yadan. 2019. *Hydra - a framework for elegantly configuring complex applications*. Last accessed: 11/30/2025.
- Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc Le, Denny Zhou, and Xinyun Chen. 2024. Large language models as optimizers. In *The Twelfth International Conference on Learning Representations (ICLR’24)*. ICLR. Published online: iclr.cc.
- Li Yin and Zhangyang Wang. 2025. LLM-AutoDiff: Auto-differentiate any LLM workflow. *arXiv:2501.16673 [cs.CL]*.
- Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Zhi Huang, Carlos Guestrin, and James Zou. 2024. *TextGrad: Automatic “Differentiation” with Text*. Last accessed: 11/30/2025.
- Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Pan Lu, Zhi Huang, Carlos Guestrin, and James Zou. 2025. Optimizing generative AI by backpropagating language model feedback. *Nature*, 639(8055):609–616.
- Tom Zehle, Moritz Schlager, Timo Hei, and Matthias Feurer. 2025. CAPO: Cost-aware prompt optimization. In *4th International Conference on Automated Machine Learning*.
- Qing Zhang, Bing Xu, Xudong Zhang, Yifan Shi, Yang Li, Chen Zhang, Yik Chung Wu, Ngai Wong, Yijie Chen, Hong Dai, Xiansen Chen, and Mian Zhang. 2025. ELPO: Ensemble learning based prompt optimization for Large Language Models. *arXiv:2511.16122 [cs.CL]*.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *Proceedings of the 38th International Conference on Machine Learning (ICML’21)*, volume 139 of *Proceedings of Machine Learning Research*, pages 12697–12706. PMLR.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-judge with MT-bench and Chatbot Arena. In *Proceedings of the 36th International Conference on Advances in Neural Information Processing Systems (NeurIPS’23)*, pages 46595–46623. Curran Associates.
- Y. Zhou, A. Ioan Muresanu, Z. Han, K. Paster, S. Pitis, H. Chan, and J. Ba. 2023. Large language models are human-level prompt engineers. In *The Eleventh International Conference on Learning Representations (ICLR’23)*. ICLR. Published online: iclr.cc.
- Kaijie Zhu, Qinlin Zhao, Hao Chen, Jindong Wang, and Xing Xie. 2024. PromptBench: A unified library for evaluation of Large Language Models. *Journal of Machine Learning Research*, 25(254):1–22.

A Appendix

A.1 Delimitation from DSPy

One might argue that the existing framework DSPy (Khattab et al., 2024) is already very similar to promptolusion, as it is also open-source and modular, supports prompt optimization, implements multiple optimizers and allows for several LLM implementations, and provides additional utilities. However, there are several important differences, both in underlying purpose and in concrete design, that clearly distinguish the two frameworks:

Different focus : DSPy includes prompt optimization only as one component within a broader compilation workflow tightly coupled to a program structure, whereas promptolusion is not a full application framework and instead focuses exclusively on the prompt optimization stage.

Abstraction level : promptolusion provides a relatively low-level abstraction, emphasizing full user control over each component and transparency via detailed logging and callback hooks. In contrast, DSPy hides much of the prompt construction behind a high-level declarative interface (though it remains inspectable).

Target user base : Thus, promptolusion is geared toward researchers and advanced ML practitioners, while DSPy primarily targets AI application developers building end-to-end workflows such as agents or RAG systems.

Experiments : promptolusion makes it easy to implement custom optimizers, tasks, and other components, and is particularly suited for systematic large-scale benchmark experiments due to the integrated config framework. Its extensibility explicitly encourages contributions of new algorithms and tasks.

Compatibility : promptolusion integrates seamlessly into existing inference pipelines, since the result of optimization is a single string, whereas DSPy essentially requires substituting parts of its program representation.

A.2 Experiment Details

We evaluate our framework against DSPy and AdalFlow. Since neither provides a straightforward way to restrict compute budgets based on token counts, we enforce the token limit by raising an exception inside the respective LLM wrappers once the limit is exceeded, returning the last sug-

gested prompt.

For evaluation, we route every LLM API call through the same interface using langchain to ensure a fair comparison of the optimized prompts, and compare exact matches between the predicted and true labels (allowing for differences in capitalization). In the case of *GEPA*, we had to manually clean the LLM outputs because they did not follow the required output format stated in the task description (encapsulating the prediction within `<final_answer>` tags). Specifically, we had to remove the `“[[ ## target ## ]]”` and `“[[ ## completed ## ]]”` tags. For AdalFlow, we had to intercept API calls due to faulty extraction of system and user prompts. The system-prompt extraction mechanism relied on tags, but these were not forwarded correctly (e.g., `<START_OF_USER_PROMPT>` was expected instead of the provided `<START_OF_USER>`). These changes are documented in the experiment repository.

Dataset	Huggingface ID	n _{dev}	n _{test}
SST5	SeetFit/sst5	500	300
GSM8K	openai/gsm8k	500	300

Table 5: Overview of the utilized HuggingFace datasets.

The utilized datasets and sample sizes are detailed in Table 5. The dev set is used for optimization, the test set for holdout evaluation of the final prompts. Few-shot examples, if considered by the optimizer, are also taken from the dev set.

We make our complete experiment code publicly available for full reproducibility at <https://github.com/finitearth/prompt-optimization-framework-comparison>.

A.3 Quality Standards

To ensure high code quality, we adopt established software engineering best practices throughout our package. We maintain a comprehensive test suite that automatically verifies expected behavior after code changes. All tests must pass before a release is published, ensuring users encounter no issues when updating. The main branch is protected, and all contributions must be submitted via pull requests, each reviewed by another main contributor. We additionally employ pre-commit hooks to automatically check code formatting, documentation, and other basic quality issues before commits are made, improving readability and maintainability. We also

maintain strict documentation standards and do not accept poorly documented pull requests. Dedicated CI and CI/CD pipelines enable an automated build, test, and release of the package and documentation.

A.4 Documentation & Tutorials

Alongside the open-source software package, we provide extensive documentation available at <https://automl.github.io/promptolution/>. It covers the major components and their functionality, and also includes tutorials that guide users through their first prompt optimization use case. This further enhances the accessibility and usability of our framework.

A.5 System Demonstration

The system demonstration under <https://youtu.be/gySdgjEhsZA> uses the following code example, through which we guide step-by-step. The implementation requires the installation of the promptolution package with API support enabled. The subsequent imports establish the necessary components for the LLM wrapper, CAPO, and the evaluation task definitions.

```
! pip install promptolution[api]
```

```
from promptolution.llms import APILLM
from promptolution.optimizers import CAPO
from promptolution.tasks import JudgeTask
from promptolution.predictors import
    MarkerBasedPredictor
from promptolution.utils import
    create_prompts_from_task_description
import pandas as pd
```

The dataset, containing raw email-generation instructions, is loaded from a CSV file using pandas:

```
df_emails = pd.read_csv("emails.csv")
```

After loading the data, we define a task description that specifies the desired persona, tone, and formatting constraints (including the closing signature). This serves as the reference specification for the entire workflow.

```
task_description = "Write_concise,_polite,_
    professional_academic_emails_for_me_as_a_
    PhD_student,_asking_clarifying_questions_
    when_my_instructions_are_vague,_avoiding_
    cliché_openings,_and_always_ending_with_
```

```
Yours_sincerely,_Tom_Zehle."
```

We then instantiate the underlying LLM. In this example, the APILLM wrapper connects to the Llama-3.3-70B-Instruct-Turbo model via the DeepInfra API, which serves as the downstream LLM, the judge, and the meta-LLM.

```
llm = APILLM(
    model_id="meta-llama/Llama-3.3-70B-\
        \"Instruct-Turbo\",
    api_url="https://api.deepinfra.com/\
        \"v1/openai\",
    api_key=open("token.txt").read(),
)
```

To initialize the optimization search space, a set of candidate prompts is derived directly from the task description using the framework's utility function. These serve as the starting population for the optimizer we use later.

```
initial_prompts =
    create_prompts_from_task_description(
        task_description,
        llm,
    )
```

Given the subjective nature of the email generation task (lacking a ground truth), a JudgeTask is configured. This setup utilizes an LLM-as-a-Judge approach to evaluate the semantic quality of the generated outputs against the input instructions. A great alternative could be the ClassificationTask in cases where ground truth labels are available, or the RewardTask when the user can define an objective function that scores the outputs.

```
task = JudgeTask(
    df_emails,
    judge_llm=llm,
    task_description=task_description,
    x_column="instruction",
)
```

Next, we instantiate the CAPO optimizer with a marker-based predictor. The optimization routine is executed for six iterations (`n_steps=6`) to refine the prompt candidates iteratively.

```
optimizer = CAPO(
    task=task,
    predictor=MarkerBasedPredictor(llm),
    meta_llm=llm,
```

```
        initial_prompts=initial_prompts,  
        check_fs_accuracy=False  
    )  
    final_prompts = optimizer.optimize(n_steps=6)
```

After optimization completes, the final optimized prompts are returned and ready for use in email generation.