

MLCopilot: Unleashing the Power of Large Language Models in Solving Machine Learning Tasks

Lei Zhang^{*1} Yuge Zhang¹ Kan Ren² Dongsheng Li¹ Yuqing Yang¹

¹Microsoft Research, ²ShanghaiTech University

isleizhang@outlook.com, yugzhan@microsoft.com, renkan@shanghaitech.edu.cn

Abstract

The field of machine learning (ML) has gained widespread adoption, leading to significant demand for adapting ML to specific scenarios, which is yet expensive and non-trivial. The predominant approaches towards the automation of solving ML tasks (*e.g.*, AutoML) are often time-consuming and hard to understand for human developers. In contrast, though human engineers have the incredible ability to understand tasks and reason about solutions, their experience and knowledge are often sparse and difficult to utilize by quantitative approaches. In this paper, we aim to bridge the gap between machine intelligence and human knowledge by introducing a novel framework MLCopilot¹, which leverages the state-of-the-art large language models to develop ML solutions for novel tasks. We showcase the possibility of extending the capability of LLMs to comprehend structured inputs and perform thorough reasoning for solving novel ML tasks. And we find that, after some dedicated design, the LLM can (*i*) observe from the existing experiences of ML tasks and (*ii*) reason effectively to deliver promising results for new tasks. The solution generated can be used directly to achieve high levels of competitiveness.

1 Introduction

Past decades have witnessed a great advance and rapid development of machine learning (ML), but ML algorithms are still notoriously hard to configure (Hutter et al., 2019). For specific tasks, configuring and conducting corresponding ML solutions is non-trivial, which thus requires extensive human labor. Many challenges arise in developing practical ML solutions. First, it is of large human efforts considering the large space of ML solutions, such

as feature engineering, model design, optimization details, etc. Second, ML algorithms are sensitive to even minor changes of the task context. As a result, even the same algorithm may need to be re-configured for different application tasks. Last but not least, transferring successful experiences across different tasks is also intractable, which demands high-level reasoning abilities of human experts to derive reasonable solutions for novel tasks.

The predominant approaches that relieve the human effort of algorithm configuration have been some automation mechanisms such as AutoML (Automated Machine Learning) (Hutter et al., 2019). One major branch of AutoML formulates the problem as black-box optimization, and resorts to some optimization approaches such as Bayesian optimization (BO) (Frazier, 2018) to solve it. Though the obtained results are shown to be promising, it is time-consuming to spawn multiple trials, especially for large datasets and complex tasks. Moreover, AutoML does not follow the natural pattern of ML development that humans are accustomed to, which leaves a huge gap for humans to understand and control the whole process. Specifically, it is either difficult to explain the behavior of auto-tuning, or intractable to incorporate human prior such as the knowledge of the model architectures into the process, making it less flexible for human developers. Furthermore, the ML solutions derived by these optimization-based methods may *only* fit to the specific domains, and the transferring ability of these results also remains an open problem (Chen et al., 2022; Yan et al., 2022).

Contrarily, we notice two tendencies in how humans approach an ML task. Instead of jumping into solving the new task directly, humans often try to comprehend the task at hand and draw from their past experiences on relevant tasks. Additionally, humans recall their knowledge, which may have come from a textbook or prior experiences. This process differs significantly from the auto-

Work done during Lei Zhang’s internship at Microsoft Research Asia. Correspondence to Kan Ren.

¹Examples and code available at <https://github.com/microsoft/CoML>

mated approach mentioned earlier, which leads us to a natural question: *can we leverage both machine intelligence and human design patterns to improve our ability to solve ML tasks?* The advances of Large Language Models (LLM) (Brown et al., 2020; Chowdhery et al., 2022; Ouyang et al., 2022) have illustrated tremendous promising performance in mimicking human behaviors on conversation-based tasks. It seems plausible to utilize the power of LLM to address ML problems in a more human-like way.

Nevertheless, several challenges remain when incorporating LLMs to achieve this goal. First, we discovered that LLMs have trouble performing ML tasks based solely on the task description, in which case the performance is no better than random generation. Attempting to leverage historical ML experience, we found that the data often reside in heterogeneous formats (*e.g.*, code, configs and logs), which need to be canonicalized into formats that are acceptable to LLMs. Moreover, the amount of information that can be incorporated into in-context learning (Brown et al., 2020) is quite limited, and thus some retrieval strategy is desired to make the best out of it. Finally, deriving a ML solution based on historical experience is in its essence a mathematical thinking and logical reasoning problem (Patel et al., 2021), which necessitates some mechanisms to reasoning over knowledge.

In this paper, we explore and present a novel framework *MLCopilot*, which leverages LLMs to suggest solutions for novel real-world ML tasks, based on the existing experiences from historical tasks. We decompose the problem into offline and online stages. In the offline stage, *MLCopilot* canonicalizes historical data and creates an experience pool. LLMs are then used to *elicit* valuable knowledge from historical experience. In the online stage, *MLCopilot* retrieves experiences from the most relevant tasks from the experience pool, given the description of the target task. It then interacts with LLMs to obtain multiple suggested ML solutions in one round. We demonstrate that, with a well-designed framework, LLMs can not only elicit meaningful knowledge from historical experiences but also provide reasonable and competitive ML solutions for novel tasks.

Our work presents a three-fold contribution, which can be summarized as follows. (i) To the best of our knowledge, we are the first to utilize LLMs as a tool to generate solutions for new ML tasks. (ii) A novel retrieve-and-prompt framework has

been proposed to solve ML tasks almost instantaneously, *without* any time-consuming searching or optimization. (iii) We leverage the text understanding and generation capabilities of LLMs to produce *interpretable*² results for ML tasks. This approach has shown comparable or even better performance on a variety of real-world ML benchmarks.

2 Related Work

2.1 Large Language Models

Large language models (LLMs) are neural networks of significant sizes (typically containing tens or hundreds of billions of parameters). They have gained the incredible ability of processing and generating natural languages, due to the training on massive amounts of text data (Radford et al., 2018, 2019; Brown et al., 2020). Studies show that LLMs beyond a certain scale have “emergent abilities” (Wei et al., 2022), and perform remarkably well in applications such as chatbots, machine translation, and text summarization (Zhao et al., 2023; Touvron et al., 2023).

While LLMs have illustrated superior performance on natural language understanding and human-like text generation, they are still quite limited for complicated tasks that require reasoning (Huang and Chang, 2022) and mathematical skills (Patel et al., 2021; Thawani et al., 2021; Han et al., 2022; Saxton et al., 2019). The stream of task automation (Lu et al., 2023; Shen et al., 2023) investigated a general approach to decompose a task into a sequence of sub-tasks, but they are orthogonal to our work, since they did not take into past experience or knowledge from other tasks when planning a new task.

2.2 Machine Learning and AutoML

Machine learning (ML) is a subfield of artificial intelligence (AI) that involves developing optimization algorithms that can learn from data and make predictions (Bishop and Nasrabadi, 2006) or decisions (Sutton and Barto, 2018). Although ML has been successful in many real-world applications, designing an effective ML solution for a new task can be challenging due to the numerous design choices required. AutoML (Hutter et al., 2019) emerges as an approach to alleviate the manual effort involved. Popular methodologies include neu-

²The concept of “interpretability” in the context of our work primarily pertains to the human-readable knowledge generated by LLMs, which serves as a transparent reference for decision-making.

Term	Definition	Example
Task T	ML problem to solve (optionally with constraints).	Find an optimizer for a ResNet on ImageNet dataset.
Solution space \mathbb{S}	Solution hypothesis space to the task.	Optimizer: {Adam, SGD}; Learning rate: $[10^{-6}, 0.1]$.
Solution S	One particular choice within solution space.	2-layer ResNet with SGD optimizer using LR 10^{-3} .
Experience E	Successful solutions on historical tasks.	SGD with lr 0.024 achieves 76.2% accuracy for ResNet on ImageNet.
Knowledge K	High-level information acquired from experiences.	Usage of small LR makes training slower but could yield better final result.

Table 1: Terminologies used throughout this paper.

ral architecture search (NAS) (Pham et al., 2018), meta-learning (Andrychowicz et al., 2016), and Bayesian optimization (Frazier, 2018).

AutoML is able to reach beyond-human levels in solving ML tasks, but it still faces a few drawbacks. First, most AutoML methods require many rounds of trial-and-error, which can be time-consuming. Second, AutoML typically searches from scratch for a new task and neglects the experience on previous tasks. Finally, most AutoML methods are not interpretable due to their black-box nature, which excludes human understanding. Some methods may address one or two of the drawbacks, but not all of them. For example, the stream of transferrable AutoML research seeks to leverage past experience to assist in searching for new tasks (Bardenet et al., 2013; Wistuba et al., 2016; Mittal et al., 2020; Yan et al., 2022; Wang et al., 2021a), but they lack interpretability and most of them only work for specific types of tasks. A recent study (Chen et al., 2022) aims to use Transformer model (Vaswani et al., 2017) with large-scale pre-training to deal with broader types of tasks, but it is still non-interpretable and a cost search remains required for new tasks. Most recently, (Zheng et al., 2023) tried to search for neural architectures using GPT-4 (OpenAI, 2023). It prompts LLM to explain rationales, but it only explores model architectures, and still requires costly trial-and-error.

3 Preliminaries

The goal of MLCopilot is to assist humans in solving complex ML problems. Generally speaking, given a *task* which is a real-world problem for ML models to tackle, the goal of ML development is to conduct a concrete *solution*. The solution can be either a pipeline, configuration, or code snippet, based upon which a concrete ML model could be learned to handle the target task. The solution is also a particular sample within a complicated *solution space* that involves various design choices. These choices are mutually correlated and the outcome of different alternatives often influences the

others and eventually affects the final performance of the overall ML solution.

To create reasonable ML solutions for new tasks, we can draw on *experiences* from previous relevant tasks. MLCopilot is designed to use historical experiences for *knowledge* elicitation and effectively conduct effective solutions for the given novel ML task (details in § 4). To improve comprehension and clarity, we summarize the terminologies with descriptions and examples in Table 1.

4 MLCopilot

In this section, we present MLCopilot, with the formulation of the main problem and the overall architecture of our method. Then, we will describe some key components of MLCopilot in detail, including target task description, retrieval, canonicalization, and knowledge elicitation.

4.1 Overall Framework

As discussed previously, to unleash the power of LLMs in solving complex ML tasks, explicitly leveraging historical experience is crucial. However, utilizing past experience is not straightforward considering the heterogeneous data format and the huge number of records. Therefore, our technical design mainly focuses on addressing two problems: (i) how to comprehend and exploit the abundant raw experiences; (ii) how to effectively solve ML tasks based on the result of (i).

The main idea behind MLCopilot is *knowledge-based reasoning*, that is to leverage LLMs to conduct reasoning and task solving based on the previous knowledge, which has been analyzed and elicited from past experiences. To this end, MLCopilot contains two stages, including offline and online parts, both of which have been visually illustrated in Figure 1. In the offline stage, LLM has been incorporated to analyze the canonicalized historical experience data and elicit useful knowledge. And in the online stage, the user will query MLCopilot, which is also built upon LLM, to obtain a suitable ML solution for the novel task.

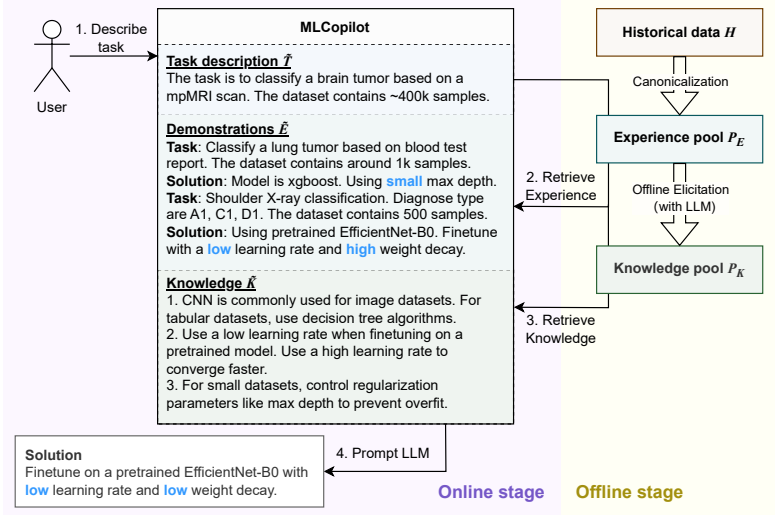


Figure 1: Overview of MLCopilot. MLCopilot has offline and online stages. During the offline stage, it creates pools of experience and knowledge. In the online stage, it retrieves experience and knowledge based on the novel task description. Finally, MLCopilot invokes LLM and returns solutions.

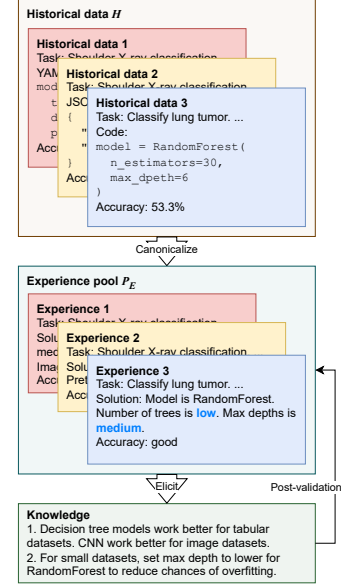


Figure 2: Offline stage: canonicalization, knowledge elicitation.

4.1.1 Offline Stage: Understanding and Reasoning

We first present the data settings and describe the corresponding preprocessing procedure briefly. Let $H = \{D_1, \dots, D_{N_H}\}$ be the raw historical data with N_H previous records. The i -th record D_i is defined as a three-element tuple $\langle T_i, S_i, M_i \rangle$ which contains a task $T_i \in \mathbb{T}$, a solution $S_i \in \mathbb{S}$, and the evaluated metric performance M_i , *e.g.*, classification accuracy.

Note that, the historical data H may have heterogeneous and diverse formats. For example, the task can be described in the natural text, while the solution can be a JSON configuration, a row of tabular data, or code snippets. An experience pool P_E is constructed, to canonicalize the data and store them as experiences $P_E = \{E_1, \dots, E_{N_E}\}$, where $E_j = \mathcal{C}(D_j)$, and $\mathcal{C}(\cdot)$ is the canonicalization function. For the simplicity of notations, we assume all the solutions within P_E come from a universal solution space. It is easy to extend the framework to scenarios with multiple solution spaces.

Knowledge, is high-level information acquired from the experience (Dictionary, 1989), and we leverage LLM to elicit knowledge from the constructed experience pool, in the offline stage. Knowledge is the easy-to-understand summarization of previous ML experiences, which will further be utilized when inferring the final solution in the online stage. To be specific, a subset of experience is first sampled from the experience pool, then a LLM is incorporated to read and understand the ex-

perience data, allowing us to “elicit” knowledge K from it. The process of elicitation is formulated as $K = \mathcal{I}_K(P_E; \text{LLM})$, which is an iterative process by interacting with LLM along with post-validation on the obtained knowledge. The detailed process of elicitation is discussed in § 4.5. All the generated knowledge is stored in a knowledge pool $P_K = \{K_1, \dots, K_{N_K}\}$ with totally N_K items.

The obtained experience pool and knowledge pool will be further utilized by MLCopilot in the online stage, to conduct reasonable, promising, and competitive ML solutions for novel tasks.

4.1.2 Online Stage: Retrieving and Solving

The online stage of MLCopilot aims to conduct reasoning and task solving based on the off-the-shelf information obtained from the offline stage. Specifically, given the user query with a task description, MLCopilot will respond with the corresponding reasonable ML solutions via retrieving relevant experiences and knowledge, and interacting with LLM by a curated prompt, in one round.

When a user comes with a novel target task \tilde{T} , which has never been seen in history, MLCopilot first retrieves the relevant experiences of other relevant tasks as demonstrations $\tilde{E} = \mathcal{R}_E(\tilde{T}, P_E)$, where $\mathcal{R}_E(\cdot)$ is the retrieval functions for the experience pool. It also retrieves knowledge $\tilde{K} = \mathcal{R}_K(\tilde{T}, P_K)$ to guide the response on the new task, where $\mathcal{R}_K(\cdot)$ is the retrieval functions for the knowledge pool. MLCopilot finally generates a recommended solution by invoking LLM once:

$$\tilde{S} = \text{LLM}(\tilde{T}, \tilde{E}, \tilde{K}).$$

The framework is shown in Figure 1, where we illustrate an example of how MLCopilot handles a task to classify a brain tumor by leveraging previous experiences and knowledge. Next we will introduce the dedicated components in detail.

4.2 Task Description in Natural Language

Firstly, we show how the target task is described in our framework, which is the input to MLCopilot. The prior works (Feurer et al., 2015; Wang et al., 2021a) usually use meta-features designed by humans for specific types of tasks (e.g., the number of samples in the dataset) to describe a task, so as to ease the difficulty of comprehending the task. However, such design might degenerate the ability of LLM to generalize to new types of tasks. We believe that task description in natural language is more straightforward to users. It is also agnostic to task types, and does not require heuristics to design meta-features. As such, we adapt the task description without any feature engineering, and users can freely describe dataset names, characteristics, domain-specific constraints, and more. Furthermore, our experiments (§ D) illustrated that incorporating a natural language user interface helps recall and leverage previous knowledge contained in the training corpus of LLMs.

4.3 Retrieval

The retrieval technique has been used to (i) gather some demonstrations of the historical ML solutions to the relevant tasks and (ii) apply useful knowledge previously to further motivate and prompt the LLM to better solve the target ML task.

We first discuss how to retrieve experience \mathcal{R}_E as demonstrations \tilde{E} . Intuitively, the most helpful experience in solving a new task should come from the most relevant tasks. The key question then becomes how relevance is defined. To this end, we first embed the task description by invoking a language model \mathcal{E} (e.g., GPT-3 (Brown et al., 2020)) to generate an embedding vector of the textual content. Given a new task \tilde{T} , MLCopilot retrieves the most relevant historical tasks from the experience pool P_E by calculating the cosine similarity between the embeddings of the new task and the stored tasks. The corresponding experience to these embeddings will serve as demonstrations \tilde{E} , as calculated as

$$\tilde{E} = \mathcal{R}_E(\tilde{T}, P_E) = \arg \text{top-}k_{\langle T, S, M \rangle \in P_E} \left(\frac{\mathcal{E}(T) \cdot \mathcal{E}(\tilde{T})}{|\mathcal{E}(T)| \cdot |\mathcal{E}(\tilde{T})|} \right),$$

where the most relevant k entries of experience will be retrieved for subsequent demonstration.

The retrieval of knowledge \mathcal{R}_K is based on matching of solution space – retrieving *all* the knowledge that are elicited from the same solution space as the new task, from the knowledge pool P_K . This simplicity of \mathcal{R}_K is due to the fact that the knowledge produced in the offline stage of MLCopilot is concise and of high quality, whose generation procedure will be discussed in § 4.5.

4.4 Canonicalization

As mentioned previously, the data of raw ML experience are heterogeneous and of diverse formats. While some of them (e.g., task descriptions) have already been in natural text format, the ML solutions and the corresponding metric performance are often expressed in structured configurations, tabular formats, or programming languages. More importantly, they might even contain a lot of numbers, which language models or even LLMs are not good at processing and understanding (Thawani et al., 2021; Han et al., 2022; Saxton et al., 2019). To better unleash the power of LLM, we canonicalize all the data to express it in natural language.

The essential part of canonicalization is to convert the raw data into a well-formed natural language, as shown in the left part of Figure 2. Other than unifying the solutions in diverse formats, a crucial technique is number discretization which avoids feeding numbers to LLM directly. We follow (Thawani et al., 2021) that discretizes continuous numerical data into several intervals, and mapping each value within each interval to the same discrete value. To minimize performance loss, we discretize each numerical value based on the corresponding distribution and percentile points. More details can be found in § 5.1.

4.5 Knowledge Elicitation

With the canonicalized experience stored in pool P_E , MLCopilot can then elicit knowledge, to better support the online stage for solving novel ML tasks. It is important to note that knowledge elicitation occurs offline, prior to serving user tasks. The approach involves the following steps: (i) constructing a prompt that consists of a *random subset* of the experience pool (to avoid bias towards certain tasks), along with an inquiry that asks for analysis and summary; (ii) sending the prompt to LLMs to generate a knowledge “candidate”; (iii) validating the candidate on experience pool. The flow of this

process is illustrated in the right part of Figure 2 (pseudo-code in § A).

We elaborate on the validation step, which we call automated *post-validation* after requesting knowledge from the LLM. This step is designed to alleviate the hallucination issue (Ji et al., 2023) and raise the quality of generated knowledge. It tests the knowledge by using it to solve a set of *validation tasks*. If the generated knowledge is found invalid (e.g., due to hallucination), it adjusts the generation settings, such as the order of experiences in the prompt, tone of hypophora questions, and parameters of LLM invocation, and let LLM regenerate knowledge. This iterative process continues until the performance on the validation tasks has been converged, or the invocation has reached the maximum number. This process can be represented formally as $K = \mathcal{I}_K(P_E; \text{LLM})$.

We argue that our knowledge elicitation is novel and different from prior works of knowledge extraction (Zhang et al., 2022) or knowledge generation (Yu et al., 2022; Lu et al., 2023) in natural language processing. Firstly, our knowledge is obtained from heterogeneous resources using general text completion models, without requiring predefined templates or complicated pipelines. Secondly, acquiring knowledge for ML tasks requires analysis, summarization, and high-level reasoning, which is significantly more challenging than simply extracting simple facts (Zhang et al., 2022). Finally, the knowledge is anchored in experience data, rather than purely based on LLM’s pre-training corpus (Yu et al., 2022), which makes the framework scalable to new scenarios.

Knowledge elicited by MLCopilot can be beneficial not only for LLMs but also for human ML practitioners. Since the knowledge is expressed in natural language, it could potentially serve as a cookbook for ML developers. In an effort to share our findings and inspire future ML research, we have released all the knowledge obtained so far (see § E). Hopefully this will reveal some of the “secret sauce” behind how ML works and promote knowledge sharing within the community.

5 Experiment

We evaluate MLCopilot on a series of benchmarks, aiming to answer the following research questions: (i) Can MLCopilot outperform traditional approaches or simple interactions with LLMs? (ii) How important are individual techniques in ML-

Copilot, e.g., knowledge and experience? (iii) Is the elicited knowledge informative and reasonable?

5.1 Experiment Setup

Implementation details. The current implementation of MLCopilot involves maintaining dedicated experience and knowledge pools for each solution space. The historical data is sourced from the benchmarks described below, while task descriptions are crawled from benchmark websites. Numerical values in the data are discretized into five levels: “very low”, “low”, “medium”, “high”, and “very high”. The precise value of each level is determined by analyzing the statistics of the best solutions within the solution space (see detailed analysis in § 5.3). We interact with the general-purpose GPT-3.5 model³ (code-named “text-davinci-003”, *without* additional fine-tuning), and “text-embedding-ada-002” to obtain embeddings for task descriptions. (Results with other LLMs can be found in § D.3.) The temperature is set to 0 to minimize randomness. Additional details regarding prompt design can be found in § C.

Benchmarks. We selected benchmarks that have established a predetermined solution space for all possible solutions and provided performance metrics for all the solutions in the solution space (either through a lookup table or surrogate). We conducted experiments using MLCopilot on three ML benchmarks: HPO-B (Arango et al., 2021), PD1 (Wang et al., 2021b), and HyperFD (Yan et al., 2022). These benchmarks comprise numerous ML tasks and datasets, covering a broad spectrum of scenarios such as tabular data classification and regression, image classification, and object detection. Details can be found in § B.1.

Evaluation metrics. In each experiment, every compared method makes three attempts to predict successful solutions for an unseen task. Solutions are evaluated in the order they were suggested. $\text{Metric}@t$, where $t \geq 1$, is defined as the best metric performance achieved among the first t suggested solutions. The reported performance metrics are averaged over at least 5 random seeds.

5.2 Main Results

We show the performance of MLCopilot in Table 2. Baselines we compared with include:

- Traditional AutoML or meta learning methods,

³<https://platform.openai.com/docs/models/gpt-3-5>

Method	HPO-B \uparrow			PD1 \uparrow			HyperFD (Rank \downarrow AP \uparrow)					
	nAcc@1	nAcc@2	nAcc@3	nAcc@1	nAcc@2	nAcc@3	Rank@1	Rank@2	Rank@3	AP@1	AP@2	AP@3
Random	54.70	60.70	64.80	-0.86	-0.08	0.39	109.55	73.16	54.79	90.76	91.20	91.38
Constant	72.85	74.61	75.02	1.27	1.56	1.59	78.00	54.33	49.25	91.13	91.23	91.28
TST-M	72.73	74.44	74.56	1.10	1.35	1.40	57.67	43.50	42.67	91.22	91.37	91.38
HyperSTAR	67.37	68.14	68.71	1.10	1.27	1.34	97.75	72.97	52.03	90.78	91.05	91.25
ASKL	77.01	81.76	85.02	1.26	1.29	1.44	92.58	64.67	51.25	90.93	91.15	91.34
FLAML	77.84	82.95	88.06	1.28	1.31	1.58	66.42	43.33	31.83	91.09	91.29	91.33
HyperFD	–	–	–	–	–	–	56.97	47.91	31.75	91.17	91.26	91.44
LLM-ZS	61.37	79.41	80.56	-1.03	1.25	1.26	119.25	90.42	41.00	90.69	90.96	91.38
LLM-FS	78.93	83.10	89.73	0.43	0.57	0.62	66.69	52.43	40.98	91.26	91.40	91.48
MLCopilot	81.59	83.23	90.72	1.48	1.54	1.62	59.74	38.67	25.58	91.38	91.60	91.66

Table 2: Main results on HPO-B, PD1 and HyperFD. **nAcc**, **AP**: the higher the better. **Rank**: the lower the better.

including Random, ASKL (Feurer et al., 2015), Constant (Bardenet et al., 2013; Kotthoff et al., 2019), TST-M (Wistuba et al., 2016), HyperSTAR (Mittal et al., 2020), HyperFD (Yan et al., 2022) and FLAML-Zero (Wang et al., 2021a). Details described in § B.2.

- **LLM-ZS** directly prompts LLM to generate a zero-shot solution based solely on the task description, which is similar to using tools such as GitHub Copilot⁴ or Amazon CodeWhisperer⁵.
- **LLM-FS** uses the few-shot prompt technique (Brown et al., 2020) by adding some demonstrations to the prompt to enable in-context learning. The demonstrations are randomly selected from our canonicalized experience pool. Unlike MLCopilot, LLM-FS does not have access to advanced techniques such as experience and knowledge retrieval.

MLCopilot achieved the highest normalized accuracy (nAcc) across all three trials. The improvement is particularly significant for the first attempt (nAcc@1). It is remarkable that LLM-FS has already surpassed all the traditional baselines, suggesting the large capability of LLMs on ML tasks.

On PD1, Normalized accuracy (nAcc) are in range $[-2, 2]$ following the setting of (Wang et al., 2021b). MLCopilot remains the best out of all methods compared. Notably, “Constant” baseline almost outcompetes all other baselines, which casts doubt on the effectiveness of the task similarities measured by other baselines. Meanwhile, both LLM-ZS and LLM-FS fail on PD1, indicating PD1 is more challenging for LLMs.

For HyperFD, Following (Yan et al., 2022), we use average precisions (AP) (the higher the better) and rankings (within $[1, 216]$, the lower are better) to measure the performance. Similar to what

was observed in HPO-B, LLM-FS achieves comparable performance to most baselines with a few demonstrations. It is expected that the performance of LLM-FS would improve with the inclusion of techniques from MLCopilot. However, it is worth noting that HyperFD is a private benchmark and its benchmark was released after the knowledge cutoff of GPT-3.5, making it unlikely that LLM has memorized the best solutions on this benchmark.

5.3 Ablation study

In the ablation study, we use nAcc@1 (or Rank@1) as the main metric for comparisons.

Study of retrieval. The first question we are trying to answer is whether retrieving experience and knowledge are necessary – what happens if either of them is missing from the prompt sent to LLM? The results are shown in Table 4. While the absence of knowledge leads to a reduction in performance, the absence of demonstrations leads to a complete collapse. Examining the knowledge generated (§ E), we found it often contains vague claims such as “*The size of the dataset can influence the configuration of eta*” (HPO-B, Space 5971). The knowledge did not clarify what is the “*influence*”, which is why experience is still much needed even with the presence of knowledge.

We then compare different retrieval methods (“Pipeline MLCopilot” columns in Table 3). MLCopilot retrieves the most relevant tasks based on the *embedding of textual description*. Alternatively, we can (i) measure similarities based on *meta-features*; (ii) simply retrieve experiences *randomly*. Shown in Table 3, both meta-feature and text embedding consistently outperform random retrieval.

When choosing between meta-features or text embedding, we believe that the latter has demonstrated advantages over manually designed meta-features. This is partly due to the fact that the performance of meta-features depends largely on

⁴<https://github.com/features/copilot>

⁵<https://aws.amazon.com/codewhisperer/>

Retrieved by	Pipeline	HPO-B \uparrow		PD1 \uparrow		HyperFD \downarrow	
		ASKL	MLCopilot	ASKL	MLCopilot	ASKL	MLCopilot
Text embedding		75.34 \pm 0.00	81.59 \pm 0.94	1.40 \pm 0.00	1.48 \pm 0.06	79.17 \pm 0.00	59.74 \pm 1.89
Meta-feature		80.61 \pm 0.00	83.29 \pm 1.46	1.02 \pm 0.00	1.41 \pm 0.09	107.67 \pm 0.00	50.49 \pm 6.38
Random		73.67 \pm 2.51	78.00 \pm 2.82	0.06 \pm 0.25	1.37 \pm 0.12	84.23 \pm 14.84	57.95 \pm 10.19

Table 3: Comparison of approaches to retrieve experience (*i.e.*, based on what measures to retrieve the experience) and to consume the retrieved experience (ASKL: directly use the solutions for retrieved tasks on the new task; MLCopilot: use the retrieved experience as demonstrations along with knowledge to prompt LLM).

Retrieve	HPO-B \uparrow	PD1 \uparrow	HyperFD \downarrow
Exp.+Know.	81.59 \pm 0.94	1.48 \pm 0.06	59.74 \pm 1.89
Know.	62.77 \pm 1.78	1.10 \pm 0.00	127.75 \pm 0.00
Exp.	76.21 \pm 0.16	1.36 \pm 0.09	63.20 \pm 3.55

Table 4: Effect of retrieving experience and knowledge.

the quality of their design. While meta-features have been shown to be promising for tabular datasets where they are well-studied and carefully designed (Feurer et al., 2015; Wang et al., 2021a), the design of meta-features for complex tasks in PD1 is non-trivial. In contrast, the text embedding approach has the additional advantage of not requiring any manual design for new types of tasks. Furthermore, text embedding is more promising for handling new and varied tasks, while meta-feature for new tasks is not easily scalable.

Nevertheless, we would like to emphasize that the key factor is not solely the method of retrieving experience, but rather how the retrieved experience is utilized. When the retrieved experience is used directly, as done in ASKL, all retrieval strategies perform poorly. In contrast, MLCopilot has the ability to not only retrieve relevant experience, but also provide guidance through elicited knowledge and leverage the power of LLMs.

Study of canonicalization. As shown in Table 6, the performance suffers considerably without discretization as sending continuous numerical values directly to LLM is not feasible. Furthermore, it is crucial to compute the split points based on the statistics of the best solutions. If the range is expanded to include all possible values, the split point may not fall on the sensitive points, resulting in sub-par performance. This is demonstrated by “On All” in Table 6, which performs even worse than no discretization at all⁶.

Study of knowledge. In Table 5, we conducted an ablation study to evaluate the impact of knowledge on our method. The results show that, removing knowledge retrieval in the online stage of our

⁶Please note that HyperFD is not included in the ablation study as it adopts a discrete solution space.

Method	HPO-B \uparrow	HyperBO \uparrow	HyperFD \downarrow
MLCopilot	81.59 \pm 0.94	1.48 \pm 0.06	59.74 \pm 1.89
w/o Post-Val.	78.34 \pm 0.71	1.44 \pm 0.05	62.41 \pm 3.66
w/o Know.	76.21 \pm 0.16	1.36 \pm 0.09	63.20 \pm 3.55

Table 5: Ablation on knowledge utilization in online stage and post-validation in offline stage.

Discretization	HPO-B \uparrow	PD1 \uparrow	HyperFD \downarrow
On Best	81.59 \pm 0.94	1.48 \pm 0.06	59.74 \pm 1.89
On All	70.82 \pm 0.01	1.41 \pm 0.02	–
\times	74.09 \pm 0.44	1.45 \pm 0.05	84.96 \pm 7.16

Table 6: Discretization in canonicalization.

method results in a significant decrease in the final performance. This is because knowledge is instrumental in helping LLMs arrive at the most effective solutions. Furthermore, the post-validation in elicitation procedure in the offline stage of MLCopilot also plays a vital role in enhancing its usefulness.

Based on our qualitative study on the generated knowledge (see § E), we found that, the knowledge serves as a helpful summary of past ML experiences while providing guidance on how to adjust parameters and settings based on task characteristics. We observe that post-validation significantly reduces the chances that trivial, vague, or hallucinated knowledge is produced, although such knowledge is still sometimes observed. For example, in the case of “UniRef50” task with “Transformer” model on PD1, the knowledge contains certain numerical examples that were not part of the demonstrations and instead the result of hallucinations.

6 Conclusion

In conclusion, this paper proposes MLCopilot, a framework that unleashes the power of LLMs to solve practical ML tasks. MLCopilot showcases the versatility of LLMs, that it can handle not only text-related tasks, but also tasks involving heterogeneous inputs and intricate reasoning. We believe this represents a significant advancement in expanding the scope of LLM applications to a broader spectrum of complex problems.

7 Ethical considerations

The architecture of MLCopilot is meticulously engineered to ensure that the solutions it recommends always remain within the bounds of the solution space provided by the user. As a result, it acts as a safeguard against the generation of unethical solutions, provided that the defined solution space adheres to ethical standards.

However, the foundational techniques outlined in this paper, including experience retrieval and knowledge elicitation, possess broader applicability across various scenarios beyond machine learning, such as task automation (Lu et al., 2023) and scientific research (Boiko et al., 2023). In these contexts where the solution space extends beyond the constraints of a strictly-defined machine learning problem and where Large Language Models (LLMs) exhibit inherent limitations, the potential for unpredictability arises. Therefore, it becomes imperative to exercise ethical prudence when deploying MLCopilot in diverse cases.

8 Limitations

Potential data leakage. Since LLMs are trained on large corpus of data from Internet, it is likely that the benchmarks (especially HPO-B based on OpenML) have already been encountered during the pre-training phase of LLMs. To mitigate this potential bias, we conducted an evaluation of MLCopilot on HyperFD (Yan et al., 2022). It is worth noting that the HyperFD dataset was introduced in a paper published after the knowledge cutoff date of GPT-3.5, and the dataset itself remains private. We empirically reveal that MLCopilot exhibits robust performance on the HyperFD dataset.

Furthermore, our findings indicate a significant performance enhancement when the data is canonicalized (Table 6). If the data were indeed memorized during the pre-training process, LLMs would likely benefit from access to unaltered, raw data. These results provide valuable supporting evidence for the assertion that the capabilities of LLMs extend beyond mere memorization. They encompass a broader spectrum of cognitive skills, including mathematical reasoning and logical thinking.

Distinction from AutoML methods. MLCopilot is not intended to serve as a replacement for established AutoML approaches. The distinction is grounded in the inherent limitations of Large Language Models (LLMs) when it comes to performing mathematical computations, as illustrated

in recent work (Imani et al., 2023). Consequently, it is improbable that MLCopilot would surpass state-of-the-art Bayesian optimization methods in the pursuit of superior solutions. In Table 2 we terminated our evaluation at $t = 3$ (i.e., three solutions), as we observed that performance reached a point of saturation with further increases in t .

We argue that the true value of MLCopilot lies in the following facets: (i) it accepts arbitrary types of task descriptions; (ii) it leverages ML experiences from diverse sources, encompassing both pretraining and prompting; (iii) it exhibits an exceptional ability to rapidly produces multiple out-of-the-box solutions for a novel task. Consequently, we envision the possibility of combining MLCopilot with existing AutoML methods, opening up an intriguing avenue for future exploration.

Robustness of MLCopilot. As MLCopilot has the ability to accommodate heterogeneous formats of inputs, it is worth discussing the robustness of MLCopilot in the wild. This consideration extends to situations where users submit poorly-formatted task descriptions and when the experience pool includes data with noisy accuracy labels or flawed canonicalization. A detailed assessment of MLCopilot’s robustness is presented in § D.

The experiments conducted shed light on the system’s robustness against certain challenges (e.g., the choice of LLMs and task description formats). But it is still important to note that its performance can degrade under specific conditions, such as when dealing with a severely limited prompt context window length.

References

- Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. 2016. Learning to learn by gradient descent by gradient descent. *Advances in neural information processing systems*, 29.
- Sebastian Pineda Arango, Hadi Samer Jomaa, Martin Wistuba, and Josif Grabocka. 2021. Hpo-b: A large-scale reproducible benchmark for black-box hpo based on openml. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Rémi Bardenet, Mátyás Brendel, Balázs Kégl, and Michele Sebag. 2013. Collaborative hyperparameter tuning. In *International conference on machine learning*, pages 199–207. PMLR.

- Christopher M Bishop and Nasser M Nasrabadi. 2006. *Pattern recognition and machine learning*, volume 4. Springer.
- Daniil A. Boiko, Robert MacKnight, and Gabe Gomes. 2023. [Emergent autonomous scientific research capabilities of large language models](#).
- Leo Breiman. 2001. Random forests. *Machine learning*, 45:5–32.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794.
- Yutian Chen, Xingyou Song, Chansoo Lee, Zi Wang, Richard Zhang, David Dohan, Kazuya Kawakami, Greg Kochanski, Arnaud Doucet, Marc’aurilio Ranzato, et al. 2022. Towards learning universal hyperparameter optimizers with transformers. *Advances in Neural Information Processing Systems*, 35:32053–32068.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, and others. 2022. Palm: Scaling language modeling with pathways. *ArXiv*, abs/2204.02311.
- Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine learning*, 20:273–297.
- Oxford English Dictionary. 1989. Oxford english dictionary. *Simpson, Ja & Weiner, Esc*, 3.
- Matthias Feurer, Aaron Klein, Katharina Eggensperger, Jost Springenberg, Manuel Blum, and Frank Hutter. 2015. Efficient and robust automated machine learning. *Advances in neural information processing systems*, 28.
- Peter I Frazier. 2018. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*.
- Hongwei Han, Jialiang Xu, Mengyu Zhou, Yijia Shao, Shi Han, and Dongmei Zhang. 2022. Luna: Language understanding with number augmentations on transformers via number plugins and pre-training. *arXiv preprint arXiv:2212.02691*.
- Jie Huang and Kevin Chen-Chuan Chang. 2022. Towards reasoning in large language models: A survey. *arXiv preprint arXiv:2212.10403*.
- Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren. 2019. *Automated machine learning: methods, systems, challenges*. Springer Nature.
- Shima Imani, Liang Du, and Harsh Shrivastava. 2023. [Mathprompter: Mathematical reasoning using large language models](#).
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38.
- Lars Kotthoff, Chris Thornton, Holger H Hoos, Frank Hutter, and Kevin Leyton-Brown. 2019. Auto-weka: Automatic model selection and hyperparameter optimization in weka. *Automated machine learning: methods, systems, challenges*, pages 81–95.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2017. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90.
- Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. 2023. Chameleon: Plug-and-play compositional reasoning with large language models. *arXiv preprint arXiv:2304.09842*.
- 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.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the Role of Demonstrations: What Makes In-Context Learning Work? In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Gaurav Mittal, Chang Liu, Nikolaos Karianakis, Victor Fragoso, Mei Chen, and Yun Fu. 2020. Hyperstar: Task-aware hyperparameters for deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8736–8745.
- Yurii Evgen’evich Nesterov. 1983. A method of solving a convex programming problem with convergence rate $O(k^{-2})$. In *Doklady Akademii Nauk*, volume 269, pages 543–547. Russian Academy of Sciences.
- OpenAI. 2023. [Gpt-4 technical report](#).

- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Arkil Patel, Satwik Bhattamishra, and Navin Goyal. 2021. Are nlp models really able to solve simple math word problems? *arXiv preprint arXiv:2103.07191*.
- Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. 2018. Efficient neural architecture search via parameters sharing. In *International conference on machine learning*, pages 4095–4104. PMLR.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. 2019. Analysing mathematical reasoning abilities of neural models. *arXiv preprint arXiv:1904.01557*.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. Hugging-gpt: Solving ai tasks with chatgpt and its friends in huggingface. *arXiv preprint arXiv:2303.17580*.
- Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- Avijit Thawani, Jay Pujara, Pedro A Szekely, and Filip Ilievski. 2021. Representing numbers in nlp: a survey and a vision. *arXiv preprint arXiv:2103.13136*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. Llama: Open and Efficient Foundation Language Models. *ArXiv*, abs/2302.13971.
- Joaquin Vanschoren, Jan N Van Rijn, Bernd Bischl, and Luis Torgo. 2014. Openml: networked science in machine learning. *ACM SIGKDD Explorations Newsletter*, 15(2):49–60.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Chi Wang, Qingyun Wu, Markus Weimer, and Erkang Zhu. 2021a. Flaml: A fast and lightweight automl library. *Proceedings of Machine Learning and Systems*, 3:434–447.
- Zi Wang, George E Dahl, Kevin Swersky, Chansoo Lee, Zelda Mariet, Zachary Nado, Justin Gilmer, Jasper Snoek, and Zoubin Ghahramani. 2021b. Pre-trained gaussian processes for bayesian optimization. *arXiv preprint arXiv:2109.08215*.
- Albert Webson and Ellie Pavlick. 2022. Do prompt-based models really understand the meaning of their prompts? In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2300–2344.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, E. Chi, Tatsunori Hashimoto, Oriol Vinyals, P. Liang, J. Dean, and W. Fedus. 2022. Emergent Abilities of Large Language Models. *ArXiv*, abs/2206.07682.
- Martin Wistuba, Nicolas Schilling, and Lars Schmidt-Thieme. 2016. Two-stage transfer surrogate model for automatic hyperparameter optimization. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2016, Riva del Garda, Italy, September 19-23, 2016, Proceedings, Part I 16*, pages 199–214. Springer.
- Chenqian Yan, Yuge Zhang, Quanlu Zhang, Yaming Yang, Xinyang Jiang, Yuqing Yang, and Baoyuan Wang. 2022. Privacy-preserving online automl for domain-specific face detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4134–4144.
- Wenhao Yu, Dan Iter, Shuohang Wang, Yichong Xu, Mingxuan Ju, Soumya Sanyal, Chenguang Zhu, Michael Zeng, and Meng Jiang. 2022. Generate rather than retrieve: Large language models are strong context generators. *arXiv preprint arXiv:2209.10063*.
- Ningyu Zhang, Xin Xu, Liankuan Tao, Haiyang Yu, Hongbin Ye, Shuofei Qiao, Xin Xie, Xiang Chen, Zhoubo Li, Lei Li, et al. 2022. Deepke: A deep learning based knowledge extraction toolkit for knowledge base population. *arXiv preprint arXiv:2201.03335*.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.
- Mingkai Zheng, Xiu Su, Shan You, Fei Wang, Chen Qian, Chang Xu, and Samuel Albanie. 2023. [Can gpt-4 perform neural architecture search?](#)

A Algorithms

We summarize our method as [algorithm 1](#) and [algorithm 2](#).

Algorithm 1: Offline Stage of MLCopilot

Input : Historical data $H = \{D_1, \dots, D_{N_H}\}$. Maximum iterations *rounds*. Stagnation patience *patience*. Candidate question list *Questions*. Validation tasks *ValTasks*.
Output : Experience Pool P_E ; Knowledge k^* .

```
1  $P_E \leftarrow \{\mathcal{C}(D_i)\}$ ; /*  $\mathcal{C}$  is canonicalization function */
2  $r^* \leftarrow -\infty$ ;
3  $stagnation \leftarrow 0$ ;
4 for  $n=1$  to rounds do
5    $E \leftarrow \text{RandomSample}(P_E)$ ;
6    $q \leftarrow \text{RandomSample}(\text{Questions})$ ; /* Sample one hypophora question */
7    $\tau \leftarrow \text{Uniform}(0, 1)$ ; /* Random temperature */
8    $k \leftarrow \text{LLM}(E, q; \tau)$ ; /* Generate knowledge candidates */
9    $S \leftarrow \text{LLM}(E, k; \text{ValTasks}; 0)$ ; /* Mock online stage on validation tasks */
10   $r \leftarrow \text{Evaluate}(S)$ ; /* Run and evaluate the solution */
11  if  $r > r^*$  then
12     $r^* \leftarrow r$ ;
13     $k^* \leftarrow k$ ;
14     $stagnation \leftarrow 0$ ;
15  else
16     $stagnation \leftarrow stagnation + 1$ ;
17    if  $stagnation > patience$  then
18      break;
19 return  $P_E, k^*$ 
```

Algorithm 2: Online Stage of MLCopilot

Input : A new task description \tilde{T} . Experience pool P_E . Knowledge pool P_K .
Output : Solution \tilde{S} .

```
1  $\tilde{E} \leftarrow \mathcal{R}_E(\tilde{T}, P_E)$ ; /* Retrieve experiences */
2  $\tilde{K} \leftarrow \mathcal{R}_K(\tilde{T}, P_K)$ ; /* Retrieve knowledge */
3  $\tilde{S} \leftarrow \text{LLM}(\tilde{T}, \tilde{E}, \tilde{K})$ ;
4 return  $\tilde{S}$ 
```

B Experiment Details

B.1 Benchmarks

HPO-B. HPO-B-v3 (Arango et al., 2021) comprises 16 solution spaces for 101 datasets obtained from the OpenML (Vanschoren et al., 2014). Each space has a fixed ML algorithm such as random forest (Breiman, 2001), SVM (Cortes and Vapnik, 1995), or XGBoost (Chen and Guestrin, 2016), and the goal is to determine the optimal configuration of the algorithm for a given dataset. HPO-B also provides successful configurations from past tasks, which are canonicalized into experience in our case. Additionally, they have released surrogate models to expedite the evaluation of solutions that have not been attempted before. The final benchmark performance is determined by averaging the normalized accuracy (nAcc) across all datasets, following the normalization protocol in (Arango et al., 2021).

PD1. The PD1 Neural Net Tuning Dataset is proposed by HyperBO (Wang et al., 2021b), consisting of 24 classification tasks, covering image classification, next token prediction, and translation. Each task is associated with a predefined neural network (CNN (Krizhevsky et al., 2017) or transformer (Vaswani et al., 2017)), and has four configurable parameters of a SGD optimizer with Nesterov momentum (Nesterov, 1983). Due to the high cost of training neural networks, evaluating the solutions suggested by MLCopilot by running them in real-time is not feasible. So we created a surrogate model to predict the performance of suggested solutions for each task (see § B.4 for details). As per (Wang et al., 2021b), we report normalized accuracy (nAcc).

HyperFD. HyperFD (Yan et al., 2022) is a benchmark designed to optimize the performance of a neural face detector on an unseen dataset by properly configuring data augmentation, neural architecture, loss function, and training recipe. For this purpose, they formulated a solution space and provided both average precision (AP) and rank for every possible solution on all datasets. The benchmark was published after the claimed knowledge cutoff of GPT-3.5 (September 2021), and it has not been publicly released yet, making it unlikely that it has appeared in the training data of LLM.

B.2 Compared Baselines

Details about the traditional baselines used in our experiment are described below.

- **Random** method randomly generates a solution.
- **ASKL** (Auto-sklearn 1.0) (Feurer et al., 2015) finds the most similar tasks based on manually selected meta-features of tasks and directly uses the best solutions on them.
- **Constant** (Bardenet et al., 2013; Kotthoff et al., 2019) (a.k.a. Average) uses a constant set of solutions for any new task. The produced set of solutions is the one with the best **average** performance on the historical tasks. This method is straightforward and has no specific literature reference. We reference the two literatures for “Constant” as they also adopt a similar baseline for comparison.
- **TST-M** (Wistuba et al., 2016) employs Gaussian processes to approximate the performance of solutions in the solution space for each task. When a new task is encountered, it combines performance predictions of solutions for different tasks and predicts the performance of each solution on the new task by averaging predictions on history tasks weighted by task similarities.
- **HyperSTAR** (Mittal et al., 2020) trains a performance predictor for a joint encoding of solution and task features. HyperSTAR is originally built for vision tasks. To adapt it for non-image tasks, we incorporate handcrafted meta-features as task features.
- **HyperFD** (Yan et al., 2022) is a method specifically designed for the HyperFD benchmark, which uses a sophisticated meta-feature extractor for neural face detection tasks. However, it is not a general-purpose method and is not designed to work with other types of tasks.
- **FLAML-Zero** (Wang et al., 2021a) is a recent method that generates a portfolio of ML solutions through offline meta-training, minimizing overall regret across meta-training tasks. It uses meta-features to link new tasks to existing ones based on their similarity.

B.3 Post-validation

The post-validation step that we have incorporated draws inspiration from established practices within machine learning, where a dedicated validation set is employed to enhance model performance. In our specific case, we allocate 10% of the training meta-dataset for validation purposes, allowing us to systematically filter and select the most valuable generated knowledge. This additional layer of validation contributes significantly to addressing hallucination-related issues.

Detailed steps of post-validation include a loop of: sampling hypophora question, sampling temperature, generating knowledge candidates, validating candidate knowledge, and an earlystopping mechanism that determines stagnation. This is described in [algorithm 1](#).

B.4 Building Surrogate Model for PD1

The metric in PD1 contains many NaN values, which correspond to network training divergence. For benchmarking purposes, it is more important to be able to distinguish the top-performing solutions, *i.e.*, solutions above medium accuracy. To accomplish this, we adopt a two-stage surrogate approach. We use a classification model to distinguish the top-performing solutions, and then two regression models: one specially optimized for the top-performing solutions, and the other one for all solutions. We utilize XGBoost (Chen and Guestrin, 2016) for building classifiers and regressors. Default parameters are used for those models.

C Prompt Design

We show two example prompts used on HPO-B. One of them is used for the online stage, as shown in Table 7, when MLCopilot receives a task description given by the user and sends it to LLM to obtain a recommended solution.

Table 8 shows an example of prompt used during the offline stage, when we generate a series of knowledge candidates. For post-validation, we use the prompt same as the online stage (example in Table 7).

Prompt
Space description
Here are some classification datasets along with best hyper-parameter configurations to train a R language model "Learner mlr.classif.svm from package(s) e1071" on them.
Demonstrations
Dataset: The dataset name is "ada_agnostic". It contains 2 classes, 4562 instances, 49 features, 48 numeric features, 1 categorical features. The majority class size is 3430 and the minority class size is 1132. Configuration 1: cost is very small. kernel is linear. Configuration 2: cost is very small. kernel is linear. Configuration 3: cost is very small. kernel is linear.
Dataset: The dataset name is "credit-g". It contains 2 classes, 1000 instances, 21 features, 7 numeric features, 14 categorical features. The majority class size is 700 and the minority class size is 300. Configuration 1: cost is medium. gamma is small. kernel is radial. Configuration 2: cost is medium. gamma is very small. kernel is radial. Configuration 3: cost is medium. gamma is small. kernel is radial.
Dataset: The dataset name is "ozone-level-8hr". It contains 2 classes, 2534 instances, 73 features, 72 numeric features, 1 categorical features. The majority class size is 2374 and the minority class size is 160. Configuration 1: cost is small. gamma is small. kernel is radial. Configuration 2: cost is very small. gamma is small. kernel is radial. Configuration 3: cost is small. gamma is small. kernel is radial.
Knowledge
Guidelines: 1. For datasets with many numeric features, larger cost values and smaller gamma values tend to be more effective. 2. For datasets with many categorical features, linear kernels tend to be more effective. 3. For datasets with few numeric features, small cost values and larger gamma values tend to be more effective. 4. For datasets with few categorical features, polynomial kernels tend to be more effective.
Instruction
Based on the examples and guidelines above, recommend 3 hyper-parameter configurations for a new classification dataset
Description for new task
Dataset: The dataset name is "gina_agnostic". It contains 2 classes, 3468 instances, 971 features, 970 numeric features, 1 categorical features. The majority class size is 1763 and the minority class size is 1705.

Table 7: Example prompt for HPO-B in online serving.

Prompt
Space description
Here are some classification datasets along with best hyper-parameter configurations to train a R language model "Learner mlr.classif.svm from package(s) e1071" on them.
Demonstrations
Dataset: The dataset name is "wilt". It contains 2 classes, 4839 instances, 6 features, 5 numeric features, 1 categorical features. The majority class size is 4578 and the minority class size is 261. Configuration 1: cost is medium. gamma is large. kernel is radial. Configuration 2: cost is medium. gamma is medium. kernel is radial. Configuration 3: cost is large. gamma is medium. kernel is radial.
Dataset: The dataset name is "ilpd". It contains 2 classes, 583 instances, 11 features, 9 numeric features, 2 categorical features. The majority class size is 416 and the minority class size is 167. Configuration 1: cost is medium. gamma is medium. kernel is radial. Configuration 2: cost is very small. gamma is very large. kernel is radial. Configuration 3: cost is medium. gamma is very large. kernel is radial.
Dataset: The dataset name is "steel-plates-fault". It contains 2 classes, 1941 instances, 34 features, 33 numeric features, 1 categorical features. The majority class size is 1268 and the minority class size is 673. Configuration 1: cost is small. kernel is linear. Configuration 2: cost is very small. kernel is linear. Configuration 3: cost is very small. kernel is linear.
Instruction
Q: From the examples above, what patterns can we observe about the relationship between dataset characteristics and the best hyper-parameter configurations? Answer MUST be concise, critical, point-by-point, line-by-line, and brief. Only include relevant observations without unnecessary elaboration.

Table 8: Example prompt for HPO-B in the offline stage.

Table 9: Comparison of different description formats.

Description Format	HPO-B \uparrow	PD1 \uparrow	HyperFD \downarrow
Original	81.59 \pm 0.94	1.48 \pm 0.06	59.74 \pm 1.89
Condense	79.66 \pm 0.06	1.52 \pm 0.01	57.33 \pm 3.21
Anonymous	77.43 \pm 0.04	1.21 \pm 0.06	68.42 \pm 10.01
Misleading names	75.80 \pm 3.01	1.43 \pm 0.05	62.52 \pm 3.92




Figure 3: Effect of the experience number and the number of solutions demonstrated for each task.

D Robustness of MLCopilot

D.1 Task description in the wild.

When MLCopilot is deployed, it may be impractical to require users to strictly adhere to a specific format when writing task descriptions. We must consider if MLCopilot is robust enough to handle various formats. To simulate diverse formats, we ask GPT-3.5 to rewrite the descriptions by: (i) condensing the original task descriptions; and (ii) anonymizing the descriptions by removing task names. The results are shown in Table 9. We observed fluctuations in performance when the description format changed, indicating that LLM is sensitive to prompt format. This aligns with the previous researches (Webson and Pavlick, 2022; Lu et al., 2022) suggesting that LLMs may not interpret inputs in the same way as humans. Performance was particularly worse when tasks were anonymized, leading us to conjecture that task names stimulate neurons in the LLM that are important for solving the relevant tasks and also leverage the previous knowledge already memorized from the training corpus of LLMs. To further verify this, we conducted an additional experiment by randomly swapping task names between tasks, and surprisingly observed performance improvement on PD1 and HyperFD. This echoes the finding in (Min et al., 2022), which suggests that “random labels are better than no labels at all”.

D.2 Length of prompt.

We study the effect of prompt length, which is mainly influenced by the number of retrieved ML experiences as demonstrations, to the performance. In our previous experiments, we retrieved as many experiences as possible, either until all the experiences had been retrieved or the maximum prompt length was reached. For each task, the three best solutions were demonstrated, which was an arbitrary choice in our early experiments. In this section, we vary these two factors. As shown in Figure 3, the performance generally improves as the number of demonstrations in the prompt increases, measured by prompt tokens. However, the performance soon saturates at around 1.5k tokens and fluctuates. Moreover, demonstrating more solutions for each task leverages more data and has higher potential, especially above 3k tokens.

D.3 Choice of LLMs.

We report performance of MLCopilot if equipped with LLMs other than GPT-3.5 (code-named “text-davinci-003”) used in our main experiments. The models we have experimented with include:

Method (with LLM)	HPO-B \uparrow	PD1 \uparrow	HyperFD \downarrow
ASKL	77.01 \pm 0.00	1.26 \pm 0.00	92.58 \pm 0.00
FLAML	77.84 \pm 0.00	1.28 \pm 0.00	66.42 \pm 0.00
MLCopilot (gpt-3.5-turbo)	81.67 \pm 1.99	1.37 \pm 0.10	72.41 \pm 10.48
MLCopilot (text-davinci-001)	82.13 \pm 2.05	1.58 \pm 0.04	71.25 \pm 10.70
MLCopilot (LLAMA-7B)	79.51 \pm 0.57	1.43 \pm 0.08	67.47 \pm 5.60
MLCopilot (text-davinci-003)	81.59 \pm 0.94	1.48 \pm 0.06	59.74 \pm 1.89

Table 10: Performance of MLCopilot equipped with different LLMs.

- GPT-3.5 Turbo⁷: a cost-efficient version of GPT-3.5 that uses chat completion as its user interface.
- GPT-3 (Brown et al., 2020) (code-named “text-davinci-001”): the original GPT-3 model trained without instruction finetuning.
- LLAMA-7B (Touvron et al., 2023): a well-known open-source model with a large user community, with a relatively loose requirement of GPU memory.

We compare the results with the original results with text-davinci-003 (GPT-3.5) and our main baselines. As shown in Table Table 10, MLCopilot is robust to choices of LLMs. It is compatible with all the LLMs we have tested and achieves competitive results under different settings. Also, we see a trend that when working with stronger and larger models, MLCopilot still achieves even better results.

D.4 Noisy accuracy and faulty canonicalization.

Method	w/ original data	w/ perturbed data
FLAML	77.84 \pm 0.00	73.53 \pm 0.00
MLCopilot	81.59 \pm 0.94	78.54 \pm 3.25

Table 11: Impact of noises in accuracy.

Method	nAcc
MLCopilot (original)	81.59 \pm 0.94
MLCopilot (faulty canonicalization)	77.11 \pm 1.77
FLAML	77.84 \pm 0.00

Table 12: Effect of faulty canonicalization.

We discuss the cases where the experience pool is polluted during the operation of the system. We evaluated the robustness under such scenarios on HPO-B (the largest solution space).

Firstly, we assessed the impact of noises in accuracy, by perturbing the accuracies in historical data. We added Gaussian noise to the accuracy values. The standard deviation of the Gaussian noise is 10% of the accuracy distribution. As a result of such disturbance, suboptimal configurations might pop up as the best configurations and serve as demonstrations in the prompt. After experimenting on HPO-B, we found (in Table 11) that MLCopilot does suffer from such disturbance (performance drop from 81.59 to 78.54). However, such a result is still competitive with the state-of-the-art baselines (FLAML 77.84). Moreover, if a similar disturbance was done to the input data of FLAML, its performance further drops to 73.53. Full results (top-1 normalized accuracy) are shown in the table below.

Secondly, we investigate the effect of faulty canonicalization. In our paper, we showed (in Table 6) that canonicalization is an important component of MLCopilot, and a misconfigured canonicalization can lead to degraded performance. Following your suggestions, we introduced random noise into the canonicalization process by replacing 10% of the canonicalized data with random discrete values. That is, each parameter of the configuration has a 10% probability to be replaced with a random choice from “very low”, “low”, “medium”, “high”, “very high”. We show the results (top-1 normalized accuracy on HPO-B) in Table 12.

Although the result is still competitive with baseline FLAML, we can see that faulty canonicalization does lead to worse performance. Notably, the impact is even more severe than the setting of perturbed accuracy. We speculate that false canonicalization can be particularly misleading for the logical reasoning of large language models. We will include a discussion of these findings in our revision.

⁷<https://openai.com/blog/introducing-chatgpt-and-whisper-apis>

E Knowledge

All contents in this section are generated by Large Language Models.

E.1 HPO-B

HPO-B contains 16 design spaces. We finalize one set of knowledge for each space.

Space: 5860

1. Generally, datasets with more **numeric features** require larger alphas and smaller lambdas for better performance.
2. Datasets with a higher ratio of minority to majority **class size** require smaller alphas and larger lambdas for better performance.
3. Datasets with **more features** require larger alphas and smaller lambdas for better performance.
4. Datasets with more **categorical** features require larger alphas and larger lambdas for better performance.

Space: 4796

1. For datasets with a large majority **class size** and a small minority **class size**, a larger cp and minbucket size tend to be better hyper-parameter configurations.
2. For datasets with a small majority **class size** and a large minority **class size**, a smaller cp and minbucket size tend to be better hyper-parameter configurations.
3. For datasets with a large number of **numeric features**, a larger cp and minbucket size tend to be better hyper-parameter configurations.
4. For datasets with a small number of **numeric features**, a smaller cp and minbucket size tend to be better hyper-parameter configurations.
5. For datasets with a large number of **categorical** features, a smaller cp and minbucket size tend to be better hyper-parameter configurations.
6. For datasets with a small number of **categorical** features, a larger cp and minbucket size tend to be better hyper-parameter configurations.

Space: 5971

1. Generally, **larger datasets** require higher nrounds and larger subsample values.
2. The majority **class size** and minority **class size** of the dataset can influence the configuration of alpha, booster, colsample bylevel, colsample bytree, eta, lambda, max depth, min child weight, nrounds, and subsample.
3. The number of numeric and **categorical** features in the dataset can determine the booster used.
4. The **size of the dataset** can influence the configuration of eta, lambda, max depth, min child weight, nrounds, and subsample.
5. The **size of the minority** class can determine the configuration of alpha, colsample bylevel, colsample bytree, eta, lambda, max depth, min child weight, nrounds, and subsample.

Space: 6766

1. For datasets with a larger majority **class size**, high values of alpha and low values of lambda tend to perform better.
2. For datasets with a smaller majority **class size**, low values of alpha and high values of lambda tend to perform better.
3. For datasets with more **numeric features**, medium values of alpha and low values of lambda tend to perform better.
4. For datasets with more **categorical features**, high values of alpha and large values of lambda tend to perform better.
5. For datasets with a **larger number of features**, high values of alpha and large values of lambda tend to perform better.

Space: 5965

1. The larger the majority **class size**, the smaller the min node size and sample fraction tend to be.
2. The larger the minority **class size**, the larger the min node size and sample fraction tend to be.
3. The **larger the number of features**, the larger the mtry tends to be.
4. The larger the number of **numeric features**, the larger the mtry tends to be.
5. The larger the number of **categorical features**, the smaller the mtry tends to be.
6. The **larger the number of trees**, the smaller the mtry tends to be.
7. The **larger the number of instances**, the larger the sample fraction tends to be.
8. The replace parameter is **usually** set to True.
9. The respect unordered factors parameter is **usually** set to False.

Space: 5906

1. **Smaller datasets** tend to have smaller alpha and eta values, while **larger datasets** tend to have larger values.
2. Datasets with **more features** tend to have larger colsample bylevel and colsample bytree values, while datasets with fewer features tend to have smaller values.
3. Datasets with more **numeric features** tend to have larger lambda and max depth values, while datasets with fewer **numeric features** tend to have smaller values.
4. **Smaller datasets** tend to have smaller nrounds and subsample values, while **larger datasets** tend to have larger values.
5. Datasets with more **categorical features** tend to have smaller min child weight values, while datasets with fewer **categorical features** tend to have larger values.

Space: 7607

1. The min node size generally decreases as the **dataset size increases**.
2. The mtry is **usually** small for datasets with few features and large for datasets with many features.
3. The num trees is **usually** small for datasets with few instances and large for datasets with many instances.
4. Replace is **usually** set to False for small datasets and True for

large datasets.

5. Respect unordered factors is **usually** set to False for datasets with few **categorical** features and True for datasets with many **categorical** features.
6. Sample fraction is **usually** set to small for datasets with few instances and large for datasets with many instances.

Space: 6794

1. For datasets with a large majority **class size**, larger min node size and sample fraction values are **usually** used, while for datasets with a smaller majority **class size**, smaller min node size and sample fraction values are **usually** used.
2. For datasets with **more features**, larger mtry values are **usually** used.
3. For datasets with more **numeric features**, replace is **usually** set to True, while for datasets with more **categorical** features, replace is **usually** set to False.
4. Respect unordered factors is **usually** set to True when the dataset has more **categorical** features.

Space: 7609

1. For datasets with **more features**, larger mtry values are preferred.
2. For datasets with **more instances**, larger sample fractions are preferred.
3. For datasets with **more majority class instances**, smaller min node sizes are preferred.
4. For datasets with more **numeric features**, replace is typically set to True.
5. For datasets with more **categorical** features, respect unordered factors is typically set to False.
6. For datasets with a more balanced **class size**, num trees is typically set to a smaller value.

Space: 5859

1. **larger datasets** tend to require smaller cp values and larger minbucket values.
2. **Smaller datasets** tend to require larger cp values and smaller minbucket values.
3. For **larger datasets**, maxdepth tends to be very large or medium, whereas for **Smaller datasets**, maxdepth tends to be very small or small.
4. For **larger datasets**, minsplit tends to be very large or large, whereas for **Smaller datasets**, minsplit tends to be very small or small.

Space: 5889

1. The **larger the dataset size**, the larger the mtry and num trees, and the smaller the sample fraction.
2. The larger the majority **class size**, the larger the mtry and num trees, and the smaller the sample fraction.

3. The **smaller the number of features**, the smaller the mtry and num trees, and the larger the sample fraction.
4. The more **numeric features**, the larger the mtry and num trees, and the smaller the sample fraction.
5. The more **categorical** features, the smaller the mtry and num trees, and the larger the sample fraction.
6. The replace parameter is **usually** set to True.

Space: 6767

1. Datasets with a larger majority **class size** tend to require larger nrounds and larger subsample values.
2. Datasets with more **numeric features** tend to require larger colsample bylevel and colsample bytree values.
3. Datasets with more **categorical** features tend to require smaller min child weight values.
4. Datasets with a smaller minority **class size** tend to require smaller eta and lambda values.
5. Datasets with **more features** tend to require larger max depth values.

Space: 5970

1. For datasets with more **numeric features**, smaller alpha and smaller lambda values tend to be the best hyper-parameter configurations.
2. For datasets with more **categorical** features, larger alpha and larger lambda values tend to be the best hyper-parameter configurations.
3. For datasets with majority **class size** significantly larger than minority **class size**, larger alpha and larger lambda values tend to be the best hyper-parameter configurations.

Space: 5527

1. The cost parameter tends to increase as the **dataset size increases**.
2. The gamma parameter tends to decrease as the number of numeric features increases.
3. The kernel parameter tends to be radial for datasets with numeric features, and polynomial or linear for datasets with **categorical** features.
4. The degree parameter tends to increase as the number of **categorical** features increases.

Space: 5636

1. The larger the majority **class size**, the smaller the cp value should be.
2. The larger the minority **class size**, the larger the cp value should be.
3. The **larger the number of features**, the smaller the maxdepth value should be.
4. The larger the number of **numeric features**, the larger the minbucket value should be.
5. The larger the number of **categorical** features, the smaller the

minbucket value should be.

6. The **larger the number of instances**, the larger the minsplit value should be.

Space: 5891

1. For datasets with many **numeric features**, larger cost values and smaller gamma values tend to be more effective.
2. For datasets with many **categorical features**, linear kernels tend to be more effective.
3. For datasets with few **numeric features**, small cost values and larger gamma values tend to be more effective.
4. For datasets with few **categorical features**, polynomial kernels tend to be more effective.

E.2 PD1

We performed leave-one-out evaluation on the PD1 benchmark, which consists of 23 tasks. However, some tasks are using the same model and dataset but only different in batch size. These tasks should not appear in training tasks and test tasks at the same time (Wang et al., 2021b). Therefore, only 13 distinct sets of training tasks were available for testing. For each set of training tasks, we generated a corresponding set of knowledge, which is presented below.

Test task: CIFAR100, Wide ResNet

1. Set the initial learning rate (LR) according to the **size of the dataset** and the complexity of the model.
2. Set the momentum parameter to a lower value for **larger datasets** and a higher value for **simpler models**.
3. Set the power parameter to a higher value for **more complex models**.
4. Set the lambda parameter to a higher value for **more complex models** and a lower value for **simpler models**.

Test task: CIFAR10, Wide ResNet

1. Adjust the initial learning rate and momentum based on the size and complexity of the dataset: higher for **large and complex datasets**, lower for **small and simple datasets**.
2. Adjust the power and lambda parameters based on the desired **speed of the learning process**: higher power and lower lambda for faster learning, lower power and higher lambda for slower learning.
3. Consider any **domain-specific constraints** when configuring the optimizer, **such as** accuracy requirements.

Test task: Fashion-MNIST, Max Pooling CNN with ReLU

1. Set the initial learning rate to a **low or medium** value.
2. Set the momentum to a **high or medium** value.
3. Set the power to a **low or medium** value.
4. Set the lambda to a **low or medium** value.
5. Adjust the initial learning rate, momentum, power, and lambda according to the characteristics of the task, **such as the dataset size, model architecture, and the complexity** of the prediction task. For example, for tasks with **larger datasets**, a higher initial learning rate may be beneficial, while for tasks with smaller

datasets, a lower initial learning rate may be more suitable. Similarly, for tasks with [more complex models](#), a higher momentum may be beneficial, while for [simpler models](#), a lower momentum may be more suitable. Additionally, for tasks with more [complex prediction tasks](#), a higher power may be beneficial, while for [simpler tasks](#), a lower power may be more suitable. Finally, for tasks with more complex models, a higher lambda may be beneficial, while for [simpler models](#), a lower lambda may be more suitable.

Test task: Fashion-MNIST, Max Pooling CNN with Tanh

1. Choose an initial LR that is appropriate for the [size of the dataset](#) and complexity of the model.
2. Set the momentum to a value that is appropriate for [the size of the dataset](#) and complexity of the model.
3. Set the power parameter to a value that is appropriate for the [size of the dataset](#) and complexity of the model.
4. Set the lambda parameter to a value that is appropriate for the [size of the dataset](#) and complexity of the model.

Test task: Fashion-MNIST, Simple CNN

1. Set the initial learning rate (LR) to a value that is appropriate for the [size of the dataset](#).
2. Set the momentum to a value that is appropriate for [the size of the dataset](#).
3. Set the power parameter to a value that is appropriate for the [size of the dataset](#).
4. Set the lambda parameter to a value that is appropriate for the [size of the dataset](#) and the desired [level of regularization](#).

Test task: ImageNet, ResNet50

1. For tasks with [larger batch sizes](#), use a higher initial learning rate and higher momentum. For tasks with [smaller batch sizes](#), use a lower initial learning rate and lower momentum.
2. For tasks with [larger vocabularies](#), use a higher lambda value. For tasks with smaller vocabularies, use a lower lambda value.
3. For tasks with [more complex models](#), use a higher power value. For tasks with [simpler models](#), use a lower power value.

Test task: LM1B, Transformer

1. Set the initial learning rate to a value that is suitable for the [size and complexity of the dataset](#).
2. Set the momentum to a value that is suitable for the [size and complexity of the dataset](#).
3. Set the power parameter to a value that is suitable for the [noise and outliers in the dataset](#).
4. Set the lambda parameter to a value that is suitable for the [noise and outliers in the dataset](#).

Test task: MNIST, Max Pooling CNN with ReLU

1. Set the initial learning rate to a **low or medium** value.
2. Set the momentum to a **high or medium** value.
3. Set the power to a **low or medium** value.
4. Set the lambda to a low or high value.
5. Consider the characteristics of the task, **such as** the **dataset size, model architecture, and the complexity** of the prediction task, when adjusting the parameters.
6. For tasks with **larger datasets**, a higher initial learning rate and lower momentum may be more suitable.
7. For tasks with **Smaller datasets**, a lower initial learning rate and higher momentum may be more suitable.
8. For tasks with **more complex models**, a higher initial learning rate and lower momentum may be more suitable.
9. For tasks with **simpler models**, a lower initial learning rate and higher momentum may be more suitable.
10. For tasks with more **complex prediction tasks**, a higher initial learning rate and lower momentum may be more suitable.
11. For tasks with **simpler prediction tasks**, a lower initial learning rate and higher momentum may be more suitable.

Test task: MNIST, Max Pooling CNN with Tanh

1. Set the initial learning rate to a high value to ensure that the model is able to learn quickly and efficiently.
2. Set the momentum to a low value to prevent the model from overfitting.
3. Set the power and/or lambda to high values to ensure that the learning rate decays slowly and the model is able to continue learning for a longer period of time.
4. For tasks **such as** training a CNN with max-pool and ReLU on Fashion MNIST, set the initial learning rate to a low value to prevent the model from overfitting.
5. For tasks **such as** training a ResNet50 on ImageNet, set the initial learning rate to a high value to ensure that the model is able to learn quickly and efficiently.
6. For tasks **such as** training a Wide ResNet on CIFAR100, set the initial learning rate to a very high value to ensure that the model is able to learn quickly and efficiently.
7. For tasks **such as** training a Transformer on UniRef50, set the initial learning rate to a low value to prevent the model from overfitting.

Test task: MNIST, Simple CNN

1. Adjust the initial learning rate and momentum according to the **size and complexity of the dataset**: higher for large and complex datasets, lower for **small and simple datasets**.
2. Adjust the power and lambda parameters according to the **size and complexity of the dataset**: higher for **large and complex datasets**, lower for **small and simple datasets**.
3. Adjust the initial learning rate and momentum according to the task requirements: higher for tasks requiring high accuracy, lower

for tasks requiring high speed.

Test task: SVHN, Wide ResNet

1. For **image classification tasks**, set the initial learning rate (LR) to a higher value and the momentum to a lower value.
2. For **language tasks**, set the initial LR to a lower value and the momentum to a higher value.
3. For tasks with **larger batch sizes**, set the initial LR to a higher value and the momentum to a lower value.
4. For tasks with **smaller batch sizes**, set the initial LR to a lower value and the momentum to a higher value.
5. For tasks with **more complex models**, set the power to a higher value and the lambda to a higher value.
6. For tasks with **simpler models**, set the power to a lower value and the lambda to a lower value.

Test task: UniRef50, Transformer

1. Set the initial learning rate, momentum, power, and lambda values according to the following guidelines:
 - For **larger batch sizes** and **larger datasets**, use a higher learning rate, higher momentum, lower power, and higher lambda.
 - For **smaller batch sizes** and **Smaller datasets**, use a lower learning rate, lower momentum, higher power, and lower lambda.
2. Examples:
 - For a CNN with max-pool and ReLU on Fashion MNIST with a batch size of 256, use an initial learning rate of 0.001, a momentum of 0.9, a power of 0.1, and a lambda of 0.01.
 - For a Wide ResNet on CIFAR10 with a batch size of 2048, use an initial learning rate of 0.01, a momentum of 0.9, a power of 0.5, and a lambda of 0.001.
 - For a Transformer on LM1B with a batch size of 2048, use an initial learning rate of 0.001, a momentum of 0.9, a power of 0.01, and a lambda of 0.001.

Test task: WMT15, xformer

1. Set the initial learning rate to a **low or medium** value.
2. Set the momentum to a **high or medium** value.
3. Set the power to a **low or medium** value.
4. Set the lambda to a **high or medium** value.
5. Adjust the initial learning rate and momentum based on the characteristics of the task, **such as** the dataset size, model architecture, and the complexity of the prediction task. For example, for tasks with **larger datasets**, a higher initial learning rate and a lower momentum may be more suitable, while for tasks with smaller datasets, a lower initial learning rate and a higher momentum may be more suitable. Additionally, for tasks with **more complex models**, a higher initial learning rate and a lower momentum may be more suitable, while for tasks with **simpler models**, a lower initial learning rate and a higher momentum may be more suitable. Finally, for tasks with more **complex prediction tasks**, a higher initial learning rate and a lower momentum may be more suitable, while for

tasks with [simpler prediction tasks](#), a lower initial learning rate and a higher momentum may be more suitable.

E.3 HyperFD

Similar to PD1, evaluation on HyperFD is also leave-one-out on 12 tasks. We show 12 sets of knowledge based on the choices of test tasks.

Test task: AFLW

1. Configure crop size and anchor matching IoU threshold based on the [number of faces in the dataset](#):
 - For [datasets with more faces](#), use larger crop sizes and higher anchor matching IoU thresholds.
 - For [datasets with fewer faces](#), use smaller crop sizes and lower anchor matching IoU thresholds.
2. Configure learning rate and negative to positive ratio based on the [number of faces in the dataset](#):
 - For [datasets with more faces](#), use higher learning rates and more negative to positive ratios.
 - For [datasets with fewer faces](#), use lower learning rates and fewer negative to positive ratios.
3. Configure location loss weight based on the [presence of facial landmarks in the dataset](#):
 - For datasets with [facial landmarks](#), use higher location loss weights.

Test task: ANIME

1. Set the crop size and anchor matching IoU threshold according to the [number of faces in the dataset](#):
 - For [datasets with more faces](#), use larger crop sizes and higher anchor matching IoU thresholds.
 - For [datasets with fewer faces](#), use smaller crop sizes and lower anchor matching IoU thresholds.
2. Set the location loss weight according to the [presence of facial landmarks in the dataset](#):
 - For datasets with [facial landmarks](#), use higher location loss weights.
 - For datasets without [facial landmarks](#), use lower location loss weights.
3. Set the learning rate and optimizer according to the [negative to positive ratio in the dataset](#):
 - For datasets with higher negative to positive ratios, use higher learning rates and optimizers [such as](#) SGD or Adam.

Test task: FaceMask

1. Set the crop size according to the [number of faces in the dataset](#): larger crop sizes for [datasets with more faces](#), and smaller crop sizes for [datasets with fewer faces](#).
2. Set the anchor matching IoU threshold according to [the number of faces](#) in the dataset: higher thresholds for [datasets with more faces](#), and lower thresholds for [datasets with fewer faces](#).
3. Set the location loss weight according to the [presence of facial](#)

landmarks in the dataset: higher weights for datasets with facial landmarks, and lower weights for datasets without **facial landmarks**.

4. Set the negative to positive ratio according to **the number of faces** in the dataset: higher ratios for **datasets with more faces**, and lower ratios for **datasets with fewer faces**.
5. Set the learning rate according to the number of faces in the dataset: higher rates for **datasets with more faces**, and lower rates for **datasets with fewer faces**.

Test task: FDDB

1. Set the crop size to be larger and the anchor matching IoU threshold to be higher for **datasets with more faces**.
2. Increase the location loss weight and decrease the negative to positive ratio for **datasets with more faces**.
3. Use a lower learning rate and an optimizer **such as** Adam or SGD for datasets with **facial landmarks**.

Test task: FDDB-360

1. For **datasets with more faces**, use a larger crop size and a higher anchor matching IoU threshold.
2. For **datasets with fewer faces**, use a smaller crop size and a lower anchor matching IoU threshold.
3. For datasets with no **facial landmarks**, use a lower location loss weight and a higher negative to positive ratio.
4. For datasets with **facial landmarks**, use a higher location loss weight and a lower negative to positive ratio.
5. For **datasets with more faces**, use a higher learning rate and an SGD optimizer.
6. For **datasets with fewer faces**, use a lower learning rate and an Adam optimizer.

Test task: MAFA

1. Set the crop size and anchor matching IoU threshold according to the number of faces per image in the dataset: larger crop sizes and higher IoU thresholds for **datasets with more faces** per image, and smaller crop sizes and lower IoU thresholds for datasets with fewer faces per image.
2. Set the location loss weight according to the **presence of facial landmarks in the dataset**: higher weights for datasets with facial landmarks, and lower weights for datasets without **facial landmarks**.
3. Set the negative to positive ratio according to the **difficulty of the dataset**: higher ratios for datasets with more challenging scenarios (e.g. weather-based degradations, motion blur, focus blur).
4. Set the learning rate and optimizer according to the **size of the dataset**: lower learning rates and optimizers **such as** Adam or SGD for datasets with more images.

Test task: PASCAL VOC

1. Set the crop size according to the **number of faces in the dataset**: larger crop sizes for **datasets with more faces**, and smaller crop

sizes for [datasets with fewer faces](#).

2. Set the anchor matching IoU threshold according to [the number of faces](#) in the dataset: higher thresholds for [datasets with more faces](#), and lower thresholds for [datasets with fewer faces](#).
3. Set the location loss weight according to the [presence of facial landmarks in the dataset](#): lower weights for datasets with no facial landmarks, and higher weights for datasets with [facial landmarks](#).
4. Set the negative to positive ratio according to [the number of faces](#) in the dataset: higher ratios for [datasets with more faces](#), and lower ratios for [datasets with fewer faces](#).
5. Set the learning rate and optimizer according to the [difficulty of the dataset](#): higher learning rates and optimizers [such as](#) SGD for more challenging datasets.

Test task: UFDD

1. Set the crop size and anchor matching IoU threshold according to the [number of faces in the dataset](#): larger crop size and higher IoU threshold for [datasets with more faces](#), smaller crop size and lower IoU threshold for [datasets with fewer faces](#).
2. Set the location loss weight and negative to positive ratio according to the [number of faces in the dataset](#): higher location loss weight and higher negative to positive ratio for [datasets with more faces](#), lower location loss weight and lower negative to positive ratio for [datasets with fewer faces](#).
3. Set the learning rate and optimizer according to the presence of [facial landmarks](#) in the dataset: lower learning rate and Adam optimizer for datasets with [facial landmarks](#), higher learning rate and SGD optimizer for datasets without [facial landmarks](#).

Test task: UMDAA-02

1. Set the crop size according to the [number of faces in the dataset](#): larger crop sizes for [datasets with more faces](#), and smaller crop sizes for [datasets with fewer faces](#).
2. Set the anchor matching IoU threshold according to [the number of faces](#) in the dataset: higher thresholds for [datasets with more faces](#), and lower thresholds for [datasets with fewer faces](#).
3. Set the location loss weight according to the [presence of facial landmarks in the dataset](#): higher weights for datasets with facial landmarks, and lower weights for datasets without [facial landmarks](#).
4. Set the negative to positive ratio according to [the number of faces](#) in the dataset: higher ratios for [datasets with more faces](#), and lower ratios for [datasets with fewer faces](#).
5. Set the learning rate and optimizer according to the [difficulty of the dataset](#): higher learning rates and optimizers [such as](#) SGD or Adam for more challenging datasets.

Test task: WIDER FACE

1. Set the crop size to a value that is proportional to [the number of faces](#) in the dataset.
2. Set the anchor matching IoU threshold to a value that is proportional to the [number of faces in the dataset](#).

3. Set the negative to positive ratio to a value that is proportional to the [number of faces in the dataset](#).
4. Set the learning rate to a value that is proportional to the [number of faces in the dataset](#).
5. If the dataset contains [facial landmarks](#), set the location loss weight to a value that is proportional to the number of faces in the dataset.

Test task: WIDER-FACE-360

1. Set the crop size and anchor matching IoU threshold according to the [number of faces in the dataset](#): larger crop size and higher IoU threshold for [datasets with more faces](#), and smaller crop size and lower IoU threshold for [datasets with fewer faces](#).
2. Set the location loss weight according to the presence of facial landmarks: higher weight for datasets with [facial landmarks](#), and lower weight for datasets without [facial landmarks](#).
3. Set the negative to positive ratio according to [the number of faces](#) in the dataset: higher ratio for [datasets with more faces](#), and lower ratio for [datasets with fewer faces](#).
4. Set the learning rate according to the number of faces in the dataset: higher rate for [datasets with more faces](#), and lower rate for [datasets with fewer faces](#).
5. Set the optimizer according to the [number of faces in the dataset](#): SGD for [datasets with more faces](#), and Adam for datasets with fewer faces.

Test task: WIKI

1. Set the crop size to a value that is proportional to [the number of faces](#) in the dataset.
2. Set the anchor matching IoU threshold to a value that is proportional to the [number of faces in the dataset](#).
3. Set the location loss weight to a value that is proportional to the presence of [facial landmarks](#) in the dataset.
4. Set the learning rate to a value that is inversely proportional to the [negative to positive ratio in the dataset](#).
5. Use an optimizer [such as](#) Adam or SGD.