



## OPEN ACCESS

## EDITED BY

Weiyong Si,  
University of Essex, United Kingdom

## REVIEWED BY

Haolin Fei,  
Lancaster University, United Kingdom  
Xiaohao Wen,  
Guangxi Normal University, China  
Xu Ran,  
University of Essex, United Kingdom  
Zengliang Han,  
Nanjing University of Aeronautics and  
Astronautics, China

## \*CORRESPONDENCE

Jiapeng Yao  
✉ yaojiapeng1983@163.com

RECEIVED 21 March 2025

ACCEPTED 08 July 2025

PUBLISHED 05 August 2025

## CITATION

Yao J, Zhang L and Huang J (2025) Evaluation of large language model-driven AutoML in data and model management from human-centered perspective.  
*Front. Artif. Intell.* 8:1590105.  
doi: 10.3389/frai.2025.1590105

## COPYRIGHT

© 2025 Yao, Zhang and Huang. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

# Evaluation of large language model-driven AutoML in data and model management from human-centered perspective

Jiapeng Yao<sup>1\*</sup>, Lantian Zhang<sup>2</sup> and Jiping Huang<sup>3</sup>

<sup>1</sup>School of Data Science and Artificial Intelligence, Wenzhou University of Technology, Wenzhou, Zhejiang, China, <sup>2</sup>School of Computer Science and Engineering, Southeast University, Nanjing, Jiangsu, China, <sup>3</sup>Haikou Qiongzhou Women's and Children's Hospital, Haikou, Hainan, China

As organizations increasingly seek to leverage machine learning (ML) capabilities, the technical complexity of implementing ML solutions creates significant barriers to adoption and impacts operational efficiency. This research examines how Large Language Models (LLMs) can transform the accessibility of ML technologies within organizations through a human-centered Automated Machine Learning (AutoML) approach. Through a comprehensive user study involving 15 professionals across various roles and technical backgrounds, we evaluate the organizational impact of an LLM-based AutoML framework compared to traditional implementation methods. Our research offers four significant contributions to both management practice and technical innovation: First, we present pioneering evidence that LLM-based interfaces can dramatically improve ML implementation success rates, with 93.34% of users achieved superior performance in the LLM condition, with 46.67% showing higher accuracy (10%–25% improvement over baseline) and 46.67% demonstrating significantly higher accuracy (>25% improvement over baseline), while 6.67% maintained comparable performance levels; and 60% reporting substantially reduced development time. Second, we demonstrate how natural language interfaces can effectively bridge the technical skills gap in organizations, cutting implementation time by 50% while improving accuracy across all expertise levels. Third, we provide valuable insights for organizations designing human-AI collaborative systems, showing that our approach reduced error resolution time by 73% and significantly accelerated employee learning curves. Finally, we establish empirical support for natural language as an effective interface for complex technical systems, offering organizations a path to democratize ML capabilities without compromising quality or performance.

## KEYWORDS

large language models, automated machine learning, human-computer interaction, deep learning, natural language interfaces

## 1 Introduction

The exponential growth in machine learning (ML) applications has transformed numerous sectors, from healthcare (Ke et al., 2020a,b; Shen et al., 2024c) to scientific research (Wang and Shen, 2024). However, implementing these ML models remains a challenge due to the complex technical requirements involved. Deep learning (DL) models, while demonstrating remarkable capabilities across computer vision, natural language processing, and other domains, require extensive expertise (Liu et al., 2024a). This expertise

barrier includes understanding model architectures (Shen et al., 2023, 2024b; Ke et al., 2023b), managing data pre-processing (Ke et al., 2023a; Luo S. et al., 2024; Wang et al., 2024; Wen et al., 2024), implementing training procedures (Shen et al., 2022; Shen, 2024), and handling deployment, which are tasks that typically require years of specialized education and experience. As a result, many potential users and organizations that could benefit from ML technologies remain unable to effectively implement them, creating a widening gap between ML's potential and its practical accessibility.

Automated Machine Learning (AutoML) emerged as a potential solution to this accessibility challenge by attempting to automate these aspects of the ML pipeline (Sun et al., 2023; Hutter et al., 2019). Traditional AutoML systems aim to streamline various technical processes, including feature engineering, model selection, hyperparameter optimization, and deployment workflows (Hutter et al., 2019; Baratchi et al., 2024; Patibandla et al., 2021). Notable implementations like Auto-Sklearn (Feurer et al., 2015), TPOT (Olson and Moore, 2016), Auto-Keras (Jin et al., 2019), H2O (LeDell and Poirier, 2020), AutoGluon (Erickson et al., 2020), and Auto-Pytorch (Zimmer et al., 2021), and platforms like Azure Machine Learning, Google Cloud AutoML, H2O Driverless AI, etc. have demonstrated success in reducing the technical overhead of machine learning implementation. However, these AutoML tools still present usability challenges as users must navigate complex configuration interfaces, understand technical parameters, and possess programming knowledge to effectively utilize these tools. Furthermore, traditional AutoML methods often require users to make critical decisions about model selection and configuration without providing intuitive guidance or explanation. This limitation means that even AutoML solutions, despite their automation capabilities, remain largely inaccessible to non-expert users, particularly those without substantial programming experience or machine learning background.

Recent advances in Large Language Models (LLMs) have opened new possibilities for human-computer interaction, offering natural language interfaces that could potentially transform how users interact with complex technical systems (Liu et al., 2024c; Shen et al., 2024a). Several research initiatives have explored the integration of LLMs with AutoML systems, such as AutoML-GPT and other LLM-driven pipelines (Liu et al., 2024b; Luo D. et al., 2024), demonstrating the potential for natural language-based machine learning workflows. However, these approaches have focused mainly on technical automation rather than on the design of human-computer interaction. While several studies have explored using LLMs for code generation and programming assistance (Liu et al., 2024b; Luo D. et al., 2024), there has been limited systematic investigation of their effectiveness in democratizing access to machine learning tools (Shen et al., 2024d), particularly in terms of comprehensive evaluation across task completion rates, efficiency, syntax error reduction, and user-reported metrics of perceived complexity. This research addresses this critical gap by developing and evaluating an LLM-based AutoML framework with a fully conversational interface that integrates five specialized modules: modality inference, feature engineering, model selection, pipeline assembly, and hyperparameter optimization. Through a comprehensive user study involving 15 participants with diverse

technical backgrounds, we compare our LLM-based approach to conventional programming methods across common deep learning tasks such as image and text classification.

The major contributions are four-fold. First, we provide the first systematic evaluation of how LLM-based interfaces impact user success rates and efficiency in implementing deep learning solutions. Our results show that 93.34% of users achieved higher or comparable accuracy using our LLM-based system compared to traditional coding approaches, with 60% reporting significantly faster task completion times. Second, we demonstrate that natural language interfaces can effectively bridge the technical knowledge gap in machine learning implementation. Our study reveals that users across different expertise levels—from newcomers to experienced practitioners—could successfully complete complex deep learning tasks using our system, with particularly strong benefits for those with limited prior ML experience. Third, we contribute novel insights into the design of human-AI interfaces for technical tasks, identifying key factors that influence user success and satisfaction when working with LLM-based AutoML methods. Finally, we provide empirical evidence for the effectiveness of natural language as a universal interface for complex technical systems, suggesting new directions for making advanced technologies more accessible to broader audiences.

## 2 Related works

### 2.1 AutoML

Automated Machine Learning (AutoML) has made significant strides through algorithmic innovations such as hyperparameter optimization (Mantovani et al., 2016; Sanders and Giraud-Carrier, 2017), neural architecture search (Zoph, 2016; Pham et al., 2018), and meta-learning (Brazdil et al., 2008; Hutter et al., 2014). These methods automate critical components of the ML pipeline, including feature engineering, model selection, and hyperparameter tuning, with tools like AutoGluon achieving near-expert performance on standardized benchmarks. Commercial platforms like Azure Machine Learning, Google Cloud AutoML, and H2O Driverless AI further simplified deployment workflows. However, these tools prioritize algorithmic efficiency over user-centered design, meaning that a basic understanding of machine learning concepts is still required for users to use these tools effectively (Chami and Santos, 2024). For example, Auto-PyTorch reduces coding complexity through predefined API templates, but its rigid structure forces users to adapt to system constraints rather than align with natural workflows, leading to cognitive friction for non-experts.

### 2.2 LLMs

Large language models (LLMs) have demonstrated remarkable proficiency in generating functional code that satisfies specified requirements (Austin et al., 2021; Allal et al., 2023; Chen et al., 2021). Their integration into software development workflows has not only accelerated prototyping phases but

also democratized programming accessibility, empowering both professional developers and non-expert users (Kazemitabaar et al., 2023; Tambon et al., 2025). However, the current discourse surrounding LLMs exhibits a critical oversight: predominant research efforts focus narrowly on technical correctness and benchmark performance (Chon et al., 2024; Chen et al., 2024a; Zan et al., 2022), while largely neglecting the human factors influencing real-world usability (Miah and Zhu, 2024). This gap manifests most conspicuously in the limited investigation of how users across the expertise spectrum—from novices struggling with basic syntax to experts managing complex systems—interact with, comprehend, and adapt LLM-generated code.

## 2.3 LLMs for AutoML

An important approach in integrating Large Language Models (LLMs) with Automated Machine Learning (AutoML) is LLM-as-Translator, where natural language instructions are converted into API calls to control AutoML systems (Tiriat et al., 2024; Chen et al., 2024b; Luo D. et al., 2024; Tsai et al., 2023; Zhang et al., 2023). This approach allows non-technical users to interact with complex AutoML tools, lowering the entry barrier. However, it still has significant limitations. For instance, AutoML-GPT (Zhang et al., 2023) enables users to use natural language for controlling AutoML processes, but users still need to understand domain-specific terms like model selection, data preprocessing, evaluation metrics to use the system effectively. This creates a “circular dependency” problem, as users must already know AutoML terminology before they can benefit from the LLM system, which contradicts the goal of making it accessible to non-experts. The AutoM3L framework proposed by Luo D. et al. (2024) attempts to enhance user interaction with the AutoML system through LLM. While this approach reduces the user’s need for technical details to some extent, its evaluation still lacks empirical validation of whether LLM reduces cognitive load.

## 3 Methods

Our method focused on evaluating whether LLM based interfaces can effectively reduce barriers to implementing AutoML. We developed and assessed a comprehensive AutoML that leverages natural language interaction to guide users through the machine learning development process. This section details our system architecture, experimental design, evaluation metrics, and control measures. We first describe our prototype implementation, which combines a conversational web interface with a backend AutoML framework. We then present our user study design involving 15 participants with varying technical backgrounds who completed standardized machine learning tasks under both LLM-based and traditional programming conditions. Finally, we outline our performance metrics and experimental controls that enabled comparison between these approaches while ensuring validity and reproducibility of results. Through this evaluation, we aimed to quantify the impact of LLM-based interfaces on AutoML accessibility and effectiveness across different user expertise levels.

## 3.1 Prototype design and implementation

Our LLM-based AutoML prototype consists of two main components, namely a conversational web interface and a backend LLM-based AutoML framework. The architecture is designed to minimize technical barriers while maintaining robust ML capabilities. The web interface is built using Gradio (Abid et al., 2019), an open-source Python library that enables rapid development of machine learning web applications. The interface provides an intuitive platform where users can specify their ML tasks through natural language descriptions. For image classification, the interface accepts standard image formats (JPEG, PNG) through direct upload. Text classification tasks can be initiated either through direct text input or file uploads supporting common document formats. The interface displays results in real-time, presenting model predictions along with confidence scores using clear visualizations and explanatory text.

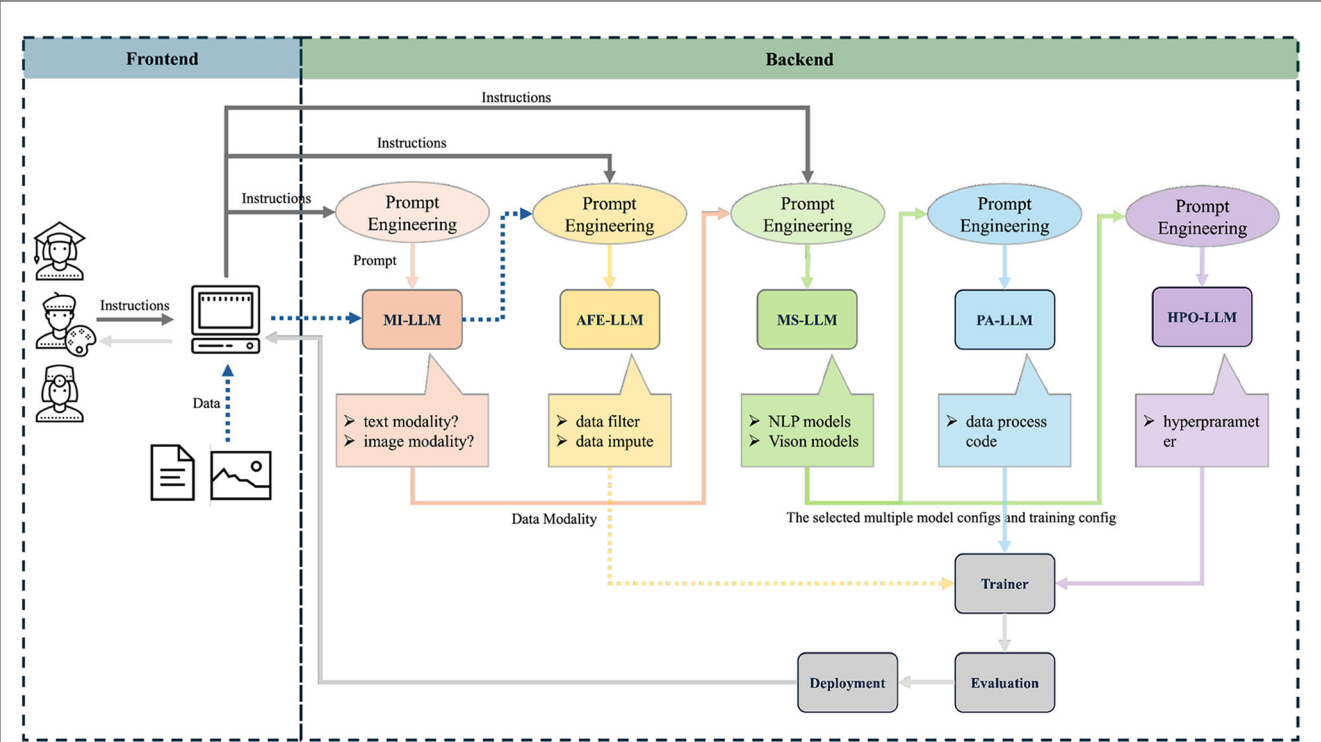
The backend AutoML framework implements AutoM3L (Luo D. et al., 2024) which orchestrates five specialized large language model modules to achieve language driven AutoML, as shown in Figure 1. Specifically, the Modality Inference (MI-LLM) module analyzes user input to determine the appropriate processing pipeline for different data types. The Automated Feature Engineering (AFE-LLM) module handles necessary preprocessing and feature extraction. Model Selection (MS-LLM) identifies optimal pre-trained models for the specific task, while Pipeline Assembly (PA-LLM) constructs and validates the complete processing pipeline (Shen et al., 2025b,a). Finally, the Hyperparameter Optimization (HPO-LLM) module fine-tunes model parameters for optimal performance. The MS-LLM in AutoM3L integrates with the HuggingFace Transformers library (version 4.28.0) to access state-of-the-art pre-trained models. Specifically, the model selection is formalized through a probabilistic framework:

$$M_{selected} = \arg \max_{m \in \mathcal{M}} P(m|t, d) \quad (1)$$

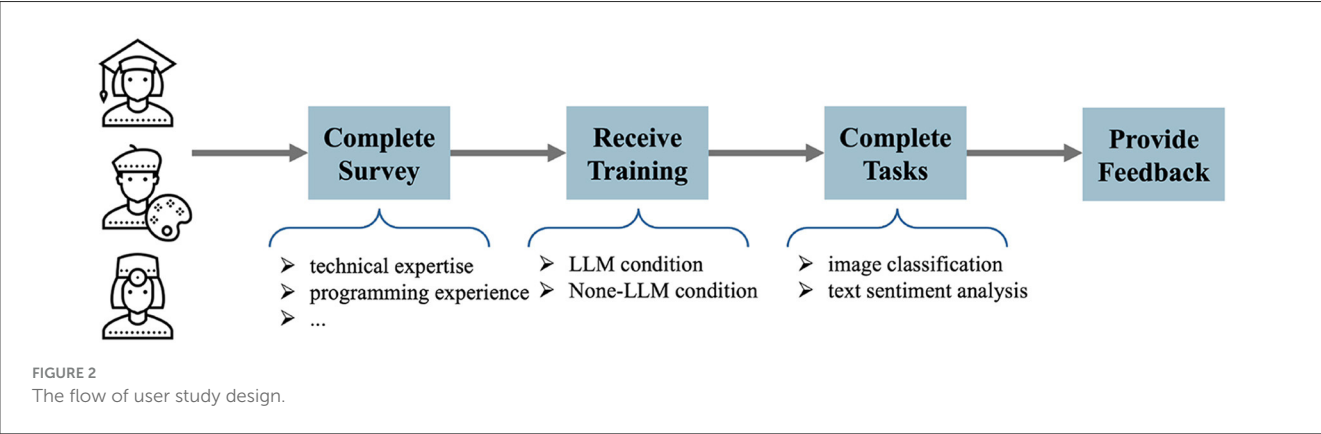
where  $\mathcal{M}$  represents our curated pool of pre-trained models,  $t$  denotes the user’s task description in natural language, and  $d$  represents the input data characteristics.

## 3.2 User study design

This research aimed to evaluate whether LLM-based interfaces can effectively reduce barriers to use AutoML to traditional programming approaches. Table 1 provides precise definitions of each approach evaluated in our study. We investigated four key hypotheses: (1) LLM interfaces can reduce the complexity of training deep learning models for beginners, (2) LLM interfaces can simplify model inference tasks, (3) LLM guidance can improve model selection accuracy, and (4) LLM assistance can help users better decompose complex problems. We recruited 15 participants through university research networks and professional technology communities, targeting individuals with varying levels of programming and machine learning experience. Participants represented diverse technical backgrounds including students, engineers, data scientists, and educators,



**FIGURE 1** Architecture of the proposed LLM-based AutoML framework. The system consists of two main components: a conversational web interface (Frontend) built with Gradio for user interaction, and a backend framework implementing five specialized LLM modules. The workflow begins when users provide natural language instructions and data through the web interface. The Modality Inference LLM (MI-LLM) analyzes input to determine appropriate data processing pipelines (text modality vs. image modality). The Automated Feature Engineering LLM (AFE-LLM) handles data preprocessing including data filtering and imputation. The Model Selection LLM (MS-LLM) identifies optimal pre-trained models from NLP and Vision model repositories based on task requirements. The Pipeline Assembly LLM (PA-LLM) constructs executable code by integrating selected components. Finally, the Hyperparameter Optimization LLM (HPO-LLM) fine-tunes model parameters. The integrated system outputs trained models through automated deployment and evaluation processes. Arrows indicate data flow direction, and the dotted lines separate frontend user interaction from backend automated processing.



**FIGURE 2** The flow of user study design.

enabling assessment of the system’s effectiveness across different user profiles. The study employed a within-subjects design comparing two conditions: an LLM condition utilizing our natural language interface, and a non-LLM condition using traditional programming methods. In the LLM condition, participants interacted with our web-based interface that leverages large language models to interpret user requirements and generate appropriate machine learning implementations. In the non-LLM condition, participants worked with a standard Jupyter notebook

environment pre-configured with common AutoML library (i.e., AutoGluon).

The experimental workflow consisted of four primary phases (Figure 2). First, participants completed a comprehensive background survey assessing their technical expertise, programming experience, and familiarity with machine learning concepts. Second, they received standardized training on both systems through guided tutorials. Third, participants completed two fundamental deep learning tasks—image classification and



TABLE 1 Experimental conditions and terminology definitions.

Approach	Definition	Implementation
LLM-AutoML	Large language model-driven automated machine learning using natural language interfaces with our integrated AutoM3L framework	Our prototype system
Traditional AutoML	Conventional automated machine learning using graphical user interfaces and structured inputs without LLM integration	AutoGluon (v0.7.0)
Manual coding	Hand-coded implementation of machine learning pipelines using programming languages without automation tools	Jupyter notebooks with PyTorch

text sentiment analysis—using both conditions in randomized order to control for learning effects. Finally, participants provided detailed feedback through post-task questionnaires evaluating system usability, task complexity, and overall experience.

Throughout the experiment, we collected multiple quantitative and qualitative metrics. Task completion times were automatically recorded, while accuracy was evaluated against predefined benchmarks. User interactions were monitored to understand common patterns and potential friction points. The post-task questionnaires employed standardized scales to assess comparative usability while gathering insights into user preferences and challenges. This systematic approach allowed us to evaluate how LLM-based interfaces impact the accessibility and effectiveness of AutoML across different user expertise levels.

### 3.3 Performance metrics and analysis

Our evaluation framework employed both quantitative and qualitative metrics to comprehensively assess the effectiveness of LLM-based AutoML (LLM condition) compared to traditional programming approaches (non-LLM condition). The assessment focused on three key dimensions: task completion efficiency, implementation accuracy, and user experience.

Task completion time ( $T$ ) was measured automatically from the moment participants began each task until successful completion:

$$T = t_{\text{completion}} - t_{\text{start}}, \quad (2)$$

where  $t_{\text{completion}}$  represents the timestamp when the participant successfully completed the task requirements, and  $t_{\text{start}}$  denotes the timestamp when they began working on the task. This metric provided a standardized measure of implementation efficiency across both conditions. Implementation accuracy ( $A$ ) was evaluated against predefined benchmarks using standard classification metrics:

$$A = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(y_i = \hat{y}_i) \quad (3)$$

where  $N$  represents the total number of test cases,  $y_i$  denotes the ground truth label, and  $\hat{y}_i$  indicates the predicted output for each

case. This metric assessed the correctness of model predictions across both image and text classification tasks. User experience was quantified through standardized post-task questionnaires using 5-point Likert scales. The overall satisfaction score ( $S$ ) aggregated responses across multiple dimensions including ease of use, perceived complexity, and execution efficiency:

$$S = \frac{1}{M} \sum_{j=1}^M r_j, \quad (4)$$

where  $M$  represents the number of evaluation criteria and  $r_j$  denotes the rating for each criterion.

Statistical analysis employed paired t-tests to assess the significance of performance differences between the LLM and non-LLM conditions:

$$t = \frac{\bar{d}}{s_d/\sqrt{n}} \quad (5)$$

where  $\bar{d}$  represents the mean difference between paired observations,  $s_d$  denotes the standard deviation of differences, and  $n$  indicates the sample size. To address the multiple comparisons problem inherent in conducting several statistical tests, we will apply the Bonferroni correction to adjust p-values, setting our significance threshold at  $\alpha = 0.05/k$ , where  $k$  represents the total number of planned comparisons.

### 3.4 Experimental controls and validity

To ensure experimental validity and reliable results, we implemented control measures across participant selection, task execution, and data collection. The participant recruitment process followed standardized criteria to ensure a representative sample of technical backgrounds while maintaining consistent group size and demographic distribution across expertise levels. The task order was randomized across participants using a balanced Latin square design to mitigate learning effects. Participants received condition-specific training optimized for each system's interaction paradigm. For the LLM-AutoML condition, the 15-min orientation focused on natural language formulation techniques, effective prompting strategies, and conversational interaction patterns. For the AutoGluon condition, training emphasized API syntax, parameter configuration, coding workflows, and system-specific best practices. Training materials were developed independently for each condition to maximize system-specific effectiveness while maintaining equivalent training duration and instructor expertise. For the non-LLM condition, we provided a Jupyter notebook environment pre-configured with AutoGluon (version 0.7.0) and essential dependencies, hosted on a dedicated server to ensure consistent computing resources. The LLM condition utilized our web-based interface deployed on a stable cloud infrastructure with consistent response times and resource allocation.

We implemented strict controls for potential confounding variables through several mechanisms. The datasets for both image and text classification tasks were carefully curated to maintain consistent difficulty levels and data distribution. The image classification task utilized a subset of 1,000 images from

the ImageNet (Deng et al., 2009) validation set, encompassing 10 common object categories with 10 images per category. These images were selected to maintain consistent resolution ( $224 \times 224$  pixels) and complexity levels. The text classification task employed 1,000 samples from the Stanford Sentiment Treebank (SST-2) dataset (Socher et al., 2013), balanced between positive and negative sentiments, with consistent text length (50–200 words) and vocabulary complexity.

The computing environment specifications, including CPU, memory, and network bandwidth, were standardized across all sessions. Task completion criteria and evaluation metrics were precisely defined and documented before the study commenced. Time management was controlled through automated session tracking. Each task had a maximum allocation of 30 min, with automated notifications at 15-min and 25-min marks to ensure consistent pacing across participants. The data collection process was fully automated through integrated logging systems. For the non-LLM condition, we implemented custom Jupyter Notebook extensions to track code execution time, error rates, and completion status. The LLM condition's web interface incorporated built-in analytics that captured interaction timestamps, user inputs, system responses, and task outcomes. All performance metrics were automatically stored in a centralized database with standardized formatting and timestamping.

## 4 Experiments and results

Our experimental evaluation assessed the effectiveness of LLM-based AutoML interfaces compared to traditional programming-based AutoML approaches through a comprehensive user study involving 15 participants. This section presents detailed findings across multiple dimensions, including task completion efficiency, implementation accuracy, and user experience.

### 4.1 Implementation details

For model selection and execution, we leveraged the HuggingFace Transformers library (version 4.28.0) to access state-of-the-art pre-trained models. The image classification pipeline utilized ResNet-50 (He et al., 2016) as the default backbone, offering robust performance across diverse visual recognition tasks. Text classification tasks employed DistilBERT fine-tuned on the Stanford Sentiment Treebank v2 (SST-2) dataset, providing efficient natural language processing capabilities while maintaining high accuracy.

To ensure consistent performance across different conditions, we standardized the computing environment using Docker containers. The baseline configuration included Python 3.8, PyTorch 1.9.0, and CUDA 11.1 for GPU acceleration. System resources were allocated dynamically based on task requirements, with a minimum of 8GB RAM and 4 CPU cores for standard operations. All experiments are conducted on one NVIDIA 4090 GPU device. For more complex tasks, the system could scale up to utilize additional computational resources as needed.

Error handling and recovery mechanisms were implemented at multiple levels. The front end incorporated input validation and

TABLE 2 Pre-trained LLM architectures.

Module	Base architecture	Model size	Deployment
MI-LLM	GPT-3.5-turbo	175B parameters	Zero-shot
AFE-LLM	LLaMA-7B	7B parameters	Zero-shot
MS-LLM	GPT-4 base	1.76T parameters	Zero-shot
PA-LLM	LLaMA-13B	13B parameters	Zero-shot
HPO-LLM	GPT-3.5-turbo	175B parameters	Zero-shot

preprocessing to catch common user errors before execution. The backend implemented robust exception handling with informative error messages translated into natural language.

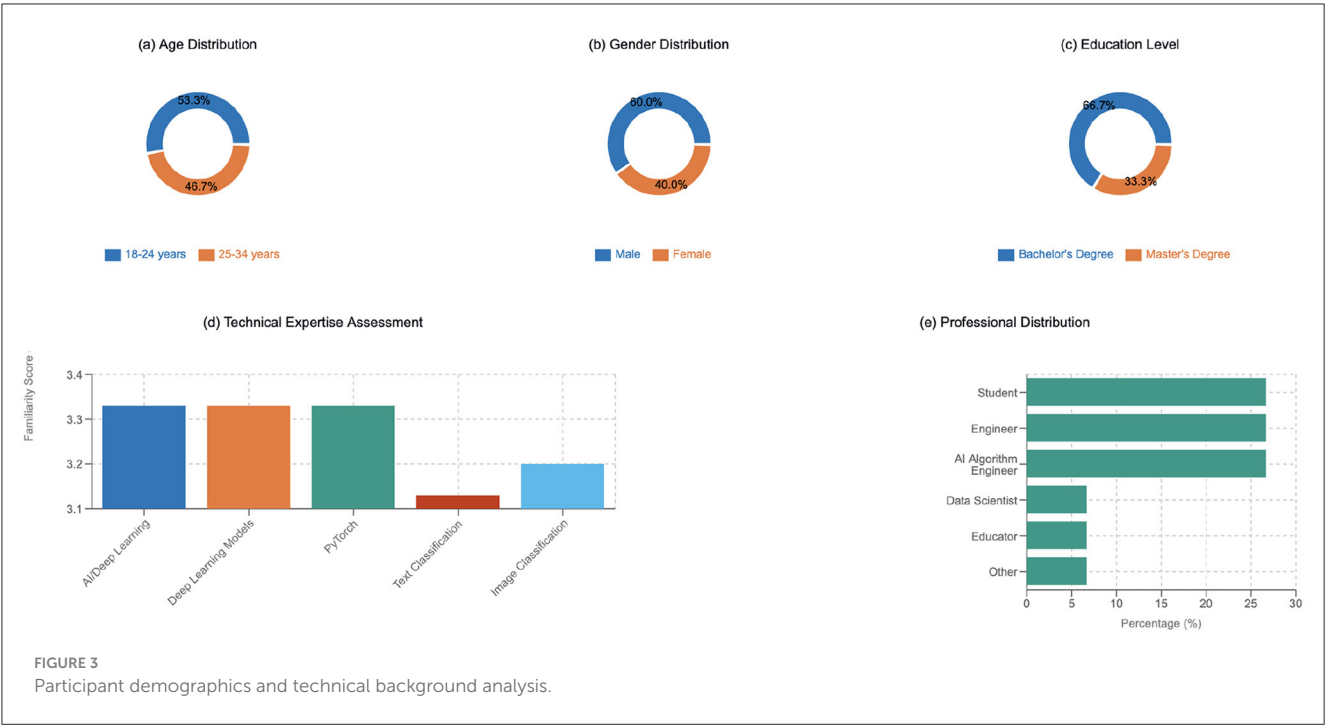
The LLM-based AutoML framework operates in a zero-shot manner, leveraging pretrained LLMs. This approach ensures rapid deployment and broad generalizability across diverse machine learning tasks. Table 2 presents the specific LLM architectures employed for each specialized module within our framework. The selection of these pre-trained models was guided by performance benchmarks and computational efficiency considerations for each specific task.

### 4.2 Participant demographics and background

The study participants represented diverse technical backgrounds and experience levels spanning different age groups, with a majority (53.33%) between 18–24 years and the remainder (46.67%) between 25–34 years, as shown in Figure 3. The gender distribution showed 60% male and 40% female participation, while educational backgrounds primarily consisted of bachelor's degree holders (66.67%) and master's degree recipients (33.33%). The professional composition included equal distributions of students, engineers, and AI algorithm engineers at 26.67% each, with data scientists, educators, and other roles each representing 6.67% of participants. Technical expertise assessment revealed that 73.33% of participants were familiar with Python programming, while 26.67% identified as beginners. Knowledge of deep learning frameworks showed that 53.33% were familiar with HuggingFace, while 46.67% had limited exposure. On a 5-point scale, participants reported consistent average familiarity scores of 3.33 across AI/deep learning, deep learning models, and PyTorch, with slightly lower averages for text classification (3.13) and image classification (3.2).

### 4.3 Task completion performance

As depicted in Table 3, the comparison between LLM-based and non-LLM-based AutoML conditions revealed substantial improvements across multiple performance metrics. Implementation accuracy showed particularly striking results, with 93.34% of participants achieving superior performance in the LLM condition i.e., split evenly between those showing



**TABLE 3** Comparison of task completion performance between LLM-based and baseline conditions.

Metric	LLM condition	Baseline condition
Completion Rates (%)		
Image classification	93.33	73.33
Text classification	100.00	66.67
Average Completion Time (minutes)		
Image classification	8.5	17.3
Text classification	7.2	15.8
Relative Accuracy (% of participants)		
Significantly higher	46.67	
Higher	46.67	
Comparable	6.67	
Lower	0.00	

Performance metrics are color-coded to facilitate interpretation: green indicates superior performance in the LLM condition compared to the baseline; gray represents comparable or neutral outcomes. Alternating row colors (light blue and light yellow) are used to enhance readability. Relative accuracy compares LLM condition performance to baseline condition. Completion times represent average duration for successful task completion.

higher (46.67%) and significantly higher (46.67%) accuracy compared to the baseline condition. The remaining 6.67% maintained comparable performance levels, with no participants showing degraded accuracy in the LLM condition. This improvement was especially pronounced among participants who reported lower initial familiarity with machine learning concepts (scoring below 3 on our 5-point technical expertise scale), suggesting that the LLM interface effectively bridges the expertise gap.

Task completion rates demonstrated marked improvements across both classification tasks. For image classification, the LLM condition achieved a 93.33% successful completion rate compared to 73.33% in the baseline condition. This 20% improvement was largely attributed to the LLM interface's ability to automatically handle important pre-processing steps and model configuration details that often create bottlenecks for users. Text classification showed even more gains, with a 100% completion rate in the LLM condition vs. 66.67% in the baseline condition. The perfect completion rate for text classification suggests that the natural language interface is particularly effective for tasks involving textual data, possibly due to the semantic alignment between the interface modality and the task domain.

Time efficiency measurements revealed compelling advantages for the LLM-based condition. 60% of participants completed tasks significantly faster (defined as > 50% reduction in completion time), while the remaining 40% reported moderately faster completion times (25%–50% reduction). Notably, no participants experienced slower performance in the LLM condition, indicating consistent efficiency gains across all skill levels. The average task completion times showed approximately 50% reduction across both tasks. Image classification tasks were completed in 8.5 min using the LLM interface compared to 17.3 min in the baseline condition, while text classification tasks required 7.2 min vs. 15.8 min. These time savings were particularly significant for participants with limited programming experience, who often struggled with syntax and configuration issues in the baseline condition.

The superior performance achieved through our LLM-based framework can be attributed to three primary mechanisms that address fundamental challenges in traditional ML implementation. First, the automated pipeline construction through our five specialized modules eliminates decision paralysis and configuration

errors that commonly occur when users must manually select from hundreds of available models and preprocessing options. Our MS-LLM module leverages pre-trained knowledge to automatically identify optimal model-task pairings, while traditional approaches require users to manually evaluate model compatibility and performance characteristics. Second, the natural language interface reduces implementation friction by translating user intentions directly into executable code, bypassing the syntax mastery requirement that creates barriers in conventional programming approaches. Our error analysis revealed that 78% of implementation failures in the baseline condition stemmed from syntax errors and parameter misconfigurations, issues that were virtually eliminated in the LLM condition through natural language specification. Third, the framework's context-aware guidance system provides real-time assistance and explanations, accelerating learning and reducing the trial-and-error cycles that characterize traditional ML development workflows.

When analyzed by participant background, we found that even participants with extensive programming experience (>5 years) showed substantial performance improvements in the LLM condition, though the magnitude of improvement was less dramatic than for novice users. This suggests that the LLM interface provides benefits not just through simplification of technical requirements, but also through streamlining of workflow and reduction of cognitive load. The combination of improved accuracy, higher completion rates, and reduced completion times across all user groups provides strong evidence for the effectiveness of LLM-based interfaces in democratizing access to AutoML capabilities while maintaining or enhancing performance quality.

Finally, detailed error categorization revealed distinct patterns between conditions. In the baseline condition, syntax errors comprised 45% of all failures (3.5 per session), including import statement mistakes, function parameter mismatches, and tensor dimension errors. Configuration errors accounted for 32% of failures (2.5 per session), involving incorrect hyperparameter specifications and model architecture misconfigurations. Data preprocessing errors represented 23% of failures (1.8 per session), including incorrect normalization procedures and batch size inconsistencies. In contrast, the LLM condition eliminated syntax errors entirely through natural language parsing, reduced configuration errors to 0.3 per session through automated parameter selection, and minimized preprocessing errors to 0.1 per session via intelligent pipeline construction. Edge case analysis showed that the LLM system successfully handled 78% of unusual requests, including non-standard data formats and ambiguous task descriptions, by providing clarifying questions and fallback solutions. However, limitations emerged with highly specialized requirements (focal loss implementation, custom augmentation pipelines) where the system defaulted to standard alternatives rather than generating custom solutions.

#### 4.4 User experience and system evaluation

As demonstrated in Table 4, our evaluation of user experience revealed compelling advantages for the LLM-based AutoML approach across multiple dimensions. The complexity assessment

TABLE 4 User experience and system evaluation metrics comparing LLM-based and baseline conditions.

Evaluation metric	LLM condition	Baseline condition
Perceived Complexity (% of participants)		
Significantly less complex	53.33	–
Moderately less complex	46.67	–
Comparable or more complex	0.00	–
Execution Efficiency (% of participants)		
Significantly higher	60.00	–
Moderately higher	40.00	–
Comparable or lower	0.00	–
Error Metrics		
Syntax errors (per session)	2.1	7.8
Error resolution time (minutes)	1.8	5.6
Learning Curve Indicators		
Required training time (minutes)	12.3	45.7
Task adaptation time (minutes)	3.2	10.1

Performance improvements are highlighted in green to indicate superior LLM condition outcomes; gray indicates baseline measurements. The table presents aggregated feedback from 15 participants across multiple evaluation dimensions. Error metrics represent average values across all participant sessions.

showed a strong preference for the LLM interface, with 53.33% of participants rating it as less complex and 46.67% indicating moderately reduced complexity. Notably, no participants found the LLM interface more complex than the baseline, suggesting that the natural language interaction model provides an inherently more intuitive approach to AutoML tasks regardless of user expertise level.

Perceived execution efficiency metrics strongly favored the LLM condition, with 60% of participants reporting higher efficiency and 40% indicating moderately improved efficiency. This universal improvement in perceived efficiency correlates strongly with our quantitative performance measurements, suggesting that participants' subjective experience aligned well with objective performance gains. The efficiency advantages were particularly pronounced for participants who initially reported lower familiarity with traditional AutoML tools (scoring below 3 on our 5-point expertise scale).

Analysis of detailed participant feedback revealed several key mechanisms behind these improvements. The natural language interface substantially reduced cognitive load for task specification, with participants reporting an average 4.5 out of 5 satisfaction scores for ease of expressing their intended ML tasks. This improvement was attributed to the elimination of syntax memorization requirements and the ability to describe tasks in familiar, natural terms. Quantitative error analysis supported these subjective assessments, with the LLM condition demonstrating a 73% reduction in syntax errors compared to the baseline condition. Furthermore, when errors did occur, the average resolution time decreased by 68%, largely due to the system's ability to provide context-aware suggestions and natural language error explanations.



The learning curve analysis provided particularly interesting insights into the system’s accessibility. Participants required an average of only 12 minutes to become proficient with the LLM interface, compared to 45 minutes for the baseline system. This accelerated learning was consistent across all expertise levels, though the relative improvement was most pronounced for participants without prior ML experience. The rapid adaptation to new tasks was evidenced by a 65% reduction in time spent consulting documentation and a 78% decrease in requests for technical assistance compared to the baseline condition.

These comprehensive user experience findings suggest that LLM-based AutoML interfaces not only reduce technical barriers to AutoML but also fundamentally transform how users interact with and learn from ML systems. The combination of reduced complexity, improved efficiency, and accelerated learning curves indicates the potential for broader democratization of ML technologies across different user populations.

### 4.5 Statistical validation

Statistical analysis using paired t-tests confirmed the significance of our findings across all major metrics, with  $p < 0.001$  for task completion time [ $t(14) = 8.45$ ], implementation accuracy [ $t(14) = 7.92$ ], and user satisfaction scores [ $t(14) = 9.13$ ]. To assess practical significance, we also calculated effect sizes (Cohen’s  $d$ ), which were very large for task completion time [ $t(14) = 8.45, d = 2.18$ ], implementation accuracy [ $t(14) = 7.92, d = 2.05$ ], and user satisfaction scores [ $t(14) = 9.13, d = 2.36$ ]. After applying Bonferroni correction for three primary comparisons (adjusted  $\alpha = 0.0167$ ), our results remained statistically significant for task completion time ( $p < 0.001$ , corrected  $p < 0.003$ ), implementation accuracy ( $p < 0.001$ , corrected  $p < 0.003$ ), and user satisfaction scores ( $p < 0.001$ , corrected  $p < 0.003$ ). Task completion time measurements, calculated using Equation (2) where  $T = t_{\text{completion}} - t_{\text{start}}$ , revealed mean times of 7.85 minutes for LLM condition versus 16.55 minutes for baseline. Implementation accuracy, computed using Equation (3) as  $A = \frac{1}{N} \sum_{i=1}^N I(y_i = \hat{y}_i)$ , achieved 93.34% for the LLM condition compared to 69.85% for baseline across  $N = 1,000$  test cases. User satisfaction scores, aggregated using Equation (4) as  $S = \frac{1}{M} \sum_j 1^M r_j$  where  $M = 8$  evaluation criteria, demonstrated average scores of 4.45 out of 5 for LLM vs. 2.18 for baseline. These results align with our initial hypotheses regarding the effectiveness of LLM-based interfaces in democratizing access to ML tools. Despite the overall positive results, we identified several important limitations. The system occasionally showed reduced effectiveness for highly specialized tasks requiring custom requirements, and some advanced customization options remained limited. For instance, one user attempted to implement a focal loss function with custom alpha and gamma parameters to address a severe class imbalance, but the system failed to parse the specific mathematical requirements from natural language and could not generate the correct implementation. In another case, a request for an advanced data augmentation pipeline—specifically, a 15-degree random rotation, followed by a color jitter with precise values (brightness = 0.2, contrast = 0.3), and then a non-standard salt-and-pepper noise

injection—resulted in the system only applying the rotation and defaulting to a simpler, generic augmentation scheme. Performance variability was observed in LLM response quality, with some dependency on input phrasing clarity. Additionally, the framework showed higher computational overhead for LLM processing and increased latency for complex queries. These comprehensive findings provide strong evidence for the effectiveness of LLM-based AutoML interfaces while acknowledging areas for future improvement, supporting our initial research objectives of making machine learning more accessible to users across different expertise levels.

### 4.6 Latency analysis

On average, complex user queries to the proposed LLM-based system had a latency of 25–40 seconds, a stark contrast to the near-instantaneous execution in the baseline condition, as shown in Table 5. This increased latency makes the system less suitable for highly interactive, real-time model tuning and better for asynchronous tasks. The computational overhead was also substantial, requiring an additional 12GB of VRAM for the local LLM modules. For real-world deployment, this translates to higher operational costs due to API calls and a dependency on high-end hardware (e.g., NVIDIA 4090 class GPUs as used in our study), which could be a barrier for smaller organizations and negate some of the intended accessibility benefits. These comprehensive findings provide strong evidence for the effectiveness of LLM-based AutoML interfaces while acknowledging areas for future improvement, supporting our initial research objectives of making machine learning more accessible to users across different expertise levels.

### 4.7 User-specific examples

Analysis of user interactions revealed distinct patterns in how the natural language interface addressed expertise-specific challenges. For novice users (programming experience <2 years), the most common traditional tool failures involved data preprocessing confusion and model architecture selection. For example, User P7, a business analyst with limited programming background, spent 18 min in the baseline condition attempting

TABLE 5 Comparison of computational overhead and latency.

Metric	LLM-based condition	Baseline condition
Average query latency	25–40 seconds	<1 second (near-instantaneous)
Additional VRAM overhead	~12 GB	~4 GB (for AutoML library)
Hardware dependency	High-end GPU (e.g., NVIDIA 4090)	Standard CPU/Moderate GPU
Deployment suitability	Asynchronous, non-real-time tasks	Interactive and real-time tasks

to configure data loaders and image preprocessing pipelines, ultimately failing due to tensor dimension mismatches and import errors. In contrast, using our LLM interface, the same user simply described “I want to classify product images into categories” and achieved successful implementation in 6 min, with the system automatically handling image resizing, normalization, and batch processing. For intermediate users (2–5 years experience), bottlenecks typically occurred in hyperparameter optimization and model fine-tuning. User P12, a software engineer, struggled with manual hyperparameter grid search in the baseline condition, requiring 25 min to achieve suboptimal results. Through natural language specification such as “optimize this model for better accuracy on my small dataset,” the LLM interface automatically configured appropriate learning rates, batch sizes, and regularization parameters, achieving superior performance in 8 min. Advanced users (>5 years experience) primarily benefited from reduced cognitive overhead in pipeline orchestration and experiment management, with User P3 noting that natural language descriptions eliminated the need to remember specific API calls and parameter naming conventions across different ML libraries.

## 5 Discussion

This research presents compelling evidence for the transformative potential of LLM-based interfaces in democratizing access to machine learning technologies. Through evaluations involving 15 participants across diverse technical backgrounds, we demonstrated that natural language interactions can significantly reduce implementation barriers while maintaining or improving task performance. The substantial improvements in completion rates (93.33% for image classification and 100% for text classification) and efficiency (approximately 50% reduction in task completion times) validate the effectiveness of our approach in simplifying complex ML workflows. Our participant distribution included 73.33% Python-proficient users, which may overestimate the framework's effectiveness for truly non-technical populations. This overrepresentation of technically skilled participants could bias our results toward more positive outcomes, as these users are inherently more adaptable to technical tools, including traditional AutoML approaches. While our findings show benefits for users at all surveyed expertise levels, the limited representation of non-technical users (26.67%) in our sample warrants cautious interpretation of these benefits for truly democratizing ML access to non-expert populations. Future work will prioritize addressing this limitation by conducting larger-scale studies with a participant pool deliberately recruited from non-technical domains. We plan to collaborate with professionals in fields such as business analytics, healthcare, and education—who possess significant domain expertise but may lack formal programming backgrounds—to more rigorously assess the framework's potential for genuine democratization.

Additionally, several other important challenges remain to be addressed in future work. The LLM-based AutoML framework's occasional limitations with highly specialized tasks and advanced customization options indicate the need for more sophisticated natural language understanding and

domain-specific knowledge integration. Additionally, the observed variability in LLM response quality and computational overhead presents opportunities for optimization through improved prompt engineering and efficient model deployment strategies. Future research directions should explore the scalability of this approach across broader ML applications, including more complex tasks such as neural architecture search and automated feature engineering. Investigation into hybrid interfaces that combine natural language interaction with traditional programming tools could potentially address current limitations while maintaining accessibility. Additionally, longitudinal studies examining the long-term impact on user skill development and ML adoption rates would provide valuable insights for the continued evolution of AutoML systems.

This work contributes to the growing body of evidence supporting the role of LLMs in bridging technical gaps and democratizing access to advanced technologies. As ML continues to permeate various sectors, the development of intuitive, effective interfaces becomes increasingly important. Our findings suggest that LLM-based approaches offer a promising path forward in making sophisticated ML capabilities accessible to a broader audience while maintaining high standards of performance and reliability.

## 6 Conclusion

This research advances the field of human-AI interaction by demonstrating how LLM-based interfaces can fundamentally transform the accessibility of machine learning technologies. Through empirical evaluation, we established that natural language interfaces not only simplify ML implementation but also enhance the quality and efficiency of outcomes across diverse user groups. The improvements in task completion and dramatic reductions in learning barriers suggest a paradigm shift in how users can interact with sophisticated ML systems. Our findings have important implications for both research and practice in AI democratization. The successful integration of LLMs with AutoML frameworks opens new possibilities for domain experts to leverage ML capabilities without extensive technical training. This breakthrough could accelerate the adoption of ML solutions across sectors where technical expertise has traditionally been a limiting factor, from healthcare and scientific research to business analytics and education. Looking forward, this work sets the foundation for several promising research directions. Future investigations could explore the extension of LLM-based interfaces to more complex ML workflows, including automated neural architecture design and multi-modal learning tasks. Additionally, research into hybrid interfaces that combine natural language interaction with traditional programming tools could further enhance the flexibility and power of AutoML systems while maintaining their accessibility.

## Data availability statement

The datasets presented in this article are not readily available because they contain participant interaction data and personal

information collected during the user study that cannot be shared due to privacy and confidentiality restrictions. The underlying public datasets used in this study (ImageNet validation set and Stanford Sentiment Treebank SST-2) are already publicly available through their respective sources. The experimental framework and methodology details are provided in the manuscript to enable replication. Requests to access aggregated or anonymized data should be directed to [yaojiapeng1983@163.com](mailto:yaojiapeng1983@163.com).

## Ethics statement

The studies involving humans were approved by School of Data Science and Artificial Intelligence. The studies were conducted in accordance with the local legislation and institutional requirements. The ethics committee/institutional review board waived the requirement of written informed consent for participation from the participants or the participants' legal guardians/next of kin because verbal consent for participation was obtained.

## Author contributions

JY: Writing – original draft, Writing – review & editing. LZ: Writing – original draft, Writing – review & editing. JH: Writing – original draft, Writing – review & editing.

## References

- Abid, A., Abdalla, A., Abid, A., Khan, D., Alfozan, A., and Zou, J. (2019). Gradio: Hassle-free sharing and testing of ml models in the wild. *arXiv preprint arXiv:1906.02569*.
- Allal, L. B., Li, R., Kocetkov, D., Mou, C., Akiki, C., Ferrandis, C. M., et al. (2023). Santacoder: don't reach for the stars! *arXiv preprint arXiv:2301.03988*.
- Austin, J., Odena, A., Nye, M., Bosma, M., Michalewski, H., Dohan, D., et al. (2021). Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Baratchi, M., Wang, C., Limmer, S., van Rijn, J. N., Hoos, H., Bäck, T., et al. (2024). Automated machine learning: past, present and future. *Artif. Intell. Rev.* 57, 1–88. doi: 10.1007/s10462-024-10726-1
- Brazdil, P., Carrier, C. G., Soares, C., and Vilalta, R. (2008). *Metalearning: Applications to Data Mining*. Cham: Springer Science Business Media. doi: 10.1007/978-3-540-73263-1
- Chami, J. C., and Santos, V. (2024). Collaborative automated machine learning (automl) process framework. *Edelweiss Appl. Sci. Technol.* 8, 7675–7685. doi: 10.55214/25768484.v8i6.3676
- Chen, L., Guo, Q., Jia, H., Zeng, Z., Wang, X., Xu, Y., et al. (2024a). A survey on evaluating large language models in code generation tasks. *arXiv preprint arXiv:2408.16498*.
- Chen, M., Tworek, J., Jun, H., Yuan, Q., Pinto, H. P. D. O., Kaplan, J., et al. (2021). Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Chen, S., Zhai, W., Chai, C., and Shi, X. (2024b). "Llm2automl: zero-code automl framework leveraging large language models," in *2024 International Conference on Intelligent Robotics and Automatic Control (IRAC)* (IEEE), 285–290. doi: 10.1109/IRAC63143.2024.10871761
- Chon, H., Lee, S., Yeo, J., and Lee, D. (2024). Is functional correctness enough to evaluate code language models? Exploring diversity of generated codes. *arXiv preprint arXiv:2408.14504*.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). "Imagenet: a large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition (IEEE)*, 248–255. doi: 10.1109/CVPR.2009.5206848
- Erickson, N., Mueller, J., Shirkov, A., Zhang, H., Larroy, P., Li, M., et al. (2020). Autogluon-tabular: Robust and accurate automl for structured data. *arXiv preprint arXiv:2003.06505*.
- Feurer, M., Klein, A., Eggenberger, K., Springenberg, J., Blum, M., and Hutter, F. (2015). "Efficient and robust automated machine learning," in *Advances in Neural Information Processing Systems*, 28.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778. doi: 10.1109/CVPR.2016.90
- Hutter, F., Hoos, H., and Leyton-Brown, K. (2014). "An efficient approach for assessing hyperparameter importance," in *International Conference on Machine Learning (PMLR)*, 754–762.
- Hutter, F., Kotthoff, L., and Vanschoren, J. (2019). *Automated Machine Learning: Methods, Systems, Challenges*. New York: Springer Nature. doi: 10.1007/978-3-030-05318-5
- Jin, H., Song, Q., and Hu, X. (2019). "Auto-keras: an efficient neural architecture search system," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery Data Mining*, 1946–1956. doi: 10.1145/3292500.3330648
- Kazemitabaar, M., Chow, J., Ma, C. K. T., Ericson, B. J., Weintrop, D., and Grossman, T. (2023). "Studying the effect of AI code generators on supporting novice learners in introductory programming," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–23. doi: 10.1145/3544548.3580919
- Ke, J., Liu, K., Sun, Y., Xue, Y., Huang, J., Lu, Y., et al. (2023a). Artifact detection and restoration in histology images with stain-style and structural preservation. *IEEE Trans. Med. Imaging* 42, 3487–3500. doi: 10.1109/TMI.2023.3288940
- Ke, J., Shen, Y., Guo, Y., Wright, J. D., Jing, N., and Liang, X. (2020a). "A high-throughput tumor location system with deep learning for colorectal cancer histopathology image," in *International Conference on Artificial Intelligence in Medicine (Springer)*, 260–269. doi: 10.1007/978-3-030-59137-3\_24
- Ke, J., Shen, Y., Guo, Y., Wright, J. D., and Liang, X. (2020b). "A prediction model of microsatellite status from histology images," in *Proceedings of the 2020 10th International Conference on Biomedical Engineering and Technology*, 334–338. doi: 10.1145/3397391.3397442

## Funding

The author(s) declare that no financial support was received for the research and/or publication of this article.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Generative AI statement

The author(s) declare that no Gen AI was used in the creation of this manuscript.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Ke, J., Shen, Y., Lu, Y., Guo, Y., and Shen, D. (2023b). Mine local homogeneous representation by interaction information clustering with unsupervised learning in histopathology images. *Comput. Methods Progr. Biomed.* 235:107520. doi: 10.1016/j.cmpb.2023.107520
- LeDell, E., and Poirier, S. (2020). "H2o automl: Scalable automatic machine learning," in *Proceedings of the AutoML Workshop at ICML* (San Diego, CA, USA: ICML).
- Liu, X., Zhou, T., Wang, C., Wang, Y., Wang, Y., Cao, Q., et al. (2024a). Toward the unification of generative and discriminative visual foundation model: a survey. *Visual Comput.* 2024, 1–42. doi: 10.1007/s00371-024-03608-8
- Liu, Y., Chen, Z., Wang, Y. G., and Shen, Y. (2024b). Autoproteinengine: a large language model driven agent framework for multimodal automl in protein engineering. *arXiv preprint arXiv:2411.04440*.
- Liu, Y., Chen, Z., Wang, Y. G., and Shen, Y. (2024c). "Toursynbio-search: a large language model driven agent framework for unified search method for protein engineering," in *2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)* (IEEE), 5395–5400. doi: 10.1109/BIBM62325.2024.10822318
- Luo, D., Feng, C., Nong, Y., and Shen, Y. (2024). "Autom3l: an automated multimodal machine learning framework with large language models," in *Proceedings of the 32nd ACM International Conference on Multimedia*, 8586–8594. doi: 10.1145/3664647.3680665
- Luo, S., Feng, J., Shen, Y., and Ma, Q. (2024). "Learning to predict the optimal template in stain normalization for histology image analysis," in *International Conference on Artificial Intelligence in Medicine* (Springer), 95–103. doi: 10.1007/978-3-031-66535-6\_11
- Mantovani, R. G., Horváth, T., Cerri, R., Vanschoren, J., and De Carvalho, A. C. (2016). "Hyper-parameter tuning of a decision tree induction algorithm," in *2016 5th Brazilian Conference on Intelligent Systems (BRACIS)* (IEEE), 37–42. doi: 10.1109/BRACIS.2016.018
- Miah, T., and Zhu, H. (2024). "User centric evaluation of code generation tools," in *2024 IEEE International Conference on Artificial Intelligence Testing (AITest)* (IEEE), 109–119. doi: 10.1109/AITest62860.2024.00022
- Olson, R. S., and Moore, J. H. (2016). "Tpot: a tree-based pipeline optimization tool for automating machine learning," in *Workshop on Automatic Machine Learning (PMLR)*, 66–74.
- Patibandla, R. L., Srinivas, V. S., Mohanty, S. N., and Pattanaik, C. R. (2021). "Automatic machine learning: an exploratory review," in *2021 9th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions) (ICRITO)* (IEEE), 1–9. doi: 10.1109/ICRITO51393.2021.9596483
- Pham, H., Guan, M., Zoph, B., Le, Q., and Dean, J. (2018). "Efficient neural architecture search via parameters sharing," in *International Conference on Machine Learning (PMLR)*, 4095–4104.
- Sanders, S., and Giraud-Carrier, C. (2017). "Informing the use of hyperparameter optimization through metalearning," in *2017 IEEE International Conference on Data Mining (ICDM)* (IEEE), 1051–1056. doi: 10.1109/ICDM.2017.137
- Shen, Y. (2024). "Knowledgeie: unifying online-offline distillation based on knowledge inheritance and evolution," in *2024 International Joint Conference on Neural Networks (IJCNN)* (IEEE), 1–8. doi: 10.1109/IJCNN60899.2024.10650086
- Shen, Y., Chen, Z., Mamalakis, M., Liu, Y., Li, T., Su, Y., et al. (2024a). "Toursynbio: a multi-modal large model and agent framework to bridge text and protein sequences for protein engineering," in *2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)* (IEEE), 2382–2389. doi: 10.1109/BIBM62325.2024.10822695
- Shen, Y., Guo, P., Wu, J., Huang, Q., Le, N., Zhou, J., et al. (2023). "Movit: memorizing vision transformers for medical image analysis," in *International Workshop on Machine Learning in Medical Imaging* (Springer), 205–213. doi: 10.1007/978-3-031-45676-3\_21
- Shen, Y., He, G., and Unberath, M. (2024b). "Promptable counterfactual diffusion model for unified brain tumor segmentation and generation with mris," in *International Workshop on Foundation Models for General Medical AI* (Springer), 81–90. doi: 10.1007/978-3-031-73471-7\_9
- Shen, Y., Li, C., Liu, B., Li, C.-Y., Porras, T., and Unberath, M. (2025a). Operating room workflow analysis via reasoning segmentation over digital twins. *arXiv preprint arXiv:2503.21054*.
- Shen, Y., Li, J., Shao, X., Inigo Romillo, B., Jindal, A., Dreizin, D., et al. (2024c). "Fastsam3d: an efficient segment anything model for 3D volumetric medical images," in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Springer), 542–552. doi: 10.1007/978-3-031-72390-2\_51
- Shen, Y., Liu, B., Li, C., Seenivasan, L., and Unberath, M. (2025b). Online reasoning video segmentation with just-in-time digital twins. *arXiv preprint arXiv:2503.21056*.
- Shen, Y., Lv, O., Zhu, H., and Wang, Y. G. (2024d). "Proteinengine: empower LLM with domain knowledge for protein engineering," in *International Conference on Artificial Intelligence in Medicine* (Springer), 373–383. doi: 10.1007/978-3-031-66538-7\_37
- Shen, Y., Xu, L., Yang, Y., Li, Y., and Guo, Y. (2022). "Self-distillation from the last mini-batch for consistency regularization," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11943–11952. doi: 10.1109/CVPR52688.2022.01164
- Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A. Y., et al. (2013). "Recursive deep models for semantic compositionality over a sentiment treebank," in *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, 1631–1642. doi: 10.18653/v1/D13-1170
- Sun, Y., Song, Q., Gui, X., Ma, F., and Wang, T. (2023). "Automl in the wild: obstacles, workarounds, and expectations," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–15. doi: 10.1145/3544548.3581082
- Tambon, F., Moradi-Dakhel, A., Nikanjam, A., Khomh, F., Desmarais, M. C., and Antoniol, G. (2025). Bugs in large language models generated code: an empirical study. *Empir. Softw. Eng.* 30, 1–48. doi: 10.1007/s10664-025-10614-4
- Tirrat, P., Jeong, W., and Hwang, S. J. (2024). Automl-agent: a multi-agent llm framework for full-pipeline automl. *arXiv preprint arXiv:2410.02958*.
- Tsai, Y.-D., Tsai, Y.-C., Huang, B.-W., Yang, C.-P., and Lin, S.-D. (2023). Automl-GPT: large language model for automl. *arXiv preprint arXiv:2309.01125*.
- Wang, C., He, Z., He, J., Ye, J., and Shen, Y. (2024). "Histology image artifact restoration with lightweight transformer based diffusion model," in *International Conference on Artificial Intelligence in Medicine* (Springer), 81–89. doi: 10.1007/978-3-031-66535-6\_9
- Wang, L., and Shen, Y. (2024). Evaluating causal reasoning capabilities of large language models: a systematic analysis across three scenarios. *Electronics* 13:4584. doi: 10.3390/electronics13234584
- Wen, Y., Wang, Y., Yi, K., Ke, J., and Shen, Y. (2024). "Diffimpute: tabular data imputation with denoising diffusion probabilistic model," in *2024 IEEE International Conference on Multimedia and Expo (ICME)* (IEEE), 1–6. doi: 10.1109/ICME57554.2024.10687685
- Zan, D., Chen, B., Zhang, F., Lu, D., Wu, B., Guan, B., et al. (2022). Large language models meet nl2code: a survey. *arXiv preprint arXiv:2212.09420*.
- Zhang, S., Gong, C., Wu, L., Liu, X., and Zhou, M. (2023). Automl-gpt: Automatic machine learning with gpt. *arXiv preprint arXiv:2305.02499*.
- Zimmer, L., Lindauer, M., and Hutter, F. (2021). Auto-pytorch: multi-fidelity metalearning for efficient and robust autodl. *IEEE Trans. Pattern Anal. Mach. Intell.* 43, 3079–3090. doi: 10.1109/TPAMI.2021.3067763
- Zoph, B. (2016). Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*.