LLM-FE: Automated Feature Engineering for Tabular Data with LLMs as Evolutionary Optimizers

Nikhil Abhyankar Virginia Tech nikhilsa@vt.edu Parshin Shojaee Virginia Tech parshinshojaee@vt.edu Chandan K. Reddy Virginia Tech reddy@cs.vt.edu

Abstract

Automated feature engineering plays a critical role in improving predictive model performance for tabular learning tasks. Traditional automated feature engineering methods are limited by their reliance on pre-defined transformations within fixed, manually designed search spaces, often neglecting domain knowledge. Recent advances using Large Language Models (LLMs) have enabled the integration of domain knowledge into the feature engineering process. However, existing LLM-based approaches use direct prompting or rely solely on validation scores for feature selection, failing to leverage insights from prior feature discovery experiments or establish meaningful reasoning between feature generation and data-driven performance. To address these challenges, we propose LLM-FE, a novel framework that combines evolutionary search with the domain knowledge and reasoning capabilities of LLMs to automatically discover effective features for tabular learning tasks. LLM-FE formulates feature engineering as a program search problem, where LLMs propose new feature transformation programs iteratively, and data-driven feedback guides the search process. Our results demonstrate that LLM-FE consistently outperforms state-of-the-art baselines, significantly enhancing the performance of tabular prediction models across diverse classification and regression benchmarks.

The code is available at: O:https://github.com/nikhilsab/LLMFE

1 Introduction

Feature engineering, the process of transforming raw data into meaningful features for machine learning models, is crucial for improving predictive performance, particularly when working with tabular data [10]. In many tabular prediction tasks, well-designed features have been shown to significantly enhance the performance of tree-based models, often outperforming deep learning models that rely on learned representations [14]. However, data-centric tasks such as feature engineering are one of the most challenging processes in the tabular learning workflow [2, 20], as they require experts and data scientists to explore many possible combinations in the vast combinatorial space of feature transformations. Classical feature engineering methods [22, 24, 23, 21, 45] construct extensive search spaces of feature processing operations, relying on various search and optimization techniques to identify the most effective features. However, these search spaces are mostly constrained by predefined, manually designed transformations and often fail to incorporate domain knowledge [45]. Domain knowledge can serve as an invaluable prior for identifying these transformations, leading to reduced complexity and more interpretable and effective features [20].

Recently, Large Language Models (LLMs) have emerged as a powerful solution to this challenge, offering access to extensive embedded domain knowledge that can be leveraged for feature engineering. While recent approaches have demonstrated promising results in incorporating this knowledge into automated feature discovery, current LLM-based methods [20, 17] rely predominantly



Figure 1: Overview of the LLM-FE Framework. For a given dataset, LLM-FE follows these steps: (a) **New Feature Generation**, where an LLM generates feature transformation hypotheses as programs for a given tabular dataset; (b) **Feature Engineering**, where the feature transformation program is applied to the underlying dataset, resulting in a modified dataset; (c) **Feature Evaluation**, where the modified dataset with the new features is evaluated using a prediction model; (d) **Experience Management**, which maintains a buffer of high-scoring programs that act as in-context samples for LLM's iterative refinement prompt. The features generated by LLM-FE are interpretable, using LLM's knowledge of the problem's domain.

on direct prompting mechanisms or validation scores to guide the feature generation process. These approaches do not leverage insights from prior feature discovery experiments, thereby falling short of establishing meaningful reasoning between feature generation and data-driven performance.

To address these limitations, we propose LLM-FE, *a novel framework integrating the capabilities of LLMs with tabular prediction models and evolutionary search to facilitate effective feature optimization.* As shown in Figure 1, LLM-FE follows an iterative process to generate and evaluate the hypothesis of the feature transformation, using the performance of the tabular prediction model as a reward to enhance the generation of effective features. Starting from an initial feature transformation program, LLM-FE leverages the LLMs' embedded domain knowledge by incorporating task-specific details, feature descriptions, and a subset of data samples to generate new feature discovery programs (Figure 1(a)). At each iteration, LLM acts as a knowledge-guided evolutionary optimizer, which mutates examples of previously successful feature transformation programs to generate new effective features [30]. The newly proposed features are then integrated with the original dataset to yield an augmented dataset (Figure 1(b)). The prediction model's performance is evaluated on a held-out validation set derived from the augmented dataset (Figure 1(c)), provides data-driven feedback that, combined with a dynamic memory of previously explored feature transformation programs (Figure 1(d)), guides the LLM to refine its feature generation iteratively.

Table 1 compares LLM-FE to several stateof-the-art classical and LLM-based feature engineering methods. Traditional methods lack adaptability and deeper contextual understanding, while LLM-based methods generate simple features [25] or use feedback to iteratively refine only a single rule. In contrast, LLM-FE supports all four aspects by leveraging LLM-based domain knowledge and feedback-driven optimization to generalize well across table prediction tasks. We evaluate LLM-FE

Table 1: Comparison of existing feature engineering methods.

Method	Domain Knowledge	Feedback Driven	Complex Features	Multi-Feature Refinement
AutoFeat [21]	×	X	1	×
OpenFE [45]	×	×	1	×
FeatLLM [17]	1	×	×	×
CAAFE [20]	1	1	×	×
OCTree [31]	✓	1	×	×
LLM-FE	 Image: A second s	1	1	1

with GPT-3.5-Turbo [34] and Llama-3.1-8B-Instruct [11] backbones on classification and regression tasks across diverse tabular datasets. LLM-FE consistently outperforms the state-of-the-art

feature engineering methods, identifying contextually relevant features that improve downstream performance. In particular, we observe improvements with tabular models like XGBoost [6], TabPFN [19], and MLP [13]. Our analysis also highlights the importance of evolutionary search in achieving effective results. The major contributions of this work can be summarized as.

• We introduce LLM-FE, a novel framework that leverages the LLM's reasoning capabilities and domain knowledge, coupled with evolutionary search, to perform automated feature engineering for tabular data.

• Our experimental results demonstrate the effectiveness of LLM-FE, showcasing its ability to outperform state-of-the-art baselines across diverse benchmark datasets.

• Through a comprehensive ablation study, we highlight the critical role of domain knowledge, evolutionary search, data-driven feedback, and data samples in guiding the LLM to efficiently explore the feature space and discover impactful features more effectively.

2 Related Works

Feature Engineering. Feature engineering involves creating meaningful features from raw data to improve predictive performance [20]. The growing complexity of datasets has driven the automation of feature engineering to reduce manual effort and optimize feature discovery. Traditional automated feature engineering methods include tree-based exploration [24], iterative subsampling [21], and transformation enumeration [22]. Learning-based methods leverage machine learning and reinforcement learning for feature transformation [33, 23, 44]. OpenFE [45] integrates a feature-boosting algorithm with a two-stage pruning strategy. These traditional approaches often fail to leverage domain knowledge for feature discovery, making LLMs well-suited for such tabular prediction tasks due to their prior contextual domain understanding.

LLMs and Optimization. Recent advances in LLMs enable them to leverage their pre-trained knowledge to handle novel tasks through techniques such as prompt engineering and in-context learning without requiring additional training [4, 40]. Despite the progress made with LLMs, they often struggle with factually incorrect or inconsistent outputs [29, 47]. Researchers have thus explored methods that use feedback or refinement mechanisms to improve LLM responses and leverage them within complex optimization tasks [29, 16]. More recent approaches involving evolutionary optimization frameworks couple LLMs with evaluators [27, 28, 41, 26], using LLMs to perform adaptive mutation and crossover operations [30]. This approach has shown success in areas such as prompt optimization [43, 15], neural architecture search [46, 5], discovery of mathematical heuristics [35], and symbolic regression [36]. Building on these concepts, our LLM-FE framework utilizes an LLM as an evolutionary optimizer by coupling its prior knowledge with data-driven refinement to discover optimal features for the underlying tabular learning task.

LLMs for Tabular Learning. LLMs have been applied to structured data by converting tables into natural language [9, 18, 39]. Approaches include table-specific tokenization for robust pretraining [42] and using fine-tuning or few-shot in-context learning for tabular prediction tasks [9, 18, 32]. Recently, LLMs have been used for feature engineering with FeatLLM [17], improving tabular predictions by generating and parsing rules to engineer binary features. CAAFE [20] introduces a context-aware approach where LLMs generate features directly from task descriptions, while OCTree [31] incorporates an additional decision tree reasoning step to iteratively improve a single rule at a time. In contrast, LLM-FE maintains multiple promising candidates and uses evolutionary search to efficiently navigate the feature space via mutation and crossover without using any additional steps. LLM-FE thus uses LLMs to efficiently navigate the optimization space of feature discovery, generating meaningful features that are informed by prior knowledge and enriched with data-driven insights and evolutionary exploration.

3 LLM-FE Approach

3.1 Problem Formulation

A tabular dataset \mathcal{D} comprises N rows (or instances), each characterized by d columns (or features). Each data instance x_i is a d-dimensional feature vector with feature names denoted by $C = \{c_j\}_{j=1}^d$. The dataset is accompanied by metadata \mathcal{M} , which contains feature descriptions and task-specific information. For supervised learning tasks, each instance x_i is associated with a corresponding label y_i , where $y_i \in \{0, 1, ..., K\}$ for classification tasks with K classes, and $y_i \in \mathbb{R}$ for regression tasks. Given a labeled tabular dataset $\mathcal{D} = (x_i, y_i)_{i=1}^N$ and prediction model f to map from the input feature space \mathcal{X} to its corresponding label space \mathcal{Y} , the feature engineering objective is to determine an optimal feature transformation \mathcal{T} , which enhances the performance of a predictive model when trained on the transformed input space. Formally, the feature engineering task can be defined as:

$$\max_{\mathcal{T}} \mathcal{E}(f^*(\mathcal{T}(\mathcal{X}_{\text{val}})), \mathcal{Y}_{\text{val}})$$
(1)

subject to:

$$f^* = \arg\min_{f} \mathcal{L}_f(f(\mathcal{T}(\mathcal{X}_{\mathrm{tr}})), \mathcal{Y}_{\mathrm{tr}})$$
(2)

where $(\mathcal{X}_{tr}, \mathcal{Y}_{tr})$ and $(\mathcal{X}_{val}, \mathcal{Y}_{val})$ are the sub-training set and validation set, respectively, that is derived from the training data $(\mathcal{X}_{train}, \mathcal{Y}_{train})$. The feature transformation \mathcal{T} generated by the LLM π_{θ} and defined as $\mathcal{T} = \pi_{\theta}(\mathcal{X}_{train})$, meaning the transformation is learned from the training data by the LLM. The predictive model f^* is then trained on the transformed training data $\mathcal{T}(\mathcal{X}_{train})$ to minimize loss. Consequently, the bilevel optimization problem seeks to identify the feature transformations \mathcal{T} that maximize the performance \mathcal{E} on $\mathcal{T}(\mathcal{X}_{val})$ while minimizing the loss function on the transformed training data, thereby efficiently exploring the potential feature space.

3.2 Feature Generation

Figure 1(a) illustrates the feature generation step that uses an LLM to create multiple new feature transformation programs, leveraging the model's prior knowledge, reasoning, and in-context learning abilities to effectively explore the feature space.

3.2.1 Input Prompt

To facilitate the creation of effective and contextually relevant feature discovery programs, we develop a structured prompting methodology. The prompt is designed to provide comprehensive data-specific information, an initial feature transformation program for the evolution starting point, an evaluation function, and a well-defined output format (see Appendix B.2 for more details). Our input prompts p are composed of the following key elements:

Instruction. The LLM is assigned the task of finding the most relevant features to help solve the given task. The task emphasizes using the LLM's prior knowledge of the dataset's domain to generate features. The LLM is explicitly instructed to generate novel features and provide clear step-by-step reasoning for their relevance to the prediction task. Moreover, since LLMs tend to generate simple features, we specifically instruct the LLM to generate complex features.

Dataset Specification. After providing the instructions, we provide LLM with the dataset-specific information from the metadata \mathcal{M} . This information encompasses a detailed description of the intended downstream task, along with the feature names C and their corresponding descriptions. In addition, we provide a limited number of representative samples from the tabular dataset. To improve the effective interpretation of the data, we adopt the serialization approach used in previous works [9, 18, 17]. We serialized the data samples as follows:

Serialize
$$(x_i, y_i, C) =$$
'If c_1 is x_i^1, \dots, c_d is x_i^d . Then Result is y_i ' (3)

By providing dataset-specific details, we guide the language model to focus on the most contextually pertinent features that directly support the dataset and task objective.

Evaluation Function. The evaluation function, incorporated into the prompt, guides the language model to generate feature transformation programs that align with performance objectives. These programs augment the original dataset with new features, which are assessed on the basis of a prediction model's performance when trained on the augmented data. The model's evaluation score on the augmented validation set serves as an indicator of feature quality. By including the evaluation function in the prompt, the LLM generates programs that are inherently aligned with the desired performance criteria.

In-Context Demonstration. Specifically, we sample the k highest-performing demonstrations from previous iterations, enabling the LLM to build on successful outputs. The iterative interaction between the LLM's generative outputs and the evaluator's feedback, informed by these examples, facilitates a systematic refinement process. With each iteration, the LLM progressively improves its outputs by leveraging patterns and insights identified in previous successful demonstrations.

3.2.2 Feature Sampling

At each iteration t, we construct the prompt p_t by sampling the previous iteration as input to the LLM π_{θ} , resulting in the output $\mathcal{T}_1, \ldots, \mathcal{T}_b = \pi_{\theta}(p_t)$ representing a set of b sampled programs. To promote diversity and maintain a balance between exploration (creativity) and exploitation (prior knowledge), we employ stochastic temperature-based sampling. Each of the sampled feature transforms (\mathcal{T}_i) is executed before evaluation to discard error-prone programs. This ensures that only valid and executable feature transformation programs are considered further in the optimization pipeline. In addition, to ensure computational efficiency, a maximum execution time threshold is enforced, discarding any programs that exceed it.

3.3 Data-Driven Evaluation

As illustrated in Figure 1(b), we use the generated features to augment the original dataset with the newly derived features. Similar to [20, 31], our feature evaluation process comprises two stages: (i) model training on the augmented dataset, and (ii) performance assessment for feature quality (Figure 1(c)). We fit a tabular predictive model f^* , to the transformed training set $\mathcal{T}(\mathcal{X}_{tr})$, by minimizing the loss \mathcal{L}_f as shown in Eq.1. Subsequently, we evaluated the LLM-generated feature transformations \mathcal{T} by evaluating the model's performance on the augmented validation set $\mathcal{T}(\mathcal{X}_{val})$ (see Eqs. 1 and 2). As explained in Section 3.1, the objective is to find optimal features that maximize the performance \mathcal{E} , i.e., accuracy for classification and error metrics for regression.

3.4 Experience Management

То promote diverse feature discovery and avoid stagnation in local optima, LLM-FE employs evolutionary multi-population experience management (Figure 1(d)) to store feature discovery programs in a dedicated database. Then, it uses samples from this database to construct in-context examples for LLM, facilitating the generation of novel features. This step consists of two components: (i) multi-population memory to maintain a longterm memory buffer, and (ii) sampling from this memory buffer to construct in-context example demonstrations. After evaluating the feature transforms in iteration t, we store the pair of feature transforms and score (\mathcal{T}, s) in the population buffer \mathcal{P}_t to iteratively refine the search process. То effectively evolve a population of programs, we adopt a multi-population model inspired by the 'island' model employed by [7, 36, 35]. The program

Algorithm 1 LLM-FE

Input: LLM π_{θ} , Dataset \mathcal{D} , Metadata \mathcal{M} , Iterations T, Model f, Metric \mathcal{E} 1: $\mathcal{P}_0 \leftarrow \texttt{BufferInit}()$ 2: $\mathcal{T}^*, s^* \leftarrow \text{null}, -\infty$ 3: $p \leftarrow \texttt{UpdatePrompt}(\mathcal{D}, \mathcal{M})$ 4: for t = 1 to T - 1 do $p_t \leftarrow p + \mathcal{P}_{t-1}.\texttt{topk}()$ 5: 6: $\{\mathcal{T}_j\}_{j=1}^b \leftarrow \pi_\theta(p_t)$ for j = 1 to b do 7: 8: $s_j \leftarrow \texttt{FeatureScore}(f, \mathcal{T}_j, \mathcal{D}, \mathcal{E})$ 9: if $s_j > s^*$ then $\check{\mathcal{T}}^*, s^* \leftarrow \mathcal{T}_j, s_j$ 10: end if 11: 12: $\mathcal{P}_t \leftarrow \texttt{UpdateBuffer}(\mathcal{P}_{t-1}, \mathcal{T}_j, s_j)$ end for 13: 14: end for 15: return \mathcal{T}^*, s^*

population is divided into m independent islands, each evolving separately and initialized with a copy of the user's initial example (see Figure 6(d)). This enables parallel exploration of the feature space, mitigating the risk of suboptimal solutions. At each iteration t, we select one of the m islands and sample programs from the memory buffer to update the prompt with new in-context examples. The newly generated feature samples b are evaluated, and if their scores s_j exceed the current best score, the feature score pair (\mathcal{T}_j, s_j) is added to the same island from which the in-context examples were sampled. To preserve diversity and ensure that programs with different performance characteristics are maintained in the buffer, we cluster programs within islands based on their signature, defined by their scores. To build refinement prompts, we follow the sampling process from [35], first sampling one of the m available islands, followed by sampling the k programs from the selected island to create k-shot in-context examples for the LLM. Cluster selection prefers high-scoring programs and follows Boltzmann sampling [8] with a score-based probability of choosing a cluster *i*: $P_i = \frac{exp(s_i/\tau_c)}{\sum_i exp(s_i/\tau_c)}$, where s_i denotes the mean score of the *i*-th cluster and τ_c is the temperature parameter. The sampled feature transformation programs from the memory buffer are then included in the prompt as examples to guide LLM toward successful feature transformations—incurring negligible computational overhead. Refer to Appendix D.1 for more details.

Algorithm 1 presents the pseudocode of LLM-FE. We begin with the initialization of a memory buffer BufferInit, incorporating an initial population that contains a simple feature transform. This initialization serves as the starting point for the evolutionary search for feature transformation programs to be evolved in the subsequent steps. At each iteration t, the function topk is used to sample k in-context examples from the population of the previous iteration \mathcal{P}_{t-1} to update the prompt. Subsequently, we prompt the LLM using this updated prompt to sample b new programs. The sampled programs are then evaluated using FeatureScore, which represents the Data-Driven Evaluation (Section 3.3). After T iterations, the best-scoring program \mathcal{T}^* from \mathcal{P}_t and its corresponding score s^* are returned as the optimal solution found for the problem. LLM-FE employs an iterative search to enhance programs, harnessing the LLM's capabilities. Learning from the evolving pool of experiences in its buffer, the LLM steers the search toward effective solutions.

4 Experimental Setup

We evaluated LLM-FE on a range of tabular datasets, encompassing classification and regression tasks. Our experimental analysis included quantitative comparisons with baselines and detailed ablation studies. Specifically, we assessed our approach using three known tabular predictive models with distinct architectures: (1) XGBoost, a tree-based model [6], (2) MLP, a neural model [13], and (3) TabPFN [19], a transformer-based foundation model [38]. The results highlight LLM-FE's capability to generate effective features that consistently enhance the performance of different prediction models across diverse datasets.

4.1 Datasets

We followed [20] to select datasets from [17, 20, 45] that include descriptive feature names and are not perfectly solvable using XGBoost. Our analysis contains 19 classification and 10 regression datasets, each containing mixed categorical and numerical features. We also include 8 large-scale, high-dimensional classification datasets to ensure comprehensive evaluation. These datasets were sourced from established machine learning repositories, including OpenML [37, 12], UCI Machine Learning Repository [3], and Kaggle. Each dataset is accompanied by metadata, which includes a natural language description of the prediction task and descriptive feature names. We partitioned each dataset into train and test sets using an 80-20 split. Following [20], we evaluated all methods over five iterations, each time using a distinct random seed and train-test splits. For more details, check Appendix A.

4.2 Baselines

We evaluated LLM-FE against state-of-the-art feature engineering approaches, including OpenFE [45] and AutoFeat [21], as well as LLM-based methods CAAFE [20], FeatLLM [17] and OCTree [31]. We used XGBoost as the default tabular data prediction model in comparison with baselines and employed GPT-3.5-Turbo as the default LLM backbone for all LLM-based methods. To ensure a fair comparison, all LLM-based baselines were configured to query the LLM backbone for a total of 20 samples until they converged to their best performance. Appendix B.1 contains additional implementation details.

4.3 LLM-FE Configuration

In our experiments, we utilized GPT-3.5-Turbo and Llama-3.1-8B-Instruct as backbone LLMs, with a sampling temperature parameter of t = 0.8 and the number of islands set to m = 3. At each iteration, the LLM generated b = 3 feature transformation programs per prompt in Python. To ensure consistency with baselines, LLM-FE was also configured with a total of 20 LLM samples for each experiment. Finally, we sampled the top m (where m denotes the number of islands) feature discovery programs based on their respective validation scores and reported the final prediction through an ensemble. More implementation details are provided in Appendix B.2.

Table 2: Performance of XGBoost on Classification Datasets using various Feature Engineering (FE) Methods, evaluated using accuracy (higher values indicate better performance). We report the mean values and standard deviation across five splits. \times : denotes execution time of greater than 12 hours or failure due to execution errors. **bold:** indicates the best performance. <u>underline</u>: indicates the second-best performance. 'n': indicates the number of samples; 'p': indicates the number of features.

Dataset	n	n	Base	Classical F	E Methods	LLM	-based FE Me	thods	LLM-FE
Duniber		Р	Dust	AutoFeat	OpenFE	CAAFE	FeatLLM	OCTree	
adult	48.8k	14	$0.873 \pm {\scriptstyle 0.002}$	×	$0.873 \pm {\scriptstyle 0.002}$	$0.872 \pm {\scriptstyle 0.002}$	$0.842 \pm {\scriptstyle 0.003}$	$0.870\pm {\scriptstyle 0.002}$	$0.874 \pm {\scriptstyle 0.003}$
arrhythmia	452	279	0.657 ± 0.019	×	×	×	×	×	0.659 ± 0.018
balance-scale	625	4	$0.856\pm {\scriptscriptstyle 0.020}$	$0.925\pm {\scriptscriptstyle 0.036}$	0.986 ± 0.009	0.966 ± 0.029	$0.800 \pm {}_{0.037}$	$0.882\pm$ 0.022	0.990 ± 0.013
bank-marketing	45.2k	16	0.906 ± 0.003	×	$\overline{0.908_{0.002}}$	$0.907 \pm {\scriptstyle 0.002}$	$0.907 \pm {\scriptstyle 0.002}$	$0.900\pm{\scriptstyle 0.002}$	$0.907 \pm {\scriptstyle 0.002}$
breast-w	699	9	$0.956\pm {\scriptscriptstyle 0.012}$	$0.956\pm {\scriptstyle 0.019}$	$0.956\pm {\scriptscriptstyle 0.014}$	$\overline{0.960\pm}_{0.009}$	0.967 ± 0.015	0.969 ± 0.009	0.970 ± 0.009
blood-transfusion	748	4	0.742 ± 0.012	$0.738 \pm {}_{0.014}$	0.747 ± 0.025	$0.749 \pm {}_{0.017}$	$0.771 \pm {\scriptstyle 0.016}$	0.755 ± 0.026	0.751 ± 0.036
car	1728	6	0.995 ± 0.003	0.998 ± 0.003	0.998 ± 0.003	0.999 ± 0.001	$0.808 \pm {}_{0.037}$	0.995 ± 0.004	0.999 ± 0.001
cdc diabetes	253k	21	$0.849\pm {\scriptstyle 0.001}$	×	$\overline{0.849\pm {}_{ m 0.001}}$	$0.849\pm{\scriptstyle 0.001}$	$0.849\pm{\scriptstyle 0.001}$	$0.849\pm{\scriptstyle 0.001}$	$0.849 \pm {\scriptstyle 0.001}$
cmc	1473	9	0.528 ± 0.029	0.505 ± 0.015	0.517 ± 0.007	$0.524\pm$ 0.016	0.479 ± 0.015	0.525 ± 0.027	$0.531 \pm {\scriptstyle 0.015}$
communities	1.9k	103	0.706 ± 0.016	×	$0.704\pm{\scriptstyle 0.009}$	0.707 ± 0.013	0.593 ± 0.012	$0.708\pm{\scriptstyle 0.016}$	$0.711 \pm \scriptscriptstyle 0.012$
covtype	581k	54	$0.870\pm {\scriptscriptstyle 0.001}$	×	$0.885\pm$ 0.007	$0.872\pm {\scriptstyle 0.003}$	$0.554\pm {\scriptstyle 0.001}$	0.832 ± 0.002	$0.882\pm{\scriptstyle 0.003}$
credit-g	1000	20	0.751 ± 0.019	0.757 ± 0.017	0.758 ± 0.017	$0.751 \pm {}_{0.020}$	$0.707 \pm {}_{0.034}$	$0.753 \pm {}_{0.021}$	0.766 ± 0.015
eucalyptus	736	19	$0.655 \pm {}_{0.024}$	0.664 ± 0.028	0.663 ± 0.033	$0.679 \pm {}_{\scriptscriptstyle 0.024}$	×	$0.658 \pm {}_{0.041}$	0.668 ± 0.027
heart	918	11	0.858 ± 0.013	$0.857 \pm {}_{0.021}$	$0.854 \pm {}_{0.020}$	0.849 ± 0.023	0.865 ± 0.030	$0.852\pm{\scriptstyle 0.022}$	0.866 ± 0.021
jungle_chess	44.8k	6	0.869 ± 0.001	×	0.900 ± 0.004	0.901 ± 0.038	0.577 ± 0.002	0.869 ± 0.002	0.969 ± 0.004
myocardial	1.7k	111	$0.784 \pm {}_{0.023}$	×	$0.787 \pm {}_{0.026}$	0.789 ± 0.023	$0.778 \pm {}_{0.023}$	$0.787 \pm {}_{0.031}$	$0.789 \pm {\scriptstyle 0.023}$
pc1	1109	21	0.931 ± 0.004	$0.931 \pm {}_{0.014}$	0.931 ± 0.009	0.929 ± 0.005	$0.933\pm ext{0.007}$	0.934 ± 0.007	0.935 ± 0.006
tic-tac-toe	958	9	0.998 ± 0.002	$1.000\pm$ 0.000	0.994 ± 0.006	0.996 ± 0.003	0.653 ± 0.037	0.997 ± 0.003	0.998 ± 0.005
vehicle	846	18	$0.754 \pm \scriptscriptstyle 0.016$	$0.788 \pm {\scriptstyle 0.018}$	$\underline{0.785 \pm {\scriptstyle 0.008}}$	$0.771 \pm {\scriptstyle 0.019}$	$0.744 \pm \scriptscriptstyle 0.035$	$0.753\pm{\scriptstyle 0.036}$	$\overline{0.761\pm \scriptscriptstyle 0.027}$
Mean Rank			4.26	4.89	3.26	3.31	4.94	3.84	1.47

Table 3: Performance of XGBoost on Regression Datasets using various Feature Engineering (FE) Methods, evaluated using normalized root mean square error (N-RMSE) (lower values indicate better performance). We report the mean and standard deviation across five splits. **bold:** indicates the best performance. <u>underline</u>: indicates the second-best performance. 'n': indicates the number of samples; 'p': indicates the number of features.

Dataset	n	р	Base	Classical F	E Methods	LLM-FE
		1		AutoFeat	OpenFE	
airfoil_self_noise	1503	6	0.013 ± 0.001	0.012 ± 0.001	0.013 ± 0.001	0.011 ± 0.001
bike	17389	11	0.216 ± 0.005	$\overline{0.223 \pm 0.006}$	0.216 ± 0.007	0.207 ± 0.006
cpu_small	8192	10	0.034 ± 0.003	0.034 ± 0.002	0.034 ± 0.002	0.033 ± 0.003
crab	3893	8	0.234 ± 0.009	$\overline{0.228\pm0.008}$	0.224 ± 0.001	0.223 ± 0.013
diamonds	53940	9	0.139 ± 0.002	0.140 ± 0.004	$\overline{0.137 \pm 0.002}$	0.134 ± 0.002
forest-fires	517	13	1.469 ± 0.080	1.468 ± 0.086	1.448 ± 0.113	1.417 ± 0.083
housing	20640	9	0.234 ± 0.009	0.231 ± 0.013	$\overline{0.224\pm0.005}$	0.218 ± 0.009
insurance	1338	7	0.397 ± 0.020	0.384 ± 0.024	$\overline{0.383\pm0.022}$	0.381 ± 0.028
plasma_retinol	315	13	0.390 ± 0.032	0.411 ± 0.036	$\overline{0.392\pm0.032}$	0.388 ± 0.033
wine	4898	10	0.110 ± 0.001	0.109 ± 0.001	0.108 ± 0.001	$\textbf{0.105} \pm \textbf{0.001}$
Mean Rank			3.40	3.10	2.20	1.00

5 Results and Discussion

5.1 Performance Comparisons

In Table 2, we compare LLM-FE against various feature engineering baselines across 19 classification datasets. The results demonstrate that LLM-FE consistently enhances predictive performance from the base model (using raw data). LLM-FE also obtains the lowest mean rank (best performance) at a lower computational cost (see Appendix D.1), showing better effectiveness in enhancing feature discovery compared to other leading baselines. To further evaluate the effectiveness of LLM-FE, we perform experiments on 10 regression datasets using the same evaluation settings employed for the classification datasets. Due to the lack of regression data implementations in the available codebases for LLM-based baselines, in Table 3, we restrict our comparison to only non-LLM methods (OpenFE and AutoFeat), which have been previously validated on regression tasks. The results indicate that LLM-FE outperforms all baseline methods, achieving the lowest mean rank and consistently improving across all datasets.

5.2 Generalizability Analysis

To evaluate the generalizability of the LLM-FE, we examine its performance across multiple tabular prediction models and various LLM backbones. Specifically, we employ two LLM backbones, Llama-3.1-8B-Instruct and GPT-3.5-Turbo, in conjunction with three distinct tabular prediction models: XGBoost [6], a widely-used tree-based algorithm for tabular tasks; Multilayer Perceptron (MLP), a simple yet common deep-learning architecture tailored to tabular datasets [13]; and TabPFN [19], a recent transformer-based foundation model specifically designed for tabular data. Table 4 summarizes our findings, demonstrating that LLM-FE effectively identifies features that enhance the performance of various prediction models and LLM backbones across different tasks. Notably, the results indicate that features generated by LLM-FE using either LLM backbone consistently improve base model prediction performance compared to scenarios without any feature engineering. Detailed analyses on additional prediction models, hyperparameter optimization, feature transferability and more are provided in Appendix C.

Table 4: Performance improvement by LLM-FE using different prediction models and LLM backbones. We report the aggregated values for accuracy on classification tasks and normalized root mean square error on regression tasks. All results represent the mean and standard deviation computed across five splits. **bold:** indicates the best performance. TabPFN* evaluations are conducted using only 10,000 samples due to its limited processing capacity.

Method	LLM	Regression \downarrow				
XGBoost						
Base	-	$0.820\pm$ 0.020	$0.324 \pm \textbf{0.016}$			
LLM-FE	Llama 3.1-8B GPT-3.5 Turbo	0.832 ± 0.021 0.840 ± 0.022	$\begin{array}{c} 0.310 \pm 0.022 \\ \textbf{0.306} \pm \textbf{0.015} \end{array}$			
MLP						
Base	-	$0.745\pm \scriptscriptstyle 0.034$	$0.871 \pm \textbf{0.027}$			
LLM-FE	Llama 3.1-8B GPT-3.5 Turbo	0.768 ± 0.032 0.791 ± 0.029	$\begin{array}{c} 0.794 \pm \textit{0.016} \\ \textbf{0.631} \pm \textit{0.043} \end{array}$			
	Tal	bPFN*				
Base	_	$0.852\pm _{0.028}$	$0.289 \pm \textbf{0.016}$			
LLM-FE	Llama 3.1-8B GPT-3.5 Turbo	$\begin{array}{c} 0.856 \pm \textbf{0.017} \\ \textbf{0.863} \pm \textbf{0.018} \end{array}$	$\begin{array}{c} 0.288 \pm 0.016 \\ \textbf{0.286} \pm \textbf{0.015} \end{array}$			

6 Analysis

6.1 Ablation Study

We perform an ablation study on the classification datasets (<10,000 samples) listed in Table 2 to assess the contribution of each component in LLM-FE. Figure 2 illustrates the impact of individual components on overall performance, using XGBoost and GPT-3.5-Turbo. We report the accuracy aggregated and normalized over all the datasets. In the 'w/o Domain Knowledge' setting, dataset and task-specific details are removed from the prompt and feature names are anonymized with generic placeholders such as C_1, C_2, \ldots, C_n . In this way, we remove any semantic meaning that could provide contextual insights about the problem. Without domain knowledge, the performance significantly drops to 0.626, underscoring its critical role in generating meaningful features. The 'w/o Evolutionary **Refinement'** setting) also leads to the greatest decline in performance (0.587), emphasizing the importance of iterative data-driven feedback in addition to domain knowledge for refining feature



Figure 2: Aggregated ablation study results across classification datasets, showcasing the impact of individual components on LLM-FE's performance: (a) Data Examples, (b) Domain Knowledge, and (c) Evolutionary Refinement. Values are normalized with respect to the base LLM-FE model to facilitate fair comparison across conditions.

transforms. Lastly, the results show that 'w/o Data Examples' variant leads to only a slight performance drop, as LLMs might struggle to comprehend the nuances and patterns within the data samples. LLM-FE benefits significantly from each component, leading to an improved performance.



Figure 3: Qualitative Analysis on Impact of Domain Knowledge. illustrating how LLM-FE (b) **utilizes** domain knowledge to create meaningful features with descriptions, in contrast to feature engineering without domain insights (a) leading to uninterpretable outputs.



Figure 4: Quantitative impact of domain knowledge on model accuracy. Using domain knowledge boosts performance compared to both the base model and LLM-FE without domain knowledge.



Figure 5: Performance Trajectory Analysis. for LLM-FE *w/o* evolutionary refinement and LLM-FE. LLM-FE demonstrates a better trajectory, highlighting the advantage of evolutionary refinement.

6.2 Impact of Domain Knowledge and Evolutionary Refinement

Figure 3 illustrates the qualitative benefits of incorporating domain knowledge into feature engineering. In this example, two approaches are contrasted: one without domain knowledge (Figure 3(a)), and LLM-FE guided by domain-specific insights through an LLM-based feature engineering (Figure 3(b)). The domain-agnostic variant creates arbitrary transformations, such as combining features C1 and C3 using a square root of their product and dropping feature C2 without clear justification. In contrast, LLM-FE leverages its embedded knowledge to derive interpretable and clinically meaningful features. Figure 4 presents a quantitative comparison of model performance on the same dataset, showing that LLM-FE with domain knowledge achieves the highest accuracy, outperforming both the base model and LLM-FE without domain knowledge. Figure 5 illustrates the validation accuracy trajectory of LLM-FE with and without evolutionary refinement across 20 iterations. The variant without refinement shows early improvement but quickly plateaus, indicating convergence to a local optimum. In contrast, LLM-FE continues to improve across iterations, achieving higher accuracy overall. This comparison highlights the effectiveness of evolutionary refinement in enhancing performance by enabling the model to escape local optima and optimize more effectively. For detailed analysis on more datasets, feature quality, and computational efficiency, see Appendix D.

7 Conclusion

In this work, we introduce a novel framework LLM-FE that leverages LLMs as evolutionary optimizers to discover new features for tabular prediction tasks. By combining LLM-driven hypothesis generation with data-driven feedback and evolutionary search, LLM-FE effectively automates the feature engineering process. Through comprehensive experiments on diverse tabular learning

tasks, we demonstrate that LLM-FE consistently outperforms state-of-the-art baselines, delivering substantial improvements in predictive performance across various tabular prediction models. Future work could explore integrating more powerful or domain-specific language models to enhance the relevance and quality of generated features for domain-specific problems. Moreover, our framework could extend beyond feature engineering to other stages of the tabular learning and data-centric pipeline, such as data augmentation, automated data cleaning (including imputation and outlier detection), and model tuning.

Acknowledgments

This research was partially supported by the U.S. National Science Foundation (NSF) under Grant No. 2416728.

References

- [1] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2623–2631, 2019.
- [2] Anaconda. The state of data science 2020. Website, 2020.
- [3] Arthur Asuncion, David Newman, et al. Uci machine learning repository, 2007.
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. Advances in neural information processing systems, 33:1877–1901, 2020.
- [5] Angelica Chen, David Dohan, and David So. Evoprompting: language models for code-level neural architecture search. *Advances in Neural Information Processing Systems*, 36, 2024.
- [6] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining, pages 785–794, 2016.
- [7] Miles Cranmer. Interpretable machine learning for science with pysr and symbolic regression. jl. *arXiv preprint arXiv:2305.01582*, 2023.
- [8] Michael De La Maza and Bruce Tidor. Increased flexibility in genetic algorithms: The use of variable boltzmann selective pressure to control propagation. In *Computer Science and Operations Research*, pages 425–440. Elsevier, 1992.
- [9] Tuan Dinh, Yuchen Zeng, Ruisu Zhang, Ziqian Lin, Michael Gira, Shashank Rajput, Jy-yong Sohn, Dimitris Papailiopoulos, and Kangwook Lee. Lift: Language-interfaced fine-tuning for non-language machine learning tasks. *Advances in Neural Information Processing Systems*, 35:11763–11784, 2022.
- [10] Pedro Domingos. A few useful things to know about machine learning. *Communications of the ACM*, 55(10):78–87, 2012.
- [11] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. arXiv preprint arXiv:2407.21783, 2024.
- [12] Matthias Feurer, Jan N Van Rijn, Arlind Kadra, Pieter Gijsbers, Neeratyoy Mallik, Sahithya Ravi, Andreas Müller, Joaquin Vanschoren, and Frank Hutter. Openml-python: an extensible python api for openml. *Journal of Machine Learning Research*, 22(100):1–5, 2021.
- [13] Yury Gorishniy, Ivan Rubachev, Valentin Khrulkov, and Artem Babenko. Revisiting deep learning models for tabular data. Advances in Neural Information Processing Systems, 34:18932– 18943, 2021.
- [14] Léo Grinsztajn, Edouard Oyallon, and Gaël Varoquaux. Why do tree-based models still outperform deep learning on typical tabular data? *Advances in neural information processing systems*, 35:507–520, 2022.

- [15] Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yujiu Yang. Connecting large language models with evolutionary algorithms yields powerful prompt optimizers. arXiv preprint arXiv:2309.08532, 2023.
- [16] Patrick Haluptzok, Matthew Bowers, and Adam Tauman Kalai. Language models can teach themselves to program better. *arXiv preprint arXiv:2207.14502*, 2022.
- [17] Sungwon Han, Jinsung Yoon, Sercan O Arik, and Tomas Pfister. Large language models can automatically engineer features for few-shot tabular learning. *arXiv preprint arXiv:2404.09491*, 2024.
- [18] Stefan Hegselmann, Alejandro Buendia, Hunter Lang, Monica Agrawal, Xiaoyi Jiang, and David Sontag. Tabllm: Few-shot classification of tabular data with large language models. In *International Conference on Artificial Intelligence and Statistics*, pages 5549–5581. PMLR, 2023.
- [19] Noah Hollmann, Samuel Müller, Katharina Eggensperger, and Frank Hutter. Tabpfn: A transformer that solves small tabular classification problems in a second. *arXiv preprint arXiv:2207.01848*, 2022.
- [20] Noah Hollmann, Samuel Müller, and Frank Hutter. Large language models for automated data science: Introducing caafe for context-aware automated feature engineering. Advances in Neural Information Processing Systems, 36, 2024.
- [21] Franziska Horn, Robert Pack, and Michael Rieger. The autofeat python library for automated feature engineering and selection. In *Machine Learning and Knowledge Discovery in Databases: International Workshops of ECML PKDD 2019, Würzburg, Germany, September 16–20, 2019, Proceedings, Part I*, pages 111–120. Springer, 2020.
- [22] James Max Kanter and Kalyan Veeramachaneni. Deep feature synthesis: Towards automating data science endeavors. In 2015 IEEE international conference on data science and advanced analytics (DSAA), pages 1–10. IEEE, 2015.
- [23] Udayan Khurana, Horst Samulowitz, and Deepak Turaga. Feature engineering for predictive modeling using reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [24] Udayan Khurana, Deepak Turaga, Horst Samulowitz, and Srinivasan Parthasrathy. Cognito: Automated feature engineering for supervised learning. In 2016 IEEE 16th international conference on data mining workshops (ICDMW), pages 1304–1307. IEEE, 2016.
- [25] Jaris Küken, Lennart Purucker, and Frank Hutter. Large language models engineer too many simple features for tabular data. *arXiv preprint arXiv:2410.17787*, 2024.
- [26] Robert Lange, Yingtao Tian, and Yujin Tang. Large language models as evolution strategies. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 579–582, 2024.
- [27] Joel Lehman, Jonathan Gordon, Shawn Jain, Kamal Ndousse, Cathy Yeh, and Kenneth O Stanley. Evolution through large models. In *Handbook of Evolutionary Machine Learning*, pages 331–366. Springer, 2023.
- [28] Tennison Liu, Nicolás Astorga, Nabeel Seedat, and Mihaela van der Schaar. Large language models to enhance bayesian optimization. *arXiv preprint arXiv:2402.03921*, 2024.
- [29] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36, 2024.
- [30] Elliot Meyerson, Mark J Nelson, Herbie Bradley, Adam Gaier, Arash Moradi, Amy K Hoover, and Joel Lehman. Language model crossover: Variation through few-shot prompting. *ACM Transactions on Evolutionary Learning*, 4(4):1–40, 2024.
- [31] Jaehyun Nam, Kyuyoung Kim, Seunghyuk Oh, Jihoon Tack, Jaehyung Kim, and Jinwoo Shin. Optimized feature generation for tabular data via llms with decision tree reasoning. *arXiv* preprint arXiv:2406.08527, 2024.
- [32] Jaehyun Nam, Jihoon Tack, Kyungmin Lee, Hankook Lee, and Jinwoo Shin. Stunt: Few-shot tabular learning with self-generated tasks from unlabeled tables. *arXiv preprint arXiv:2303.00918*, 2023.

- [33] Fatemeh Nargesian, Horst Samulowitz, Udayan Khurana, Elias B Khalil, and Deepak S Turaga. Learning feature engineering for classification. In *Ijcai*, volume 17, pages 2529–2535, 2017.
- [34] R OpenAI. Gpt-4 technical report. arxiv 2303.08774. View in Article, 2(5), 2023.
- [35] Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog, M Pawan Kumar, Emilien Dupont, Francisco JR Ruiz, Jordan S Ellenberg, Pengming Wang, Omar Fawzi, et al. Mathematical discoveries from program search with large language models. *Nature*, 625(7995):468–475, 2024.
- [36] Parshin Shojaee, Kazem Meidani, Shashank Gupta, Amir Barati Farimani, and Chandan K Reddy. Llm-sr: Scientific equation discovery via programming with large language models. arXiv preprint arXiv:2404.18400, 2024.
- [37] Joaquin Vanschoren, Jan N Van Rijn, Bernd Bischl, and Luis Torgo. Openml: networked science in machine learning. ACM SIGKDD Explorations Newsletter, 15(2):49–60, 2014.
- [38] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- [39] Zifeng Wang, Chufan Gao, Cao Xiao, and Jimeng Sun. Anypredict: Foundation model for tabular prediction. *CoRR*, 2023.
- [40] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems, 35:24824–24837, 2022.
- [41] Xingyu Wu, Sheng-hao Wu, Jibin Wu, Liang Feng, and Kay Chen Tan. Evolutionary computation in the era of large language model: Survey and roadmap. *arXiv preprint arXiv:2401.10034*, 2024.
- [42] Jiahuan Yan, Bo Zheng, Hongxia Xu, Yiheng Zhu, Danny Z Chen, Jimeng Sun, Jian Wu, and Jintai Chen. Making pre-trained language models great on tabular prediction. arXiv preprint arXiv:2403.01841, 2024.
- [43] Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou, and Xinyun Chen. Large language models as optimizers, 2024.
- [44] Jianyu Zhang, Jianye Hao, Françoise Fogelman-Soulié, and Zan Wang. Automatic feature engineering by deep reinforcement learning. In *Proceedings of the 18th International Conference* on Autonomous Agents and MultiAgent Systems, pages 2312–2314, 2019.
- [45] Tianping Zhang, Zheyu Aqa Zhang, Zhiyuan Fan, Haoyan Luo, Fengyuan Liu, Qian Liu, Wei Cao, and Li Jian. Openfe: automated feature generation with expert-level performance. In International Conference on Machine Learning, pages 41880–41901. PMLR, 2023.
- [46] Mingkai Zheng, Xiu Su, Shan You, Fei Wang, Chen Qian, Chang Xu, and Samuel Albanie. Can gpt-4 perform neural architecture search? arXiv preprint arXiv:2304.10970, 2023.
- [47] Zhaocheng Zhu, Yuan Xue, Xinyun Chen, Denny Zhou, Jian Tang, Dale Schuurmans, and Hanjun Dai. Large language models can learn rules. *arXiv preprint arXiv:2310.07064*, 2023.

Impact Statement

The introduction of LLM-FE as a framework for leveraging LLMs in automated feature engineering has the potential to significantly impact the field of machine learning by improving predictive performance and reducing the manual effort involved in feature generation. This is particularly beneficial in resource-intensive domains, where efficient feature extraction and transformation are crucial for accelerating model development. By combining domain expertise with evolutionary optimization, LLM-FE improves over existing methods, which often struggle to identify optimal feature representations. While LLM-FE currently targets feature engineering, its potential extends to broader data-centric AI applications. Future extensions could include automated data cleaning (such as imputation and outlier detection), exploratory data analysis (with techniques like clustering or dimensionality reduction), and data augmentation, further strengthening the quality and robustness of training datasets. Moreover, LLM-FE could be adapted to support model tuning and hyperparameter optimization, contributing to a more streamlined and interpretable machine learning pipeline, and improving overall model generalization.

Reproducibility Statement

To ensure the reproducibility of our work, we provide comprehensive implementation details of LLM-FE. Section 3 outlines the full methodology, while Appendix B.2 offers an in-depth description of the framework, including the specific LLM prompts used. The datasets employed in our experiments are also detailed in Appendix A. Additionally, we release our code and data at this to facilitate further research.

A Dataset Details

Table 5 describes the diverse collection of datasets spanning three major categories: (1) binary classification, (2) multi-class classification, and (3) regression problems used in our evaluation. The datasets were primarily sourced from established platforms, including OpenML [37, 12], UCI [3], and Kaggle. We specifically selected datasets with descriptive feature names, excluding those with merely numerical identifiers. Each dataset includes a task description, enhancing contextual understanding for users. Our selection encompasses not only small datasets but also larger ones, featuring extensive data samples and high-dimensional datasets with over 50 features. This diverse and comprehensive selection of datasets represents a broad spectrum of real-world scenarios, varying in both feature dimensionality and sample size, thereby providing a robust framework for evaluating feature engineering works.

Dataset	#Features	#Samples	Source	ID/Name		
		Binary Cla	assification			
adult	14	48842	OpenML	1590		
blood-transfusion	4	748	OpenML	1464		
bank-marketing	16	45211	OpenML	1461		
breast-w	9	699	OpenML	15		
credit-g	20	1000	OpenML	31		
tic-tac-toe	9	958	OpenML	50		
pc1	21	1109	OpenML	1068		
pima-indian-diabetes	8	768	OpenML	43582		
Multi-class Classification						
arrhythmia	279	452	OpenML	5		
balance-scale	4	625	OpenML	11		
car	6	1728	OpenML	40975		
cmc	9	1473	OpenML	23		
eucalyptus	19	736	OpenML	188		
jungle_chess	6	44819	OpenML	41027		
vehicle	18	846	OpenML	54		
cdc diabetes	21	253680	Kaggle	diabetes-health-indicators-dataset		
heart	11	918	Kaggle	heart-failure-prediction		
communities	103	1994	UCI	communities-and-crime		
myocardial	111	1700	UCI	myocardial-infarction-complications		
		Regr	ession			
airfoil self noise	6	1503	OpenML	44957		
cpu small	12	8192	OpenML	562		
diamonds	9	53940	OpenML	42225		
plasma retinol	13	315	OpenML	511		
forest-fires	13	517	OpenML	42363		
housing	9	20640	OpenML	43996		
crab	8	3893	Kaggle	crab-age-prediction		
insurance	7	1338	Kaggle	us-health-insurancedataset		
bike	11	17389	UČI	bike-sharing-dataset		
wine	10	4898	UCI	wine-quality		

Table 5: Dataset statisti	ICS.
----------------------------------	------

B Implementation Details

B.1 Baselines

We implement and evaluate various state-of-the-art feature engineering baselines, spanning traditional methods to recent LLM-based approaches, for comparison with LLM-FE. After generating features with each baseline, we apply a unified preprocessing pipeline to prepare the data for training and evaluation in the machine learning model. We implement the following baselines:

AutoFeat. AutoFeat is a classical feature engineering approach that uses iterative feature subsampling with beam search to select informative features. We utilize the open-source autofeat¹ package, retaining the default parameter settings. For parameter settings, we refer to the example '.ipynb' files provided in their official repository.

OpenFE. OpenFE is another state-of-the-art traditional feature engineering method using feature boosting and pruning algorithms. We employ the open-source openfe² package with standard parameter settings.

FeatLLM. FeatLLM uses an LLM to generate rules to binarize features that are then used as input to a simple model, such as linear regression. We adapt the open-source featllm³ implementation, modifying the pipeline to use an XGBoost model for inference. To ensure a fair comparison with other methods, we provide the entire training dataset to train the XGBoost model while using only a subset of the dataset (10 samples) to the LLM to generate binary features. We report the results through an ensemble over three samples to maintain consistency with LLM-FE.

CAAFE. We utilize the official implementation of CAAFE,⁴, maintaining all parameter settings as specified in the original repository. Following their workflow, we preprocess the data using their pipeline before inputting it into the prediction model after the feature engineering process.

OCTree. The official OCTree implementation⁵ was modified to keep the data loading and model initialization part common. We implemented OCTree only for classification datasets, as the official implementation is limited to classification datasets, and running for regression datasets on our own could have resulted in incorrect implementation.

B.2 LLM-FE

Feature Generation. Figure 6 presents an example prompt for the balance-scale dataset. The prompt begins with general instructions, followed by dataset-specific details, such as task descriptions, feature descriptions, and a subset of data instances serialized and expressed in natural language. To introduce diversity in prompting, we randomly sample between this approach and an alternative set of instructions, encouraging the LLM to explore a wider range of operators from OpenFE [45], as prior LLMs tend to favor simpler operators [25] (see Figure 7). The quality of features generated has been detailed in Appendix D.2. By providing this structured context, the model can leverage its domain knowledge to generate semantically and contextually meaningful hypotheses for new feature optimization programs.

Data-Driven Evaluation. After prompting the LLM, we sample b = 3 outputs. Based on preliminary experiments, we set the temperature for LLM output generation to t = 0.8 to balance creativity (exploration) and adherence to problem constraints, as well as reliance on prior knowledge (exploitation). The data modification process is illustrated in Figure 6(c), where the outputs are used to modify the features via modify_features(input). These modified features are then input into a prediction model, and the resulting validation score is calculated. To ensure efficiency, our evaluation is constrained by time and memory limits set at T = 30 seconds and M = 2GB, respectively.

¹https://github.com/cod3licious/autofeat.git

²https://github.com/IIIS-Li-Group/OpenFE.git

³https://github.com/Sungwon-Han/FeatLLM

⁴https://github.com/noahho/CAAFE

⁵https://github.com/jaehyun513/0CTree

Programs exceeding these limits are disqualified and assigned None scores, ensuring timely progress and resource efficiency in the search process.

Memory Management. Following the 'islands' model used by [7, 36, 35], we maintain the generated hypotheses along with their evaluation scores in a memory buffer comprising multiple islands (m = 3) that evolve independently. Each island is initialized with a basic feature transformation program specific to the dataset. Each island is initialized with a simple feature transformation program specific to the dataset (def modify_features_v0() in Figure 6(d)). In each iteration, novel hypotheses and their validation metrics are incorporated into their respective islands only if they exceed the island's current best score. Within each island, we additionally cluster feature discovery programs based on their signature, characterized by their validation score. Feature transformation programs that produce identical scores are consolidated together, creating distinct clusters. This clustering approach helps preserve diversity by ensuring that programs with varying performance characteristics remain in the population. We leverage this island model to construct prompts for the LLM. After an initial update of the prompt template with dataset-specific information, we integrate in-context demonstrations from the buffer. Following [36, 35], we randomly select one of the m available islands. Within the chosen island, we sample k = 2 programs to serve as in-context examples. To sample programs, we first select clusters based on their signatures using the Boltzmann selection strategy [8] to sample clusters based on their signatures with a preference for clusters with higher scores. Let s_i be the score of the i-th cluster, and the probability P_i for selecting the i-th cluster is given as:

$$P_i = \frac{exp(\frac{s_i}{\tau_c})}{\sum_i \left(\frac{s_i}{\tau_c}\right)}, \text{ where } \tau_c = T_0(1 - \frac{u \mod N}{N})$$
(4)

where τ_c is the temperature parameter, u is the current number of programs on the island, and $T_0 = 0.1$ and N = 10,000 are hyperparameters. Once a cluster is selected, we sample the programs from it.

C Additional Results

C.1 Transferability of Generated Features

While traditional approaches typically use the same model for both feature generation and inference, we demonstrate that the features generated by one model can be utilized by other models. Following [31], we use XGBoost, a computationally efficient decision tree-based model, to generate features to be used by more complex architectures for inference. As demonstrated in Table 6, XGBoost-generated features show an improvement in the performance of MLP and TabPFN over their base versions. This cross-architecture performance improvement suggests that the generated features capture meaningful data characteristics that are valuable across different modeling paradigms.

Table 6: Comparative analysis of LLM-FE using feature transfer. We use XGBoost to perform feature engineering and apply these features to MLP and TabPFN (indicated as LLM-FE _{XGB}). We report the accuracy for classification tasks and the normalized root mean square error for regression tasks. We report the mean and standard deviation across five random splits. **bold:** indicates the best performance.

Method	LLM	Classification \uparrow	Regression \downarrow				
MLP							
Base	-	$0.745\pm \scriptstyle 0.034$	$0.871 \pm \textbf{0.027}$				
LLM-FE _{XGB} LLM-FE	GPT-3.5-Turbo GPT-3.5-Turbo	$\begin{array}{c} 0.763 \pm 0.030 \\ \textbf{0.791} \pm \textbf{0.029} \end{array}$	$\begin{array}{c} 0.848 \pm 0.017 \\ \textbf{0.631} \pm \textbf{0.043} \end{array}$				
	TabPFN						
Base	-	$0.852 \pm \textbf{0.028}$	$0.289 \pm \textbf{0.016}$				
LLM-FE _{XGB} LLM-FE	GPT-3.5-Turbo GPT-3.5-Turbo	$\begin{array}{c} 0.861 \pm \textbf{0.017} \\ \textbf{0.863} \pm \textbf{0.018} \end{array}$	$\begin{array}{c} 0.287 \pm 0.015 \\ \textbf{0.286} \pm \textbf{0.015} \end{array}$				



Figure 6: Example of an input prompt for balance-scale dataset containing (a) instruction, (b) dataset specification containing the details about the task, features, and data samples, (c) evaluation function, (d) initial in-context demonstration, and (e) function to complete.



Figure 7: An example of the alternate set of instructions used to direct the model to use a complex set of operations over simple operators for generating features.

C.2 LLM-FE and Hyperparameter Optimization (HPO)

To assess hyperparameter optimization (HPO) effectiveness, we establish a comprehensive search space for our baseline models and employ the Optuna library [1] for tuning. Our methodology incorporates 400 trials with random sampling across multiple dataset splits. We concentrate our analysis on five classification datasets where baseline models achieve accuracies below 0.8. Table 7 demonstrates that HPO consistently enhances performance across all evaluated datasets for the Base model. Significantly, our proposed method LLM-FE delivers additional performance improvements even in the HPO-optimized setting. These findings confirm that while HPO provides meaningful performance gains, LLM-FE offers complementary and substantial enhancements, reinforcing its effectiveness as a feature engineering approach independent of hyperparameter tuning benefits.

Table 7: Comparison of XGBoost classification accuracy across d	latasets using Base and LLM-FE models,
evaluated under (a) without hyperparameter optimization (HPO) and	d (b) with HPO.

Dataset	w/o]	HPO	w/ HPO		
	Base	LLM-FE	Base	LLM-FE	
eucalyptus credit-g cmc blood-transfusion vehicle	$\begin{array}{c} 0.655 \pm 0.024 \\ 0.751 \pm 0.019 \\ 0.528 \pm 0.030 \\ 0.674 \pm 0.017 \\ 0.754 \pm 0.016 \end{array}$	$\begin{array}{c} \textbf{0.668} \pm \textbf{0.027} \\ \textbf{0.766} \pm \textbf{0.025} \\ \textbf{0.531} \pm \textbf{0.015} \\ \textbf{0.782} \pm \textbf{0.017} \\ \textbf{0.761} \pm \textbf{0.027} \end{array}$	$\begin{array}{c} 0.659 \pm 0.022 \\ 0.761 \pm 0.022 \\ 0.554 \pm 0.026 \\ 0.777 \pm 0.021 \\ 0.776 \pm 0.035 \end{array}$	$\begin{array}{c} 0.678 \pm 0.020 \\ 0.784 \pm 0.017 \\ 0.578 \pm 0.021 \\ 0.805 \pm 0.009 \\ 0.801 \pm 0.033 \end{array}$	

C.3 Additional Results

We extend the results from Section 5, showcasing the performance improvements achieved by LLM-FE across various prediction models. Specifically, we employ XGBoost, MLP, and TabPFN to generate features and subsequently use the same models for inference. As shown in Table 8, the features using GPT-3.5-Turbo by LLM-FE consistently enhance model performance across different datasets, outperforming the base versions trained without feature engineering. To further assess the generalizability of LLM-FE, we conducted experiments on smaller prediction models like CatBoost and Logistic Regression. From Table 9 that LLM-FE outperforms the respective base models for most of the datasets.

Table 8: Performance improvement with LLM-FE. We report the mean and standard deviation over five splits. We use Normalized Root Mean Square Error for all regression datasets, with a lower value indicating better performance, and Accuracy for classification datasets, with a higher value indicating better performance. **bold:** indicates the best performance.

Dataset	XGE	Boost	М	LP	TabPFN	
Butuber	Base	LLM-FE	Base	LLM-FE	Base	LLM-FE
Classification Datasets						
balance-scale	$0.856\pm \scriptscriptstyle 0.020$	0.990 ± 0.013	$0.933\pm$ 0.008	$0.997 \pm \scriptscriptstyle 0.004$	$0.970\pm$ 0.016	$1.000\pm$ 0.000
breast-w	$0.956\pm ext{0.012}$	$0.970\pm$ 0.009	$0.957\pm$ 0.010	0.964 ± 0.005	$0.971\pm$ 0.006	$0.971\pm$ 0.007
blood-transfusion	$0.742\pm ext{0.012}$	$0.751 \pm ext{0.036}$	$0.674 \pm \scriptscriptstyle 0.071$	$0.782 \pm$ 0.017	$0.790\pm$ 0.012	$0.791 \pm {\scriptstyle 0.011}$
car	$0.995\pm ext{0.003}$	0.999 ± 0.001	$0.929\pm$ 0.019	$0.950\pm$ 0.009	$0.984 \pm \scriptscriptstyle 0.007$	0.996 ± 0.006
cmc	$0.528\pm ext{0.030}$	$0.531 \pm ext{0.015}$	$0.559 \pm \scriptscriptstyle 0.028$	0.566 ± 0.028	$0.563\pm$ 0.030	0.566 ± 0.036
credit-g	$0.751\pm$ 0.019	0.766 ± 0.025	$0.558\pm {\scriptscriptstyle 0.144}$	$0.633 \pm {\scriptstyle 0.101}$	0.728 ± 0.008	$0.794 \pm \scriptscriptstyle 0.022$
eucalyptus	0.655 ± 0.024	0.668 ± 0.027	$0.414 \pm \scriptscriptstyle 0.064$	0.456 ± 0.062	$0.712\pm ext{0.016}$	$0.715 \pm \scriptscriptstyle 0.021$
heart	$0.858\pm$ 0.013	$0.866 \pm \scriptscriptstyle 0.021$	$0.840\pm$ 0.010	$0.844 \pm$ 0.006	$0.882 \pm \scriptscriptstyle 0.025$	$0.880\pm$ 0.021
pc1	$0.931\pm \scriptscriptstyle 0.004$	$0.935\pm$ 0.006	$0.931 \pm \scriptscriptstyle 0.002$	$0.904\pm ext{0.055}$	$0.936\pm$ 0.007	0.937 ± 0.003
tic-tac-toe	0.998 ± 0.004	0.998 ± 0.005	$0.816\pm ext{0.029}$	$0.854 \pm \scriptscriptstyle 0.052$	0.984 ± 0.005	0.986 ± 0.009
vehicle	$0.754\pm$ 0.016	$0.761 \pm ext{0.027}$	$0.583 \pm \scriptscriptstyle 0.062$	$0.673 \pm {\scriptstyle 0.043}$	$0.852\pm ext{0.016}$	$0.856 \pm ext{0.028}$
		Regr	ession Datasets			
airfoil_self_noise	$0.013 \pm \scriptscriptstyle 0.001$	$0.011 \pm ext{0.001}$	$0.275\pm ext{0.008}$	$0.108 \pm ext{0.001}$	$0.008\pm ext{0.001}$	$0.007 \pm \scriptscriptstyle 0.001$
bike	0.216 ± 0.005	0.207 ± 0.005	$0.636\pm$ 0.015	$0.551 \pm \scriptscriptstyle 0.022$	0.200 ± 0.005	0.199 ± 0.006
cpu_small	$0.034\pm ext{0.003}$	0.033 ± 0.003	3.793 ± 0.731	$\textbf{2.360} \pm \textbf{1.263}$	$0.036\pm ext{0.001}$	$0.035\pm$ 0.001
crab	$0.234\pm ext{0.009}$	$0.223 \pm \scriptscriptstyle 0.014$	$0.214\pm$ 0.010	$0.212 \pm \scriptscriptstyle 0.011$	$0.208\pm$ 0.013	$0.207 \pm \scriptscriptstyle 0.014$
diamond	$0.139\pm \scriptscriptstyle 0.002$	$0.134 \pm \scriptscriptstyle 0.002$	$0.296\pm ext{0.018}$	$0.265 \pm \scriptscriptstyle 0.011$	0.132 ± 0.005	0.130 ± 0.005
forest-fires	1.469 ± 0.080	1.417 ± 0.083	$1.423\pm ext{0.104}$	1.344 ± 0.091	$1.270\pm ext{0.101}$	1.269 ± 0.114
housing	$0.234\pm ext{0.009}$	$0.218 \pm ext{0.009}$	$0.505\pm ext{0.009}$	0.444 ± 0.036	$0.210\pm {\scriptstyle 0.004}$	0.202 ± 0.003
insurance	$0.397\pm {\scriptscriptstyle 0.144}$	$0.381 \pm \scriptscriptstyle 0.142$	$0.896\pm ext{0.053}$	$0.487 \pm \scriptscriptstyle 0.026$	$0.351\pm$ 0.018	0.346 ± 0.020
plasma_retinol	$0.390\pm \scriptscriptstyle 0.032$	0.388 ± 0.033	$0.440\pm$ 0.070	0.411 ± 0.053	$0.348 \pm ext{0.048}$	0.348 ± 0.055
wine	0.110 ± 0.001	$0.105\pm ext{0.001}$	$0.125\pm ext{0.001}$	$0.125\pm$ 0.013	$0.117\pm \scriptscriptstyle 0.004$	$0.116 \pm ext{0.004}$

Table 9: Performance improvement with LLM-FE on CatBoost and Logistic Regression. We report the mean and standard deviation over five splits. We use Accuracy for classification datasets, with a higher value indicating better performance. **bold:** indicates the best performance.

Dataset	Logistic F	Regression	CatBoost		
	Base	LLM-FE	Base	LLM-FE	
balance-scale breast-w blood-transfusion car cmc credit-g eucalyptus heart pc1 tic-tac-toe vehicle	$\begin{array}{c} 0.874 \pm 0.024 \\ 0.955 \pm 0.014 \\ 0.799 \pm 0.014 \\ 0.690 \pm 0.017 \\ 0.520 \pm 0.019 \\ 0.764 \pm 0.006 \\ 0.671 \pm 0.036 \\ 0.877 \pm 0.021 \\ 0.931 \pm 0.003 \\ 0.997 \pm 0.004 \\ 0.772 \pm 0.028 \end{array}$	$\begin{array}{c} \textbf{0.994} \pm \textbf{0.012} \\ \textbf{0.962} \pm \textbf{0.008} \\ \textbf{0.799} \pm \textbf{0.009} \\ \textbf{0.696} \pm \textbf{0.031} \\ \textbf{0.525} \pm \textbf{0.012} \\ \textbf{0.780} \pm \textbf{0.015} \\ \textbf{0.667} \pm \textbf{0.042} \\ \textbf{0.872} \pm \textbf{0.025} \\ \textbf{0.935} \pm \textbf{0.003} \\ \textbf{0.996} \pm \textbf{0.005} \\ \textbf{0.769} \pm \textbf{0.015} \end{array}$	$\begin{array}{c} 0.870 \pm 0.020 \\ 0.957 \pm 0.009 \\ 0.742 \pm 0.012 \\ 0.999 \pm 0.001 \\ 0.518 \pm 0.028 \\ 0.714 \pm 0.046 \\ 0.436 \pm 0.027 \\ 0.845 \pm 0.015 \\ 0.929 \pm 0.005 \\ 0.673 \pm 0.027 \\ 0.719 \pm 0.045 \end{array}$	$\begin{array}{c} \textbf{1.000} \pm \textbf{0.000} \\ \textbf{0.962} \pm \textbf{0.008} \\ \textbf{0.751} \pm \textbf{0.036} \\ \textbf{0.999} \pm \textbf{0.001} \\ \textbf{0.548} \pm \textbf{0.027} \\ \textbf{0.700} \pm \textbf{0.021} \\ \textbf{0.509} \pm \textbf{0.050} \\ \textbf{0.839} \pm \textbf{0.018} \\ \textbf{0.932} \pm \textbf{0.012} \\ \textbf{0.747} \pm \textbf{0.152} \\ \textbf{0.725} \pm \textbf{0.033} \end{array}$	

C.4 Robustness to Noise

Noise is an inherent challenge in real-world data, arising from various sources, including sensor errors, human mistakes, environmental factors, and equipment limitations. Such noise can mask underlying patterns and impair machine learning models' ability to learn true relationships in the data. To evaluate how effectively LLM-FE leverages prior knowledge and evolutionary search to handle noisy data, we introduced Gaussian noise ($\sigma = 0, 0.01, 0.05, 0.1$) into numerical classification datasets. As shown in Figure 8, we compared XGBoost's performance across different feature engineering approaches, using GPT-3.5-Turbo as the LLM backbone for both the LLM-based approaches. The results demonstrate that LLM-FE maintains superior accuracy and robustness even under increasing noise conditions.



Figure 8: Impact of Noise Levels on XGBoost model performance across different feature engineering approaches, under increasing noise conditions ($\sigma = 0.0$ to 0.1). We report the mean accuracy across six classification datasets containing only numerical features.

D Qualitative Analysis

D.1 Computational Efficiency

Automated feature engineering methods, both classical and LLM-based, universally employ model training and validation to evaluate feature relevance. This evaluation strategy represents standard methodology across all automated feature engineering approaches rather than an additional computational burden specific to LLM-FE. We conduct our efficiency-performance trade-off analysis on the datasets with higher sample counts from Section 5, as these datasets present greater complexity with their substantial number of samples and features. Our comparative Pareto analysis (Figure 9) presents the base model alongside various feature engineering Our proposed method, LLM-FE, baselines. demonstrates Pareto optimality by achieving superior performance with substantially reduced computational requirements compared to existing methods, which either exhibit longer execution



Figure 9: Pareto Plot: comparing trade-off between performance (accuracy) vs time (in seconds) for LLM-FE and other feature engineering baselines.

times or yield inferior performance metrics on these datasets. Only the base method requires less computation time than LLM-FE, but at a significant performance cost. This positions LLM-FE as the optimal solution in the efficiency-performance space, delivering state-of-the-art results with reasonable computational demands even when handling datasets of considerable complexity.

D.2 Complexity of Generated Features

Despite a tendency of LLMs to prefer simpler operators [25] (e.g., addition, subtraction, absolute value), resulting in a less diverse feature set compared to specialized feature-engineering tools, LLM-FE demonstrates the valuable capability of identifying certain complex and informative transformations that are rarely generated by conventional LLM-based automated methods. Specifically, complex operators like groupbythenmean, groupbythenmin, groupbythenmax, residual, and sigmoid are also recommended frequently by LLM-FE, as illustrated in Figure 10. Such complex operations have the potential to capture meaningful patterns involving group-based aggregation that simpler transformations may miss. Thus, while further refinement is needed to balance operator selection, the ability of LLM-FE to discover nuanced, aggregation-based features emphasizes its promising role as a complementary technique in the broader automated feature engineering toolkit.



Figure 10: Frequency of Feature Engineering Operators. We compare the operators for LLM-FE with simple LLM-based methods.

D.3 Impact of Domain Knowledge



Figure 11: Quantitative and Qualitative Analysis on Impact of Domain Knowledge for LLM-FE on Heart and Breast-W datasets. (a) Comparison of XGBoost performance for LLM-FE against its domain-agnostic variant and traditional methods, such as OpenFE and AutoFeat, which do not integrate domain knowledge and exhibit reduced performance. (b) Features generated using the *w/o* Domain Knowledge variant of LLM-FE. (c) Feature discovery program generated by LLM-FE. The generated programs emphasize how incorporating domain expertise leads to more interpretable features that improve model performance.

Figure 11 highlights the qualitative and quantitative benefits of domain-specific feature transforms. We demonstrate this using two datasets: the Breast-W dataset, which focuses on distinguishing between benign and malignant tumors, and the Heart dataset, which predicts cardiovascular disease risk based on patient attributes. These tasks underscore the crucial role of domain knowledge in identifying meaningful features. Using embedded domain knowledge, LLM-FE not only significantly improves accuracy but also provides the reasoning for choosing the given feature, leading to more interpretable feature engineering. For example, in the Heart dataset, LLM-FE suggests the feature 'Log_Cholesterol', recognizing cholesterol's critical role in heart health and applying a logarithmic transformation to reduce the impact of outliers and stabilize the variance. In contrast, the 'w/o Domain Knowledge' variant arbitrarily combines existing features, leading to uninterpretable transformations and reduced overall performance (Figure 11(a)). Similarly, for breast cancer prediction, LLM-FE identifies 'proliferation_activity' a biologically relevant metric leading

to performance improvement, whereas the absence of domain knowledge results in a simple mean of all features, lacking interpretability and clinical significance (Figures 11(b) and 11(c)).

D.4 Impact of Evolutionary Refinement

Figure 12 shows the detailed performance trajectory of LLM-FE compared with its 'w/o Evolutionary Refinement' variant on PC1 and Balance-Scale datasets. The graph demonstrates that LLM-FE, using evolutionary search, consistently improves validation accuracy, while the non-refinement variant stagnates due to local optima. On the PC1 dataset, the non-refinement variant plateaus after seven iterations, and on the Balance-Scale dataset, it stagnates after five iterations. LLM-FE's evolutionary refinement helps it escape local optima with more robust optimization, leading to better validation accuracy on both datasets.



Figure 12: Performance Trajectory Analysis. Validation Accuracy progression for LLM-FE *w/o* evolutionary refinement and LLM-FE. LLM-FE demonstrates better validation accuracy, highlighting the advantage of evolutionary iterative refinement.