

Evolve Cost-aware Acquisition Functions Using Large Language Models

Yiming Yao^{1,2}, Fei Liu², Ji Cheng², and Qingfu Zhang²

¹ City University of Hong Kong (Dongguan), Dongguan523000, China

² Department of Computer Science, City University of Hong Kong, Hong Kong
 {yimingyao3-c,flui36-c,J.Cheng}@my.cityu.edu.hk,
 qingfu.zhang@cityu.edu.hk

Abstract. Many real-world optimization scenarios involve expensive evaluation with unknown and heterogeneous costs. Cost-aware Bayesian optimization stands out as a prominent solution in addressing these challenges. To approach the global optimum within a limited budget in a cost-efficient manner, the design of cost-aware acquisition functions (AFs) becomes a crucial step. However, traditional manual design paradigm typically requires extensive domain knowledge and involves a labor-intensive trial-and-error process. This paper introduces EvolCAF, a novel framework that integrates large language models (LLMs) with evolutionary computation (EC) to automatically design cost-aware AFs. Leveraging the crossover and mutation in the algorithm space, EvolCAF offers a novel design paradigm, significantly reduces the reliance on domain expertise and model training. The designed cost-aware AF maximizes the utilization of available information from historical data, surrogate models and budget details. It introduces novel ideas not previously explored in the existing literature on acquisition function design, allowing for clear interpretations to provide insights into its behavior and decision-making process. In comparison to the well-known EIpu and EI-cool methods designed by human experts, our approach showcases remarkable efficiency and generalization across various tasks, including 12 synthetic problems and 3 real-world hyperparameter tuning test sets.

Keywords: Cost-aware Bayesian optimization · Acquisition functions · Large language models · Evolutionary computation.

1 Introduction

Bayesian optimization is a powerful tool for solving expensive optimization problems and has found wide application in many real-world scenarios [28,8,23,7]. It typically employs a surrogate model to approximate the expensive function and well-designed acquisition functions (AFs) to select potential solutions in a sample-efficient manner. Popular acquisition functions include probability of improvement (PI) [14], expected improvement (EI) [21], upper confidence bound (UCB) [27], knowledge gradient (KG) [6], etc.

Vanilla BO typically sets the number of evaluations as the budget constraint, implicitly assuming a uniform evaluation cost in the design space [16]. However, this is rarely the case in many real-world applications, leading to the concept of cost-aware BO. For example, in hyperparameter optimization (HPO) tasks for machine learning models, the costs with different hyperparameter configurations may even differ in the order of magnitudes [16]. Under the budget constraint of accumulated costs, approaching the global optimum is a challenge for traditional AFs due to their unawareness of heterogeneous evaluation costs, which highlights the importance of designing efficient cost-aware AFs.

Previous works have proposed several heuristics to take the cost information into account [26,16,15]. Representative ones include EI per unit cost (EIpu) [26] which divides EI by the cost function to promote solutions with both low cost and high improvement, and EI-cool [16] that introduces the cost-cooling strategy to make EIpu adapt to problems with expensive global optimum. Based on EIpu and EI-cool, several enhanced approaches have been suggested [19,24,9]. However, designing these AFs typically necessitates significant involvement of domain experts and extensive trial-and-error testing to refine and improve upon previous methods. This manual design paradigm is labor-intensive and cannot be automated. Furthermore, the information currently used to define cost-aware AFs is inadequate, as it only considers the EI metric and budget details while overlooking the complete historical data, which severely restricts the exploration in the algorithm space. Simply integrating the EI metric with budget information does not inherently stimulate innovative ideas, thereby greatly limiting the performance and generalization of the designed AFs. While some model-based methods have been proposed to automatically learn AFs parameterized by neural networks in meta-BO community [30,11,20], they often require substantial effort in complex framework design and model training. Additionally, the resulting AFs are represented by network parameters, leading to poor interpretability compared to widely used AFs that have explicit mathematical expressions. Besides, these methods are designed for problems with uniform costs and are not applicable to many real-world applications with heterogeneous costs.

In the past three years, large language models (LLMs) have been widely used in code generation, mathematical reasoning, and automatic algorithm design [25,17]. While recent studies have explored the use of LLMs to enhance vanilla BO [32,18], these approaches rely on querying LLMs to directly suggest candidate solutions. The absence of a clearly defined search strategy results in inadequate explainability. Moreover, whenever a new problem arises, it needs to conduct a substantial number of queries for LLMs from scratch, which can be expensive and impractical for real-world applications.

We propose a novel paradigm, named EvolCAF, which integrates LLMs in an evolutionary framework to automatically design explicit AFs to enhance cost-aware BO. Different from the existing works, it enjoys good automation and explainability outperforming existing human-crafted methods. Our main contributions are as follows:

- We introduce EvolCAF, which integrates large language models (LLMs) with evolutionary computation (EC) to automatically design cost-aware AFs. It enables crossover and mutation in the algorithm space to iteratively search for elite AFs, significantly reducing the reliance on expert knowledge and domain model training.
- We leverage EvolCAF to design a cost-aware AF that fully utilizes the available history information. Remarkably, the designed AF introduces novel ideas that have not been explored in existing literature on acquisition function design. The designed AF can be expressed explicitly, allowing for clear interpretations to provide insights into its behavior and decision-making process.
- We evaluate the designed AF on diverse synthetic functions as well as practical hyperparameter optimization (HPO) problems. Compared to the popular EIpu and EI-cool methods designed by domain experts, our approach demonstrates remarkable efficiency and generalization, which highlights the promising potential in addressing many related real-world applications.

2 Background and Related Works

2.1 Background

In vanilla Bayesian optimization (BO), we consider finding the optimal solution $\mathbf{x}^*$ that maximizes the black-box objective function f : $\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$, where $\mathcal{X}$ is a compact subset of $\mathbb{R}^d$, we assume $f : \mathcal{X} \rightarrow \mathbb{R}$ is continuously differentiable and expensive to evaluate.

Gaussian Processes To approximate the expensive objective function, BO typically employs a Gaussian process (GP) model [1] as the surrogate. GP is an infinite distribution over functions f specified by a prior mean function $\mu(\cdot)$ and covariance function $k(\cdot, \cdot)$: $f(\mathbf{x}) \sim \mathcal{GP}_f(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$. Suppose in iteration t , the historical data set $\mathcal{D}_t = \{(\mathbf{x}_i, y_i)\}_{i=1}^t$ are obtained from the observation model $y_i = f(\mathbf{x}_i) + \epsilon_i$ with observation noise $\epsilon_i \sim \mathcal{N}(0, \sigma_\epsilon^2)$. Given the test point $\mathbf{x}^*$, the predictive distribution $p(y|\mathbf{x}^*, \mathcal{D}_t)$ is also Gaussian with mean $\mu(\mathbf{x}^*) = K_{\mathbf{x}^*, \mathbf{X}}(K_{\mathbf{X}, \mathbf{X}} + \sigma_\epsilon^2 \mathbf{I})^{-1} \mathbf{y}$ and variance $\sigma^2(\mathbf{x}^*) = k(\mathbf{x}^*, \mathbf{x}^*) - K_{\mathbf{x}^*, \mathbf{X}}(K_{\mathbf{X}, \mathbf{X}} + \sigma_\epsilon^2 \mathbf{I})^{-1} K_{\mathbf{X}, \mathbf{x}^*}$, where $\mathbf{y} = [y_1, y_2, \dots, y_t]^T$ are noisy output values observed from the latent functions $\mathbf{f} = [f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_t)]^T$ at training points $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t]^T$, $K_{\mathbf{X}, \mathbf{X}} = [k(\mathbf{x}_i, \mathbf{x}_j)]_{\mathbf{x}_i, \mathbf{x}_j \in \mathbf{X}}$ is the covariance matrix and $K_{\mathbf{X}, \mathbf{x}^*} = [k(\mathbf{x}_i, \mathbf{x}^*)]_{\mathbf{x}_i \in \mathbf{X}}$ is the correlation vector for all training and test points.

Acquisition Functions The acquisition function (AF) defines a utility that measures the benefit of evaluating an unknown point $\mathbf{x}$. We denote the definition of AF as $\alpha(\mathbf{x})$, which may contain the historical data set $\mathcal{D}_t$, model information

$\mu(\mathbf{x})$, $\sigma^2(\mathbf{x})$, etc. One of the popular AFs is expected improvement (EI) [21], which quantifies the expected amount of improvement over the current best observation $y^* = \max_i y_i$ at a given point $\mathbf{x}$ in the search space:

$$\alpha_{\text{EI}}(\mathbf{x}) = \mathbb{E}_{f(\mathbf{x})}[[f(\mathbf{x}) - y^*]_+] = \sigma(\mathbf{x}) h\left(\frac{\mu(\mathbf{x}) - y^*}{\sigma(\mathbf{x})}\right), \quad (1)$$

where $[\cdot]_+$ is the $\max(0, \cdot)$ operation, $h(z) = \phi(z) + z\Phi(z)$, ϕ and Φ are the standard normal density and distribution functions, respectively.

2.2 Cost-aware Bayesian Optimization

Problem Setting Cost-aware Bayesian optimization assumes evaluating the objective function is expensive, and the evaluation in different regions will incur heterogeneous costs, we denote the unknown cost function as $c(\mathbf{x})$. For every query $\mathbf{x}_i$, we can obtain the noisy observation y_i with cost $z_i = c(\mathbf{x}_i) + \eta_i$, where $\eta_i \sim \mathcal{N}(0, \sigma_\eta^2)$. Similar to the objective function f , the black-box cost function is modeled as a draw from the Gaussian process $c(\mathbf{x}) \sim \mathcal{GP}_c$. We use the posterior predictive mean of $\mathcal{GP}_c$ for calculating the cost function as [26] and [19] do.

Given the historical data set $\bar{\mathcal{D}}_t = \{(\mathbf{x}_i, y_i, z_i)\}_{i=1}^t$ with t evaluated samples and the limited total budget B_{total} , we can only find a near-optimal solution with the constraint of cumulative cost $\sum_{i=1}^T z_i \leq B_{\text{total}}$, where T is the maximum number of evaluated samples satisfies the budget constraint. The general framework followed by the vast majority of existing cost-aware BO methods is shown in Algorithm 1, the main difference lies in the different definitions of cost-aware AFs, which will be introduced below.

Algorithm 1 Cost-aware BO

Input:

B_{total} : total budget, $\bar{\mathcal{D}}_t$: initial data set with t evaluated samples

- 1: Initialize used budget $B_{\text{used}} = \sum_{i=1}^t z_i$
- 2: Train objective and cost models $\mathcal{GP}_f$ and $\mathcal{GP}_c$ using $\bar{\mathcal{D}}_t$
- 3: **while** $B_{\text{used}} < B_{\text{total}}$ **do**
- 4: Query candidate: $\mathbf{x}_{t+1} = \arg \max_{\mathbf{x}} \alpha(\mathbf{x})$
- 5: Evaluate candidate: $y_{t+1}, z_{t+1} \leftarrow f(\mathbf{x}_{t+1}), c(\mathbf{x}_{t+1})$
- 6: Update data set $\bar{\mathcal{D}}_{t+1} \leftarrow \bar{\mathcal{D}}_t \cup \{(\mathbf{x}_{t+1}, y_{t+1}, z_{t+1})\}$
- 7: Update models $\mathcal{GP}_f$ and $\mathcal{GP}_c$ using $\bar{\mathcal{D}}_{t+1}$
- 8: $B_{\text{used}} \leftarrow B_{\text{used}} + z_{t+1}$
- 9: $t \leftarrow t + 1$

10: **end while**

- 11: Total number of evaluated samples $T \leftarrow t$

Output: Best configuration $\arg \max_{(\mathbf{x}_i, y_i) \in \bar{\mathcal{D}}_T} y_i$

EI per Unit Cost (EIpu) To balance the cost and quality of evaluations, Snoek et al. [26] proposed EI per unit cost (EIpu), which normalizes the EI

metric by the cost function $c(\mathbf{x})$:

$$\alpha_{\text{EIpu}}(\mathbf{x}) = \frac{\alpha_{\text{EI}}(\mathbf{x})}{c(\mathbf{x})}. \quad (2)$$

By using the cost function to penalize EI, EIpu tends to carefully select candidate points with low cost and high improvement, making the search process cost-aware. Benefiting from the preference for cheaper regions, EIpu can be consistently improved when the optimum is cheap to evaluate since higher EI and lower cost are both encouraged. However, the preference becomes a drawback when the optimum lies in the expensive regions of the design space, which is common in many real-world applications. The cost penalty term prevents EIpu from exploring near-optimal regions that are expensive to evaluate, experiments have shown that sometimes EIpu performs even worse than EI [16].

EI-cool To alleviate the above problem, Lee et al. [16] introduced a cost-cooling factor α in EIpu called EI-cool:

$$\alpha_{\text{EI-cool}}(\mathbf{x}) = \frac{\alpha_{\text{EI}}(\mathbf{x})}{c(\mathbf{x})^\alpha}, \quad (3)$$

where $\alpha = (B_{\text{total}} - B_{\text{used}}) / (B_{\text{total}} - B_{\text{init}})$, B_{init} is the budget spent in evaluating the initial sample points, B_{used} is the budget already used and B_{total} is the given total budget. As B_{used} increases from B_{init} to B_{total} during the search process, the factor α gradually decays from 1 to 0, resulting in the transition of EI-cool from EIpu to EI. Intuitively, the cost-cooling strategy diminishes the significance of the cost model as the budget is consumed, making EI-cool to operate in an *early and cheap, late and expensive* fashion to encourage exploring expensive regions when the remaining budget is tight.

Although EI-cool alleviates the problem of performance degradation when searching for the expensive optimum, it always uses EIpu as the starting strategy, which is not flexible and may not adapt well to different problems [19]. Besides, previous analysis and experiments have shown that when the remaining budget is gradually tight, although deemphasizing the cost can increase the likelihood of exploring expensive regions, it is still possible to miss the optimum in high-cost regions before the budget is exhausted [19], as the exploration or exploitation in very cheap regions can still result in very large values of EI-cool metric, which is called low-cost-preference weakness in [24].

Variants Based on EIpu and EI-cool Based on EIpu and EI-cool, some improved methods have been proposed such as using a multi-armed bandit algorithm to automatically select either EI or EIpu [19], developing more aggressive methods to alleviate the low-cost-preference weakness of EI-cool[24], and cost-aware EI based on Pareto optimality to achieve the trade-off between cost and improvement [9]. It is evident that the design process typically requires significant involvement of domain knowledge and extensive trial-and-error testing

based on the flaws of previous methods, which are labor-intensive and cannot be automated. Besides, the existing AFs just combine the EI metric with budget information in different ways, which severely restricts the exploration in the algorithm space and does not inherently foster the generation of innovative ideas, so the performance and generalization are greatly limited.

2.3 Automatic Design for Acquisition Functions

In the meta-BO community, which focuses on meta-learning or learning to learn to enhance vanilla BO [4,29,5,1], there has been research dedicated to automatically generating efficient and generalizable AFs via a learning model. While these works are not tailored for cost-aware contexts, we will review them to emphasize the strengths and potential of our framework.

To learn a meta-acquisition function, Volpp et al. [30] replaced the hand-designed AF with a neural network named neural acquisition function (NAF), which is meta-trained on related source tasks by policy-based reinforcement learning. Hsieh et al. [11] utilized a deep Q-network (DQN) as a surrogate differentiable AF to achieve a few-shot fast adaptation of AFs. Maraval et al. [20] introduced an end-to-end differentiable framework based on transformer architectures called neural acquisition process (NAP) to meta-learn acquisition functions with the surrogate model jointly. Nevertheless, despite achieving promising results, these model-based methods often demand substantial effort in framework design and model training. Moreover, the resulting AFs are represented by network parameters, resulting in poor interpretability compared to widely used AFs that have explicit mathematical expressions.

3 EvolCAF: Evolve Cost-aware Acquisition Functions with LLMs

3.1 Framework

The proposed EvolCAF framework embraces the basic components of evolutionary computing (EC), including initialization, crossover, mutation, and population management. In EvolCAF, each individual represents an acquisition function solving a branch of synthetic instances, which is represented with an algorithm description and a code block implementation instead of an encoded vector in traditional EC. During the evolution process, the initialization, crossover, and mutation operations on the individuals are all performed by prompting LLM in the algorithm space. The entire process is completely automated without any intervention from human experts.

Fig. 1 illustrates the detailed flowchart of EvolCAF. At each generation, we maintain a population of N AFs, each AF is evaluated on a set of synthetic instances in a cost-aware BO loop to calculate the fitness value, which is the optimal gap between the true optimal value and the optimal value obtained by the AF. After new individuals are added to the population, the worst individuals are deleted according to the fitness values.




Fig. 1. Flowchart of EvolCAF framework. The left box presents the evolution of cost-aware AFs enabled by EvolCAF, wherein each individual in the population is an AF, represented with an algorithm description and a code block implementation. The initialization, crossover, and mutation are facilitated by LLMs. The middle box shows the cost-aware BO loop, each AF is evaluated on a set of synthetic instances to calculate the optimal gap as its fitness value.

3.2 General Definition for Evolved AFs

The general definition for evolved cost-aware AFs can be formulated as follows:

$$\begin{aligned}\alpha_{\text{EvolCAF}}(\mathbf{x}) &= \alpha(\mathbf{x}; \boldsymbol{\theta}_{\text{data}}, \boldsymbol{\theta}_{\text{model}}, \boldsymbol{\theta}_{\text{budget}}) \\ &= \alpha(\mathbf{x}; \mathbf{X}, \mathbf{y}, \mathbf{x}^*, y^*, \mu(\mathbf{x}), \sigma(\mathbf{x}), c(\mathbf{x}), B_{\text{used}}, B_{\text{total}}).\end{aligned}\quad (4)$$

The inputs of cost-aware AFs incorporate three groups of information: (1) historical data $\boldsymbol{\theta}_{\text{data}} = \{\mathbf{X}, \mathbf{y}, \mathbf{x}^*, y^*\}$, (2) prediction and uncertainty provided by the model $\boldsymbol{\theta}_{\text{model}} = \{\mu(\mathbf{x}), \sigma(\mathbf{x})\}$, and (3) budget information during the optimization $\boldsymbol{\theta}_{\text{budget}} = \{c(\mathbf{x}), B_{\text{used}}, B_{\text{total}}\}$. The first two groups include the data and model information for searching with uncertainties, while the last group informs the AF of the budget constraints in optimization. During the evolutionary process, EvolCAF is encouraged to explore the algorithm space to generate and refine elite cost-aware AFs.

3.3 Prompt Engineering

The general format of prompt engineering used to inform LLMs consists of four parts: (1) a general description of the task, (2) code instructions for implementing algorithms, including the function name, inputs and output, (3) interpretations for the inputs and output, including their detailed meanings in our task, the variable formats and dimensions implemented in the code, (4) helpful hints to inform LLMs to generate executable codes and utilize input information as much

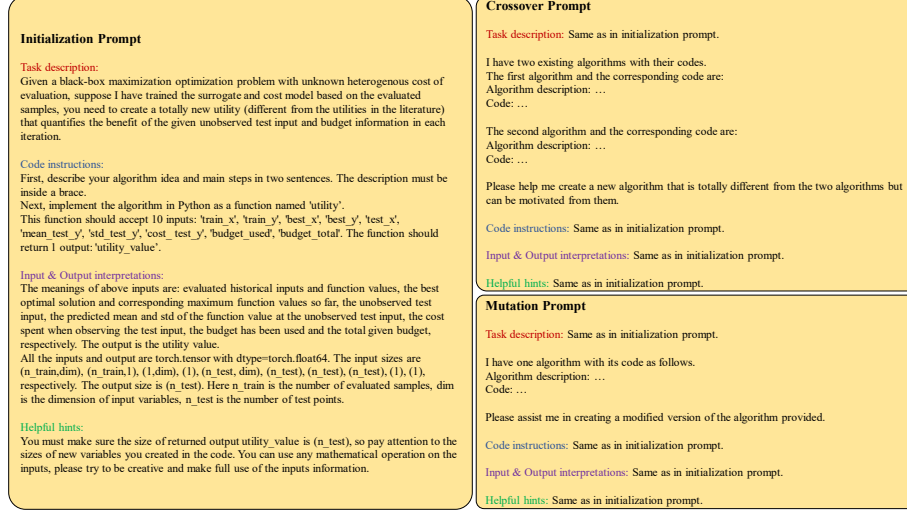


Fig. 2. Prompts used in EvolCAF for initialization, crossover, and mutation.

as possible to create novel ideas. Following the general format, in initialization, we instruct LLMs to create a completely new AF to promote population diversity. In crossover, we suggest combining the selected parent AFs to facilitate the preservation of high-performing components in the following generations. While in mutation, we aim to encourage the exploration of better AFs based on parent AFs in the algorithm space. The details of prompts for initialization, crossover, and mutation are shown in Fig. 2.

4 Experimental Studies

4.1 Experimental Settings

Settings for AF Evolution In the evolutionary process, EvolCAF maintains 10 AFs and evolves over 20 generations. We generate 1 offspring individual in each generation based on 2 parent individuals, with the crossover probability set to 1.0 and mutation probability set to 0.5. The GPT-3.5-turbo pre-trained LLM is used for generating AFs.

To generate AFs that can be efficiently optimized, we set a time threshold of 60 seconds for completing the cost-aware BO loop, which serves as a selection pressure together with the fitness value. Any AF that exceeds this time limit will be automatically eliminated during the evolutionary process.

Settings for Cost-aware BO To calculate the fitness value, we evaluate each evolved AF on 2D Ackley and 2D Rastigin functions with 10 different random seeds in the experimental design, resulting in a total of 20 instances. The aim is




Fig. 3. The evolutionary process of acquisition functions.

to achieve improved generalization results across various initial surrogate landscapes with an acceptable evaluation time during evolution, as the training and inference of GP models on a large number of instances in each generation can be expensive. The fitness value for each evolved AF is calculated by averaging the optimal gaps obtained from the 20 instances.

To simulate the scenarios in many real-world applications, we carefully design a cost function that is most expensive to evaluate at the global optimum $\mathbf{x}^*$ of the synthetic function, the formulation is similar to that used in [15] and can be expressed as:

$$c(\mathbf{x}) = \exp(-\|\mathbf{x} - \mathbf{x}^*\|_2), \quad (5)$$

where each dimension of $\mathbf{x}$ and $\mathbf{x}^*$ is normalized to $[0,1]$. In order to achieve good results for the evolved AF given a small budget, we set the total budget B_{total} as 30 in the evolutionary process, indicating that the smallest number of evaluations is 30, we will further verify the generalization using sufficient budget in the following experiments. We initialize $2d$ random samples using experimental design, where d is the dimension of the decision variable.

All BO methods are implemented using BoTorch [2], the acquisition functions are optimized through multi-start optimization using scipy’s L-BFGS-B optimizer, using 20 restarts seeded from 100 pseudo-random samples through BoTorch’s initialization heuristic for efficient optimization.

4.2 Evolution Results

Fig. 3 demonstrates the evolutionary process. In each generation, we maintain 10 AFs represented by blue dots. The mean and optimal fitness values of the population are represented with orange and red lines, respectively. With a population size of 10 and 20 generations, the fitness value of the evolving AFs can converge to a notably low level. The results show the capability of our framework to automatically generate and evolve elite AFs.

Fig. 4 shows the optimal AF with the minimum fitness value, including a general description of the algorithmic idea in defining the AF and a detailed code implementation.

After converting the code into easily understandable mathematical expressions, we observe that the optimal AF consists of three parts:

$$\alpha_{\text{EvolCAF}}(\mathbf{x}) = \alpha_1(\mathbf{x}) + \alpha_2(\mathbf{x}) + \alpha_3(\mathbf{x}). \quad (6)$$

Specifically, $\alpha_1(\mathbf{x})$ combines a modified EI with uncertainty information:

$$\begin{aligned} \alpha_1(\mathbf{x}) = & [(\mu(\mathbf{x}) - y^*)\Phi(\mathbf{z}) + \sqrt{\sigma^2(\mathbf{x}) + \sigma^2(\mathbf{y})} \cdot \phi(\mathbf{z})] \cdot \\ & (1 - \log \sqrt{\frac{\sigma^2(\mathbf{x}) + \sigma^2(\mathbf{y})}{\sigma^2(\mathbf{y})}}), \end{aligned} \quad (7)$$

where $\mathbf{z} = \frac{\mu(\mathbf{x}) - y^*}{\sqrt{\sigma^2(\mathbf{x}) + \sigma^2(\mathbf{y})}}$, ϕ and Φ are the standard normal density and distribution functions, respectively, $\sigma^2(\mathbf{y})$ represents the variance of the current historical observations.

Similar to EI, $\alpha_1(\mathbf{x})$ encourages searching regions close to the current best observation with uncertainties. However, an improvement is that $\alpha_1(\mathbf{x})$ incorporates the information of all historical observations rather than solely focusing on the current best one.

$\alpha_2(\mathbf{x})$ mainly focuses on the current remaining budget and the cost of evaluating the unknown point $\mathbf{x}$:

$$\alpha_2(\mathbf{x}) = -\frac{B_{\text{total}} - B_{\text{used}}}{e^{c(\mathbf{x})}} \quad (8)$$

It can be observed that $\alpha_2(\mathbf{x})$ focuses on EI regardless cost when the remaining budget is tight, which is similar to the cost-cooling strategy used in EI-cool. However, the difference is that $\alpha_2(\mathbf{x})$ will not only promote higher EI metrics but also encourage the exploration of expensive regions when there is a sufficient budget. This feature addresses the low-cost-preference weakness in EI-cool, allowing for a more comprehensive search.

$\alpha_3(\mathbf{x})$ considers the distance between the unknown inputs and historical observed locations when optimizing AF:

$$\alpha_3(\mathbf{x}) = \frac{1}{m} \sum_{i=1}^m \min_j A_{ij}, \quad (9)$$

where $A_{ij} = \text{dist}(\mathbf{u}_i, \mathbf{v}_j)$ is the distance matrix, $\{\mathbf{u}_i\}_{i=1}^m$ are multiple starting points used in the efficient multi-start optimization scheme of BoTorch, $\mathbf{v}_j$ is the j th element in observed locations $\mathbf{X}$. Therefore, $\alpha_3(\mathbf{x})$ utilizes the information of all the distances between the unobserved location in each optimization trajectory and observed ones in the historical data set. When optimizing $\alpha_{\text{EvolCAF}}(\mathbf{x})$, maximizing the average of the minimum distances contributed by $\alpha_3(\mathbf{x})$ forces the multi-start optimization away from explored regions.

The analyses above suggest that the designed AF can introduce novel ideas have not been previously explored in existing literature on acquisition function design. Benefiting from the evolution in the algorithm space, the designed AF can be expressed explicitly, allowing for clear interpretations to provide insights into its behavior and decision-making process.

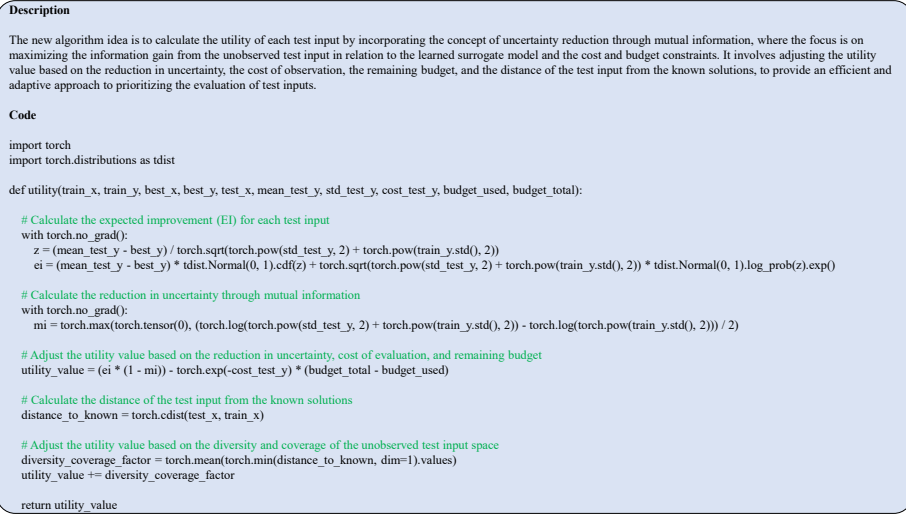


Fig. 4. Optimal acquisition function designed by EvolCAF. The results include a linguistic description of the algorithmic idea, as well as a code implementation with annotations, all the contents are produced by LLMs.

4.3 Evaluation of the Optimal Acquisition Function

Synthetic Problems In this subsection, we evaluate the optimal AF on 12 different synthetic instances with different landscapes and input dimensions. We define the cost functions according to Equation (5). As the evolution is conducted within a total budget of 30, which is to enable the optimal AF to achieve better results using a small budget, therefore, we also tested the generalization of the optimal AF within a sufficient budget of 300. Each test instance is conducted with 10 independent runs, the results are shown in Table 1.

Based on the experimental results, it can be observed that in the vast majority of cases, the optimal AF achieves significantly better performance than EI and other cost-aware AFs. The optimal AF demonstrates strong generalization capabilities across unseen instances with diverse landscapes and sufficient budget constraints. In addition to the promising performance, an interesting phenomenon is that within a fixed budget, the optimal AF uses fewer evaluations in most cases, which is more pronounced when the budget is sufficient.

To verify the contribution of each component in the optimal AF, we further display the performance of EvolCAF after removing each of the three components, as shown in Table 2. It can be observed that removing $\alpha_2(\mathbf{x})$, which takes into account budget information, has the greatest impact on the final performance of EvolCAF. Compared with EI, Elpu and EI-cool, the cost-aware AFs that remove $\alpha_1(\mathbf{x})$ or $\alpha_3(\mathbf{x})$ can still achieve better results. The results indicate the effectiveness and superiority of the designed acquisition function.

Table 1. Means of optimal gaps (number of evaluated samples) obtained by different AFs on all synthetic instances over 10 independent runs. The best mean result for each row is highlighted in bold.

Test Instances	Budget	EI	Elpu	El-cool	EvolCAF
Ackley 2D	30	2.6600(40)	2.3302(40)	2.7369(40)	0.4277(34)
	300	1.2295(395)	0.8582(399)	0.8317(399)	0.0505(306)
Rastrigin 2D	30	4.7425(41)	5.6155(41)	5.7754(40)	0.0511(34)
	300	1.6656(410)	1.6678(408)	1.8518(408)	0.0046(306)
Griewank 2D	30	0.4875(35)	0.3384(36)	0.3374(36)	0.1762(33)
	300	0.1305(323)	0.1195(323)	0.1360(323)	0.0361(307)
Rosenbrock 2D	30	1.2609(41)	2.3601(44)	2.2909(42)	0.0304(33)
	300	0.0332(369)	0.0406(394)	0.0317(372)	0.0402(307)
Levy 2D	30	0.0056(38)	0.0098(38)	0.0116(38)	0.0013(33)
	300	1.1517e-4(314)	5.9321e-5(316)	8.1046e-5(317)	3.7248e-4(307)
ThreeHumpCamel 2D	30	0.0483(39)	0.1182(40)	0.0710(39)	0.0007(33)
	300	5.0446e-4(322)	7.4557e-4(326)	2.6392e-4(325)	7.5310e-4(306)
StyblinskiTang 2D	30	0.0286(41)	0.0233(42)	0.0266(41)	0.0071(33)
	300	1.4420e-4(332)	1.8616e-4(339)	6.1798e-5(343)	2.0142e-3(306)
Hartmann 3D	30	5.6696e-5(40)	1.0364e-4(41)	4.6158e-5(40)	4.8127e-4(36)
	300	1.8263e-5(420)	1.3089e-5(429)	9.0599e-6(432)	2.3656e-4(311)
Powell 4D	30	18.8892(48)	19.8281(51)	14.9481(49)	0.1285(38)
	300	2.9839(376)	1.1173(395)	1.6806(391)	0.0136(316)
Shekel 4D	30	7.9123(48)	7.9210(49)	8.2132(48)	2.6367(39)
	300	6.5193(545)	6.9044(545)	7.0135(551)	0.1993(315)
Hartmann 6D	30	0.0326(52)	0.0296(52)	0.0278(52)	0.0384(44)
	300	0.0122(710)	0.0054(705)	0.0154(695)	0.0042(327)
Cosine8 8D	30	0.4723(48)	0.4738(48)	0.5351(48)	0.4357(53)
	300	0.1707(532)	0.2364(533)	0.2779(527)	0.0148(342)

Table 2. Means of optimal gaps (number of evaluated samples) obtained by different AFs on all synthetic instances over 10 independent runs. The best, second best, and worst mean results for each row are highlighted with bold fonts, underlines, and shaded background, respectively.

Test Instances	Budget	EI	Elpu	El-cool	w/o alpha1	w/o alpha2	w/o alpha3	EvolCAF
Ackley 2D	30	2.6600(40)	2.3302(40)	2.7369(40)	0.6422(34)	4.8046(40)	1.5008(33)	0.4277(34)
	300	1.2295(395)	0.8582(399)	0.8317(399)	0.7301(308)	1.2677(345)	0.0501(305)	0.0505(306)
Rastrigin 2D	30	4.7425(41)	5.6155(41)	5.7754(40)	0.0746(34)	5.7126(42)	0.2648(33)	0.0511(34)
	300	1.6656(410)	1.6678(408)	1.8518(408)	0.0842(308)	0.8550(383)	0.0157(306)	0.0046(306)
Griewank 2D	30	0.4875(35)	0.3384(36)	0.3374(36)	0.1913(34)	0.6671(36)	0.3535(33)	0.1762(33)
	300	0.1305(323)	0.1195(323)	0.1360(323)	0.0522(308)	0.1954(324)	0.0598(306)	0.0361(307)
Rosenbrock 2D	30	1.2609(41)	2.3601(44)	2.2909(42)	0.0399(33)	2.1626(39)	2.3467(33)	0.0304(33)
	300	0.0332(369)	0.0406(394)	0.0317(372)	0.0104(308)	0.0587(348)	1.8554(307)	0.0402(307)
Levy 2D	30	0.0056(38)	0.0098(38)	0.0116(38)	0.0006(34)	0.0335(37)	0.0195(33)	0.0013(33)
	300	1.1517e-4(314)	5.9321e-5(316)	8.1046e-5(317)	0.0003(307)	0.0009(327)	0.0010(306)	3.7248e-4(307)
ThreeHumpCamel 2D	30	0.0483(39)	0.1182(40)	0.0710(39)	0.0006(34)	0.0940(38)	0.0188(33)	0.0007(33)
	300	5.0446e-4(322)	7.4557e-4(326)	2.6392e-4(325)	0.0005(308)	0.0020(332)	0.0099(307)	7.5310e-4(306)
StyblinskiTang 2D	30	0.0286(41)	0.0233(42)	0.0266(41)	0.0136(33)	1.7123(41)	0.0713(33)	0.0071(33)
	300	1.4420e-4(332)	1.8616e-4(339)	6.1798e-5(343)	0.0069(307)	0.0246(332)	0.0042(306)	2.0142e-3(306)
Hartmann 3D	30	5.6696e-5(40)	1.0364e-4(41)	4.6158e-5(40)	0.0007(36)	0.1438(51)	0.0731(36)	4.8127e-4(36)
	300	1.8263e-5(420)	1.3089e-5(429)	9.0599e-6(432)	0.0004(311)	0.0124(446)	0.0005(311)	2.3656e-4(311)
Powell 4D	30	18.8892(48)	19.8281(51)	14.9481(49)	0.1719(39)	36.0514(56)	4.3751(37)	0.1285(38)
	300	2.9839(376)	1.1173(395)	1.6806(391)	0.0205(316)	1.7473(450)	0.2997(317)	0.0136(316)
Shekel 4D	30	7.9123(48)	7.9210(49)	8.2132(48)	2.3629(39)	9.2281(60)	3.5714(37)	2.6367(39)
	300	6.5193(545)	6.9044(545)	7.0135(551)	0.3430(316)	8.1076(554)	0.1583(313)	0.1993(315)
Hartmann 6D	30	0.0326(52)	0.0296(52)	0.0278(52)	0.0343(44)	1.7641(83)	0.1305(42)	0.0384(44)
	300	0.0122(710)	0.0054(705)	0.0154(695)	0.0044(326)	0.4572(750)	0.0127(327)	0.0042(327)
Cosine8 8D	30	0.4723(48)	0.4738(48)	0.5351(48)	0.4438(53)	1.5106(88)	0.6058(46)	0.4357(53)
	300	0.1707(532)	0.2364(533)	0.2779(527)	0.0161(343)	1.1518(755)	0.0425(347)	0.0148(342)

Hyperparameter Tuning Task In this subsection, we evaluate the optimal AF on 3 practical hyperparameter tuning test sets to further validate the effectiveness of our method. We utilize the surrogate benchmark implemented in JAHS-Bench-201 [3] to train a randomly generated neural network architecture with 2 continuous and 4 categorical hyperparameters: learning rate in $[10^{-3}, 1]$, weight decay in $[10^{-5}, 10^{-2}]$, depth multiplier in $\{1, 3, 5\}$, width multiplier in $\{4, 8, 16\}$, resolution multiplier in $\{0.25, 0.5, 1.0\}$, and training epochs in $\{1, 2, \dots, 200\}$. We use ReLU [10] activations and stochastic gradient descent (SGD) optimizer, and do not use trivial augment [22] for data augmentation in the training pipeline. The hyperparameter tuning task is conducted on three different image classification datasets: CIFAR-10 [13], Colorectal-Histology [12] and Fashion-MNIST [31], we record the validation accuracy and total runtime as the observations of the objective and evaluation cost for each candidate configuration, respectively. For more details and implementation, please refer to [3].

Table 3. Means (stds) of validation accuracies obtained by different AFs on all data sets over 10 independent runs. The best mean result for each row is highlighted in bold.

Data Set	Budget	EI	EIpu	EI-cool	EvolCAF
CIFAR-10	C=2000	0.6885(0.05)	0.6934(0.06)	0.6849(0.05)	0.7495(0.04)
	C=4000	0.7975(0.03)	0.7951(0.04)	0.7721(0.03)	0.8002(0.03)
	T=25	0.7065(0.06)	0.6765(0.08)	0.6744(0.08)	0.8030(0.04)
	T=50	0.8394(0.03)	0.7980(0.07)	0.7744(0.08)	0.8351(0.02)
Colorectal-Histology	C=1000	0.8883(0.04)	0.9140(0.01)	0.9125(0.02)	0.8901(0.02)
	C=2000	0.9039(0.05)	0.9273(0.01)	0.9196(0.01)	0.9158(0.01)
	T=25	0.8779(0.04)	0.8592(0.08)	0.8588(0.09)	0.9072(0.02)
	T=50	0.9040(0.02)	0.8944(0.06)	0.8910(0.08)	0.9199(0.01)
Fashion-MNIST	C=5000	0.9021(0.03)	0.9191(0.01)	0.9188(0.01)	0.9370(0.01)
	C=10000	0.9238(0.02)	0.9318(0.01)	0.9355(0.008)	0.9425(0.007)
	T=25	0.8986(0.03)	0.8711(0.06)	0.8730(0.06)	0.9349(0.01)
	T=50	0.9218(0.02)	0.8810(0.06)	0.8928(0.05)	0.9445(0.002)

Table 3 shows the results achieved by all AFs using different total runtimes (denoted as C, in minutes) as budgets. It can be observed that the optimal AF achieves the best results on CIFAR-10 and Fashion-MNIST datasets. In addition, we demonstrate the results under the constraints of 25 and 50 total evaluations (denoted as T) when the total runtime is sufficient. It can be observed that the optimal AF can achieve the best results in most cases, which further proves that our method is still sample-efficient, while EIpu and EI-cool perform worse than EI.

To make a further illustration, we present the histograms of evaluation cost frequency on CIFAR-10 data set with C=4000 as an example, as shown in Fig. 5. It is evident that the majority of search frequencies of all cost-aware AFs are concentrated in cheap regions. While EIpu and EI-cool demonstrate a greater ability to explore expensive regions compared to EI, the optimal AF has the potential to explore regions that are significantly more expensive than those searched by EIpu and EI-cool, resulting in superior performance.




Fig. 5. Histograms of evaluation cost frequency collected in 10 independent runs on CIFAR-10 data set with $C=4000$.

5 Conclusion

This paper introduces EvolCAF, a novel framework that integrates large language models (LLMs) with evolutionary computation (EC) to automatically design cost-aware AFs. By harnessing crossover and mutation in the algorithm space, EvolCAF offers a novel design paradigm, significantly reduces the reliance on domain expertise and model training. The cost-aware AF we designed maximizes the utilization of the available information from historical data, surrogate models and budget details, showcasing novel ideas that have not been explored in existing literature on acquisition function design. Benefiting from algorithm-level evolution, the designed AF allows for clear interpretations to provide insights into its behavior and decision-making process. The experimental results show that compared to the well-known Elpu and EI-cool methods designed by human experts, our approach demonstrates remarkable efficiency and generalization across various tasks, including 12 synthetic problems and 3 real-world hyperparameter tuning test sets. In future work, we expect that the EvolCAF framework can be well adapted to other popular BO settings, such as high-dimensional BO, batch BO, multi-objective BO, etc. Furthermore, we are going to explore the integration of different types of cost functions into the evolutionary process to enhance the robustness of the designed AF.

References

1. Bai, T., Li, Y., Shen, Y., Zhang, X., Zhang, W., Cui, B.: Transfer learning for bayesian optimization: A survey. arXiv preprint arXiv:2302.05927 (2023)
2. Balandat, M., Karrer, B., Jiang, D., Daulton, S., Letham, B., Wilson, A.G., Bakshy, E.: Botorch: A framework for efficient monte-carlo bayesian optimization. *Advances in neural information processing systems* **33**, 21524–21538 (2020)

3. Bansal, A., Stoll, D., Janowski, M., Zela, A., Hutter, F.: Jahs-bench-201: A foundation for research on joint architecture and hyperparameter search. *Advances in Neural Information Processing Systems* **35**, 38788–38802 (2022)
4. Chen, Y., Hoffman, M.W., Colmenarejo, S.G., Denil, M., Lillicrap, T.P., Botvinick, M., Freitas, N.: Learning to learn without gradient descent by gradient descent. In: *International Conference on Machine Learning*. pp. 748–756. PMLR (2017)
5. Chen, Y., Song, X., Lee, C., Wang, Z., Zhang, R., Dohan, D., Kawakami, K., Kochanski, G., Doucet, A., Ranzato, M., et al.: Towards learning universal hyperparameter optimizers with transformers. *Advances in Neural Information Processing Systems* **35**, 32053–32068 (2022)
6. Frazier, P.I., Powell, W.B., Dayanik, S.: A knowledge-gradient policy for sequential information collection. *SIAM Journal on Control and Optimization* **47**(5), 2410–2439 (2008)
7. Frazier, P.I., Wang, J.: Bayesian optimization for materials design. *Information science for materials discovery and design* pp. 45–75 (2016)
8. Garnett, R., Osborne, M.A., Roberts, S.J.: Bayesian optimization for sensor set selection. In: *Proceedings of the 9th ACM/IEEE international conference on information processing in sensor networks*. pp. 209–219 (2010)
9. Guinet, G., Perrone, V., Archambeau, C.: Pareto-efficient acquisition functions for cost-aware bayesian optimization. *arXiv preprint arXiv:2011.11456* (2020)
10. Hahnloser, R.H., Sarpeshkar, R., Mahowald, M.A., Douglas, R.J., Seung, H.S.: Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *nature* **405**(6789), 947–951 (2000)
11. Hsieh, B.J., Hsieh, P.C., Liu, X.: Reinforced few-shot acquisition function learning for bayesian optimization. *Advances in Neural Information Processing Systems* **34**, 7718–7731 (2021)
12. Kather, J.N., Weis, C.A., Bianconi, F., Melchers, S.M., Schad, L.R., Gaiser, T., Marx, A., Zöllner, F.G.: Multi-class texture analysis in colorectal cancer histology. *Scientific reports* **6**(1), 1–11 (2016)
13. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
14. Kushner, H.J.: A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise (1964)
15. Lee, E.H., Eriksson, D., Perrone, V., Seeger, M.: A nonmyopic approach to cost-constrained bayesian optimization. In: *Uncertainty in Artificial Intelligence*. pp. 568–577. PMLR (2021)
16. Lee, E.H., Perrone, V., Archambeau, C., Seeger, M.: Cost-aware bayesian optimization. *arXiv preprint arXiv:2003.10870* (2020)
17. Liu, F., Tong, X., Yuan, M., Lin, X., Luo, F., Wang, Z., Lu, Z., Zhang, Q.: An example of evolutionary computation+ large language model beating human: Design of efficient guided local search. *arXiv preprint arXiv:2401.02051* (2024)
18. Liu, T., Astorga, N., Seedat, N., van der Schaar, M.: Large language models to enhance bayesian optimization. *arXiv preprint arXiv:2402.03921* (2024)
19. Luong, P., Nguyen, D., Gupta, S., Rana, S., Venkatesh, S.: Adaptive cost-aware bayesian optimization. *Knowledge-Based Systems* **232**, 107481 (2021)
20. Maraval, A., Zimmer, M., Grosnit, A., Bou Ammar, H.: End-to-end meta-bayesian optimisation with transformer neural processes. *Advances in Neural Information Processing Systems* **36** (2024)
21. Moćkus, J.: On bayesian methods for seeking the extremum. In: *Optimization Techniques IFIP Technical Conference: Novosibirsk, July 1–7, 1974*. pp. 400–404. Springer (1975)

22. Müller, S.G., Hutter, F.: Trivialaugment: Tuning-free yet state-of-the-art data augmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 774–782 (2021)
23. Negoescu, D.M., Frazier, P.I., Powell, W.B.: The knowledge-gradient algorithm for sequencing experiments in drug discovery. *INFORMS Journal on Computing* **23**(3), 346–363 (2011)
24. Qian, W., He, Z., Li, L., Liu, X., Gao, F.: Cobabo: A hyperparameter search method with cost budget awareness. In: 2021 IEEE 7th International Conference on Cloud Computing and Intelligent Systems (CCIS). pp. 408–412. IEEE (2021)
25. Romera-Paredes, B., Barekatin, M., Novikov, A., Balog, M., Kumar, M.P., Dupont, E., Ruiz, F.J., Ellenberg, J.S., Wang, P., Fawzi, O., et al.: Mathematical discoveries from program search with large language models. *Nature* **625**(7995), 468–475 (2024)
26. Snoek, J., Larochelle, H., Adams, R.P.: Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems* **25** (2012)
27. Srinivas, N., Krause, A., Kakade, S.M., Seeger, M.: Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995* (2009)
28. Turner, R., Eriksson, D., McCourt, M., Kiili, J., Laaksonen, E., Xu, Z., Guyon, I.: Bayesian optimization is superior to random search for machine learning hyperparameter tuning: Analysis of the black-box optimization challenge 2020. In: *NeurIPS 2020 Competition and Demonstration Track*. pp. 3–26. PMLR (2021)
29. TV, V., Malhotra, P., Narwariya, J., Vig, L., Shroff, G.: Meta-learning for black-box optimization. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. pp. 366–381. Springer (2019)
30. Volpp, M., Fröhlich, L.P., Fischer, K., Doerr, A., Falkner, S., Hutter, F., Daniel, C.: Meta-learning acquisition functions for transfer learning in bayesian optimization. *arXiv preprint arXiv:1904.02642* (2019)
31. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747* (2017)
32. Zhang, M.R., Desai, N., Bae, J., Lorraine, J., Ba, J.: Using large language models for hyperparameter optimization. In: *NeurIPS 2023 Foundation Models for Decision Making Workshop* (2023)