7-2023

Seeking better Sharpe ratio via Bayesian optimization

Peng LIU
Singapore Management University, liupeng@smu.edu.sg

Seeking Better Sharpe Ratio via Bayesian Optimization

Peng Liu

assistant professor of quantitative finance (practice) at Singapore Management University, Singapore. liupeng@smu.edu.sg

Abstract

Developing an excellent quantitative trading strategy to obtain a high Sharpe ratio requires optimizing several parameters at the same time. Example parameters include the window length of a moving average sequence, the choice of trading instruments, and the thresholds used to generate trading signals. Simultaneously optimizing all these parameters to seek a high Sharpe ratio is a daunting and time-consuming task, partly because of the unknown mechanism determining the Sharpe ratio. This article proposes using Bayesian optimization to systematically search for the optimal parameter configuration that leads to a high Sharpe ratio. The author shows that the proposed intelligent search strategy performs better than manual search, a common practice that proves to be inefficient. The author's framework also can easily be extended to other parameter selection tasks in portfolio optimization and risk management.

Key Findings

- ▪ To find the best settings for a trading strategy, we need to solve a complex problem that involves finding the highest (or lowest) value of a function, like the Sharpe ratio, without knowing how it behaves. As a nontrivial exercise, this requires adjusting many parameters simultaneously, which takes a lot of time, so doing it manually is not easy.
- ▪ Bayesian optimization is a methodical approach to searching for the best parameter settings that use probabilistic models. It can identify the most promising parameter configurations in far fewer iterations than the random search. This can result in substantial improvements in trading performance.
- ▪ Bayesian optimization is a versatile technique that can be used with a broad spectrum of trading strategies, including rule-based approaches that use technical analysis and machine learning models. It also holds great promise for future research and application in the financial industry.

To maximize the profitability of a trading strategy, we need to identify the best parameter settings that generate the highest profit in multiple backtesting periods. However, backtesting performance can vary depending on the selected period. To address this, a robust approach is to back test a specific set of parameters over multiple periods that cover different scenarios and report the average performance. The optimal set of parameters is the one that consistently produces the highest average terminal return. Unfortunately, manually fine-tuning a trading strategy by testing different parameter values is a very time-consuming process.

Two challenges exist in the search for the best parameter settings for a trading strategy. First, there may be too many possible parameter values to test, making it computationally impractical to loop through every unique configuration, especially if each parameter has multiple alternatives. This is particularly problematic for continuous parameters, which have infinite potential values to test. Second, evaluating the performance of each parameter set takes time and may be slow, which makes the search for the best strategy even more difficult.

A trading strategy usually has one or more parameters that are set to specific values within a given range. Some common parameters include the window length for a moving average, the risk-return ratio, stop-loss order thresholds, trading volume, and entry/exit points. Once these input parameters are set, the strategy generates trading signals and the resulting terminal return over a specific backtesting period is used to evaluate the performance of the parameters.

To illustrate, a trend-following strategy relies on two moving averages to generate a trading signal. The short-term and long-term moving averages are the input parameters, and their optimal values are determined by maximizing a performance metric, such as the Sharpe ratio. However, the objective function that returns the Sharpe ratio given a set of input parameters is often unknown, meaning we have no access to its explicit expression or derivative information. This makes it challenging to use modern optimization tools to identify the optimal parameter settings, a common task in many other domains such as risk management (Liu 2023a) and derivative hedging (Liu 2023b).

The performance of a trading strategy is governed by a black-box function, which is sensitive to the choice of the backtesting period for a set of input parameters. Although a more robust approach is to test the parameters over several representative periods, including live tests, doing so for multiple periods can be excessively time consuming. Instead, a common alternative is to evaluate a single period as a proxy, which is faster but produces an approximate result. However, these noisy evaluations make the global optimization problem of maximizing the Sharpe ratio even more challenging.

In this article, we suggest using Bayesian optimization (BO) as a systematic approach to finding the optimal parameters that yield the highest Sharpe ratio. We employ various BO policies that use different acquisition functions to make sequential decisions in uncertain scenarios and show that BO can identify better parameters in much fewer iterations than the random search method. In the next section, we provide an introduction to BO, including the use of a Gaussian process (GP) as a surrogate model for the unknown objective function and the acquisition function to navigate the search process. We then demonstrate how BO can be applied to identify better parameters, using pairs trading as an example and optimizing different input parameter sets. Last, we discuss the potential of BO for other quantitative trading and portfolio management tasks.

BAYESIAN OPTIMIZATION

BO is a field that focuses on solving optimization problems using the Bayesian approach. Optimization aims to find the best objective value, whether it is a global maximum or minimum, and its corresponding location within a given search domain, also known as the environment. The search process begins at a specific initial location and uses a particular policy to iteratively guide the sampling locations, collect new observations, and update the search policy.

The optimization process involves repeated interactions between the policy (the optimizer) and the environment (the unknown objective function). The policy is a function that takes in a new input parameter and outputs the next value to try out in a principled way. The policy is continually learned and improved as the search continues. A good policy directs the search toward the global optimum faster than a bad policy by spending the limited sampling budget on promising candidate values. On the other hand, the environment contains the unknown objective function that needs to be learned by the policy within a specific boundary. When probing the functional value as requested by the policy, the actual observation revealed by the environment to the policy is often corrupted by noise because of the choice of a single backtesting period, making the learning even more challenging. Therefore, BO, which is a specific approach for global optimization, aims to learn a policy that can help us navigate toward the global optimum of an unknown, noise-corrupted objective function as efficiently and effectively as possible.

When selecting which parameter value to try next, most search strategies must balance the exploration and exploitation tradeoff. Exploration means searching within an unknown and distant area, and exploitation refers to searching within the previously explored neighborhood in the hopes of finding a better functional evaluation, that is, a higher Sharpe ratio. BO also faces this tradeoff. Ideally, we should explore more at the initial phase to increase our understanding of the environment (the black-box function that determines the Sharpe ratio for a specific set of parameters) and then gradually shift toward an exploitation mode that uses existing knowledge to explore promising regions. Exhibit 1 provides the schematic overview of the iterative loop involved in the BO in the context of optimizing trading strategies.

EXHIBIT 1

Schematic Overview of BO Used to Seek a Higher Sharpe Ratio of a Given Trading Strategy

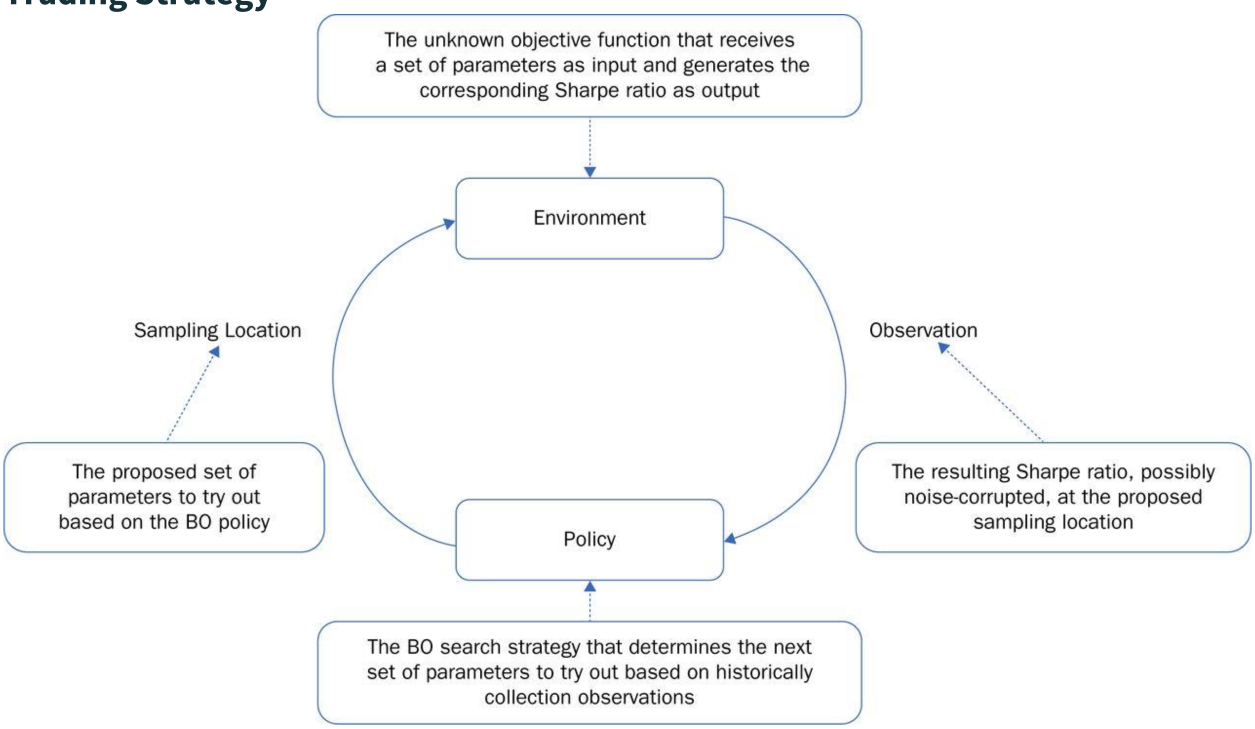

BO balances exploration and exploitation using two components: a GP to approximate the unknown objective function and an acquisition function that incorporates the exploration–exploitation tradeoff into a scalar value that reflects the sampling utility of all candidates in the domain. We will discuss each component in more detail in the following sections.

GP

GPs are a popular tool used in trading strategies to help identify optimal parameter values for a given trading algorithm. In essence, a GP is a mathematical framework used to model and approximate complex functions, where the function is assumed to follow a normal distribution for each specific input parameter. The framework is used to estimate the most likely values of parameters for a given trading strategy, which can then be used to optimize the performance of the strategy.

The GP is a commonly used stochastic process that can model an unknown black-box function and the corresponding uncertainties of modeling. It extends finite-dimensional probability distributions into a continuous search domain that contains an infinite number of variables. Any finite set of points in the domain jointly forms a multivariate Gaussian distribution. The GP is a flexible framework that can model a broad family of functions and quantify their uncertainties. It is a powerful surrogate model used to approximate the true underlying function that governs the working mechanism of the Sharpe ratio for a specific trading strategy.

The basic idea behind GPs is that they provide a flexible and powerful means of modeling functions that are too complex to be represented using simple mathematical formulas. In a trading context, this means that GPs can be used to model the relationships between various input parameters and the performance (i.e., Sharpe ratio) of a given trading strategy. By analyzing these relationships, a GP can help identify the optimal parameter values that lead to the highest possible Sharpe ratio (or other user-defined performance metric).

Overall, GPs represent a valuable tool in the development of quantitative trading strategies, allowing researchers and practitioners to more accurately model complex financial data and optimize trading performance. We refer readers to the supplementary material for a technical introduction of GPs.

Acquisition Function

Bayesian inference and GPs provide a principled approach to understanding the distribution of the objective function. However, to find the global maximum (such as the highest Sharpe ratio) of the objective function, we need to incorporate probabilistic information into our decision-making process. We must build a policy that maximizes the acquisition function and incorporates the most recent information on the objective function, recommending the most promising parameter setting for further sampling in the face of uncertainty across the domain.

The optimization policy, which is guided by maximizing the acquisition function, plays a crucial role in connecting the GP to the ultimate goal of BO. The posterior predictive distribution obtained from the updated GP provides an outlook on the objective value and the associated uncertainty for locations not yet explored. This information can be used by the optimization policy to quantify the utility of any alternative location within the domain.

After obtaining posterior knowledge about candidate locations, such as the mean and variance of the Gaussian distribution at each location (parameter setting), we need to convert this information into a single scalar utility score for the specific trading strategy. This is where the acquisition function comes in. The acquisition function is a mechanism designed to evaluate the relative potential of each candidate location, assigning each location a scalar score. The location with the maximum score is then used as the next sampling choice.

The acquisition function is a critical component of BO because it assesses how valuable a candidate location is when we acquire or sample it. It also is relatively cheap to evaluate, making it an efficient tool for locating the maximum utility score. However, evaluating the acquisition function requires solving another optimization problem, which can add additional computational complexity to the overall optimization process.

When determining the next sampling location, the acquisition function must balance exploration of uncertain regions and exploitation of promising regions based on current information. The acquisition function can be myopic, considering only the immediate impact of the current sampling decision, or nonmyopic, taking into account long-term gain versus immediate reward. Standard acquisition functions could be either single-stage myopic (e.g., Thomson sampling [TS] [Thompson 1933], probability of improvement [PI] [Kushner 1963], expected improvement [EI] [Jones, Schonlau, and Welch 1998], and UCB [Srinivas et al. 2010]), single-stage nonmyopic (e.g., knowledge gradient [KG] [Frazier, Powell, and Dayanik 2008], entropy search [ES] [Hennig and Schuler 2011], and predictive entropy search [PES] [Hernández-Lobato, Hoffman, and Ghahramani 2014]), and multistage nonmyopic (e.g., two-step lookahead EI [Wu and Frazier 2019] and multistep lookahead acquisition functions [Jiang et al. 2020]).

Acquisition functions can vary in different ways, such as the type of utility function used and the number of steps taken to look ahead. The level of risk aversion or preference also can be taken into account when designing the acquisition function. When using GP regression, the level of risk is quantified by the covariance function, which expresses the level of uncertainty about the possible values of the objective function. By incorporating risk aversion or preference into the acquisition function, we can guide the search toward regions that have lower risk or higher potential rewards. The choice of acquisition function depends on the specific goals of the optimization process. In the following, we test a few commonly used acquisition functions, including EI, q-Expected Improvement (qEI), Upper Confidence Bound (UCB), and q-Knowledge Gradient (qKG), and defer a technical introduction to these acquisition functions in the supplementary material of the article.

OPTIMIZING TRADING STRATEGIES VIA BO

In this section, we use the pairs-trading strategy to demonstrate how BO can be applied to optimize the input parameters. The strategy has three input parameters: δ_1, δ_2, and l, which determine the thresholds for entering and exiting positions and the window length used for calculating the Z-score of the spread time series. We test four different acquisition functions, including EI, UCB, qEI, and qKG, to maximize the Sharpe ratio as a black-box function $S = f(\delta_1, \delta_2, l)$. We calculate the Sharpe ratio only once over the entire investment period to simulate the noisy evaluation. In addition, we set the UCB policy parameter β to 0.8 to encourage more exploration.

We conduct experiments in three groups, optimizing δ_1, δ_1, and δ_2, and all three parameters. We start all policies with the same training set of three observations and compare them with the random policy. Because the window size parameter, l, can take only an integer value, we round the resulting proposals to the nearest integer. We also use a technique called continuous relaxation over integer-valued variables, which is shown to speed up the optimization process without affecting the performance, especially in the case of mixed search space (Daulton et al. 2022).

Exhibit 2 displays the performance of different policies over 20 iterations by comparing their maximum Sharpe ratio. There are three experiment groups, and the qKG policy performs the best in group 1 (optimizing over one parameter as shown in the left panel) but has a weaker performance in group 2 (optimizing over two parameters as shown in the middle panel) and group 3 (optimizing over three parameters as shown in the right panel). The qKG policy encourages more exploration, so it is important to appropriately balance the tradeoff between exploration and exploitation for the specific trading strategy being studied. The best tradeoff is attained by UCB in group 2 and EI in group 3.

EXHIBIT 2

Comparing the Cumulative Maximum Sharpe Ratio of the Pairs-Trading Strategy Identified Using Different Search Policies for 20 Iterations

NOTES: When optimizing over one parameter as shown in the left panel, the qKG policy takes the lead, possibly because of its nonmyopic nature in the search. However, all other BO policies (EI, qEI, and UCB) dominate qKG when optimizing over two parameters as shown in the middle panel, highlighting the dynamic nature of the exploration–exploitation tradeoff in different settings. EI reports the highest Sharpe ratio in the third group as shown in the right panel, the highest among the three groups. Most BO policies have a higher Sharpe ratio than the baseline random policy.

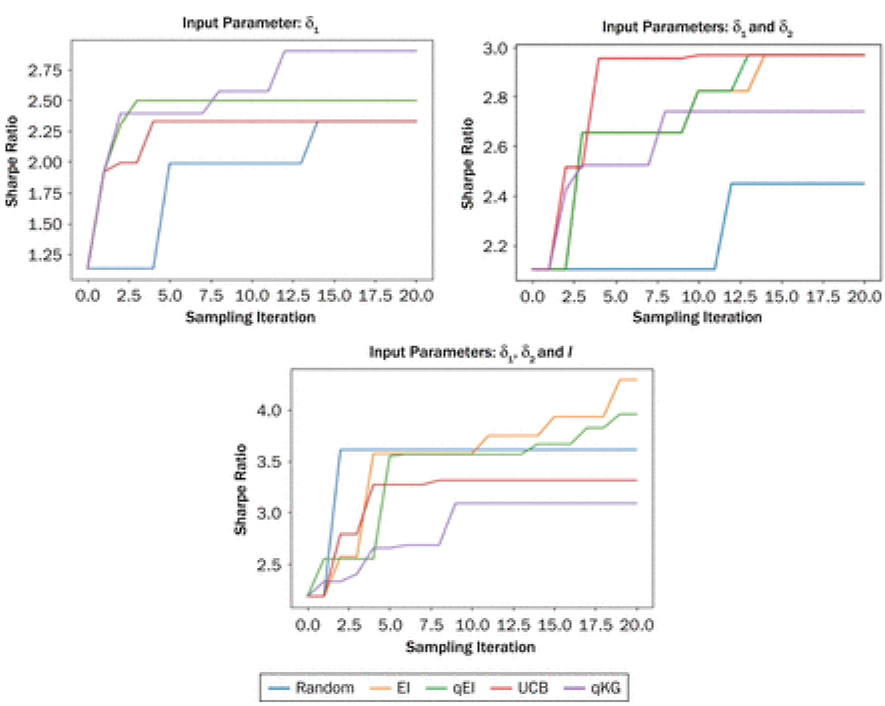

We also note that not all policies perform better than the random strategy in the third group; however, the additional performance boost from EI and qEI at the later iterations demonstrates the potential of BO in seeking a higher Sharpe ratio, which shows the effectiveness of the BO-based search strategy in finding the global optimum as more input parameters are considered. In addition, optimizing over more parameters generally leads to a higher Sharpe ratio, as shown by the scale of the *y*-axis in Exhibit 2.

The optimal trading parameters identified by BO also tend to lead to a higher wealth curve, as shown in Exhibit 3. Although a single BO strategy is not guaranteed to deliver a higher cumulative wealth at each time point, in practice we could adopt a portfolio of BO strategies and ensemble individual models to achieve a higher gain at the group level, similar to the approach proposed in Hoffman, Brochu, and De Freitas (2011). Note that two wealth curves in Exhibit 3 will overlap given the same trading parameters.

EXHIBIT 3

Comparing the Cumulative Wealth Curve for each Search Strategy across Three Experiment Groups

NOTE: BO strategies generally lead to a better-performing wealth curve than random search.

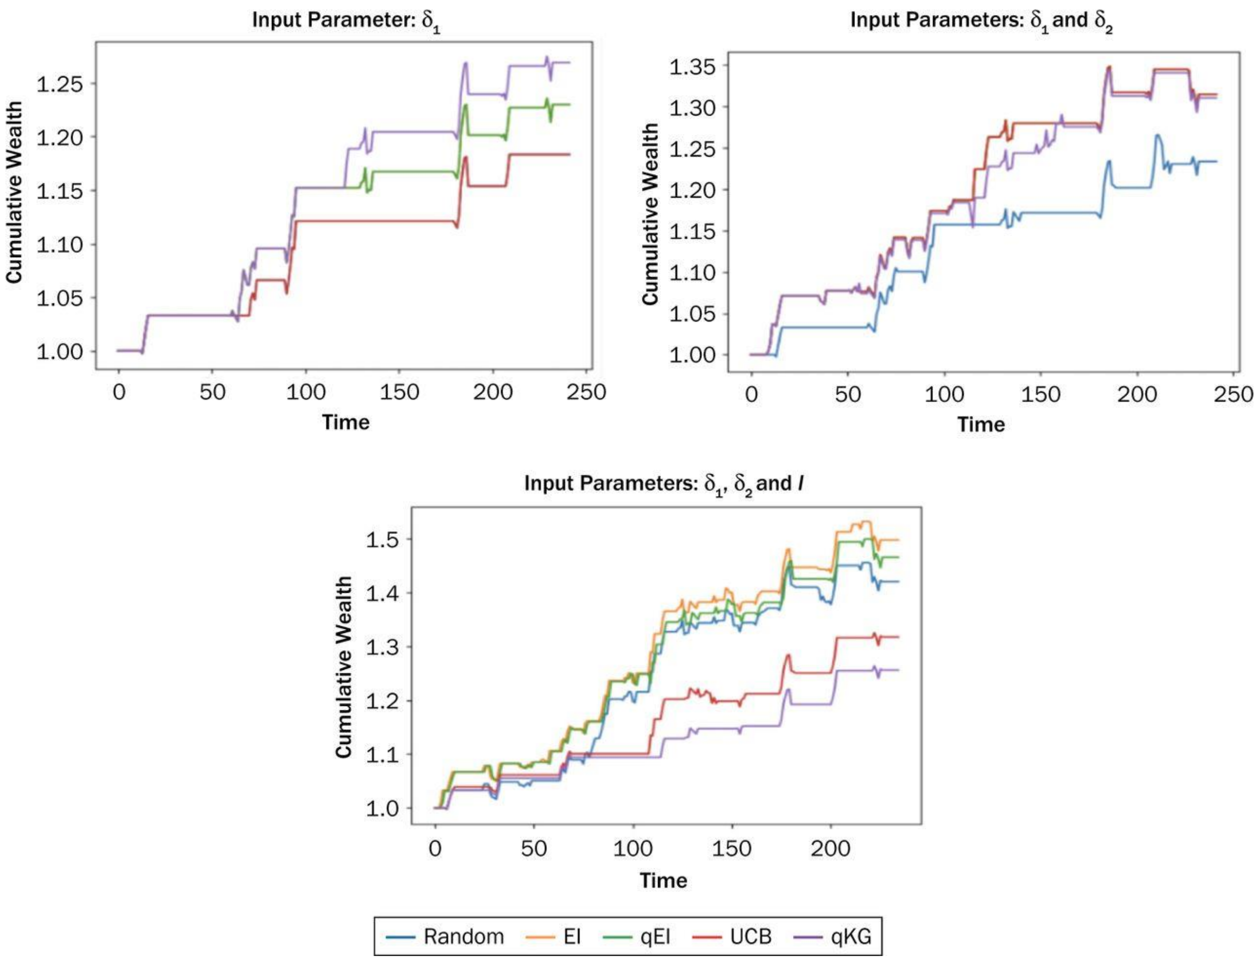

We also conducted the same experiments four times, each with a different training set, to strengthen our findings. Exhibit 4 presents the mean and standard deviation (in brackets) of the terminal Sharpe ratio achieved by various search policies after 20 iterations. The qEI acquisition function achieved the highest Sharpe ratio (highlighted in bold) in groups one and three, and UCB was the best in group two. All BO-based policies consistently outperformed the random policy, indicating the superiority of model-based BO search strategies in solving global optimization problems.

EXHIBIT 4

The Mean and Standard Deviation (in brackets) of the Terminal Sharpe Ratio Using Different Search Policies

NOTES: This exhibit shows the mean and standard deviation (in brackets) of the terminal Sharpe ratio of the pairs-trading strategy identified by different search policies after 20 iterations across three experimental groups, with the highest Sharpe ratio in boldface for each group. All BO-based polices consistently report a higher Sharpe ratio than the random policy, thus showing the advantage of BO in locating a better parameter configuration for a given trading strategy.

Optimization Parameters	Search Policy				
	Random	EI	UCB	qEI	qKG
δ_1	2.431363 (0.106445)	2.553307 (0.032761)	2.446872 (0.334184)	**2.572222** (0.000000)	2.473420 (0.089279)
δ_1 and δ_2	2.470426 (0.247212)	2.736916 (0.116130)	**2.786065** (0.106940)	2.734416 (0.146371)	2.706545 (0.041464)
Δ_1, δ_2, and *l*	3.420223 (0.219696)	3.708989 (0.170408)	3.639995 (0.396545)	**3.872040** (0.226758)	3.681820 (0.142673)

CONCLUSION

This article has demonstrated the efficacy of employing BO as a robust and efficient mechanism for enhancing trading strategies by targeting a high Sharpe ratio. Utilizing GP models and optimizing carefully crafted acquisition functions enables traders, analysts, and researchers to effectively navigate the extensive parameter landscape of trading strategies and pinpoint the most advantageous parameter setups. Our empirical analysis using the pairs-trading strategy revealed substantial improvements in trading performance, encompassing both profitability and risk management aspects. In addition, BO can be applied to a diverse array of trading strategies, including those based on technical analysis or machine learning models. Ultimately, the adoption of BO in the trading domain offers a promising avenue for future investigation and implementation within the financial sector.

ACKNOWLEDGMENTS

I thank Tee Chyng Wen for the helpful discussion on trading strategies.

REFERENCES

1. Daulton, S., Wan, X., Eriksson, D., Balandat, M., Osborne, M. A., and Bakshy, E. "Bayesian Optimization over Discrete and Mixed Spaces via Probabilistic Reparameterization." In *Advances in Neural Information Processing Systems*, edited by A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho. 35: 12760–12774, NY: Curran Associates, Inc., 2022.
   Google Scholar

2. Frazier, P., Powell, W., and Dayanik, S. 2008. “A Knowledge-Gradient Policy for Sequential Information Collection.” *SIAM Journal Control and Optimization* 47: 2410–2439.
   Google Scholar
3. Hennig, P., and Schuler, C. 2011. “Entropy Search for Information-Efficient Global Optimization.” *Journal of Machine Learning Research* 13: 6.
   Google Scholar
4. Hernández-Lobato, J., Hoffman, M., and Ghahramani, Z. 2014. “Predictive Entropy Search for Efficient Global Optimization of Black-Box Functions.” *Advances in Neural Information Processing Systems* 27 (1).
   Google Scholar
5. Hoffman, M., Brochu, E., and De Freitas, N. “Portfolio Allocation for Bayesian Optimization.” In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence (UAI’11)*, pp. 327–336. Arlington, VA: AUAI Press, 2011.
   Google Scholar
6. Jiang, S., Jiang, D., Balandat, M., Karrer, B., Gardner, J., and Garnett, R. “Efficient Nonmyopic Bayesian Optimization via One-Shot Multi-Step Trees.” In *Advances in Neural Information Processing Systems*, Vol. 33, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, pp. 18,039–18,049. NY: Curran Associates, Inc., 2020.
   Google Scholar
7. Jones, D., Schonlau, M., and Welch, W. 1998. “Efficient Global Optimization of Expensive Black-Box Functions.” *Journal of Global Optimization* 13: 455–492.
   Google ScholarCrossRef
8. Kushner, H. J. 1963. “A New Method of Locating the Maximum Point of an Arbitrary Multipeak Curve in the Presence of Noise.” *Journal of Basic Engineering* 86: 97–106.
   Google Scholar
9. Liu, P. 2023a. “An Integrated Framework on Human-in-the-Loop Risk Analytics.” *The Journal of Financial Data Science* 5 (1): 58–64.
   Google ScholarCrossRefAbstract/Full Text
10. Liu, P. 2023b. “A Review on Derivative Hedging Using Reinforcement Learning.” *The Journal of Financial Data Science* 5 (2): 136–145.
    Google ScholarCrossRefAbstract/Full Text
11. Srinivas, N., Krause, A., Kakade, S., and Seeger, M. 2010. “Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design.” International Conference on Machine Learning (ICML), 1015–1022.

Google Scholar

12. Thompson, W. R. 1933. “On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples.” *Biometrika* 25 (3/4): 285–294.
Google ScholarCrossRef

13. Wu, J., and Frazier, P. “Practical Two-Step Lookahead Bayesian Optimization.” In *Advances in Neural Information Processing Systems*, Vol. 32, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett. NY: Curran Associates, Inc., 2019.
Google Scholar