

TESO: TABU-ENHANCED SIMULATION OPTIMIZATION FOR NOISY BLACK-BOX PROBLEMS

Bulent Soykan¹, Sean Mondesire², and Ghaith Rabadi²

¹Institute for Simulation and Training, University of Central Florida, Orlando, FL, USA

²School of Modeling, Simulation, and Training, University of Central Florida, Orlando, FL, USA

ABSTRACT

Simulation optimization (SO) is frequently challenged by noisy evaluations, high computational costs, and complex, multimodal search landscapes. This paper introduces Tabu-Enhanced Simulation Optimization (TESO), a novel metaheuristic framework integrating adaptive search with memory-based strategies. TESO leverages a short-term Tabu List to prevent cycling and encourage diversification, and a long-term Elite Memory to guide intensification by perturbing high-performing solutions. An aspiration criterion allows overriding tabu restrictions for exceptional candidates. This combination facilitates a dynamic balance between exploration and exploitation in stochastic environments. We demonstrate TESO's effectiveness and reliability using a queue optimization problem, showing improved performance compared to benchmarks and validating the contribution of its memory components.

1 INTRODUCTION

Simulation Optimization (SO) is now widely recognized as a powerful methodology for decision-making across a diverse spectrum of fields, including engineering design (Deng 2007), supply chain management (Jung et al. 2004), healthcare (Lal, Roh, and Huschka 2015) and increasingly, the tuning of hyperparameters in complex machine learning models (Amaran et al. 2016). The prevalence of SO stems from its power to identify optimal parameter settings or system configurations by directly interacting with simulation models, which often represent the only viable way to capture the intricacies and stochastic nature of real-world systems when analytical models are intractable (Nezami and Anahideh 2023).

However, the practical application of SO is frequently confronted by significant challenges. Many high-fidelity simulations are *computationally expensive*, demanding substantial time or resources for each evaluation. These simulations are often treated as *black-box* functions, meaning their internal analytical structure is either unknown or too complex to be exploited directly for optimization, thus preventing the direct calculation of derivatives (Tekin and Sabuncuoglu 2004). Also, the outputs derived from *stochastic simulations* are inherently noisy due to the underlying random components within the model. This noise complicates the assessment of a candidate solution's true performance, making it difficult to reliably distinguish genuine improvements from mere statistical fluctuations based on a limited number of simulation runs (Fu 1994). In addition, the search landscapes encountered in SO problems are frequently highly complex. They can be high-dimensional, contain numerous local optima (i.e., be *multimodal*), and lack readily available or meaningful gradient information. This complexity renders standard gradient-based optimization techniques ineffective and necessitates the use of algorithms capable of global search without relying on derivative information (Figueira and Almada-Lobo 2014).

A variety of approaches have been developed to tackle the challenges inherent in SO. Among the most prominent are Surrogate-Based Optimization (SBO) methods, which aim to mitigate the high computational cost of simulations by constructing cheaper-to-evaluate approximation models, or surrogates (Hong and Zhang 2021). Gaussian Processes (GPs), often employed within a Bayesian Optimization (BO) framework, are frequently used for this purpose due to their ability to capture uncertainty in the approximation (Pearce,

Poloczek, and Branke 2019). However, constructing and tuning effective surrogate models, especially for high-dimensional or highly complex functions, can be computationally intensive itself and sensitive to the choice of kernel functions and hyperparameters (Fan et al. 2024).

Another significant class of methods, particularly relevant when comparing a finite number of system designs, falls under the umbrella of Ranking and Selection (R&S) procedures (Boesel, Nelson, and Kim 2003). These statistical methods focus on efficiently allocating simulation budget to identify the best among a discrete set of alternatives with a certain level of statistical confidence (Hong, Fan, and Luo 2021). While powerful in their domain, R&S procedures are primarily suited for problems with discrete or a relatively small number of candidate solutions and are less directly applicable to the continuous or high-dimensional search spaces often encountered in optimization.

Metaheuristics represent a third major category, offering general-purpose stochastic search strategies applicable to complex black-box problems. Common examples used in SO include Genetic Algorithms (GA) (de Carvalho Miranda, Montevechi, and de Pinho 2015), Simulated Annealing (SA) (Haddock and Mittenthal 1992), and Particle Swarm Optimization (PSO) (Kuo and Yang 2011). These methods often excel at global exploration but can suffer from premature convergence to local optima, particularly in the presence of significant simulation noise (Brahim et al. 2024). Also, their performance can be highly sensitive to the choice of algorithmic hyperparameters (e.g., population size, mutation rates, cooling schedules), often requiring substantial tuning effort, which can itself become an optimization problem (Nezami and Anahideh 2023). Distinct from these, Tabu Search (TS) is a metaheuristic optimization strategy characterized by its intelligent use of memory structures (Glover 1990). Central to TS is the concept of a *Tabu List*, which records recently visited solutions or moves, temporarily forbidding them to prevent the search from cycling and to encourage exploration of unvisited areas of the solution space. This memory-based mechanism is specifically designed to help search processes escape the confines of local optima, a common pitfall for many other heuristic methods. The integration of such memory-based strategies offers a promising direction for enhancing the robustness and exploratory capabilities of SO algorithms.

Addressing the limitations of existing approaches and benefiting from the strengths of memory-based heuristics, this paper introduces Tabu-Enhanced Simulation Optimization (TESO), a novel algorithm specifically designed for complex and noisy SO problems. The core idea behind TESO is to dynamically balance the fundamental trade-off between *exploration*, which involves discovering potentially new and promising regions of the vast solution space, and *exploitation*, which focuses on carefully refining solutions within regions already identified as high-performing. This research is motivated by the central research question: *Can the integration of TS principles, Elite Memory structures, and adaptive perturbation strategies within a unified SO framework effectively balance exploration and exploitation to reliably converge towards high-quality solutions, particularly when dealing with noisy, potentially multimodal, and computationally expensive simulation environments?*

While TS and elite set strategies are known in deterministic optimization, their application to stochastic SO presents unique challenges. Standard TS assumes deterministic objective evaluations, making concepts like 'best move' and 'aspiration' straightforward. In SO, these must be redefined to operate on noisy estimates. TESO's novelty is not in inventing these memory structures, but in adapting and integrating them into a cohesive framework for the SO context. Specifically, it uses a dual-memory approach which is explicitly designed to balance the search process under the uncertainty inherent to simulation.

To address the research questions, this paper makes the following primary contributions to the field of SO: (i) The proposal of TESO, a *novel SO algorithm* that synergistically integrates Tabu Search principles (tabu list, aspiration criterion) with Elite Memory and adaptive noise control, specifically tailored for noisy, complex black-box simulation environments. (ii) An empirical *demonstration of TESO's effective balancing of exploration and exploitation*. We show how tabu-driven diversification and random sampling interact with elite-guided intensification and adaptive perturbation to navigate challenging stochastic landscapes, leading to improved solution quality and reliability.

The remainder of this paper is structured as follows. Section 2 provides a detailed review of relevant background concepts. Section 3 presents the methodology of the proposed TESO algorithm, elaborating on its components. Section 4 details the numerical experiments conducted to evaluate TESO's performance. Finally, Section 5 concludes the paper by summarizing the key findings, acknowledging limitations, and suggesting potential directions for future research in this area.

2 BACKGROUND AND RELATED WORK

2.1 Simulation Optimization (SO)

SO refers to the process of finding optimal input parameters or decision variables, denoted by a vector x , for a system whose performance is evaluated using a computer simulation model (Andradóttir 1998). The goal is typically to minimize or maximize an objective function, $J(x)$, which represents the expected value of some performance measure obtained from the simulation output, $f(x, \omega)$ (Gosavi et al. 2015). Formally, the problem can often be expressed as:

$$\min_{x \in \mathcal{X}} J(x) = \mathbb{E}_{\omega}[f(x, \omega)] \quad (1)$$

or equivalently for maximization, where x represents the vector of decision variables, \mathcal{X} is the feasible set of decisions, $f(x, \omega)$ is the performance measure output from a single simulation run given input x and realization ω of the random elements within the simulation, and $\mathbb{E}_{\omega}[\cdot]$ denotes the expectation taken over the probability space governing the stochastic components ω . A key characteristic of many SO problems is that the simulation model acts as a *black box*, meaning an explicit analytical expression for $J(x)$ or its gradients is often unavailable (Cao, Hu, and Lian 2024).

The application of optimization techniques within this simulation context presents several distinct and significant challenges. First and foremost is the need for *handling stochastic responses* (Hong and Zhang 2021). Since the simulation output $f(x, \omega)$ for any given x is a random variable due to the inherent randomness ω within the simulation model, a single simulation run provides only a noisy estimate of the true expected performance $J(x)$. Relying on single, noisy observations can lead to incorrect assessments of solution quality and poor optimization decisions. Accurately estimating $J(x)$ typically requires multiple simulation replications for each candidate solution x , which directly impacts the overall computational effort. This inherent noise complicates tasks such as estimating gradients (if applicable), comparing candidate solutions reliably, and determining convergence.

Also, SO problems are often constrained by *computational budget limitations* (Chen and Lee 2011). High-fidelity simulations representing complex real-world systems can be computationally intensive, requiring significant time or resources for each execution. Consequently, the total number of simulation runs (across all candidate solutions and all replications per solution) that can be performed is often severely limited. This budget constraint forces a careful allocation of computational effort and exacerbates the challenge posed by stochastic responses, as obtaining highly precise estimates of $J(x)$ for many different x values may be infeasible.

In addition, SO algorithms must effectively balance the fundamental *exploration versus exploitation dilemma* (Amaran et al. 2016). Given the limited budget and noisy feedback, the algorithm must decide whether to spend resources investigating new, unexplored regions of the decision space \mathcal{X} (*exploration*) in the hope of discovering globally superior solutions, or to focus resources on refining solutions within regions already known to yield good performance (*exploitation*). Over-emphasizing exploration can waste valuable computational budget on evaluating unpromising areas, while over-emphasizing exploitation risks premature convergence to a locally optimal solution that may be significantly inferior to the true global optimum. Achieving an effective balance between these two competing objectives is critical for the success of any SO algorithm, particularly when dealing with the complex, potentially multimodal landscapes frequently encountered in simulation-based problems (Nsiye et al. 2024; Soykan and Rabadi 2024). TESO is designed specifically to address these intertwined challenges.

2.2 Metaheuristics in Simulation Optimization

Given the black-box nature and the frequent absence of readily available gradient information that characterize many SO problems, metaheuristic algorithms offer a compelling and widely adopted set of tools (Amaran et al. 2016). A primary strength of these population-based or trajectory-based methods lies in their inherent ability to perform a *global search* across the decision space, rather than being confined to local improvements like gradient-based methods (Soykan and Rabadi 2022). They operate directly on candidate solutions and their observed (potentially noisy) objective function values, bypassing the need for derivative information or strong assumptions about the objective function's structure (e.g., convexity). This makes them particularly suitable for solving the complex, potentially multimodal, and rugged landscapes often characteristic of simulation models (Thengvall, Hall, and Deskevich 2025).

However, the effective application of metaheuristics in SO is not without its difficulties. A significant challenge lies in *parameter tuning*; algorithms like GA (requiring settings for mutation and crossover rates, population size, selection mechanisms), SA (requiring an effective cooling schedule), and PSO (requiring inertia weights, cognitive and social parameters) possess numerous control parameters. The optimal settings for these parameters can be highly problem-dependent and non-trivial to determine *a priori*, often requiring extensive preliminary experimentation or specific expertise (Nezami and Anahideh 2023). Also, most metaheuristics are stochastic in nature and often lack formal *convergence guarantees* to the global optimum, especially within a finite computational budget typically available in SO. While some theoretical convergence results exist under specific assumptions (e.g., for SA under infinitely slow cooling), practical guarantees of finding the true optimum in a finite number of steps are rare.

Despite their global search capabilities, these algorithms can still suffer from *premature convergence*, becoming trapped in sub-optimal regions of the search space. The inherent *stochasticity* of simulation outputs adds another layer of complexity; noise can obscure the true underlying landscape, making fitness or objective evaluations unreliable and potentially misleading the search direction, which can exacerbate issues like premature convergence or difficulty in discerning truly superior solutions. Despite these weaknesses, metaheuristics remain a popular and often effective approach for SO due to their flexibility, general applicability to challenging black-box problems, and potential for finding high-quality solutions when gradient-based or analytical methods are infeasible. TESO aims to retain the exploratory strengths while mitigating some weaknesses, particularly premature convergence, through its specific memory structures.

2.3 Memory-Based Heuristics: Tabu Search and Elite Memory

TS is a prominent metaheuristic that exemplifies the strategic use of adaptive memory, particularly for escaping local optima in complex search spaces (Glover 1990). Its core mechanism is the *Tabu List*, a short-term memory that records recent moves or solution attributes, temporarily forbidding their reversal for a set duration (Tabu Tenure). This prevents cycling and promotes diversification by forcing the search into new areas. TS iteratively explores solution neighborhoods, selecting the best permissible move (non-tabu, or tabu but meeting an *aspiration criterion*, such as finding a new global best) even if it results in a temporary worsening of the objective function. This non-improving move capability, combined with longer-term strategies for exploitation and exploration, has made TS highly effective (Soykan and Rabadi 2016).

The broader concept of leveraging memory extends beyond TS, with many advanced heuristics incorporating mechanisms to learn from search history. A common and powerful form is the use of long-term memory to store a collection of *elite solutions*—the best candidates found so far. This elite set serves as a basis for exploitation. Other forms of memory might track historical attribute frequencies or operator performance to adaptively bias the search. By incorporating memory, whether the short-term restrictive memory of a Tabu List or the long-term guiding memory of an elite set, heuristics can achieve a more effective balance between exploration and exploitation, leading to faster convergence and increased robustness, particularly in the noisy and complex landscapes characteristic of SO (Yu, Lahrichi, and Matta 2023).

3 TESO: ALGORITHM METHODOLOGY

3.1 Framework Overview

TESO is positioned as a direct search metaheuristic specifically adapted for noisy, black-box SO. It improves upon standard metaheuristics by integrating adaptive memory structures inspired by Tabu Search (TS) to better navigate complex, uncertain landscapes. TESO accommodates simulation noise by operating on mean estimates derived from multiple replications. Key adaptations include: a *Tabu List* storing recent candidate representations to prevent re-evaluation and encourage diversification despite noise; a stochastic *aspiration criterion* allowing tabu candidates to be accepted if their estimated performance surpasses the current best; and an *Elite Memory* storing top-performing candidates based on their mean estimates, which provides robust starting points for intensification via perturbation.

When contrasted with *standard TS*, TESO differs significantly in its neighborhood exploration and evaluation. Standard TS often relies on explicitly defined, discrete neighborhood structures (e.g., adjacent swaps, bit flips). TESO, on the other hand, employs an implicit neighborhood structure defined by adaptive random perturbation around elite solutions. The search step is probabilistic rather than a deterministic evaluation of all neighbors. Also, TESO evaluates candidates using multiple simulation replications to handle noise, whereas standard TS typically assumes deterministic function evaluations. This distinction is critical. In a deterministic setting, a move is unambiguously improving or not. In TESO, all comparisons are based on mean estimates from multiple replications, which are themselves random variables. Therefore, the aspiration criterion and improvement checks (detailed later) are inherently stochastic, a key adaptation for the SO domain.

Compared to *other metaheuristics like GA, SA, or PSO* commonly used in SO, TESO's explicit use of *both* short-term restrictive memory (Tabu List) and long-term guiding memory (Elite Memory) is a key differentiator. While GAs maintain diversity through populations and PSO uses particle/global bests (a form of Elite Memory), they typically lack the systematic short-term avoidance mechanism provided by a Tabu List. Basic SA is generally memoryless. Also, TESO should be clearly distinguished from *SBO and BO* approaches. SBO/BO methods construct an explicit statistical model (the surrogate, e.g., a Gaussian Process) of the underlying simulation response. They use this surrogate, along with an acquisition function (like Expected Improvement or Upper Confidence Bound), to intelligently select the next point(s) to evaluate with the expensive simulation, aiming to improve the surrogate model and find the optimum efficiently. TESO, in contrast, is a *direct search* method. While it uses memory, it does not build an explicit global model of the objective function; it operates directly on the (replicated) outputs of the simulation model itself to guide its search trajectory.

3.2 Algorithm Methodology

TESO algorithm (Algorithm 1) begins by initializing essential components (line 2): the Tabu List (\mathcal{T}) and Elite Memory (\mathcal{E}) are cleared, the best objective value found so far (f_{best}) is set to an appropriate initial value (e.g., infinity for minimization), the adaptive noise level (η) is set to its starting value (η_{init}), and iteration counters (t , Δt) are initialized. The core logic resides within a loop iterating up to the total trial budget T (line 4). Each iteration commences with *Candidate Generation* (lines 5-9). During the initial n_{init} iterations or with a small probability p_{div} thereafter, the algorithm performs diversification by generating a candidate $x^{(t)}$ randomly (line 6). Otherwise, it engages in intensification: an elite solution x_e is selected from \mathcal{E} , and a new candidate $x^{(t)}$ is created by perturbing x_e using the current noise level η (line 8). A fallback to random generation occurs if \mathcal{E} is empty. A unique, hashable representation $h^{(t)}$ of the generated candidate is then created (line 9). Since our test problem uses a continuous decision variable, this representation can be created by discretizing the variable into bins or by applying a hash function to its string representation.

Before proceeding to potentially costly evaluation, the candidate undergoes a *Tabu Check and Aspiration* step (lines 11-14). If the candidate's representation $h^{(t)}$ is found in the Tabu List \mathcal{T} and it does not satisfy

the predefined aspiration criterion (i.e., it is not expected to significantly improve upon f_{best}), the candidate is discarded, and the algorithm proceeds directly to the next iteration via the ‘continue’ statement (line 13). If the candidate $x^{(t)}$ is not tabu or satisfies the aspiration criterion, it is subjected to *Stochastic Evaluation* (line 16). This involves running n_{rep} independent simulation replications using $x^{(t)}$ as input. The mean performance $\mu^{(t)}$ and its standard deviation $\sigma^{(t)}$ are calculated from these replications, providing a statistically robust estimate of the candidate’s objective value. Note that while the standard deviation $\sigma^{(t)}$ is calculated, it is not directly used in the decision logic of the current TESO implementation. It is recorded for analysis and could be used in more advanced versions of the algorithm, for instance, to dynamically adjust the number of replications or to inform a risk-based aspiration criterion.

Following evaluation, the algorithm performs *Updates to the Best Solution and Memory Structures* (lines 18-22). It checks if the evaluated mean $\mu^{(t)}$ constitutes an improvement over the current f_{best} . If an improvement is found, f_{best} and the corresponding best solution x_{best} are updated, and the no-improvement counter Δt is reset (line 20). If no improvement occurred (and the algorithm is past the initialization phase), Δt is incremented (line 21). Regardless of improvement, the candidate’s representation $h^{(t)}$ is added to the Tabu List \mathcal{T} , and the candidate-performance pair $(x^{(t)}, \mu^{(t)})$ is added to the Elite Memory \mathcal{E} , maintaining their respective capacity limits (line 22). Within the loop, *Adaptive Noise Control and Termination Check* occur (lines 24-25). The noise level η used for perturbation is updated according to a predefined schedule (e.g., decaying over iterations). The algorithm then checks if the termination criteria are met: either the total trial budget T is exhausted (implicit in the ‘For’ loop) or the number of iterations without improvement Δt has reached the maximum allowed Δt_{max} . If the latter occurs, the loop terminates early via the ‘break’ statement (line 25). Once the loop finishes, the algorithm returns the best solution x_{best} and its corresponding estimated objective value f_{best} (line 27).

Algorithm 1 TESO Algorithm

Require: $f(x, \omega)$, \mathcal{X} , T , n_{init} , n_{rep} , $\eta_{\text{init}} \rightarrow \eta_{\text{final}}$, $C_{\mathcal{T}}, C_{\mathcal{E}}$, p_{div} , Δt_{max} , direction.

Ensure: Best solution x_{best} , Best objective f_{best} .

```

1: Init  $\mathcal{T}, \mathcal{E}, f_{\text{best}}, \eta \leftarrow \eta_{\text{init}}, \Delta t \leftarrow 0, x_{\text{best}} \leftarrow \text{null}$ .           ▷ Initialize memory, best value, noise, counters
2: for  $t = 1$  to  $T$  do
3:                                                                                       ▷ Generate Candidate  $x^{(t)}$ 
4:   if  $t \leq n_{\text{init}}$  or  $\text{Random}() < p_{\text{div}}$  then
5:      $x^{(t)} \leftarrow \text{RandCand}(\mathcal{X})$                                                                                        ▷ Diversify
6:   else
7:      $x_e \leftarrow \text{SelectElite}(\mathcal{E}); x^{(t)} \leftarrow \text{Perturb}(x_e \text{ or } \text{RandCand}(), \eta)$            ▷ Intensify
8:   end if
9:    $h^{(t)} \leftarrow \text{Represent}(x^{(t)})$ 
10:                                                                                       ▷ Check Tabu & Aspiration
11:  if  $h^{(t)} \in \mathcal{T}$  and not  $\text{AspirCritMet}(x^{(t)}, f_{\text{best}}, \text{direction})$  then
12:    continue                                                                                       ▷ Skip tabu, non-aspirated
13:  end if
14:                                                                                       ▷ Evaluate Candidate
15:   $(\mu^{(t)}, \sigma^{(t)}) \leftarrow \text{Evaluate}(x^{(t)}, f, n_{\text{rep}})$            ▷ Run  $n_{\text{rep}}$  sims, get mean/std
16:                                                                                       ▷ Update Best & Memory
17:  if  $\text{IsImprovement}(\mu^{(t)}, f_{\text{best}}, \text{direction})$  then
18:     $f_{\text{best}} \leftarrow \mu^{(t)}; x_{\text{best}} \leftarrow x^{(t)}; \Delta t \leftarrow 0$ 
19:  else if  $t > n_{\text{init}}$  then  $\Delta t \leftarrow \Delta t + 1$ 
20:  end if
21:  Add  $h^{(t)}$  to  $\mathcal{T}$ ; Add  $(x^{(t)}, \mu^{(t)})$  to  $\mathcal{E}$                                                                                        ▷ Update memories
22:                                                                                       ▷ Adapt Noise & Check Termination
23:   $\eta \leftarrow \text{UpdateNoise}(t, T, \eta_{\text{init}}, \eta_{\text{final}})$ 
24:  if  $\Delta t \geq \Delta t_{\text{max}}$  then break                                                                                       ▷ Stop if no improvement
25:  end if
26: end for
27: Return:  $x_{\text{best}}, f_{\text{best}}$ .

```

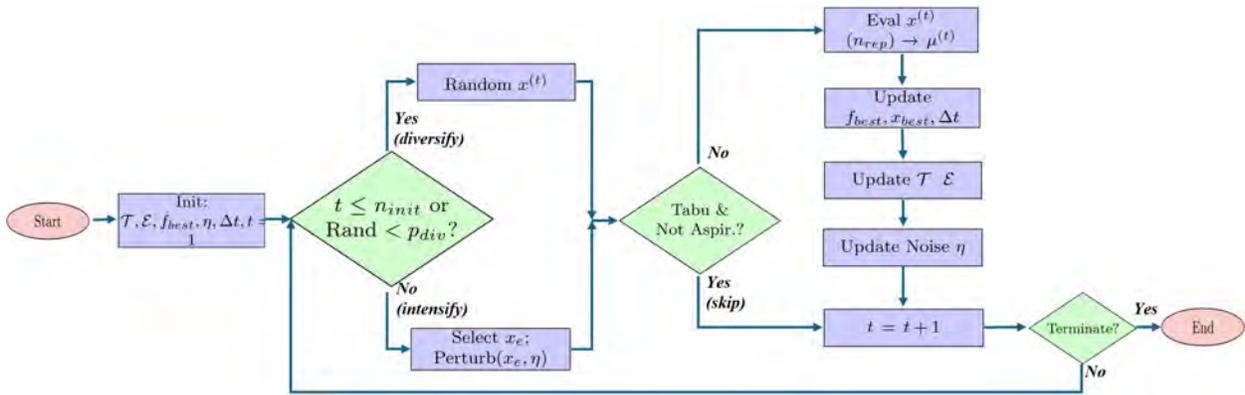


Figure 1: Flow Diagram of the TESO Algorithm, illustrating the iterative cycle including candidate generation, tabu/aspiration filtering, evaluation, and memory updates.

Figure 1 provides a visual summary of the TESO methodology detailed in Algorithm 1. The diagram illustrates the flow starting from initialization, proceeding through the main iterative loop. Key stages depicted include the adaptive candidate generation (switching between diversification via random sampling and intensification via elite perturbation), the critical filtering step based on the Tabu List (\mathcal{T}) and aspiration criteria, the stochastic evaluation using multiple replications, and the subsequent updates to the best solution, both memory structures (\mathcal{T} , \mathcal{E}), and control parameters like noise level (η) before checking termination conditions.

4 NUMERICAL EXPERIMENTS

4.1 Experimental Objectives

The numerical experiments herein aim to empirically validate the proposed TESO algorithm’s effectiveness for noisy SO problems. Our goals are fourfold: first, to demonstrate TESO’s capability to converge reliably to high-quality solutions within typical budget constraints despite stochastic evaluations; second, to quantify the impact of its core memory components (Tabu List, Elite Memory) by comparing the full algorithm against ablation variants lacking one of these structures; third, to benchmark TESO’s final solution quality, convergence speed, and reliability against baseline methods; and fourth, to analyze how TESO’s mechanisms dynamically balance exploration and exploitation during the search. These experiments collectively seek to establish TESO as an effective framework for challenging SO problems.

4.2 Test Problem: M/M/k Queue Optimization

To evaluate the performance of the TESO algorithm on a *multiple-server M/M/k queueing system*. This system is characterized by Poisson arrivals (M), exponentially distributed service times (M) shared across k identical parallel servers (k), and typically assumes an infinite buffer capacity operating under a First-In, First-Out (FIFO) discipline. The key parameters governing its behavior are the mean arrival rate, λ , the number of servers, k , and the mean service rate, μ , for *each individual server*. For stability, the system requires the total service capacity to exceed the arrival rate, i.e., $k \cdot \mu > \lambda$.

In this optimization context, the *decision variable*, x , remains the *service rate* μ provided by each of the k identical servers. We assume this rate can be controlled within a predefined feasible range $[\mu_{min}, \mu_{max}]$, where $k \cdot \mu_{min} > \lambda$ must hold for stability. The *objective function*, $J(\mu)$, aims to minimize a combination of operational performance (customer waiting time) and the cost associated with providing service capacity. Specifically, we seek to minimize the sum of the expected steady-state average sojourn time (W) experienced by customers (time spent waiting in queue plus time in service) and a cost penalty proportional to the

square of the individual server service rate, aggregated across all servers. The quadratic cost reflects the increasing marginal expense of providing faster service per server. Therefore, the objective function to be minimized is:

$$J(\mu) = \mathbb{E}[W(k, \mu, \omega)] + C \cdot k \cdot \mu^2 \quad (2)$$

where $\mathbb{E}[W(k, \mu, \omega)]$ is the true expected average sojourn time when k servers operate at rate μ , and C is a positive cost coefficient applied per server. For our experiments, we set the system parameters to create a non-trivial scenario: *Arrival Rate*: $\lambda = 2.5$ customers per unit time, *Number of Servers*: $k = 3$, *Cost Coefficient*: $C = 0.5$, *Feasible Service Rate Range*: The stability condition is $3\mu > 2.5$, or $\mu > 0.833\dots$. We define the feasible range for the individual server service rate as $\mu \in [1.0, 4.0]$.

This M/M/3 objective function retains the key characteristics relevant for testing TESO:

Stochastic / Noisy Evaluation: While complex analytical formulas exist for M/M/k steady-state probabilities (Erlang C formula), directly calculating the mean sojourn time $\mathbb{E}[W(k, \mu, \omega)]$ can be cumbersome or impractical within an optimization loop. It is typically estimated by running a discrete-event simulation of the M/M/3 system and averaging observed sojourn times. Random arrival and service processes (ω) ensure each simulation yields a noisy estimate $\hat{W}(k, \mu, \omega)$, making the evaluated objective $\hat{J}(\mu) = \frac{1}{n_{rep}} \sum_{j=1}^{n_{rep}} \hat{W}(k, \mu, \omega_j) + C \cdot k \cdot \mu^2$ inherently stochastic.

Black-Box Nature (from Optimizer's Perspective): TESO treats the simulation providing $\hat{W}(k, \mu, \omega)$ as a black box, using only the input μ and the noisy output $\hat{J}(\mu)$.

Complex Landscape: Although the underlying trade-off between reducing waiting time (higher μ) and increasing cost (higher μ) likely results in a generally convex-like shape for $J(\mu)$. Also, the noise present in the simulation-based evaluation creates a rugged landscape for the optimizer, posing a significant challenge for convergence and requiring robust handling of noisy feedback.

It is important to note that this M/M/3 problem is expected to be unimodal. While it effectively tests the algorithms' ability to handle noise and balance exploration and exploitation in a simple setting, it does not challenge their capacity to escape distinct local optima. The primary benefit of the Tabu List observed here will relate to preventing cycling in noisy regions rather than overcoming multimodality

4.3 Benchmark Algorithms for Comparison

We compare TESO's results against several benchmark algorithms to evaluate the performance and understand the contributions of the different components within the proposed TESO framework. These benchmarks are chosen to represent both a baseline level of performance and to facilitate an ablation study isolating the effects of TESO's core memory structures. The algorithms used for comparison are:

Pure Random Sampling (PRS): This serves as the most basic benchmark. In PRS, candidate solutions $x^{(t)}$ are generated purely by sampling uniformly at random from the feasible decision space \mathcal{X} for the entire duration of the optimization run (T trials). No memory or adaptive search strategy is employed. PRS provides a baseline against which the added value of any intelligent search mechanism, including TESO's, can be measured. Its performance indicates the difficulty of finding good solutions simply by chance within the given budget.

TESO without Tabu List (TESO-noTabu): This variant is identical to the full TESO algorithm (Algorithm 1) except that the Tabu List mechanism is disabled. Comparing full TESO against TESO-noTabu allows us to directly assess the impact and benefit of the short-term tabu memory on preventing cycling, encouraging diversification, and potentially improving the final solution quality or convergence speed.

TESO without Elite Memory (TESO-noElite): This variant mirrors the full TESO algorithm but removes the influence of the Elite Memory on candidate generation during the intensification phase. Instead of selecting an elite candidate x_e from \mathcal{E} and perturbing it, this variant might, for example, always perturb the single *current* best-known solution x_{best} or simply rely more heavily on the diversification mode (random generation). The comparison between full TESO and TESO-noElite aims to quantify the contribution of

the long-term Elite Memory in effectively guiding the intensification process and exploiting promising regions identified during the search.

We have intentionally omitted comparisons to other common metaheuristics like SA or GA in this study. A fair comparison would require extensive hyperparameter tuning for each method, which itself constitutes a difficult optimization problem and could obscure the specific contributions of TESO’s memory components. Our ablation study, therefore, provides a more controlled analysis of the direct impact of the Tabu List and Elite Memory, which is the primary focus of this paper.

4.4 Experimental Design

To conduct a fair and comprehensive evaluation of TESO, we define a consistent experimental setup for optimizing the M/M/3 queue problem ($k = 3, \lambda = 2.5, C = 0.5, \mu \in [1.0, 4.0]$). Each algorithm (TESO and benchmarks) is run for a total budget of $T = 300$ candidate evaluations, with the first $n_{\text{init}} = 20$ dedicated to random initialization. Candidate solutions are evaluated using $n_{\text{rep}} = 30$ independent simulation replications to estimate the mean objective and its standard deviation. For statistical robustness, all results are averaged over $N_{\text{macro}} = 30$ independent macro-replications. TESO utilizes specific parameter settings: a linearly decaying noise schedule for perturbations ($\eta_{\text{init}} = 0.2$ to $\eta_{\text{final}} = 0.01$), a Tabu List capacity $C_{\mathcal{T}} = 15$, an Elite Memory capacity $C_{\mathcal{E}} = 10$, a diversification probability $p_{\text{div}} = 0.2$, and termination if no improvement occurs in $\Delta t_{\text{max}} = 50$ iterations. The benchmark algorithms adapt this setup: *PRS* omits memory and adaptive search; *TESO-noTabu* disables the Tabu List ($C_{\mathcal{T}} = 0$); and *TESO-noElite* modifies intensification to not rely on the elite set \mathcal{E} . All algorithms operate on the same queue problem defined in Section 4.2.

Performance is assessed using several key metrics across the 30 macro-replications. We report the average and standard deviation of the *Best Mean Objective Value Found* (f_{best}^*) upon termination to gauge final solution quality and reliability. The stability of performance near the end of the search is measured by the *Average Objective Value over the Last 50 Trials*. *Convergence Plots* illustrate the progression of the average best-found solution over the trial budget, visually comparing speed and consistency (often including standard error bands). Finally, we record the average *Computational Time* per macro-replication as a practical measure of efficiency. These metrics collectively allow for a thorough comparison of the effectiveness, reliability, and efficiency of the algorithms.

4.5 Results and Discussion

The performance of the proposed TESO algorithm and the selected benchmarks are evaluated based on the queue optimization problem described in Section 4.2. The results, averaged over $N_{\text{macro}} = 30$ independent macro-replications, are presented numerically in Table 1 and visually through the convergence plots in Figure 2.

Comparison against Benchmarks As expected, PRS demonstrates the weakest performance, achieving the highest (worst) final best mean objective value (approximately 4.11) with the largest standard deviation (0.20), indicating low solution quality and poor reliability. The convergence plot for PRS shows slow improvement, confirming its inefficiency as a search strategy. In contrast, all TESO variants significantly outperform PRS. The full TESO algorithm achieves the best overall performance, converging to the lowest final best mean objective value (2.53), which is very close to the assumed true optimum (2.52). Furthermore, TESO exhibits the lowest standard deviation for the final best value (0.07) across the macro-replications, highlighting its superior reliability and consistency in finding high-quality solutions compared to all benchmarks. While TESO incurs a slightly higher average computational time compared to its ablation variants, this is justifiable given the significant improvement in solution quality and reliability. The convergence plot (Figure 2) confirms TESO’s faster convergence towards better solutions compared to the other methods.

Impact of Memory Components (Ablation Study) Comparing the full TESO algorithm with its variants allows us to analyze the specific contributions of the Tabu List and Elite Memory.

Effect of Tabu List: Comparing TESO with TESO-noElite (which has Tabu but no Elite guidance for intensification) against PRS shows the baseline benefit of using memory. Comparing TESO (full) with TESO-noTabu reveals the advantage conferred by the tabu mechanism. TESO-noTabu achieves a respectable final objective value (2.72), substantially better than PRS, demonstrating the effectiveness of elite-guided intensification. However, it performs worse than the full TESO (2.53) and exhibits higher variability (std dev 0.16 vs 0.07). This suggests that without the Tabu List, the search, while guided by elite solutions, may indeed spend unnecessary effort revisiting recently explored areas or get temporarily stuck near elite candidates, hindering diversification and slowing convergence to the very best solutions. The Tabu List effectively mitigates this.

Effect of Elite Memory: Comparing TESO with TESO-noTabu (which has Elite but no Tabu) against PRS shows the benefit of intensification. Comparing TESO (full) with TESO-noElite highlights the role of guided intensification. TESO-noElite performs better than PRS but worse than both TESO-noTabu and the full TESO (final mean obj. 2.89, std dev 0.21). This indicates that while the Tabu List helps diversification, the lack of targeted intensification based on a pool of elite solutions (instead potentially perturbing only the single best or relying more on random jumps) makes the exploitation phase less efficient. The Elite Memory is crucial for effectively focusing the search and refining high-potential regions identified during exploration. The results strongly suggest that both the short-term tabu memory and the long-term Elite Memory play vital and complementary roles in TESO’s performance. The Tabu List enhances exploration and prevents stagnation, while the Elite Memory effectively guides exploitation, leading to a synergistic effect in the full TESO algorithm that outperforms either component used in isolation (within the TESO framework).

Exploration and Exploitation Balance The convergence plots provide insights into the exploration-exploitation dynamics. PRS represents pure exploration with no learning. TESO-noElite, lacking strong guidance for exploitation, shows slower convergence after the initial phase compared to methods using Elite Memory. TESO-noTabu shows faster initial convergence due to elite guidance but potentially plateaus earlier or at a slightly higher level than full TESO, possibly due to insufficient diversification caused by revisiting areas near elite solutions. The full TESO algorithm exhibits a desirable pattern: a period of exploration (similar slope to others initially, possibly slightly slower than TESO-noTabu if elite guidance is very strong early on), followed by a phase of rapid improvement (steeper slope) as intensification guided by Elite Memory takes over, and finally converging reliably to the best solution region with low variance. This behavior suggests that the combination of tabu diversification, elite intensification, and adaptive noise successfully manages the exploration-exploitation trade-off for this noisy SO problem.

Table 1: Performance Comparison of TESO and Benchmark Algorithms.

Algorithm	Final Best Mean Obj.	Final Best Std Dev	Avg Obj. (Mean \pm Std Dev)	Avg Comp. Time (s)
PRS	4.11	0.23	4.11 \pm 0.20	53.34
TESO-noElite	2.89	0.21	2.89 \pm 0.18	101.43
TESO-noTabu	2.72	0.16	2.72 \pm 0.13	113.52
TESO	2.53	0.07	2.53 \pm 0.06	132.81

5 CONCLUSION AND FUTURE WORK

This paper introduced TESO, a novel metaheuristic framework tailored for optimizing noisy, expensive black-box simulation models. *TESO’s methodology* combines adaptive random search with a short-term Tabu List to prevent cycling and a long-term Elite Memory to guide intensification around promising solutions, employing an aspiration criterion and adaptive noise control. Our *key experimental results* on a queue optimization problem demonstrated TESO’s effectiveness, showing consistent convergence to high-quality solutions with better reliability (lower variance) compared to PRS and ablation variants

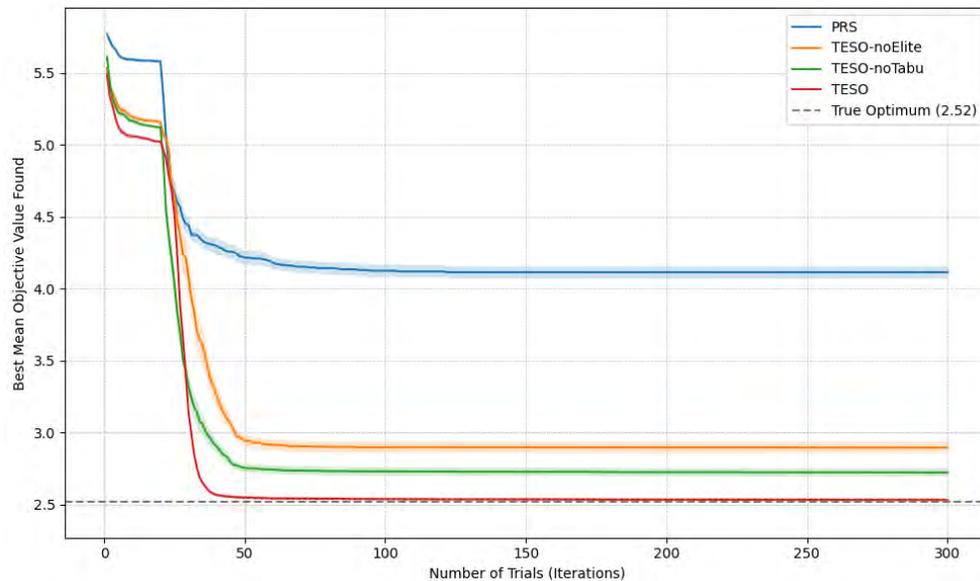


Figure 2: Convergence Plot for Queue Optimization

lacking either memory component. The study confirmed the synergistic benefit of integrating both tabu restrictions and elite guidance, leading to a *successful balance between exploration and exploitation*. The primary *contributions* include this novel integration of TS principles and Elite Memory for stochastic optimization, the empirical demonstration of improved performance, and a structured approach to managing the exploration-exploitation trade-off. However, we acknowledge *limitations* such as validation on a single, unimodal problem class, potential sensitivity to TESO's internal parameters, and untested scalability in very high-dimensional settings. Furthermore, while we criticize other metaheuristics for their sensitivity to hyperparameters, TESO is not immune. Its performance depends on parameters and the noise decay schedule. A full sensitivity analysis was beyond the scope of this initial study but is a critical direction for future work. Finally, the algorithm's scalability in very high-dimensional settings remains untested.

Future work will focus on three key areas. First, we will test TESO's generalizability on a wider range of benchmark problems with higher dimensions and multimodality. Second, we aim to enhance the algorithm by developing adaptive mechanisms for its internal parameters and exploring hybridization with surrogate models. Finally, pursuing a theoretical analysis of its convergence properties would provide important formal grounding.

REFERENCES

- Amaran, S., N. V. Sahinidis, B. Sharda, and S. J. Bury. 2016. "Simulation optimization: a review of algorithms and applications". *Annals of Operations Research* 240:351–380.
- Andradóttir, S. 1998. "Simulation optimization". *Handbook of simulation*:307–334.
- Boesel, J., B. L. Nelson, and S.-H. Kim. 2003. "Using ranking and selection to "clean up" after simulation optimization". *Operations Research* 51(5):814–825.
- Brahim, B., M. Kobayashi, M. Al Ali, T. Khatir, and M. E. A. E. Elmeliiani. 2024. "Metaheuristic optimization algorithms: An overview". *HCMCOU Journal of Science–Advances in Computational Structures*.
- Cao, H., J.-Q. Hu, and T. Lian. 2024. "Blackbox Simulation Optimization". *Journal of the Operations Research Society of China*:1–27.
- Chen, C.-H., and L. H. Lee. 2011. *Stochastic simulation optimization: an optimal computing budget allocation*, Volume 1. World scientific.
- de Carvalho Miranda, R., J. A. B. Montevechi, and A. F. de Pinho. 2015. "Development of an adaptive genetic algorithm for simulation optimization". *Acta Scientiarum. Technology* 37(3):321–328.

- Deng, G. 2007. *Simulation-based optimization*. Ph. D. thesis, University of Wisconsin–Madison.
- Fan, W., L. J. Hong, G. Jiang, and J. Luo. 2024. “Review of Large-Scale Simulation Optimization”. *arXiv preprint arXiv:2403.15669*.
- Figueira, G., and B. Almada-Lobo. 2014. “Hybrid simulation–optimization methods: A taxonomy and discussion”. *Simulation modelling practice and theory* 46:118–134.
- Fu, M. 1994. “A tutorial review of techniques for simulation optimization”. In *Proceedings of Winter Simulation Conference*, 149–156 <https://doi.org/10.1109/WSC.1994.717096>.
- Glover, F. 1990. “Tabu search: A tutorial”. *Interfaces* 20(4):74–94.
- Gosavi, A. et al. 2015. *Simulation-based optimization*, Volume 62. Springer.
- Haddock, J., and J. Mittenhal. 1992. “Simulation optimization using simulated annealing”. *Computers & industrial engineering* 22(4):387–395.
- Hong, L. J., W. Fan, and J. Luo. 2021. “Review on ranking and selection: A new perspective”. *Frontiers of Engineering Management* 8(3):321–343.
- Hong, L. J., and X. Zhang. 2021. “Surrogate-based simulation optimization”. In *Tutorials in Operations Research: Emerging Optimization Methods and Modeling Techniques with Applications*, 287–311. INFORMS.
- Jung, J. Y., G. Blau, J. F. Pekny, G. V. Reklaitis, and D. Eversdyk. 2004. “A simulation based optimization approach to supply chain management under demand uncertainty”. *Computers & chemical engineering* 28(10):2087–2106.
- Kuo, R., and C. Yang. 2011. “Simulation optimization using particle swarm optimization algorithm with application to assembly line design”. *Applied Soft Computing* 11(1):605–613.
- Lal, T. M., T. Roh, and T. Huschka. 2015. “Simulation based optimization: Applications in healthcare”. In *2015 winter simulation conference (WSC)*, 1261–1271. IEEE.
- Nezami, N., and H. Anahideh. 2023. “Hyperparameter Adaptive Search for Surrogate Optimization: a Self-Adjusting Approach”. In *2023 Winter Simulation Conference (WSC)*, 3484–3495 <https://doi.org/10.1109/WSC60868.2023.10408315>.
- Nsiye, E., T. Wright, B. Tse, and S. Mondesire. 2024, May. “A Micro-Discrete Event Simulation Environment for Production Scheduling in Manufacturing Digital Twins”. In *MODSIM World*. Norfolk, Virginia.
- Pearce, M., M. Poloczek, and J. Branke. 2019. “Bayesian Simulation Optimization with Common Random Numbers”. In *2019 Winter Simulation Conference (WSC)*, 3492–3503 <https://doi.org/10.1109/WSC40007.2019.9004680>.
- Soykan, B., and G. Rabadi. 2016. “A Tabu Search Algorithm for the Multiple Runway Aircraft Scheduling Problem”. In *Heuristics, Metaheuristics and Approximate Methods in Planning and Scheduling, International Series in Operations Research & Management Science*, edited by G. Rabadi, 165–186. Cham: Springer International Publishing.
- Soykan, B., and G. Rabadi. 2022. “A simulation-based optimization approach for multi-objective runway operations scheduling”. *SIMULATION: Transactions of the Society for Modeling and Simulation International* 98(11):991–1012.
- Soykan, B., and G. Rabadi. 2024. “Optimizing Job Shop Scheduling Problem Through Deep Reinforcement Learning and Discrete Event Simulation”. In *Proceedings of the 2024 Winter Simulation Conference (WSC’24)*, 2607–2618.
- Tekin, E., and I. Sabuncuoglu. 2004. “Simulation optimization: A comprehensive review on theory and applications”. *IIE Transactions* 36(11):1067–1081 <https://doi.org/10.1080/07408170490500654>.
- Thengvall, B. G., S. N. Hall, and M. P. Deskevich. 2025. “Measuring the effectiveness and efficiency of simulation optimization metaheuristic algorithms”. *Journal of Heuristics* 31(1):1–21.
- Yu, C., N. Lahrichi, and A. Matta. 2023. “Optimal budget allocation policy for tabu search in stochastic simulation optimization”. *Computers & Operations Research* 150:106046.

AUTHOR BIOGRAPHIES

BULENT SOYKAN is an Associate Research Scientist in the Institute for Simulation and Training at the University of Central Florida (UCF). He received his Ph.D. from Old Dominion University. His research interests include real-time decision-making under uncertainty, combinatorial optimization, digital twins, reinforcement learning. His email address is bulent.soykan@ucf.edu and his website is <https://www.ist.ucf.edu/faculty/bulent-soykan-ph-d/>.

SEAN MONDESIRE is an Assistant Professor in the School of Modeling, Simulation, and Training at UCF. His research interests include real-time decision-making through machine learning, big data analytics, scalable digital twins, and simulations. His email address is sean.mondesire@ucf.edu and his website is <https://www.ist.ucf.edu/faculty/sean-mondesire-ph-d/>.

GHAITH RABADI is a Professor and Director of the School of Modeling, Simulation, and Training at UCF. His research interests include Modeling & Simulation, Operation Research, and AI in different areas of applications including production, supply chains, transportation, and disaster response logistics and informatics. His email address is ghaith.rabadi@ucf.edu and his website is <https://www.ist.ucf.edu/faculty/ghaith-rabadi-ph-d/>.