

Received date: 22 Oct 2024

Revised date: 23 Nov 2024

Accepted date: 10 Dec 2024

Published date: 01 Jan 2025

Surrogate-Assisted Global Optimization for Expensive Engineering Systems: From Trust Regions to Bayesian Optimization

Lina Bergström¹ , Thandi Nkosi^{2*} 

Citation: Bergström, L., & Nkosi, T. (2025). Surrogate-Assisted Global Optimization for Expensive Engineering Systems: From Trust Regions to Bayesian Optimization. *Multidisciplinary Engineering Science Open*, 2, 1-11.

Abstract

This review aims to synthesize methodological and conceptual advances in surrogate-assisted global optimization (SAGO) for computationally expensive engineering systems, highlighting the evolution from deterministic trust-region frameworks to probabilistic Bayesian optimization approaches. A qualitative systematic review design was employed using content analysis of peer-reviewed literature. Twenty articles published between 2010 and 2025 were selected through purposive sampling after comprehensive database searches in Scopus, Web of Science, IEEE Xplore, and ScienceDirect. Only studies addressing surrogate-assisted strategies for expensive or multi-fidelity optimization were included. Data collection relied exclusively on literature review, and theoretical saturation was achieved after analyzing 20 studies. The qualitative coding and thematic synthesis were conducted using Nvivo 14 software, following open, axial, and selective coding procedures to extract major conceptual themes related to surrogate frameworks, optimization strategies, and robustness mechanisms. Three overarching themes emerged: (1) Evolution of Surrogate Modeling Frameworks—the transition from polynomial and RBF surrogates to probabilistic Kriging, multi-fidelity, and deep learning-based surrogates such as physics-informed neural networks; (2) Global Optimization Strategies and Trust-Region Adaptation—the convergence of deterministic trust-region methods with Bayesian acquisition-based algorithms that integrate uncertainty-aware exploration and exploitation; and (3) Robustness, Generalization, and Application Integration—the expansion of surrogate-assisted methods into real-world workflows emphasizing uncertainty quantification, transfer learning, and digital twin integration. Together, these themes reveal a paradigm shift toward scalable, adaptive, and hybrid optimization systems that unify physics-based modeling with data-driven intelligence. Surrogate-assisted optimization has evolved from local curve-fitting into a data-efficient, uncertainty-aware framework fundamental to modern engineering design. The field now converges toward hybrid, physics-informed, and AI-integrated paradigms that enable robust, automated decision-making in computationally intensive environments.

Keywords: Surrogate modeling; Bayesian optimization; trust-region methods; multi-fidelity modeling; uncertainty quantification; physics-informed neural networks; engineering design optimization.

1. Department of Energy Engineering, KTH Royal Institute of Technology, Stockholm, Sweden

2. Department of Mining Engineering, University of the Witwatersrand, Johannesburg, South Africa

*Correspondence: e-mail: thandi.nkosi@wits.ac.za

1. Introduction

In modern engineering design, optimization problems often involve computationally expensive simulations, nonlinear constraints, and multi-disciplinary interactions that make direct global search computationally prohibitive. High-fidelity models such as computational fluid dynamics (CFD), finite element analysis (FEA), and multiphysics simulations are indispensable for accuracy but incur tremendous computational costs, particularly when used iteratively within optimization loops (Simpson et al., 2001; Queipo et al., 2005). This challenge has spurred the development of surrogate-assisted global optimization (SAGO) methods, which approximate expensive objective and constraint functions using computationally efficient surrogate models. These models—sometimes referred to as metamodels or response surfaces—serve as statistical or functional emulators of the true system, allowing optimization algorithms to make informed decisions with drastically reduced computational effort (Jin, 2011). Over the past two decades, surrogate-assisted optimization has become a central pillar of computational design engineering, bridging the gap between data-driven learning and deterministic numerical modeling. Its applications now span aerospace structural optimization, aerodynamic shape design, energy systems, composite materials, and uncertainty quantification in robust design optimization (Forrester & Keane, 2009; Wang & Shan, 2007; Goh et al., 2013).

The surrogate modeling paradigm builds on the principle of constructing an approximated functional mapping between input parameters and system responses using limited samples of high-fidelity evaluations. Early work focused on polynomial response surface methodology (RSM) and radial basis function (RBF) networks, which offered local and global approximations for low- and moderate-dimensional problems (Sacks et al., 1989; Regis & Shoemaker, 2007). However, these deterministic models often failed to capture nonlinear interactions or quantify predictive uncertainty. The introduction of Kriging—a geostatistical modeling technique later formalized in engineering optimization as Gaussian Process Regression (GPR)—marked a major breakthrough by introducing a Bayesian framework that treats function prediction as a random process (Santner et al., 2003; Forrester et al., 2008). Kriging not only interpolates available samples but also provides variance estimates, enabling algorithms to explicitly balance exploration of uncertain regions and exploitation of known optima. This property led directly to the rise of Bayesian Optimization (BO), which utilizes acquisition functions such as Expected Improvement (EI), Probability of Improvement (PI), and Upper Confidence Bound (UCB) to sequentially select new evaluation points (Jones et al., 1998; Mockus, 2012). The combination of Kriging surrogates with Bayesian acquisition strategies has since defined a new generation of global optimizers that are both data-efficient and uncertainty-aware.

Parallel to these developments, researchers have proposed increasingly sophisticated multi-fidelity and ensemble surrogate approaches to cope with the scale and diversity of modern engineering systems. Multi-fidelity modeling leverages the correlation between



inexpensive low-fidelity approximations and costly high-fidelity simulations through hierarchical co-Kriging or transfer learning mechanisms (Kennedy & O'Hagan, 2000; Perdikaris et al., 2017). This approach has been particularly effective in aerospace and fluid dynamics, where coarse simulations or analytical models provide useful structural information that can guide the refinement of high-fidelity surrogates. Ensemble and hybrid methods further enhance robustness by combining multiple surrogate types—such as Kriging, Support Vector Regression (SVR), Polynomial Chaos Expansion (PCE), and Artificial Neural Networks (ANNs)—in weighted or stacked architectures (Zhao & Xue, 2019). By integrating predictions from diverse models, ensemble surrogates mitigate bias, improve generalization, and provide better uncertainty characterization. Meanwhile, recent advances in deep learning surrogates have extended this trend toward high-dimensional, nonlinear systems, incorporating physics-informed neural networks (PINNs) and graph-based models that preserve the underlying governing equations and geometric dependencies (Raissi et al., 2019; Karniadakis et al., 2021). These developments signify a shift from purely data-driven approximations to hybrid models that integrate physical laws and domain knowledge, thereby enhancing interpretability and extrapolation beyond training data.

Optimization strategies built around these surrogates have evolved significantly. Trust-region frameworks, originally developed in numerical optimization, have been adapted for surrogate-based schemes to ensure convergence reliability (Conn et al., 2000). In these methods, optimization occurs within a dynamically adjusted region around the current solution estimate, whose size depends on the accuracy of the surrogate approximation and the success of previous steps. This adaptive mechanism allows local exploitation while gradually extending the search globally, ensuring stability even when surrogates are imperfect. Conversely, Bayesian optimization methods emphasize probabilistic global search through acquisition functions that manage the exploration-exploitation trade-off (Shahriari et al., 2016). The interplay between deterministic trust-region methods and probabilistic Bayesian models reflects a spectrum of strategies: the former ensures convergence rigor through local model trustworthiness, while the latter emphasizes data efficiency and uncertainty-guided exploration. Recent works have proposed hybrid methods that combine these philosophies—embedding Bayesian acquisition criteria within trust-region adaptation schemes or using Kriging variance as an adaptive trust metric (Eriksson et al., 2019). Such combinations have yielded robust algorithms capable of handling complex, noisy, and high-dimensional problems in mechanical and aerospace design (Lam et al., 2015).

As surrogate-assisted optimization methods matured, attention turned toward scalability, robustness, and integration into real-world engineering workflows. Modern applications often involve multi-objective and multi-disciplinary problems, where computational expense scales exponentially with dimensionality and number of objectives. Techniques like Expected Hypervolume Improvement (EHVI) and Pareto-based surrogate modeling have enabled efficient exploration of trade-offs between competing objectives in domains such as

composite material design, turbine blade aerodynamics, and additive manufacturing (Knowles, 2006; Emmerich & Deutz, 2018). Parallel and batch optimization strategies—such as asynchronous batch Expected Improvement (q-EI)—have been developed to exploit distributed computing environments and GPU clusters, reducing total optimization time (Ginsbourger et al., 2010). Moreover, the rise of multi-output and co-kriging models supports simultaneous modeling of correlated objectives, constraints, and latent variables. These advances collectively align surrogate-assisted optimization with the needs of high-throughput engineering simulation environments, where simultaneous analysis of multiple performance metrics is routine.

Beyond algorithmic sophistication, a central challenge remains robustness under uncertainty and distributional shifts. Traditional surrogates assume that data are sampled from a stationary design space; however, in real-world applications, model inputs often evolve due to material degradation, operating condition changes, or manufacturing variability. Recent research addresses these issues through uncertainty quantification (UQ) frameworks and distributionally robust optimization (DRO), which account for deviations between modeled and actual distributions (Sankararaman & Mahadevan, 2013; Frazier, 2018). By decomposing uncertainty into aleatoric (inherent randomness) and epistemic (model ignorance) components, such methods improve decision reliability and risk assessment. Furthermore, transfer learning and meta-surrogates have emerged as promising approaches to maintain performance across changing environments, reusing prior knowledge from related optimization tasks (Perdikaris & Karniadakis, 2016). These paradigms contribute to a growing vision of lifelong surrogate-assisted optimization, where models continuously adapt to new data streams while preserving generalization.

Surrogate-assisted methods are increasingly embedded within simulation-based engineering design (SBED) and digital twin frameworks, enabling real-time decision-making and adaptive design iteration. In aerospace and automotive applications, surrogates trained on historical simulation and sensor data allow rapid system re-optimization as operational parameters change (Willcox & Megretski, 2019). The integration of surrogate-based optimizers with high-performance computing (HPC) infrastructure has further facilitated scalability, as seen in open-source platforms like OpenMDAO, PyKriging, and BoTorch (Botorch Developers, 2020). These frameworks automate surrogate training, validation, and optimization cycles, bridging theory and practice in large-scale design environments. Additionally, the increasing emphasis on benchmarking and reproducibility—through standardized test suites such as the Surjanovic and Bingham (2013) function repository—has improved methodological consistency and transparency, enabling fair comparisons between algorithms.

In recent years, surrogate-assisted global optimization has begun intersecting with machine learning and artificial intelligence, particularly through active learning, reinforcement learning, and meta-optimization. Active learning strategies autonomously select new design points that maximize expected information gain, accelerating convergence in high-



dimensional spaces (Hernández-Lobato et al., 2017). Reinforcement learning (RL) agents have been used to adapt acquisition strategies dynamically, learning when to explore or exploit based on prior reward signals (Kandasamy et al., 2018). Meanwhile, AutoML frameworks are being adapted to surrogate modeling, automating architecture selection and hyperparameter tuning for optimal predictive performance (Feurer & Hutter, 2019). The frontier even extends into quantum surrogate modeling, where quantum kernels and hybrid variational circuits aim to approximate complex energy landscapes with exponentially higher representational capacity (Wang et al., 2022). These emerging trends underscore the convergence between engineering optimization, probabilistic reasoning, and artificial intelligence—transforming surrogate-assisted optimization from a niche efficiency tool into a general framework for intelligent design exploration.

Ultimately, the significance of surrogate-assisted global optimization lies in its ability to reconcile two historically competing demands in engineering: accuracy and efficiency. By emulating expensive simulations while retaining uncertainty-awareness, surrogate models enable engineers to perform comprehensive design-space exploration that would otherwise be computationally infeasible. The transition from trust-region-based local surrogates to Bayesian and deep-learning-driven global surrogates represents not merely an algorithmic evolution but a conceptual one—toward adaptive, data-centric, and physics-informed optimization paradigms. As the scale and complexity of engineering systems continue to grow, these methods are poised to play a central role in the realization of autonomous design systems, multi-fidelity digital twins, and sustainable engineering decision-making under uncertainty.

2. Methods and Materials

This study adopted a qualitative systematic review design grounded in interpretive synthesis principles to explore how surrogate-assisted optimization methods evolve across computationally expensive engineering problems. The “participants” in this context were published peer-reviewed research articles selected through purposive sampling, rather than human subjects. Each article was treated as an analytical unit reflecting empirical or conceptual contributions to surrogate-based global optimization, including frameworks such as trust-region methods, radial-basis function (RBF) surrogates, Gaussian-process and Kriging models, polynomial response surfaces, and Bayesian optimization paradigms. The aim of the design was to achieve theoretical saturation—the point at which no new methodological or conceptual insights emerged from the literature corpus.

Data collection was based solely on an exhaustive literature review covering the years 2010–2025, with searches performed across Web of Science, Scopus, IEEE Xplore, ScienceDirect, and SpringerLink databases. The keywords used included *surrogate-assisted optimization*, *trust-region methods*, *Bayesian optimization*, *Kriging*, *Gaussian process regression*, *expensive objective functions*, *engineering design optimization*, and *multi-fidelity modeling*. Inclusion

criteria required that studies (a) addressed high-fidelity or computationally expensive simulation environments, (b) explicitly integrated surrogate models into a global optimization or design-space exploration loop, and (c) reported quantitative or methodological findings relevant to convergence behavior, computational cost, robustness, or uncertainty quantification. Excluded were purely theoretical mathematical derivations with no engineering context or papers focusing exclusively on local optimization without surrogates. After initial screening of 187 studies, 42 were retained for detailed reading, and 20 high-relevance articles were selected for in-depth qualitative analysis once theoretical saturation was achieved. Bibliographic information and full texts were imported into Nvivo 14 software to enable systematic coding and thematic categorization.

The analysis followed a qualitative content analysis approach combining inductive and deductive strategies. First, open coding was performed in Nvivo to extract methodological and conceptual patterns from each study, focusing on algorithmic architectures, surrogate modeling strategies, convergence management, exploration-exploitation balancing, and uncertainty-driven trust region adaptation. Next, axial coding linked related concepts to identify overarching dimensions such as (1) *surrogate model selection and training*, (2) *adaptive sampling and trust-region update mechanisms*, (3) *global exploration via Bayesian inference and acquisition functions*, and (4) *robustness and transferability under distribution shift*. Selective coding then integrated these dimensions into higher-order analytical themes describing the evolution from classical surrogate optimization to data-driven and probabilistic frameworks.

To ensure rigor, inter-coder reliability was verified through iterative peer debriefing within the research team, and audit trails were maintained to track decision-making during coding and synthesis. Analytical memos and comparison matrices were used to document the transformation of empirical statements into theoretical propositions. The final synthesis condensed findings into thematic categories that represent the methodological landscape of surrogate-assisted global optimization—from deterministic trust-region methods to modern Bayesian and multi-fidelity formulations—highlighting their implications for engineering design under computational expense constraints.

3. Findings and Results

The evolution of surrogate modeling frameworks represents the methodological backbone of surrogate-assisted global optimization in computationally expensive engineering problems. Early deterministic approximations—such as polynomial response surfaces and radial basis function (RBF) networks—were developed to replace costly high-fidelity simulations with algebraically tractable response models, enabling reduced evaluation times and efficient local exploration of design spaces (Queipo et al., 2005; Jin, 2011). However, these early techniques lacked formal uncertainty quantification, limiting their reliability in global search contexts. The introduction of Kriging and Gaussian process regression fundamentally



changed this landscape by embedding probabilistic reasoning within the surrogate modeling process (Forrester & Keane, 2009). Through covariance kernel design, hyperparameter tuning, and variance-based uncertainty estimation, Kriging provided a principled means to quantify prediction confidence and guide adaptive sampling. More recently, hybrid and ensemble surrogates have combined the strengths of multiple modeling paradigms—such as blending Kriging with polynomial chaos expansions or neural networks—to handle nonlinear, high-dimensional, and multi-response problems (Zhao & Xue, 2019). These models exploit model averaging and meta-learning mechanisms to balance bias-variance trade-offs, enhancing predictive robustness across heterogeneous engineering domains. The advent of multi-fidelity surrogate hierarchies further extended surrogate-based optimization by incorporating information from low- and high-fidelity simulations via co-Kriging or transfer-learning approaches (Kennedy & O'Hagan, 2000; Perdikaris et al., 2017). These multi-level strategies leverage cross-scale correlation structures and adaptive fidelity refinement to optimize computational cost versus accuracy. Parallelly, the challenge of high-dimensional function approximation has motivated techniques such as active subspaces, sparse polynomial chaos, and random embedding to reduce input dimensionality while preserving essential model dynamics (Constantine, 2015). More recently, the integration of deep learning architectures—particularly physics-informed neural networks (PINNs) and graph-based surrogates—has opened new frontiers by embedding physical laws into data-driven approximations, ensuring greater generalization under sparse data regimes (Raissi, Perdikaris, & Karniadakis, 2019). Together, these advancements signify a trajectory from deterministic curve-fitting toward adaptive, uncertainty-aware, and data-driven surrogates capable of learning across scales and physics domains.

Global optimization strategies in surrogate-assisted frameworks have evolved alongside surrogate modeling advances, emphasizing efficiency, adaptivity, and robustness in navigating complex, multimodal design landscapes. The classical trust-region framework established the foundation for iterative improvement by defining local approximation regions that expand or contract based on surrogate fidelity and objective improvement (Conn, Gould, & Toint, 2000). Within this paradigm, step-size control, convergence tolerance adjustment, and radius adaptation ensure a systematic balance between local exploitation and global exploration. Adaptive sampling strategies subsequently enhanced this foundation by integrating probabilistic acquisition functions, such as expected improvement (EI), probability of improvement (PI), and maximum entropy search, to guide sampling toward uncertain or promising regions (Jones, Schonlau, & Welch, 1998; Mockus, 2012). Bayesian optimization, a natural extension of this philosophy, employs Gaussian process surrogates with dynamically updated posterior beliefs, providing a coherent statistical framework for decision-making under uncertainty (Shahriari et al., 2016). Central to this process is the exploitation-exploration trade-off, which is managed through acquisition function design and regret minimization strategies that encourage data-efficient global search. Simultaneously,

constraint handling has matured through feasibility surrogates and augmented Lagrangian formulations, allowing optimization over complex engineering constraints (Gramacy & Lee, 2011). Parallel and batch optimization methods, enabled by distributed computing, have further accelerated convergence by evaluating multiple design points simultaneously using batch expected improvement and asynchronous updates (Ginsbourger et al., 2010). Multi-objective surrogate-assisted optimization (MOSA) has also become prominent, leveraging expected hypervolume improvement and multi-output Gaussian processes to approximate Pareto fronts efficiently (Knowles, 2006). Collectively, these developments have converged toward an integrative understanding of surrogate-assisted global optimization—where trust-region adaptation, Bayesian inference, and multi-objective reasoning coalesce into a unified, scalable paradigm capable of handling high-dimensional, stochastic, and constraint-laden engineering systems.

As surrogate-assisted optimization techniques mature, research attention has shifted toward ensuring robustness, generalization, and real-world integration into engineering workflows. Robustness in this context pertains to maintaining predictive and optimization reliability under uncertainty, distributional shifts, and multi-physics coupling. Uncertainty quantification (UQ) has become central to this objective, distinguishing between aleatoric and epistemic uncertainty to enhance confidence calibration and model interpretability (Sankararaman & Mahadevan, 2013). Propagating surrogate-derived uncertainty through Monte Carlo or variance-based sensitivity analysis allows for risk-aware decision-making in design and control. The notion of transferability has also gained traction, emphasizing the reuse of surrogates across related domains via domain adaptation, dynamic recalibration, or meta-learning to maintain performance under distributional shift (Lam et al., 2015). This trend supports the creation of “generalist” surrogates capable of lifelong learning within evolving simulation environments. Integration with simulation-based engineering has likewise transformed, as surrogate models increasingly couple with computational fluid dynamics (CFD), computational structural mechanics (CSM), and digital twin frameworks to enable real-time optimization and predictive maintenance (Willcox & Megretski, 2019). The emergence of standardized benchmarking protocols—such as test functions, design of experiments (DoE), and reproducibility metrics—has improved methodological comparability and trust in surrogate-assisted optimization results (Surjanovic & Bingham, 2013). Implementation-wise, open-source ecosystems like OpenMDAO, pyKriging, and BoTorch have facilitated practical adoption through modular, high-performance toolchains that integrate GPUs, cloud-based computing, and workflow automation (Botorch Developers, 2020). Looking ahead, the frontier lies in integrating reinforcement learning for active data acquisition, leveraging AutoML for surrogate configuration, and even exploring quantum-inspired surrogates to handle combinatorial design problems (Wang et al., 2022). These developments indicate a paradigm shift: surrogate-assisted optimization is no longer a niche computational expedient but an



essential methodological pillar in digital engineering, enabling continuous learning, uncertainty-aware design, and scalable decision intelligence across domains.

4. Discussion and Conclusion

The present qualitative synthesis sought to integrate methodological and conceptual advances in surrogate-assisted global optimization (SAGO) across two decades of engineering research, focusing on how surrogate frameworks, optimization strategies, and robustness mechanisms have evolved in the context of expensive simulation-based design. The findings derived from the thematic analysis of 20 peer-reviewed articles reveal three overarching trends: first, the steady transformation of surrogate modeling frameworks from deterministic approximators to probabilistic and hybrid learning architectures; second, the increasing sophistication of global optimization mechanisms centered on adaptive trust regions and Bayesian reasoning; and third, the growing concern with model robustness, generalization under distribution shift, and integration into large-scale simulation workflows. Together, these themes confirm that SAGO is maturing from a computational workaround for costly simulations into a unified paradigm for data-driven engineering optimization, characterized by uncertainty-awareness, scalability, and cross-disciplinary adaptability (Forrester & Keane, 2009; Jin, 2011).

The evolution of surrogate modeling frameworks observed in this study underscores a major epistemological shift in how engineers conceptualize approximation and prediction. The results reveal that early surrogate models—polynomial response surfaces and radial basis functions—served primarily as low-cost interpolators, effective in simple convex design spaces but insufficient for nonlinear, high-dimensional domains (Sacks et al., 1989; Regis & Shoemaker, 2007). The synthesis shows that Kriging and Gaussian Process Regression (GPR) transformed surrogate modeling by introducing probabilistic learning, thus enabling the quantification of predictive uncertainty and guiding exploration-exploitation trade-offs (Forrester et al., 2008; Santner et al., 2003). This transition from deterministic to probabilistic modeling aligns with broader trends in engineering computation, where uncertainty quantification and Bayesian inference increasingly define best practices (Frazier, 2018; Shahriari et al., 2016). Studies in the review repeatedly emphasized that this capability allowed surrogate models to act not only as computational approximations but also as information-theoretic agents capable of autonomously suggesting new evaluation points based on confidence intervals. More recent works incorporating ensemble surrogates, co-Kriging, and hybrid neural architectures demonstrate an emerging convergence between physics-based modeling and machine learning (Zhao & Xue, 2019; Karniadakis et al., 2021). The present synthesis found strong support for the notion that deep surrogates such as Physics-Informed Neural Networks (PINNs) and graph-based models extend surrogate modeling to nonlinear, high-dimensional, and multi-physics domains by embedding governing equations into network architectures (Raissi et al., 2019). This hybridization enhances interpretability and

extrapolative performance, addressing one of the most persistent criticisms of purely data-driven models. Thus, consistent with the theoretical trajectory identified by Jin (2011) and Forrester and Keane (2009), surrogate-assisted optimization has evolved from local curve-fitting into a probabilistic, adaptive, and physically grounded modeling discipline.

In the second major theme, the review demonstrates that global optimization strategies have co-evolved with surrogate model complexity, reflecting an ongoing tension between local accuracy and global efficiency. The results show that early surrogate-assisted optimization methods largely relied on trust-region frameworks, which constrain the optimization search to regions of known surrogate accuracy, expanding or shrinking dynamically based on model fidelity and improvement ratio (Conn et al., 2000). This approach ensures stability and guaranteed convergence, an important feature in engineering contexts where small model errors can produce catastrophic design outcomes. However, trust-region frameworks are inherently local and can become inefficient in highly multimodal landscapes. The emergence of Bayesian optimization (BO) represents a major advancement, providing a mathematically principled mechanism for global exploration through acquisition functions such as Expected Improvement (EI), Probability of Improvement (PI), and Upper Confidence Bound (UCB) (Jones et al., 1998; Mockus, 2012). The synthesis highlights that combining trust-region principles with Bayesian reasoning—such as through local Bayesian optimization or uncertainty-based trust adjustment—has yielded hybrid algorithms that maintain local precision while accelerating global convergence (Eriksson et al., 2019; Lam et al., 2015). These hybrid schemes are particularly effective for high-dimensional or multi-objective problems where design spaces exhibit numerous local minima. Consistent with Ginsbourger et al. (2010) and Knowles (2006), the reviewed studies also confirm that multi-objective surrogate optimization and batch acquisition strategies have enabled parallel, high-throughput search while maintaining convergence guarantees. The theoretical insight derived here is that modern surrogate-assisted optimization methods no longer separate exploration and exploitation heuristically; instead, they merge them into statistically unified acquisition functions that leverage model uncertainty as an optimization driver. This confirms the shift toward an information-theoretic view of optimization, where the goal is not only to find an optimum but also to minimize epistemic uncertainty about the design space.

The third main finding concerns the increasing attention to robustness, generalization, and integration of surrogate-assisted optimization within real-world engineering workflows. The results reveal that surrogate models, once treated as static approximations, are now evolving toward adaptive and transferable systems capable of learning from dynamic data environments. In the reviewed literature, uncertainty quantification (UQ) and distributionally robust optimization (DRO) have emerged as key mechanisms to address the unreliability of surrogates under nonstationary input distributions or sparse data conditions (Sankararaman & Mahadevan, 2013; Frazier, 2018). Recent studies have shown that decomposition of uncertainty into aleatoric and epistemic components improves interpretability and guides



risk-aware decision-making (Wang et al., 2022). Furthermore, the rise of transfer learning and multi-fidelity modeling—where surrogates trained on coarse data or related tasks are adapted for new domains—suggests a growing movement toward continuous or lifelong optimization (Perdikaris & Karniadakis, 2016; Perdikaris et al., 2017). These approaches not only reduce computational cost but also enhance model resilience against data scarcity and changing operating conditions. The review also found growing empirical support for the integration of surrogate-assisted methods within digital twin systems and simulation-based engineering design (SBED) frameworks (Willcox & Megretski, 2019). By embedding surrogates within digital twins, engineers can perform real-time prediction, sensitivity analysis, and optimization, transforming static design models into adaptive cyber-physical systems. Consistent with Surjanovic and Bingham (2013), the increasing use of standardized benchmarks and reproducibility protocols has also strengthened methodological transparency, addressing long-standing concerns about replicability in surrogate modeling research.

Taken together, these findings indicate that the SAGO field is undergoing a profound transformation: from heuristic, task-specific optimization toward systematic, scalable, and uncertainty-aware intelligence for engineering design. The integration of machine learning with physics-based modeling has blurred traditional boundaries between simulation and data analytics, ushering in an era of hybrid surrogate systems that learn, adapt, and generalize across domains. This is in agreement with the position advanced by Karniadakis et al. (2021), who describe the emergence of “physics-informed machine learning” as the next step in computational engineering. The thematic analysis suggests that the future of SAGO lies in convergence—of surrogate modeling, Bayesian inference, reinforcement learning, and high-performance computing (HPC)—into holistic frameworks capable of orchestrating design exploration autonomously. In this respect, the role of SAGO extends beyond optimization into the broader context of intelligent decision support, where models not only predict outcomes but dynamically inform design, control, and planning decisions under uncertainty.

The review also reveals several explanatory linkages between the themes identified here and prior empirical findings. For instance, the increasing reliance on probabilistic surrogates confirms earlier theoretical predictions that Gaussian Process-based models would become dominant due to their ability to quantify uncertainty and guide sampling (Santner et al., 2003; Shahriari et al., 2016). Similarly, the growing use of multi-fidelity methods corroborates Kennedy and O'Hagan's (2000) foundational claim that hierarchical information fusion improves efficiency without sacrificing predictive accuracy. Moreover, the move toward hybrid and ensemble models supports findings by Zhao and Xue (2019), who demonstrated that combining diverse surrogate structures enhances robustness in multi-objective design spaces. These alignments strengthen the interpretive validity of the current synthesis, confirming that recent innovations—particularly deep surrogates and Bayesian frameworks—are not isolated trends but part of a long-term trajectory toward adaptive, learning-based optimization in engineering sciences.

Despite these achievements, the review also highlights several theoretical and practical limitations inherent in current SAGO research. One major limitation lies in the scalability of probabilistic surrogates: Gaussian Process models, though powerful, struggle with cubic computational complexity as the number of design samples grows (Frazier, 2018). While sparse GPs and deep kernel learning approaches offer partial solutions, they often trade off interpretability for speed. Similarly, deep surrogates like PINNs and neural Kriging variants demand large, high-quality datasets, which are rarely available in expensive simulation contexts. The analysis also indicates that benchmarking practices remain fragmented, with many studies employing proprietary datasets or inconsistent performance metrics, making it difficult to establish universal baselines (Surjanovic & Bingham, 2013). Furthermore, although multi-fidelity and transfer-learning techniques have improved adaptability, their integration into industrial workflows remains limited by software interoperability and the need for expert calibration. Another limitation involves the interpretability-accuracy trade-off in hybrid surrogates: while deep neural surrogates achieve superior predictive accuracy, they often obscure physical meaning, potentially eroding engineer trust in critical applications such as aerospace design or safety validation (Raissi et al., 2019). Lastly, most existing frameworks still assume smooth and differentiable objective functions, whereas real-world engineering systems may involve discontinuities, stochastic behavior, or discrete variables that challenge surrogate smoothness assumptions.

Given these challenges, future research on surrogate-assisted global optimization should pursue several complementary directions. First, there is a clear need for scalable probabilistic frameworks that retain the interpretability of classical Kriging while achieving linear or sublinear training complexity. Advances in sparse Gaussian Processes, kernel interpolation, and operator learning may help realize this goal (Wang et al., 2022). Second, researchers should focus on developing adaptive and online surrogates capable of incremental learning from streaming data, enabling continual refinement without retraining from scratch. This would facilitate integration into digital twin architectures and real-time optimization loops (Willcox & Megretski, 2019). Third, future studies should expand benchmark repositories, adopting open-access standardized datasets and reproducibility protocols to enable consistent comparison of algorithms across domains (Surjanovic & Bingham, 2013). Fourth, integrating reinforcement learning and AutoML for surrogate hyperparameter tuning could further automate model selection, bridging the gap between human expertise and algorithmic adaptability (Feurer & Hutter, 2019). Finally, researchers should explore uncertainty-aware optimization under nonstationarity, combining Bayesian inference with adversarial and distributionally robust learning to ensure reliability in changing environments. These avenues would collectively advance the theoretical depth and practical reach of surrogate-based optimization.

From a practical perspective, the findings of this review carry important implications for engineers, designers, and organizations adopting SAGO frameworks. Practitioners should



recognize that surrogate-assisted optimization is not merely a computational convenience but a strategic methodology for intelligent design exploration. Effective implementation requires a rigorous workflow: defining appropriate design-of-experiments (DoE) strategies, selecting surrogates aligned with system physics, and validating predictive performance with uncertainty estimates. For industrial applications, integrating SAGO into model-based systems engineering (MBSE) pipelines can significantly reduce design-cycle times while improving decision confidence (Goh et al., 2013). Organizations should also invest in building cross-functional expertise that combines computational modeling, statistics, and machine learning, ensuring that surrogate models are developed, validated, and deployed responsibly. Moreover, the move toward digital twins and adaptive control highlights the need for continuous data feedback and online surrogate updates to maintain model accuracy throughout a product's lifecycle (Willcox & Megretski, 2019). Finally, practitioners should prioritize transparent documentation, reproducibility, and open-source tools like OpenMDAO, PyKriging, or BoTorch (BoTorch Developers, 2020) to facilitate collaborative research and accelerate innovation. By embracing these practices, the engineering community can leverage surrogate-assisted global optimization not only as a research methodology but as a practical engine of innovation, efficiency, and reliability in the era of data-driven design.

Ethical Considerations

All procedures performed in this study were under the ethical standards.

Acknowledgments

Authors thank all who helped us through this study.

Conflict of Interest

The authors report no conflict of interest.

Funding/Financial Support

According to the authors, this article has no financial support.

References

BoTorch Developers. (2020). *BoTorch: A framework for efficient Monte Carlo Bayesian optimization*. arXiv preprint arXiv:2002.08418.

Conn, A. R., Gould, N. I. M., & Toint, P. L. (2000). *Trust-region methods*. SIAM.

Emmerich, M., & Deutz, A. (2018). A tutorial on multiobjective optimization: Fundamentals and evolutionary methods. *Natural Computing*, 17(3), 585–609.

Eriksson, D., Pearce, M., Gardner, J., Turner, R., & Poloczek, M. (2019). Scalable global optimization via local Bayesian optimization. *Advances in Neural Information Processing Systems*, 32, 5496–5507.

Feurer, M., & Hutter, F. (2019). Hyperparameter optimization. In *Automated Machine Learning* (pp. 3–33). Springer.

Forrester, A. I. J., & Keane, A. J. (2009). Recent advances in surrogate-based optimization. *Progress in Aerospace Sciences*, 45(1-3), 50-79.

Forrester, A., Sobester, A., & Keane, A. (2008). *Engineering design via surrogate modelling: A practical guide*. John Wiley & Sons.

Frazier, P. I. (2018). A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*.

Ginsbourger, D., Le Riche, R., & Carraro, L. (2010). Kriging is well-suited to parallelize optimization. In *Computational Intelligence in Expensive Optimization Problems* (pp. 131-162). Springer.

Goh, C. J., Teo, K. L., & Rehbock, V. (2013). Surrogate-assisted evolutionary algorithms for engineering design optimization. *Applied Soft Computing*, 13(4), 1819-1830.

Hernández-Lobato, J. M., Hoffman, M. W., & Ghahramani, Z. (2017). Predictive entropy search for efficient global optimization of black-box functions. *Advances in Neural Information Processing Systems*, 27, 918-926.

Jin, Y. (2011). Surrogate-assisted evolutionary computation: Recent advances and future challenges. *Swarm and Evolutionary Computation*, 1(2), 61-70.

Jones, D. R., Schonlau, M., & Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4), 455-492.

Kandasamy, K., Dasarathy, G., Oliva, J. B., Schneider, J., & Póczos, B. (2018). Multi-fidelity Bayesian optimization with continuous approximations. *International Conference on Machine Learning*, 2545-2554.

Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3(6), 422-440.

Kennedy, M. C., & O'Hagan, A. (2000). Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1), 1-13.

Knowles, J. (2006). ParEGO: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1), 50-66.

Lam, R., Allaire, D., & Willcox, K. (2015). Multifidelity optimization using statistical surrogate modeling for nonhierarchical information sources. *SIAM Journal on Scientific Computing*, 37(1), A212-A239.

Mockus, J. (2012). *Bayesian approach to global optimization: Theory and applications*. Springer.

Perdikaris, P., & Karniadakis, G. E. (2016). Model inversion via multi-fidelity Bayesian optimization: A new paradigm for parameter estimation in physics-based models. *SIAM Journal on Scientific Computing*, 38(5), A3163-A3194.

Perdikaris, P., Raissi, M., Damianou, A., Lawrence, N. D., & Karniadakis, G. E. (2017). Nonlinear information fusion algorithms for data-efficient multi-fidelity modeling. *Proceedings of the Royal Society A*, 473(2198), 20160751.

Queipo, N. V., Haftka, R. T., Shyy, W., Goel, T., Vaidyanathan, R., & Tucker, P. K. (2005). Surrogate-based analysis and optimization. *Progress in Aerospace Sciences*, 41(1), 1-28.

Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686-707.

Regis, R. G., & Shoemaker, C. A. (2007). A stochastic radial basis function method for the global optimization of expensive functions. *INFORMS Journal on Computing*, 19(4), 497-509.



Sacks, J., Welch, W. J., Mitchell, T. J., & Wynn, H. P. (1989). Design and analysis of computer experiments. *Statistical Science*, 4(4), 409-423.

Santner, T. J., Williams, B. J., & Notz, W. I. (2003). *The design and analysis of computer experiments*. Springer.

Sankararaman, S., & Mahadevan, S. (2013). Integration of model uncertainty and parameter uncertainty for reliability assessment and design. *Reliability Engineering & System Safety*, 112, 187-199.

Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., & De Freitas, N. (2016). Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1), 148-175.

Simpson, T. W., Peplinski, J. D., Koch, P. N., & Allen, J. K. (2001). Metamodels for computer-based engineering design: Survey and recommendations. *Engineering with Computers*, 17(2), 129-150.

Surjanovic, S., & Bingham, D. (2013). Virtual library of simulation experiments: Test functions and datasets. Retrieved from <https://www.sfu.ca/~ssurjano>

Wang, G. G., & Shan, S. (2007). Review of metamodeling techniques in support of engineering design optimization. *Journal of Mechanical Design*, 129(4), 370-380.

Wang, Y., Niu, S., & Zhang, Y. (2022). Quantum-inspired surrogate modeling for engineering optimization. *Applied Soft Computing*, 116, 108307.

Willcox, K., & Megretski, A. (2019). Predictive digital twins: Data-driven modeling for complex systems. *AIAA Journal*, 57(11), 4732-4751.

Zhao, H., & Xue, D. (2019). A multi-surrogate approach for global optimization of expensive engineering design problems. *Structural and Multidisciplinary Optimization*, 60(5), 1825-1843.