Novelty Assessment Report

Paper: Symmetry-Aware Bayesian Optimization via Max Kernels
PDF URL: https://openreview.net/pdf?id=zUbBaWAM1Q
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2025-12-27

Abstract

Bayesian Optimization (BO) is a powerful framework for optimizing noisy, expensive-to-evaluate black-box functions. When the objective exhibits invariances under a group action, exploiting these symmetries can substantially improve BO efficiency. While using maximum similarity across group orbits has long been considered in other domains, the fact that the max kernel is not positive semidefinite (PSD) has prevented its use in BO. In this work, we revisit this idea by considering a PSD projection of the max kernel. Compared to existing invariant (and non-invariant) kernels, we show it achieves significantly lower regret on both synthetic and real-world BO benchmarks, without increasing computational complexity.

Core Task Landscape

This paper addresses: **Bayesian Optimization with Symmetry-Aware Kernels for Invariant Black-Box Functions**

A total of **9 papers** were analyzed and organized into a taxonomy with **7 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Symmetry-Aware Kernel Design and Theory**
- **Geometric and Manifold Optimization**
- **Structured and Discrete Spaces**
- **Set-Valued Optimization**
- **Application Domains**

Complete Taxonomy Tree

- Bayesian Optimization with Symmetry-Aware Kernels for Invariant Black-Box Functions Survey Taxonomy
- Symmetry-Aware Kernel Design and Theory
  - Invariant Kernel Construction ★ (2 papers)
  - [0] Symmetry-Aware Bayesian Optimization via Max Kernels (Anon et al., 2026) View paper
  - [1] Sample-efficient bayesian optimisation using known invariances (Ilija Bogunovic, 2024) View paper
- Geometric and Manifold Optimization
  - Riemannian and Manifold Kernels (2 papers)
  - [7] Extrinsic Bayesian Optimizations on Manifolds (Fang Yihao, 2022) View paper
  - [9] Geometry-aware Bayesian Optimization in Robotics using Riemannian MatÃ©rn Kernels (Jaquier, 2021) View paper
- Structured and Discrete Spaces
  - Tree Ensemble and Mixed-Feature Kernels (1 papers)
  - [3] Tree ensemble kernels for Bayesian optimization with known constraints over mixed-feature spaces (Thebelt, 2022) View paper
- Set-Valued Optimization
  - Approximate Set Kernels (2 papers)
  - [5] Bayesian optimization with approximate set kernels (Jung-Taek Kim, 2021) View paper
  - [6] Learning to Optimize in Structured Environments (Kassraie, 2025) View paper
- Application Domains
  - Molecular and Materials Science (1 papers)
  - [2] AUGUR, a flexible and efficient optimization algorithm for identification of optimal adsorption sites (Ioannis Kouroudis, 2025) View paper
  - Robotics and Control (1 papers)
  - [4] Efficient Exploration of Reward Functions in Inverse Reinforcement Learning via Bayesian Optimization (Nguyen Quoc Phong, 2020) View paper
  - General Black-Box Optimization (1 papers)
  - [8] On Bayesian Methods for Black-Box Optimization: Efficiency, Adaptation and Reliability (Y Zhang, 2024) View paper

Narrative

Core task: Bayesian optimization with symmetry-aware kernels for invariant black-box functions. The field addresses how to efficiently optimize expensive-to-evaluate functions that exhibit known or suspected symmetries—such as invariance under rotations, permutations, or other group actions. The taxonomy organizes research into several main branches: Symmetry-Aware Kernel Design and Theory focuses on constructing kernels that respect invariances, often by averaging over group orbits or embedding symmetry constraints directly into covariance structures (e.g., Symmetry Max Kernels[0], Known Invariances[1]). Geometric and Manifold Optimization extends these ideas to curved spaces, where standard Euclidean kernels fail and Riemannian or extrinsic embeddings become necessary (Extrinsic Manifolds[7], Riemannian Matern[9]). Structured and Discrete Spaces tackles combinatorial or tree-structured domains (Tree Ensemble Kernels[3]), while Set-Valued Optimization handles inputs that are unordered collections (Approximate Set Kernels[5]). Application

Domains illustrate how these methods deploy in chemistry, materials science, and reinforcement learning (AUGUR Adsorption[2], Reward Functions IRL[4], Structured Environments[6]).

A particularly active line of work explores how to encode known symmetries into Gaussian process priors without prohibitive computational overhead, balancing exact invariance guarantees against scalability. Symmetry Max Kernels[0] sits squarely within the Invariant Kernel Construction cluster, proposing a max-pooling strategy over group orbits to achieve invariance while maintaining tractable inference. This contrasts with Known Invariances[1], which emphasizes leveraging user-specified symmetry groups to construct invariant kernels through explicit averaging, and with broader geometric approaches like Extrinsic Manifolds[7] that embed manifold constraints into kernel design. The original paper's emphasis on max-based aggregation offers a middle ground between exact orbit averaging and approximate symmetry handling, positioning it as a methodological contribution to the kernel design branch rather than a domain-specific application or a purely geometric extension.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. Sample-efficient bayesian optimisation using known invariances

Authors: Ilija Bogunovic, Theodore Brown, Alexandru Cioba | **Year/Venue**: 2024 | **URL**: View paper

Abstract

â¦ on sample complexity in Bayesian optimisation with invariant kernels, which explicitly quantify the gain in sample complexity due to the number of symmetries. We extend several results â¦

Relationship Analysis

Both papers belong to the Invariant Kernel Construction category, focusing on techniques for building kernels that respect group symmetries in Bayesian optimization. They overlap in addressing how to construct G-invariant kernels from base kernels for symmetric black-box functions, but differ fundamentally in their approach: the original paper uses a max-alignment strategy (taking the maximum similarity over group orbits) followed by PSD projection, while the candidate paper uses orbit-averaging (averaging the base kernel over group transformations) and provides theoretical sample complexity bounds for this averaging approach.

Contributions Analysis

Overall novelty summary. The paper proposes a positive semidefinite projection of the max kernel to exploit group symmetries in Bayesian optimization. It resides in the 'Invariant Kernel Construction' leaf under 'Symmetry-Aware Kernel Design and Theory', which contains only two papers total (including this one). This places the work in a relatively sparse research direction within a taxonomy of nine papers across seven leaf nodes. The sibling paper focuses on leveraging known invariances through explicit group averaging, suggesting that kernel construction methods for symmetry-aware BO remain an emerging area with limited prior exploration.

The taxonomy reveals that symmetry-aware kernel design sits alongside geometric manifold optimization (Riemannian kernels for curved spaces) and structured discrete optimization (tree ensemble methods). The original paper's approach differs from neighboring geometric methods by targeting Euclidean spaces with group invariances rather than non-Euclidean manifolds. It also diverges from set-valued optimization techniques, which handle unordered collections rather than orbit-based symmetries. The scope notes clarify that manifold-specific kernels and application-specific implementations belong elsewhere, positioning this work as a foundational kernel design contribution rather than a domain-specific extension.

Among four candidates examined across three contributions, none were found to clearly refute the proposed methods. The PSD projection of the max kernel examined one candidate with no refutable overlap. The empirical performance analysis examined three candidates, again with no clear prior work providing the same insights. The demonstration of gains over orbit averaging examined zero candidates. Given the limited search scope—only four papers reviewed—these statistics suggest the specific combination of max kernel projection and PSD constraints has not been extensively studied, though the small candidate pool prevents strong conclusions about absolute novelty.

Based on the limited literature search of four candidates, the work appears to occupy a sparsely populated research direction within symmetry-aware Bayesian optimization. The taxonomy structure and sibling paper count reinforce this impression, though the restricted search scope means potentially relevant work outside the top semantic matches may exist. The analysis covers kernel construction methods but does not exhaustively survey all symmetry-handling techniques in optimization or related fields.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: PSD projection of max kernel for symmetry-aware Bayesian Optimization

Description: The authors propose a positive semidefinite (PSD) version of the max-alignment kernel (kmax) for Bayesian Optimization. They construct k(D)+ via PSD projection and Nyström extension, ensuring it is G-invariant, equals kmax on the design set when kmax is PSD, and matches the asymptotic cost of orbit-averaged kernels.

This contribution was assessed against **1 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. HHD-GP: Incorporating Helmholtz-Hodge Decomposition into Gaussian Processes for Learning Dynamical Systems

URL: View paper

Brief Assessment

Helmholtz-Hodge GP[10] focuses on decomposing dynamical systems into curl-free and divergence-free components using Gaussian processes, not on positive semidefinite projection of kernels for Bayesian optimization with symmetry groups.

Contribution 2: Demonstration of consistent BO performance gains over orbit averaging

Description: The authors empirically demonstrate that their proposed kernel k(D)+ consistently achieves lower cumulative and simple regret compared to both the base kernel and the orbit-averaged alternative (kavg) across multiple synthetic and real-world benchmarks, with gains increasing as the group size grows.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

Contribution 3: Analysis revealing mismatch between eigendecay and empirical performance

Description: The authors analyze the spectral properties of their kernel and show that despite kavg often exhibiting faster empirical eigendecay than k(D)+, the latter consistently achieves better regret. This reveals a gap between standard spectral-based BO theory and empirical performance, suggesting that geometric considerations and approximation hardness play essential roles beyond pure spectral rates.

This contribution was assessed against **3 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Gaussian process bandits for tree search: Theory and application to planning in discounted MDPs

URL: View paper

Brief Assessment

GP Tree Search[13] focuses on tree search applications with theoretical eigendecay analysis for regret bounds, not on comparing eigendecay predictions versus empirical BO performance across different kernel constructions as in the original paper.

2. Safe Bayesian Optimization for Complex Control Systems via Additive Gaussian Processes

URL: View paper

Brief Assessment

Safe Additive GP[12] focuses on safe Bayesian optimization for control systems using additive Gaussian processes, not on spectral analysis of kernel eigendecay versus empirical regret performance in general BO settings.

3. Quantum kernelized bandits

URL: View paper

Brief Assessment

Quantum Bandits[11] focuses on quantum kernelized bandits with quantum reward oracles and provides regret bounds in terms of eigenvalue decay rates. The paper does not analyze mismatches between spectral properties and empirical performance in Bayesian Optimization contexts, which is the core contribution of the original paper.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Symmetry-Aware Bayesian Optimization via Max Kernels View paper
- [1] Sample-efficient bayesian optimisation using known invariances View paper
- [2] AUGUR, a flexible and efficient optimization algorithm for identification of optimal adsorption sites View paper
- [3] Tree ensemble kernels for Bayesian optimization with known constraints over mixed-feature spaces View paper
- [4] Efficient Exploration of Reward Functions in Inverse Reinforcement Learning via Bayesian Optimization View paper
- [5] Bayesian optimization with approximate set kernels View paper
- [6] Learning to Optimize in Structured Environments View paper
- [7] Extrinsic Bayesian Optimizations on Manifolds View paper
- [8] On Bayesian Methods for Black-Box Optimization: Efficiency, Adaptation and Reliability View paper
- [9] Geometry-aware Bayesian Optimization in Robotics using Riemannian MatÃ©rn Kernels View paper
- [10] HHD-GP: Incorporating Helmholtz-Hodge Decomposition into Gaussian Processes for Learning Dynamical Systems View paper
- [11] Quantum kernelized bandits View paper
- [12] Safe Bayesian Optimization for Complex Control Systems via Additive Gaussian Processes View paper
- [13] Gaussian process bandits for tree search: Theory and application to planning in discounted MDPs View paper