Novelty Assessment Report

Paper: Revisiting Matrix Sketching in Linear Bandits: Achieving Sublinear Regret via Dyadic Block Sketching
PDF URL: https://openreview.net/pdf?id=FKEHiHU4bN
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2025-12-30

Abstract

Linear bandits have become a cornerstone of online learning and sequential decision-making, providing solid theoretical foundations for balancing exploration and exploitation. Within this domain, matrix sketching serves as a critical component for achieving computational efficiency, especially when confronting high-dimensional problem instances. The sketch-based approaches reduce per-round complexity from $\Omega(d^2)$ to $O(dl)$, where D is the dimension and $l<d$ is the sketch size. However, this computational efficiency comes with a fundamental pitfall: when the streaming matrix exhibits heavy spectral tails, such algorithms can incur vacuous linear regret. In this paper, we revisit the regret bounds and algorithmic design for sketch-based linear bandits. Our analysis reveals that inappropriate sketch sizes can lead to substantial spectral error, severely undermining regret guarantees. To overcome this issue, we propose Dyadic Block Sketching, a novel multi-scale matrix sketching approach that dynamically adjusts the sketch size during the learning process. We apply this technique to linear bandits and demonstrate that the new algorithm achieves sublinear regret bounds without requiring prior knowledge of the streaming matrix properties. It establishes a general framework for efficient sketch-based linear bandits, which can be integrated with any matrix sketching method that provides covariance guarantees. Comprehensive experimental evaluation demonstrates the superior utility-efficiency trade-off achieved by our approach.

Core Task Landscape

This paper addresses: **Matrix Sketching in Linear Bandits with Sublinear Regret Guarantees**

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:

- **Adaptive and Multi-Scale Sketching Methods**
- **Static Sketching Approaches and Pitfall Analysis**
- **Reward Imputation with Sketching**
- **Theoretical Foundations and Related Problems**

Complete Taxonomy Tree

- Matrix Sketching in Linear Bandits with Sublinear Regret Guarantees Survey Taxonomy
- Adaptive and Multi-Scale Sketching Methods
  - Dyadic Block Sketching for Linear Bandits ★ (1 papers)
  - [0] Revisiting Matrix Sketching in Linear Bandits: Achieving Sublinear Regret via Dyadic Block Sketching (Anon et al., 2026) View paper
  - Gaussian Process Optimization with Adaptive Sketching (2 papers)
  - [3] Gaussian process optimization with adaptive sketching: Scalable and no regret (Calandriello, 2019) View paper
  - [5] Subspace Learning with Gaussian Processes for Sparse Contextual Bandits (Chizi, 2025) View paper
- Static Sketching Approaches and Pitfall Analysis
  - Pitfall Analysis and Framework Redesign (1 papers)
  - [2] Matrix Sketching in Bandits: Current Pitfalls and New Framework (Dongxie Wen, 2024) View paper
  - Algorithm Design and Enhancement (2 papers)
  - [1] Augmenting Subspace Optimization Methods with Linear Bandits (Menickelly, 2024) View paper
  - [6] Efficient and Robust High-Dimensional Linear Contextual Bandits. (Cheng Chen, 2020) View paper
- Reward Imputation with Sketching
  - Contextual Batched Bandits with Imputation (2 papers)
  - [4] Reward imputation with sketching for contextual batched bandits (Zhang Xiao, 2023) View paper
  - [9] Partial Information as Full: Reward Imputation with Sketching in Bandits (X Zhang, 2022) View paper
- Theoretical Foundations and Related Problems
  - Uniform Bounds for Sketched Bilinear Forms (1 papers)
  - [7] Beyond Johnson-Lindenstrauss: Uniform Bounds for Sketched Bilinear Forms (Deb, 2025) View paper
  - Matrix Completion in Online Quadratic Optimization (1 papers)
  - [8] Estimate to Decide: Matrix Completion driven Smoothed Online Quadratic Optimization (N Bhuyan, 2025) View paper

Narrative

Core task: matrix sketching in linear bandits with sublinear regret guarantees. The field addresses computational and memory challenges in linear contextual bandits by employing dimensionality reduction techniques that preserve statistical guarantees. The taxonomy reveals four main branches: Adaptive and Multi-Scale Sketching Methods develop dynamic compression schemes that adjust sketch dimensions over time or across problem scales; Static Sketching Approaches and Pitfall Analysis examines fixed-dimension sketches and identifies failure modes, as illustrated by Matrix Sketching Pitfalls[2]; Reward Imputation with Sketching explores techniques like those in Reward

Imputation Sketching[4] that handle partial feedback through imputation combined with compression; and Theoretical Foundations and Related Problems investigates underlying mathematical principles and connections to matrix completion and robust estimation, exemplified by works such as Beyond Johnson Lindenstrauss[7] and Matrix Completion Smoothed[8]. These branches collectively span the spectrum from practical algorithmic design to rigorous theoretical characterization.

A central tension emerges between static versus adaptive sketching strategies: while static methods offer simplicity, they risk catastrophic regret as shown in pitfall analyses, whereas adaptive approaches like Adaptive Sketching GP[3] and Augmenting Subspace Optimization[1] dynamically refine their compression to balance computational savings with statistical accuracy. The original paper, Dyadic Block Sketching[0], sits squarely within the Adaptive and Multi-Scale Sketching Methods branch, employing a dyadic blocking structure that progressively adjusts sketch granularity. This positions it close to other adaptive schemes but with a distinctive multi-scale flavor that contrasts with the continuous subspace refinement seen in Augmenting Subspace Optimization[1] or the Gaussian process extensions in Adaptive Sketching GP[3]. The work navigates the trade-off between maintaining sublinear regret and achieving computational efficiency, a recurring theme across adaptive methods that seek to avoid the pitfalls of overly aggressive or static compression.

Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

Taxonomy-Level Summary

Both subtopics address computational efficiency in sequential decision-making through matrix sketching techniques. The original leaf focuses on linear bandits with multi-scale dyadic sketching that adapts to unknown matrix properties, while the sibling focuses on Gaussian process optimization with adaptive sketching for scalability. They represent parallel approaches to handling high-dimensional problems in different bandit frameworks.

Similarities: - Both employ adaptive sketching techniques to reduce computational complexity - Both target sublinear regret or near-optimal performance guarantees - Both address scalability challenges in high-dimensional sequential optimization - Both avoid requiring complete prior knowledge of problem structure

Differences: - Original leaf uses dyadic (multi-scale) block structure for linear contextual bandits; sibling uses adaptive sketching for Gaussian process kernels - Original leaf focuses on linear reward models; sibling handles non-linear GP-based optimization - Original leaf emphasizes dynamic sketch size adjustment without matrix knowledge; sibling emphasizes runtime scalability in GP inference - Original leaf explicitly excludes GP approaches; sibling explicitly excludes linear bandit methods

Suggested Search Directions: - Hybrid methods combining linear and kernel-based sketching for bandits - Theoretical connections between dyadic sketching and GP approximation guarantees - Comparative empirical studies of sketching strategies across linear vs. kernel bandit settings

Sibling Subtopics

- **Gaussian Process Optimization with Adaptive Sketching** (leaves: 1, papers: 2)
- Scope: Approximate GP algorithms using adaptive sketching to achieve near-optimal runtime and scalability in high-dimensional optimization.
- Exclude: Linear bandit methods and subspace learning approaches belong to their respective categories.

Contributions Analysis

Overall novelty summary. The paper proposes Dyadic Block Sketching, a multi-scale matrix sketching method for linear bandits that dynamically adjusts sketch size to achieve sublinear regret. Within the taxonomy, it occupies the 'Dyadic Block Sketching for Linear Bandits' leaf under 'Adaptive and Multi-Scale Sketching Methods', where it is currently the sole paper. This placement reflects a relatively sparse research direction focused specifically on multi-scale adaptive sketching for linear bandits, distinguishing it from the broader adaptive sketching literature that includes Gaussian process methods and subspace learning approaches in neighboring leaves.

The taxonomy reveals that the paper sits within a branch emphasizing dynamic compression strategies, contrasting with 'Static Sketching Approaches and Pitfall Analysis' which examines fixed-dimension methods and their failure modes. Neighboring work includes Gaussian process optimization with adaptive sketching and reward imputation techniques for batched bandits. The scope notes clarify that while related adaptive methods exist (e.g., continuous subspace refinement, GP extensions), this work's dyadic blocking structure and focus on spectral tail handling in linear bandits represent a distinct methodological direction within the adaptive sketching paradigm.

Among the three contributions analyzed, all show evidence of prior work within the limited search scope of 24 candidates. The 'Dyadic Block Sketching for constrained global error' contribution examined 8 candidates with 1 appearing to provide overlapping work. The 'Sublinear regret guarantee' contribution found 3 potentially refuting papers among 8 examined, suggesting more substantial prior work in this area. The 'General framework' contribution identified 2 refuting candidates from 8 examined. These statistics indicate that while some aspects have precedent in the examined literature, the specific combination and multi-scale approach may offer distinguishing features.

Based on the limited search scope of 24 semantically similar papers, the work appears to occupy a methodologically distinct position within adaptive sketching for linear bandits. However, the contribution-level statistics suggest that individual technical elements have precedent in the examined literature. The analysis does not cover the full breadth of matrix sketching or online learning research, and a more exhaustive search might reveal additional related work in adjacent areas such as streaming algorithms or randomized numerical linear algebra.

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

Contribution 1: Dyadic Block Sketching for constrained global error

Description: A multi-scale matrix sketching method that adaptively adjusts sketch sizes across multiple scales to control approximation error. The method provably bounds global error by a predetermined parameter epsilon and tracks the optimal rank-k approximation in streaming settings.

This contribution was assessed against **8 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Dynamical Sketching for Enhanced Communication Efficiency in Federated Learning

URL: View paper

Brief Assessment

Dynamical Sketching Federated[12] focuses on federated learning with dynamic sketch size selection for communication efficiency, not on multi-scale matrix sketching with global error bounds for linear bandits or streaming settings.

---

2. Seeing the Forest from the Trees in Two Looks: Matrix Sketching by Cascaded Bilateral Sampling

URL: View paper

Brief Assessment

Cascaded Bilateral Sampling[15] focuses on bilateral sampling (rows and columns) for general rectangular matrices with a two-stage cascading framework, not on multi-scale dyadic block sketching with adaptive sketch size adjustment across multiple scales for streaming matrices.

---

3. Matrix Sketching in Bandits: Current Pitfalls and New Framework

URL: View paper

Prior Art Analysis

Matrix Sketching Pitfalls[2] demonstrates that a multi-scale matrix sketching method with adaptive sketch size adjustment already exists. The candidate paper proposes 'dyadic block sketching' that adaptively adjusts sketch sizes across multiple scales to control approximation error, with the global error bounded by a predetermined parameter epsilon. This directly overlaps with the original paper's claimed contribution of a multi-scale sketching method that 'adaptively adjusts sketch sizes across multiple scales to control approximation error' and 'provably bounds global error by a predetermined parameter epsilon.'

Evidence

Evidence 1 - **Rationale**: Both papers propose the same 'dyadic block sketching' method with identical multi-scale adaptive sketch size adjustment mechanism and error control properties. - **Original**: we propose dyadic block sketching, a novel multi-scale matrix sketching approach that dynamically adjusts the sketch size during the learning process - **Candidate**: we propose dyadic block sketching, a multi-scale matrix sketching method that imposes a constraint on the global spectral error by managing the error bound within each block. we prove that the cumulative spectral error upper bound from dyadic block sketching conforms to a specified error ϵ.

Evidence 2 - **Rationale**: Both papers describe the same dyadic block sketching method with adaptive sketch size management and tracking of rank-k approximation, indicating the original paper was not the first to propose this contribution. - **Original**: we propose dyadic block sketching, a novel multi-scale matrix sketching approach that dynamically adjusts the sketch size during the learning process. we apply this technique to linear bandits and demonstrate that the new algorithm achievessublinear regretbounds - **Candidate**: we propose dyadic block sketching, an innovative streaming matrix sketching approach that adaptively manages sketch size to constrain global spectral loss. this approach effectively tracks the best rank-k approximation in an online manner

Evidence 3 - **Rationale**: Both papers provide identical theoretical guarantees that the global error is bounded by a predetermined parameter epsilon (ϵ), demonstrating that the candidate paper already established this result. - **Original**: we propose dyadic block sketching, a novel multi-scale matrix sketching method that adaptively adjusts sketch sizes across multiple scales to control error. we prove that the global error is bounded by a predetermined errorϵ. - **Candidate**: theorem 2. suppose a streaming matrix sketch, denoted as κ, achieves a covariance error $x^\top x$ - $s^\top s$ 2 $\leq \eta \cdot \|x\|2$ f with $\ell\eta$ rows and $\mu\eta$ update time. applying κ as the sketching method for each block in the dyadic block sketching and l0 is the initial sketch size, we generate a matrix sketch s for the e...

---

4. Randomized Matrix Sketching for Neural Network Training and Gradient Monitoring

URL: View paper

Brief Assessment

Sketching Neural Training[10] focuses on neural network activation compression using EMA-based sketching for gradient monitoring, not on multi-scale matrix sketching with adaptive sketch size adjustment for general streaming matrices in linear bandits.

---

5. Effective dimension adaptive sketching methods for faster regularized least-squares optimization

URL: View paper

Brief Assessment

Effective Dimension Sketching[11] focuses on adaptive sketching for regularized least-squares optimization with effective dimension, not on multi-scale dyadic block sketching for linear bandits with global error bounds controlled by epsilon.

---

6. Optimal Matrix Sketching over Sliding Windows

URL: View paper

Brief Assessment

Optimal Sliding Sketching[14] addresses matrix sketching over sliding windows with dump snapshots, not general streaming settings. The original paper's dyadic block sketching applies to linear bandits with adaptive sketch size adjustment across learning rounds, which is a different problem domain.

---

7. Scalable Batched Iterative Matrix Sketching: Theory and Practice

URL: View paper

Brief Assessment

Batched Iterative Sketching[16] focuses on batched parametrized frequent directions for matrix sketching with different error bounds and batching mechanisms, not the multi-scale dyadic framework with predetermined epsilon-bounded global error control described in the original paper.

---

8. Matrix sketching over sliding windows

URL: View paper

Brief Assessment

Sketching Sliding Windows[13] focuses on matrix sketching in sliding window models for streaming data, not on multi-scale adaptive sketching for linear bandits with regret guarantees. The candidate addresses a different problem domain (sliding windows) rather than the original's focus on controlling approximation error in bandit settings.

Contribution 2: Sublinear regret guarantee for sketch-based linear bandits

Description: By applying Dyadic Block Sketching to linear bandits, the authors achieve sublinear regret even when the streaming matrix exhibits heavy spectral tails, addressing the linear regret pitfall observed in prior sketch-based methods like SOFUL.

This contribution was assessed against **8 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Efficient Linear Bandits through Matrix Sketching

URL: View paper

Prior Art Analysis

Efficient Linear Sketching[19] demonstrates that sketch-based linear bandits can achieve sublinear regret under heavy spectral tails, predating the original paper's claim. The candidate paper proves that SOFUL with Frequent Directions sketching achieves regret bound $\tilde{O}((1+\varepsilon_m)^{(3/2)}\sqrt{t})$, where ε_m is bounded by spectral tail eigenvalues. This directly addresses the same problem of achieving sublinear regret with matrix sketching when the streaming matrix exhibits heavy spectral properties, which the original paper claims as novel through Dyadic Block Sketching.

Evidence

Evidence 1 - **Rationale**: This pair shows that Efficient Linear Sketching[19] already proved sublinear regret bounds for sketch-based linear bandits with spectral error dependence, addressing the same fundamental problem the original paper claims to solve. - **Original**: when the streaming matrix exhibits heavy spectral tails, such algorithms can incur vacuouslinear regret. in this paper, we revisit the regret bounds and algorithmic design for sketch-based linear bandits. our analysis reveals that inappropriate sketch sizes can lead to substantial spectral error, se... - **Candidate**: theorem 3. the regret of soful with fd-sketching of sizemw.h.p. satisfies rt ˜o= mλ (1 + εm)3 2 ˜m (r+ s √ λ (1 + 1 λ))√ t . similarly to abbasi-yadkori et al. (2011), we also prove a distribution dependent regret bound for soful. this bound is polylogarithmic in time and depends on the smallest...

Evidence 2 - **Rationale**: Both papers identify the same computational challenge and the same pitfall of linear regret under heavy spectral tails, showing that Efficient Linear Sketching[19] already recognized this problem. - **Original**: the sketch-based approaches reduce perround complexity fromω(d 2)too(dl), wheredis the dimension andl < dis the sketch size. however, this computational efficiency comes with a fundamental pitfall: when the streaming matrix exhibits heavy spectral tails, such algorithms can incur vacuouslinear regre... - **Candidate**: while exhibiting good theoretical and empirical performances, both algorithms require ω (d2) time to update their model after each round. in this work we investigate whether it is possible to significantly reduce this update time while ensuring that the regret remains nicely bounded. the quadratic d...

Evidence 3 - **Rationale**: This demonstrates that Efficient Linear Sketching[19] already achieved sublinear regret with sketch-based methods by controlling regret through spectral properties, without requiring prior knowledge of matrix characteristics. - **Original**: we apply this technique to linear bandits and demonstrate that the new algorithm achievessublinear regretbounds without requiring prior knowledge of the streaming matrix properties - **Candidate**: with a sketch size of m, a rank-one update of the correlation matrix takes only time o(md), which is linear in dfor a constant sketch size. however, this speed-up comes at a price, as sketching reduces the matrix rank causing a loss of information which -in turnaffects the least squares estimates us...

Evidence 4 - **Rationale**: Both papers explicitly acknowledge the linear regret pitfall under heavy spectral tails, showing that Efficient Linear Sketching[19] already identified this challenge and provided conditions under which it occurs. - **Original**: observation 1.assume the decision set is drawn from a locally convex arm spacex. if the sketch size of soful satisfiesl < d-t 3 -q, then soful incurs vacuous linear regret. consequently, when the geometry constantq≥1/3, soful suffers linear regret for any sketch sizel < d - **Candidate**: when the correlation matrix has rank larger than the sketch size, the regret of soful remains small to the extent the spectral tail of the matrix grows slowly with t. in the worst case of a spectrum with heavy tails, soful may incur linear regret. in this respect, sketching is only justified when the...

---

2. Efficient and Robust High-Dimensional Linear Contextual Bandits.

URL: View paper

Prior Art Analysis

Robust Linear Contextual[6] demonstrates that prior work achieved sublinear regret bounds for sketch-based linear bandits under heavy spectral tails. The candidate paper presents a regret bound of $\tilde{O}((\sqrt{m} + d \log(1 + \Delta_t) + \sqrt{\Delta_t})\sqrt{mt})$ where Δ_t represents spectral error, and explicitly addresses the challenge of heavy spectral tails through their SCFD method. This shows that achieving sublinear regret with matrix sketching under heavy spectral tails was already accomplished before the original paper's submission.

Evidence

Evidence 1 - **Rationale**: Both papers address the fundamental problem of matrix approximation error in sketch-based linear bandits and propose novel sketching methods to achieve better regret bounds, indicating prior work on this contribution. - **Original**: when the streaming matrix exhibits heavy spectral tails, such algorithms can incur vacuouslinear regret. in this paper, we revisit the regret bounds and algorithmic design for sketch-based linear bandits. our analysis reveals that inappropriate sketch sizes can lead to substantial spectral error, se... - **Candidate**: however, the matrix approximation error will bring additional terms to the regret bound. in this paper we first propose a novel matrix sketching method which is called spectral compensation frequent directions (scfd). then we propose an efficient approach for contextual bandits by adopting scfd to a...

Evidence 2 - **Rationale**: Both papers claim to achieve improved regret bounds through novel sketching approaches, with the candidate explicitly stating better regret bounds in high-dimensional cases. - **Original**: the new algorithm achievessublinear regretbounds without requiring prior knowledge of the streaming matrix properties - **Candidate**: theoretical analysis reveals that our method has better regret bounds than previous methods in highdimensional cases. experimental results demonstrate the effectiveness of our algorithm and verify our theoretical guarantees.

Evidence 3 - **Rationale**: The candidate paper explicitly addresses reducing the impact of spectral error (Δ_t) on regret bounds, demonstrating prior work on achieving sublinear regret under challenging spectral conditions. - **Original**: by applying the proposed sketching framework to linear bandits, we effectively address the issue of linear regret observed in prior works. our method ensures a sublinear regret, even when the streaming matrix is heavy-tailed. - **Candidate**: our method reduces the order of Δt from 3/2 to 1/2. in addition, our regret bound decouples the dimension d and Δt , which further reduces the influence ofΔt . finally, our bound is less sensitive to the parameter λ (which is usually a small number in practice) because the term Δt /λ is in the logar...

---

3. Matrix Sketching in Bandits: Current Pitfalls and New Framework

URL: View paper

Prior Art Analysis

Matrix Sketching Pitfalls[2] demonstrates that applying dyadic block sketching to linear bandits to achieve sublinear regret under heavy spectral tails was already accomplished. The candidate paper explicitly states they 'apply the proposed dyadic block sketching to linear bandits and demonstrate that the resulting bandit algorithm can achieve sublinear regret without prior knowledge of the covariance matrix, even under the worst case.' This directly refutes the novelty claim that the original paper was first to achieve sublinear regret in sketch-based linear bandits under heavy spectral tails.

Evidence

Evidence 1 - **Rationale**: Both papers explicitly claim to achieve sublinear regret in linear bandits under worst-case/heavy-tailed scenarios using dyadic block sketching, indicating the original paper was not first. - **Original**: by applying the proposed sketching framework to linear bandits, we effectively address the issue of linear regret observed in prior works. our method ensures a sublinear regret, even when the streaming matrix is heavy-tailed. - **Candidate**: then, we apply the proposed dyadic block sketching to linear bandits and demonstrate that the resulting bandit algorithm can achieve sublinear regret without prior knowledge of the covariance matrix, even under the worst case.

---

4. Augmenting Subspace Optimization Methods with Linear Bandits

URL: View paper

Brief Assessment

Augmenting Subspace Optimization[1] addresses a different problem domain (unconstrained optimization via subspace methods) rather than linear bandits with matrix sketching for online learning.

5. Efficient sparse linear bandits under high dimensional data

URL: View paper

Brief Assessment

Sparse Linear Bandits[18] addresses sparse linear bandits using LASSO and random projection for dimension reduction, not matrix sketching methods like Frequent Directions. The candidate focuses on sparsity-driven feature selection rather than controlling spectral error from streaming matrix sketches under heavy-tailed spectra.

6. Fast Second-Order Online Kernel Learning through Incremental Matrix Sketching and Decomposition

URL: View paper

Brief Assessment

Incremental Matrix Sketching[20] addresses second-order online kernel learning with matrix sketching, not linear bandits. The candidate focuses on kernel approximation and feature mapping in RKHS settings, while the original contribution specifically targets linear bandits with heavy spectral tails using Dyadic Block Sketching.

7. Gaussian process optimization with adaptive sketching: Scalable and no regret

URL: View paper

Brief Assessment

Adaptive Sketching GP[3] focuses on Gaussian process optimization with inducing point selection via ridge leverage score sampling, not on dyadic block sketching for linear bandits under heavy spectral tails.

8. Beyond Johnson-Lindenstrauss: Uniform Bounds for Sketched Bilinear Forms

URL: View paper

Brief Assessment

Beyond Johnson Lindenstrauss[7] focuses on sketched bilinear forms in bandits with a different technical approach (generic chaining, Gaussian width), not on addressing heavy spectral tails in streaming matrices via dyadic block sketching.

Contribution 3: General framework for efficient sketch-based linear bandits

Description: A flexible framework that can incorporate various matrix sketching methods (such as FD or RFD) to achieve efficient linear bandits with theoretical guarantees. The framework is robust, scalable, and achieves diverse regret bounds through different sketching approaches.

This contribution was assessed against **8 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Efficient and Robust High-Dimensional Linear Contextual Bandits.

URL: View paper

Prior Art Analysis

Robust Linear Contextual[6] presents a framework for sketch-based linear bandits that incorporates matrix sketching methods with theoretical guarantees. The paper proposes SCFD (Spectral Compensation Frequent Directions) and demonstrates how it can be integrated into linear bandits (CBSCFD algorithm) with O(md) complexity and theoretical regret bounds. This establishes prior work on creating flexible frameworks for sketch-based linear bandits with covariance guarantees.

Evidence

Evidence 1 - **Rationale**: Both papers describe frameworks that integrate matrix sketching methods into linear bandits with theoretical guarantees, showing prior work on this contribution. - **Original**: it establishes a general framework for efficient sketch-based linear bandits, which can be integrated with any matrix sketching method that provides covariance guarantees. - **Candidate**: we propose cbscfd algorithm for high-dimensional linear contextual bandits based on scfd. our method is much more efficient than linucb and require less space. compared with previous methods which use matrix sketching to accelerate linear contextual bandits, our method has better upper regret bound ...

Evidence 2 - **Rationale**: The candidate demonstrates a framework that achieves efficiency and theoretical guarantees through matrix sketching integration, establishing prior work on flexible sketch-based frameworks. - **Original**: it is robust, scalable, and flexible, achieving diverse regret bounds through various matrix sketching approaches. - **Candidate**: we present a variant of fd, which is called scfd, for matrix sketching. we then propose cbscfd algorithm for high-dimensional linear contextual bandits based on scfd. our method is much more efficient than linucb and require less space.

Evidence 3 - **Rationale**: Both papers provide comprehensive experimental evaluations demonstrating utility-efficiency tradeoffs of their sketch-based frameworks. - **Original**: comprehensive experimental evaluation demonstrates the superior utility-efficiency tradeoff achieved by our approach. - **Candidate**: we conduct experiments on both synthetic data and real-world data sets. the baseline approaches include linucb [abbasi-yadkori et al., 2011], cbrap [yu et al., 2017] and soful [kuzborskij et al., 2019]. for cbrap algorithm, we use the gaussian random matrix as the projection matrix because it has be...

2. Matrix Sketching in Bandits: Current Pitfalls and New Framework

URL: View paper

Prior Art Analysis

Matrix Sketching Pitfalls[2] already presents a general framework for efficient sketch-based linear bandits that can incorporate various matrix sketching methods. The candidate paper explicitly states 'our method is a general framework for efficient sketch-based linear bandits, applicable to all existing sketch-based approaches' and demonstrates integration with both FD and RFD methods. This directly overlaps with the original paper's claim of providing 'a flexible framework that can incorporate various matrix sketching methods (such as FD or RFD).'

Evidence

Evidence 1 - **Rationale**: Both papers describe the same framework (DBSLinUCB) that can integrate different sketching methods to achieve various regret bounds, indicating the original paper was not first. - **Original**: it establishes a general framework for efficient sketch-based linear bandits, which can be integrated with any matrix sketching method that provides covariance guarantees. - **Candidate**: dbslinucb represents a scalable framework for efficient sketch-based linear bandits that are capable of incorporating various streaming sketching techniques. we now explore two deterministic sketching techniques that provide different regret bounds of linear bandits.

Evidence 2 - **Rationale**: Both papers use nearly identical language to describe the framework's properties (robust, scalable, achieving diverse/various regret bounds through different sketching techniques). - **Original**: Furthermore, it is robust, scalable, and flexible, achieving diverse regret bounds through various matrix sketching approaches. - **Candidate**: our framework is robust, scalable, and capable of achieving various regret bounds through different sketching techniques.

3. Augmenting Subspace Optimization Methods with Linear Bandits

URL: View paper

Brief Assessment

Augmenting Subspace Optimization[1] proposes a framework for subspace optimization methods augmented with linear bandits, not a framework for sketch-based linear bandits with covariance guarantees and regret bounds.

4. Partial Information as Full: Reward Imputation with Sketching in Bandits

URL: View paper

Brief Assessment

Partial Information Full[9] focuses on reward imputation in contextual batched bandits using sketching for computational efficiency, not on developing a general framework for sketch-based linear bandits with covariance guarantees and theoretical regret bounds as described in the original contribution.

5. CONTEXTUAL BANDITS WITH PARTIALLY MAPPED LATENT FEATURES

URL: View paper

Brief Assessment

Partially Mapped Features[17] addresses contextual bandits with latent features and partial observability, not sketch-based linear bandits with matrix sketching methods for computational efficiency.

6. Subspace Learning with Gaussian Processes for Sparse Contextual Bandits

URL: View paper

Brief Assessment

Gaussian Process Subspace[5] focuses on sparse contextual bandits with Gaussian process-based learning and subspace identification, not on sketch-based linear bandits with covariance guarantees and matrix sketching methods like FD or RFD.

7. Gaussian process optimization with adaptive sketching: Scalable and no regret

URL: View paper

Brief Assessment

Adaptive Sketching GP[3] addresses GP optimization with adaptive subset selection, not a general framework for linear bandits with multiple sketching methods and covariance guarantees.

8. Beyond Johnson-Lindenstrauss: Uniform Bounds for Sketched Bilinear Forms

URL: View paper

Brief Assessment

Beyond Johnson Lindenstrauss[7] develops sketched bandit algorithms (sk9linucb, sk9lints) but does not present a multi-scale sketching framework with covariance guarantees like the original paper's dyadic block sketching approach.

Appendix: Text Similarity Detection

Textual similarity detection checked 18 papers and found 8 similarity segment(s) across 4 paper(s).

The following **4 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. Efficient and Robust High-Dimensional Linear Contextual Bandits.

Detected in: Contribution: contribution_2, Contribution: contribution_3

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

2. Matrix Sketching in Bandits: Current Pitfalls and New Framework

Detected in: Contribution: contribution_1, Contribution: contribution_2, Contribution: contribution_3

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

3. Efficient Linear Bandits through Matrix Sketching

Detected in: Contribution: contribution_2

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

4. Fast Second-Order Online Kernel Learning through Incremental Matrix Sketching and Decomposition

Detected in: Contribution: contribution_2

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] Revisiting Matrix Sketching in Linear Bandits: Achieving Sublinear Regret via Dyadic Block Sketching View paper
- [1] Augmenting Subspace Optimization Methods with Linear Bandits View paper
- [2] Matrix Sketching in Bandits: Current Pitfalls and New Framework View paper
- [3] Gaussian process optimization with adaptive sketching: Scalable and no regret View paper
- [4] Reward imputation with sketching for contextual batched bandits View paper
- [5] Subspace Learning with Gaussian Processes for Sparse Contextual Bandits View paper
- [6] Efficient and Robust High-Dimensional Linear Contextual Bandits. View paper
- [7] Beyond Johnson-Lindenstrauss: Uniform Bounds for Sketched Bilinear Forms View paper
- [8] Estimate to Decide: Matrix Completion driven Smoothed Online Quadratic Optimization View paper

- [9] Partial Information as Full: Reward Imputation with Sketching in Bandits View paper
- [10] Randomized Matrix Sketching for Neural Network Training and Gradient Monitoring View paper
- [11] Effective dimension adaptive sketching methods for faster regularized least-squares optimization View paper
- [12] Dynamical Sketching for Enhanced Communication Efficiency in Federated Learning View paper
- [13] Matrix sketching over sliding windows View paper
- [14] Optimal Matrix Sketching over Sliding Windows View paper
- [15] Seeing the Forest from the Trees in Two Looks: Matrix Sketching by Cascaded Bilateral Sampling View paper
- [16] Scalable Batched Iterative Matrix Sketching: Theory and Practice View paper
- [17] CONTEXTUAL BANDITS WITH PARTIALLY MAPPED LATENT FEATURES View paper
- [18] Efficient sparse linear bandits under high dimensional data View paper
- [19] Efficient Linear Bandits through Matrix Sketching View paper
- [20] Fast Second-Order Online Kernel Learning through Incremental Matrix Sketching and Decomposition View paper