Novelty Assessment Report

Paper: Consistent Low-Rank Approximation

PDF URL: https://openreview.net/pdf?id=3sJ4zKToW6

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-29

Abstract

We introduce and study the problem of consistent low-rank approximation, in which rows of an input matrix $\mathbf{A}\in\mathbb{R}^{n\times d}$ arrive sequentially and the goal is to provide a sequence of subspaces that well-approximate the optimal rank-K approximation to the submatrix $\mathbf{A}^{(t)}$ that has arrived at each time T, while minimizing the recourse, i.e., the overall change in the sequence of solutions. We first show that when the goal is to achieve a low-rank cost within an additive $\varepsilon\cdot||\mathbf{A}^{(t)}||_F^2$ factor of the optimal cost, roughly $\mathcal{O}\left(\frac{k}{\varepsilon}\log(nd)\right)$ recourse is feasible. For the more challenging goal of achieving a relative $(1+\varepsilon)$-multiplicative approximation of the optimal rank-K cost, we show that a simple upper bound in this setting is $\frac{k^2}{\varepsilon^2}\cdot\text{poly}\log(nd)$ recourse, which we further improve to $\frac{k^{3/2}}{\varepsilon^2}\cdot\text{poly}\log(nd)$ for integer-bounded matrices and $\frac{k}{\varepsilon^2}\cdot\text{poly}\log(nd)$ for data streams with polynomial online condition number. We also show that $\Omega\left(\frac{k}{\varepsilon}\log\frac{n}{k}\right)$ recourse is necessary for any algorithm that maintains a multiplicative $(1+\varepsilon)$-approximation to the optimal low-rank cost, even if the full input is known in advance. Finally, we perform a number of empirical evaluations to complement our theoretical guarantees, demonstrating the efficacy of our algorithms in practice.

Core Task Landscape

This paper addresses: **Consistent Low-Rank Approximation in Streaming Data**

A total of **50 papers** were analyzed and organized into a taxonomy with **22 categories**.

Taxonomy Overview

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

- **Matrix Low-Rank Approximation Methods**
- **Tensor Low-Rank Approximation Methods**
- **Application-Driven Low-Rank Methods**

Complete Taxonomy Tree

- Consistent Low-Rank Approximation in Streaming Data Survey Taxonomy
- Matrix Low-Rank Approximation Methods
  - Randomized Sketching and Single-Pass Algorithms
  - Nyström and Structured Sketching Approaches (3 papers)
    - [4] Streaming low-rank matrix approximation with an application to scientific simulation (Tropp, 2019) View paper
    - [7] Fixed-rank approximation of a positive-semidefinite matrix from streaming data (Joel A. Tropp, 2017) View paper
    - [12] Randomized techniques for low-rank approximation of matrix and tensors with applications (Alberto, 2025) View paper
  - Frequent Directions and Deterministic Streaming (2 papers)
    - [32] Low rank approximation lower bounds in row-update streams (David P. Woodruff, 2014) View paper
    - [38] Relative errors for deterministic low-rank matrix approximations (Mina Ghashami, 2014) View paper
  - Sparsity-Exploiting and Efficient Decompositions (3 papers)
    - [43] Improved Algorithms for Low Rank Approximation from Sparsity (David P. Woodruff, 2022) View paper
    - [45] An efficient, sparsity-preserving, online algorithm for low-rank approximation (David G. Anderson, 2017) View paper
    - [50] Single-pass randomized QLP decomposition for low-rank approximation (H. Ren, 2022) View paper
  - Online Matrix Factorization and Incremental Updates
  - Stochastic and Recursive Factorization (4 papers)
    - [6] Online nonnegative matrix factorization with robust stochastic approximation (Naiyang Guan, 2012) View paper
    - [16] Online projective nonnegative matrix factorization for large datasets (Zhirong Yang, 2012) View paper
    - [22] Online Algorithms for Recovery of Low-Rank Parameter Matrix in Non-stationary Stochastic Systems (Zhou, 2025) View paper
    - [27] Online Matrix Factorization via Broyden Updates (AkyÄ±ldÄ±z, 2022) View paper
  - Nonnegative and Constrained Factorization (3 papers)
    - [15] Online matrix factorization hashing for large-scale image retrieval (Lei Wang, 2018) View paper
    - [47] Visual tracking via online nonnegative matrix factorization (Yi Wu, 2013) View paper
    - [48] Check for updates Online Matrix Factorization Hashing for Large-Scale Image Retrieval (L Wang, 2018) View paper
  - Manifold-Based and Low-Rank Constrained Learning (2 papers)
    - [37] Online optimization for max-norm regularization (Jie Shen, 2017) View paper
    - [49] Online learning in the manifold of low-rank matrices (Uri Shalit, 2010) View paper
  - Specialized Matrix Approximation Techniques

- Robust and Outlier-Resistant Methods (2 papers)
    - [34] Differentially private robust low-rank approximation (Raman Arora, 2018) View paper
    - [46] Compressive Online Robust Principal Component Analysis via - Minimization (Huynh Van Luong, 2018) View paper
  - Subset Selection and Column Sampling (1 papers)
    - [3] New subset selection algorithms for low rank approximation: Offline and online (David P. Woodruff, 2023) View paper
  - Positive-Semidefinite and Structured Matrices (1 papers)
    - [19] Stochastic approximation and memory-limited subspace tracking for Poisson streaming data (Liming Wang, 2017) View paper
  - Consistent and Recourse-Minimizing Approximation ★ (1 papers)
  - [0] Consistent Low-Rank Approximation (Anon et al., 2026) View paper
- Tensor Low-Rank Approximation Methods
  - Tucker and Tree Tensor Network Decompositions (3 papers)
  - [1] Randomized algorithms for streaming low-rank approximation in tree tensor network format (Bucci Alberto, 2024) View paper
  - [18] A sequential multilinear NystrÃ¶m algorithm for streaming low-rank approximation of tensors in Tucker format (Bucci Alberto, 2024) View paper
  - [28] Low-Rank Tucker Approximation of a Tensor From Streaming Data (Yiming Sun, 2022) View paper
  - Tensor-Train and CP Factorization (2 papers)
  - [11] Tensor Train Factorization with Spatio-temporal Smoothness for Streaming Low-rank Tensor Completion (Gaohang Yu, 2024) View paper
  - [20] Tracking Online Low-Rank Approximations of Higher-Order Incomplete Streaming Tensors (LÃª Trung ThÃ nh, 2023) View paper
  - Online Tensor Factorization and Tracking (3 papers)
  - [2] Online tensor low-rank representation for streaming data (Tong Wu, 2024) View paper
  - [5] Streaming tensor factorization for infinite data sources (Shaden Smith, 2018) View paper
  - [21] Online Tensor Low-Rank Representation for Streaming Data Clustering (Tong Wu, 2022) View paper
  - Robust and Multi-Aspect Tensor Methods (2 papers)
  - [36] Outlier-Robust Multi-Aspect Streaming Tensor Completion and Factorization. (Mehrnaz Najafi, 2019) View paper
  - [39] Variational Bayesian inference for robust streaming tensor factorization and completion (Zheng Zhang, 2018) View paper
  - Tensor Sketching and Frequent Directions (2 papers)
  - [23] Effective streaming low-tubal-rank tensor approximation via frequent directions (Qianxin Yi, 2022) View paper
  - [41] Tensor-Based Sketching Method for the Low-Rank Approximation of Data Streams (Liu Cuiyu, 2022) View paper
- Application-Driven Low-Rank Methods
  - Forecasting and Time-Series Modeling (3 papers)
  - [9] Subspace learning and imputation for streaming big data matrices and tensors (Morteza Mardani, 2015) View paper
  - [13] Online forecasting matrix factorization (San Gultekin, 2018) View paper
  - [42] Streaming Low-Rank Matrix Data Assimilation and Change Identification (Henry Shaowu Yuchi, 2023) View paper
  - Recommender Systems and Collaborative Filtering (2 papers)
  - [29] Similarity based regularization for online matrix-factorization problem: An application to course recommender systems (Dhruv Shah, 2017) View paper
  - [40] Online Low Rank Matrix Completion (Jain, 2022) View paper
  - Contextual Bandits and Sequential Decision-Making (2 papers)
  - [10] Hyperbandit: Contextual bandit with hypernewtork for time-varying user preferences in streaming recommendation (Chenglei Shen, 2023) View paper
  - [24] Online statistical inference for matrix contextual bandit (Qiyu Han, 2022) View paper
  - Neural Network Training and Compression (2 papers)
  - [8] Low-rank extended Kalman filtering for online learning of neural networks from streaming data (Peter Chang, 2023) View paper
  - [14] Low-rank gradient descent for memory-efficient training of deep in-memory arrays (Siyuan Huang, 2023) View paper
  - Hardware Acceleration and Co-Design (2 papers)
  - [17] StreamSVD: Low-rank Approximation and Streaming Accelerator Co-design (Zhewen Yu, 2021) View paper
  - [31] FlashSVD: Memory-Efficient Inference with Streaming for Low-Rank Models (Shao Zishan, 2025) View paper
  - Domain-Specific Applications (5 papers)
  - [26] SOLD: Sub-optimal low-rank decomposition for efficient video segmentation (Chenglong Li, 2015) View paper
  - [30] Persona analytics: Analyzing the stability of online segments and content interests over time using non-negative matrix factorization (B. Jansen, 2021) View paper
  - [33] Hierarchical online NMF for detecting and tracking topic hierarchies in a text stream (Ding Tu, 2018) View paper
  - [35] MFDROO: Matrix Factorization-Based Deep Reinforcement Learning Approach for Stable Online Offloading in Mobile Edge Networks (Engy A. Abdelazim, 2024) View paper
  - [44] Clustering event streams with low rank Hawkes processes (Ali Caner TÃ¼rkmen, 2020) View paper
  - Data-Driven and Few-Shot Learning (1 papers)
  - [25] Few-shot data-driven algorithms for low rank approximation (Piotr Indyk, 2021) View paper

Narrative

Core task: consistent low-rank approximation in streaming data. The field addresses how to maintain compact, low-dimensional representations of large-scale data that arrive sequentially, while ensuring that successive approximations remain stable and interpretable over time. The taxonomy organizes research into three main branches: Matrix Low-Rank Approximation Methods focus on classical streaming SVD, sketching, and robust factorizations for evolving matrices; Tensor Low-Rank Approximation Methods extend these ideas to higher-order data structures using Tucker and tensor-train decompositions; and Application-Driven Low-Rank Methods tailor approximation strategies to specific domains such as recommendation systems, visual tracking, and online learning. Representative works like Streaming Matrix Approximation[4] and StreamSVD[17] illustrate efficient matrix sketching, while Streaming Tensor Factorization[5] and Tucker Streaming[28] demonstrate how tensor methods handle multi-way streaming data. Across these branches, a recurring theme is the trade-off between computational efficiency, memory footprint, and approximation accuracy.

Several active lines of work explore contrasting priorities: some emphasize speed and scalability through randomized sketching (e.g., Randomized Low-Rank Techniques[12], FlashSVD[31]), while others prioritize robustness to outliers and missing entries (e.g., Online Robust NMF[6], Outlier-Robust Streaming Tensor[36]). A smaller cluster of studies addresses consistency and recourse minimization—ensuring that updates to the approximation do not drastically alter previously computed results. Consistent Low-Rank[0] sits squarely within this latter cluster, emphasizing stability across streaming updates in a way that contrasts with purely accuracy-driven approaches like Fixed-Rank Streaming[7] or speed-focused methods such as Frequent Directions Tubal[23]. By prioritizing consistency, Consistent Low-Rank[0] addresses an emerging concern in deployment scenarios where abrupt changes in learned representations can disrupt downstream tasks, distinguishing it from neighbors that optimize traditional error metrics without explicit recourse constraints.

Related Works in Same Category

No sibling papers and no sibling subtopics were found under the same parent taxonomy node; the paper appears structurally isolated in the taxonomy.

Contributions Analysis

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

Contribution 1: Formalization of consistent low-rank approximation problem

Description: The authors formalize a new problem variant of low-rank approximation in the streaming setting that explicitly accounts for consistency by minimizing recourse (the cumulative change in output solutions over time) while maintaining approximation quality at each time step. This problem captures the practical need for stable feature representations in dynamic data environments.

This contribution was assessed against **5 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Online learning in the manifold of low-rank matrices

URL: View paper

Brief Assessment

Manifold Low-Rank Learning[49] focuses on online learning via Riemannian optimization on the manifold of low-rank matrices, without explicitly addressing consistency constraints or recourse minimization in streaming settings.

2. MFDROO: Matrix Factorization-Based Deep Reinforcement Learning Approach for Stable Online Offloading in Mobile Edge Networks

URL: View paper

Brief Assessment

MFDROO[35] focuses on mobile edge network offloading using matrix factorization for resource modeling, not on streaming low-rank approximation with consistency constraints or recourse minimization in dynamic data environments.

3. Robust Pca Via Adaptive Weighted Least Squares and Low-Rank Matrix Factorization

URL: View paper

Brief Assessment

Adaptive Weighted PCA[62] focuses on robust PCA methods using adaptive weighted least squares and low-rank matrix factorization for handling outliers and noise, not on the streaming consistency problem with recourse minimization that the original paper formalizes.

4. Low-rank incremental methods for computing dominant singular subspaces

URL: View paper

Brief Assessment

Low-Rank Incremental Methods[61] focuses on incremental SVD updates for maintaining low-rank factorizations in streaming settings, but does not explicitly formalize the consistency/recourse minimization problem that the original paper introduces.

5. Similarity based regularization for online matrix-factorization problem: An application to course recommender systems

URL: View paper

Brief Assessment

Similarity Regularization Matrix[29] focuses on course recommendation using low-rank matrix factorization with similarity regularization, not on streaming low-rank approximation with consistency constraints or recourse minimization.

Contribution 2: Algorithm achieving additive error with $O(k/\varepsilon \log(nd))$ recourse

Description: The authors develop an algorithm that maintains an additive approximation to the optimal low-rank cost at all times while achieving recourse that is sublinear in the number of rows and linear in the rank parameter k. This demonstrates that consistency is achievable even with strong approximation guarantees.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Rank overspecified robust matrix recovery: Subgradient method and exact recovery

URL: View paper

Brief Assessment

Rank Overspecified Recovery[69] addresses robust matrix recovery from corrupted measurements using subgradient methods, not the consistent low-rank approximation problem with streaming data and recourse minimization studied in the original paper.

2. On algorithms for weighted low rank approximation

URL: View paper

Brief Assessment

Weighted Low-Rank[52] focuses on weighted low-rank approximation with row sampling in a static setting, not on consistent/online low-rank approximation with recourse minimization in streaming data.

3. Adaptive sampling and fast low-rank matrix approximation

URL: View paper

Brief Assessment

Adaptive Sampling Fast[66] focuses on finding a subset of rows whose span contains a low-rank approximation, not on maintaining consistency/recourse in a streaming setting. The candidate achieves additive error but does not address the recourse metric central to the original paper's contribution.

4. Robust and sample optimal algorithms for PSD low rank approximation

URL: View paper

Brief Assessment

Robust PSD Approximation[67] addresses low-rank approximation for PSD matrices with different error guarantees (relative error) and does not focus on the consistent/recourse framework that is central to the original paper's contribution.

---

5. CUR Low Rank Approximation at Sub-linear Cost

URL: View paper

Brief Assessment

CUR Sublinear Cost[68] focuses on low-rank approximation via CUR decomposition with different error metrics and computational goals, not on the consistent low-rank approximation problem with recourse minimization in streaming settings.

---

6. Sublinear Time Low-Rank Approximation of Toeplitz Matrices

URL: View paper

Brief Assessment

Sublinear Toeplitz[70] addresses low-rank approximation of Toeplitz matrices with sublinear query complexity, not the consistent low-rank approximation problem with recourse bounds studied in the original paper.

---

7. A sublinear-time randomized algorithm for column and row subset selection based on strong rank-revealing QR factorizations

URL: View paper

Brief Assessment

Sublinear QR Selection[65] addresses column/row subset selection for low-rank approximation using SRRQR factorizations, not the dynamic streaming model with recourse constraints that the original paper studies.

---

8. Refinement of Low Rank Approximation of a Matrix at Sub-linear Cost

URL: View paper

Brief Assessment

Refinement Sublinear Cost[71] focuses on refining existing low-rank approximations through iterative methods, not on maintaining consistent approximations in streaming settings with recourse guarantees.

---

9. Low-Rank Approximation with Matrix-Vector Products

URL: View paper

Brief Assessment

Matrix-Vector Approximation[64] focuses on low-rank approximation using matrix-vector products with query complexity bounds, not on recourse or consistency in streaming settings. The candidate addresses a fundamentally different problem setting.

---

10. Sublinear Algorithms for Matrices: Theory and Applications

URL: View paper

Brief Assessment

Sublinear Matrix Algorithms[63] focuses on eigenvalue and singular value approximation via random sampling and spectral approximation, not on consistent low-rank approximation with recourse guarantees in streaming settings.

Contribution 3: Algorithm achieving relative $(1+\varepsilon)$-approximation with $k^{(3/2)}/\varepsilon^2 \cdot \mathrm{polylog}(ndM)$ recourse

Description: The authors present an algorithm that achieves multiplicative $(1+\varepsilon)$-approximation to the optimal low-rank cost while improving recourse from quadratic to sub-quadratic in k for integer-bounded matrices. This is accomplished through careful case analysis on singular value contributions and online ridge leverage score sampling.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. On algorithms for weighted low rank approximation

URL: View paper

Brief Assessment

Weighted Low-Rank[52] addresses weighted low-rank approximation using row sampling for static matrices, not the dynamic streaming setting with recourse bounds that the original paper studies.

---

2. Provably useful kernel matrix approximation in linear time

URL: View paper

Brief Assessment

Kernel Matrix Approximation[58] focuses on kernel matrix approximation in linear time using ridge leverage score sampling, not on consistent low-rank approximation with recourse bounds for streaming data.

---

3. Low Rank Approximation Directed by Leverage Scores and Computed at Sub-linear Cost

URL: View paper

Brief Assessment

Leverage Score Directed[60] focuses on low-rank approximation via leverage score sampling for static matrices, not the dynamic consistent low-rank approximation problem with recourse minimization that the original paper addresses.

---

4. Subquadratic kronecker regression with applications to tensor decomposition

URL: View paper

Brief Assessment

Subquadratic Kronecker Regression[54] focuses on Kronecker regression for tensor decomposition, not consistent low-rank approximation with recourse guarantees in streaming settings.

---

5. One-pass additive-error subset selection for subspace approximation

URL: View paper

Brief Assessment

One-Pass Subset Selection[55] addresses ℓp subspace approximation with additive error guarantees in streaming settings, while the original paper focuses on consistent low-rank approximation with multiplicative (1+ε)-approximation and recourse bounds in dynamic data streams. The technical problems and algorithmic approaches differ fundamentally.

---

6. SPALS: Fast alternating least squares via implicit leverage scores sampling

URL: View paper

Brief Assessment

SPALS[56] focuses on tensor CP decomposition via alternating least squares with leverage score sampling for Khatri-Rao products, not on consistent low-rank matrix approximation with recourse bounds in streaming settings.

---

7. Computing Approximate Sensitivities

URL: View paper

Brief Assessment

Computing Sensitivities[53] focuses on approximating ℓp sensitivities for regression tasks using leverage score sampling, not on consistent low-rank approximation with recourse bounds for streaming matrices.

---

8. Input sparsity time low-rank approximation via ridge leverage score sampling

URL: View paper

Brief Assessment

Ridge Leverage Sampling[51] focuses on input sparsity time algorithms for low-rank approximation using sampling schemes, not on consistent low-rank approximation with recourse bounds in streaming settings.

---

9. Simpler is better: A comparative study of randomized algorithms for computing the cur decomposition

URL: View paper

Brief Assessment

Simpler CUR[57] focuses on CUR/interpolative decompositions using randomized pivoting schemes (sketching + pivoting on sketches), not on consistent low-rank approximation in streaming settings with recourse bounds. The technical approaches and problem settings are fundamentally different.

---

10. Fast Low-Rank Tensor Decomposition by Ridge Leverage Score Sampling

URL: View paper

Brief Assessment

Fast Tensor Leverage[59] focuses on Tucker tensor decomposition using ridge leverage score sampling for accelerating alternating least squares, not on consistent low-rank matrix approximation with recourse guarantees in streaming settings.

Appendix: Text Similarity Detection

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

References

- [0] Consistent Low-Rank Approximation View paper
- [1] Randomized algorithms for streaming low-rank approximation in tree tensor network format View paper
- [2] Online tensor low-rank representation for streaming data View paper
- [3] New subset selection algorithms for low rank approximation: Offline and online View paper
- [4] Streaming low-rank matrix approximation with an application to scientific simulation View paper
- [5] Streaming tensor factorization for infinite data sources View paper
- [6] Online nonnegative matrix factorization with robust stochastic approximation View paper
- [7] Fixed-rank approximation of a positive-semidefinite matrix from streaming data View paper
- [8] Low-rank extended Kalman filtering for online learning of neural networks from streaming data View paper
- [9] Subspace learning and imputation for streaming big data matrices and tensors View paper
- [10] Hyperbandit: Contextual bandit with hypernewtork for time-varying user preferences in streaming recommendation View paper
- [11] Tensor Train Factorization with Spatio-temporal Smoothness for Streaming Low-rank Tensor Completion View paper
- [12] Randomized techniques for low-rank approximation of matrix and tensors with applications View paper
- [13] Online forecasting matrix factorization View paper
- [14] Low-rank gradient descent for memory-efficient training of deep in-memory arrays View paper
- [15] Online matrix factorization hashing for large-scale image retrieval View paper
- [16] Online projective nonnegative matrix factorization for large datasets View paper
- [17] StreamSVD: Low-rank Approximation and Streaming Accelerator Co-design View paper
- [18] A sequential multilinear NystrÃ¶m algorithm for streaming low-rank approximation of tensors in Tucker format View paper
- [19] Stochastic approximation and memory-limited subspace tracking for Poisson streaming data View paper
- [20] Tracking Online Low-Rank Approximations of Higher-Order Incomplete Streaming Tensors View paper
- [21] Online Tensor Low-Rank Representation for Streaming Data Clustering View paper
- [22] Online Algorithms for Recovery of Low-Rank Parameter Matrix in Non-stationary Stochastic Systems View paper
- [23] Effective streaming low-tubal-rank tensor approximation via frequent directions View paper
- [24] Online statistical inference for matrix contextual bandit View paper
- [25] Few-shot data-driven algorithms for low rank approximation View paper
- [26] SOLD: Sub-optimal low-rank decomposition for efficient video segmentation View paper
- [27] Online Matrix Factorization via Broyden Updates View paper
- [28] Low-Rank Tucker Approximation of a Tensor From Streaming Data View paper
- [29] Similarity based regularization for online matrix-factorization problem: An application to course recommender systems View paper

- [30] Persona analytics: Analyzing the stability of online segments and content interests over time using non-negative matrix factorization View paper
- [31] FlashSVD: Memory-Efficient Inference with Streaming for Low-Rank Models View paper
- [32] Low rank approximation lower bounds in row-update streams View paper
- [33] Hierarchical online NMF for detecting and tracking topic hierarchies in a text stream View paper
- [34] Differentially private robust low-rank approximation View paper
- [35] MFDROO: Matrix Factorization-Based Deep Reinforcement Learning Approach for Stable Online Offloading in Mobile Edge Networks View paper
- [36] Outlier-Robust Multi-Aspect Streaming Tensor Completion and Factorization. View paper
- [37] Online optimization for max-norm regularization View paper
- [38] Relative errors for deterministic low-rank matrix approximations View paper
- [39] Variational Bayesian inference for robust streaming tensor factorization and completion View paper
- [40] Online Low Rank Matrix Completion View paper
- [41] Tensor-Based Sketching Method for the Low-Rank Approximation of Data Streams View paper
- [42] Streaming Low-Rank Matrix Data Assimilation and Change Identification View paper
- [43] Improved Algorithms for Low Rank Approximation from Sparsity View paper
- [44] Clustering event streams with low rank Hawkes processes View paper
- [45] An efficient, sparsity-preserving, online algorithm for low-rank approximation View paper
- [46] Compressive Online Robust Principal Component Analysis via - Minimization View paper
- [47] Visual tracking via online nonnegative matrix factorization View paper
- [48] Check for updates Online Matrix Factorization Hashing for Large-Scale Image Retrieval View paper
- [49] Online learning in the manifold of low-rank matrices View paper
- [50] Single-pass randomized QLP decomposition for low-rank approximation View paper
- [51] Input sparsity time low-rank approximation via ridge leverage score sampling View paper
- [52] On algorithms for weighted low rank approximation View paper
- [53] Computing Approximate Sensitivities View paper
- [54] Subquadratic kronecker regression with applications to tensor decomposition View paper
- [55] One-pass additive-error subset selection for subspace approximation View paper
- [56] SPALS: Fast alternating least squares via implicit leverage scores sampling View paper
- [57] Simpler is better: A comparative study of randomized algorithms for computing the cur decomposition View paper
- [58] Provably useful kernel matrix approximation in linear time View paper
- [59] Fast Low-Rank Tensor Decomposition by Ridge Leverage Score Sampling View paper
- [60] Low Rank Approximation Directed by Leverage Scores and Computed at Sub-linear Cost View paper
- [61] Low-rank incremental methods for computing dominant singular subspaces View paper
- [62] Robust Pca Via Adaptive Weighted Least Squares and Low-Rank Matrix Factorization View paper
- [63] Sublinear Algorithms for Matrices: Theory and Applications View paper
- [64] Low-Rank Approximation with Matrix-Vector Products View paper
- [65] A sublinear-time randomized algorithm for column and row subset selection based on strong rank-revealing QR factorizations View paper
- [66] Adaptive sampling and fast low-rank matrix approximation View paper
- [67] Robust and sample optimal algorithms for PSD low rank approximation View paper
- [68] CUR Low Rank Approximation at Sub-linear Cost View paper
- [69] Rank overspecified robust matrix recovery: Subgradient method and exact recovery View paper
- [70] Sublinear Time Low-Rank Approximation of Toeplitz Matrices View paper
- [71] Refinement of Low Rank Approximation of a Matrix at Sub-linear Cost View paper