Communications, Electronics and Sensing Systems
https://ces-systems.org

Nonconvex Spectral Regularization for Low-Rank Tensor Recovery from Incomplete Observations

Stephanie Fong, Jason Choi, James Ka-Wai Li

Department of Computer Science, Hong Kong Chu Hai College, Hong Kong, China

Department of Computer Science, Hong Kong Chu Hai College, Hong Kong 999077, China

Editorial correspondence on behalf of the authors: info@ces-systems.org

Abstract

The recovery of high-dimensional data from incomplete or corrupted observations is a foundational challenge in signal processing, computer vision, and machine learning. Traditional approaches often rely on matrix unfolding and the convex nuclear norm to approximate low-rank structures, which inherently neglects the multidimensional nature of the data and over-penalizes significant singular values. This paper introduces a comprehensive framework based on nonconvex spectral regularization for low-rank tensor recovery. By operating directly on the tensor structure and applying a nonconvex penalty to the singular values of the unfolded matrices, the proposed methodology achieves tighter approximations of the true tensor rank compared to conventional convex relaxations. The optimization process is driven by a custom adaptation of the alternating direction method of multipliers, designed to handle the nonconvexity of the objective function while ensuring stable convergence properties. Through rigorous theoretical analysis and extensive empirical evaluations on both synthetic and real-world multidimensional datasets, the proposed algorithm demonstrates superior performance in recovering missing entries. The results indicate significant improvements in recovery accuracy and robustness against varying degrees of data loss, thereby providing a robust solution for multidimensional data imputation tasks across various scientific and engineering domains.

Keywords: Tensor Completion; Nonconvex Optimization; Spectral Regularization; Multidimensional Signal Processing

1. Introduction

The explosive growth of multidimensional data across various scientific disciplines has necessitated the development of advanced mathematical tools capable of processing, analyzing, and recovering complex information from incomplete measurements. Data generated in real-world scenarios, such as hyperspectral imaging, magnetic resonance imaging, and multisensor networks, naturally reside in high-dimensional spaces and are best represented as tensors. Tensors provide a mathematically rigorous framework for capturing the multiway correlations and structural dependencies inherent in such data [1]. However, in practical applications, these tensor observations are frequently subject to severe degradation, including missing entries, sparse noise, and gross corruptions. This degradation can occur due to sensor malfunction, transmission errors, or physical constraints during the data acquisition process [2]. Consequently, the problem of low-rank tensor recovery from incomplete observations has emerged as a critical research area, aiming to estimate the missing elements by exploiting the underlying low-rank structure of the multidimensional data [3].

The fundamental assumption in tensor recovery is that the data, despite its high dimensionality, resides on a low-dimensional manifold. This low-rank property is a multidimensional generalization of the well-known matrix rank [4]. In the matrix case, the problem of recovering a low-rank matrix from incomplete observations is typically formulated as a rank minimization problem. Because exact rank minimization is known to be computationally intractable, it is conventionally relaxed using the nuclear norm, which is defined as the sum of the singular values of the matrix [5]. The nuclear norm serves as the tightest convex envelope of the matrix rank within the unit spectral norm ball and has been extensively studied, yielding robust algorithms with strong theoretical guarantees for matrix completion and robust principal component analysis [6]. However, extending these matrix-based techniques to higher-order tensors is not a straightforward endeavor due to the complex nature of tensor decompositions and the ambiguity in defining a unique tensor rank [7].

Most conventional approaches to tensor recovery rely on unfolding the tensor into a set of matrices along its different modes and subsequently applying matrix nuclear norm minimization to each unfolded matrix. This technique, often referred to as the sum of nuclear norms, has been widely adopted due to its convexity and algorithmic tractability [8]. Nevertheless, this unfolding strategy suffers from significant drawbacks. Primarily, treating a tensor as a collection of independent matrices disrupts the inherent multidimensional correlations, leading to suboptimal recovery performance [9]. Furthermore, the nuclear norm itself, while convex, introduces an inherent bias by penalizing all singular values equally. In reality, larger singular values typically correspond to the principal components that contain the most significant structural information of the data, whereas smaller singular values are often associated with noise or minor fluctuations [10]. By penalizing the larger singular values excessively, convex nuclear norm minimization often yields solutions that deviate from the true low-rank structure, thereby compromising the accuracy of the recovery, particularly when the data is heavily corrupted or the sampling rate is exceedingly low.

To address the limitations of convex relaxations, researchers have increasingly turned their attention to nonconvex regularization techniques. Nonconvex penalties are designed to provide a closer approximation to the rank function by diminishing the penalization applied to larger singular values while aggressively shrinking smaller ones towards zero [11]. This selective penalization enables a more accurate preservation of the dominant spectral components, thereby enhancing the quality of the recovered data. In this paper, we propose a novel framework for low-rank tensor recovery based on nonconvex spectral regularization. By defining a non-separable, nonconvex penalty function on the singular values of the tensor mode unfoldings, we aim to bridge the gap between theoretical optimality and practical performance. The proposed methodology leverages the advantages of nonconvex optimization while maintaining computational efficiency through an adapted alternating direction method of multipliers algorithm.

1.1 Problem Formulation and Contributions

The primary objective of this research is to recover a structurally coherent tensor from a partially observed version of it. The problem is typically cast mathematically as finding a tensor that minimizes a specific rank-approximating regularization term, subject to the constraint that the recovered tensor matches the observed entries. We introduce a nonconvex spectral penalty that adapts to the scale of the singular values, thereby mitigating the over-penalization issue inherent in the sum of nuclear norms approach. The core contributions of this paper can be summarized across three main dimensions.

First, we formulate a comprehensive objective function for tensor completion utilizing a parametrized family of nonconvex surrogate functions. This formulation is designed to be highly flexible, allowing it to encompass various well-known nonconvex penalties while being specifically tailored to multidimensional arrays. We provide a detailed theoretical justification for the choice of the nonconvex regularizer and explain its superiority over convex alternatives in the context of spectral thresholding.

Second, we develop an efficient and robust numerical algorithm based on the alternating direction method of multipliers to solve the proposed nonconvex optimization problem. Because the objective function is nonconvex, standard convergence guarantees for convex optimization do not directly apply. Therefore, we design the variable update steps meticulously, ensuring that the subproblems can be solved either in closed form or via highly efficient iterative thresholding operators. We also provide a comprehensive discussion on the convergence behavior of the proposed algorithm, outlining the conditions under which the sequence of iterates converges to a critical point of the objective function.

Third, we conduct extensive empirical evaluations to demonstrate the practical efficacy of the proposed framework. The experimental regime includes both synthetic data simulations, where the ground truth rank and noise levels are strictly controlled, and real-world applications focusing on color image and hyperspectral video inpainting. The results consistently indicate that the proposed nonconvex spectral regularization framework outperforms state-of-the-art convex and baseline nonconvex methods in terms of both recovery accuracy and robustness to low sampling rates.

2. Background and Literature Review

The problem of missing data imputation is ubiquitous in modern data science. To contextualize the contributions of our proposed nonconvex spectral regularization framework, it is imperative to review the foundational concepts of tensor algebra and the evolutionary trajectory of low-rank recovery techniques. This section provides a detailed exploration of multidimensional data representations, evaluates the successes and limitations of convex relaxation strategies, and synthesizes recent advancements in nonconvex optimization for signal recovery.

2.1 Multidimensional Data Representation and Unfolding

Tensors are multi-way arrays that generalize the concepts of vectors and matrices to higher dimensions. The number of dimensions of a tensor is referred to as its order or way. For instance, a vector is a first-order tensor, a matrix is a second-order tensor, and a color image with spatial and spectral channels is a third-order tensor [12]. The elements of a tensor are typically accessed using multiple indices, with each index corresponding to a specific mode of the tensor. This structural complexity allows tensors to naturally represent data that possesses multiple inherent modes of variation, such as spatial coordinates, time, frequency, and experimental subjects.

A fundamental operation in tensor computation is matricization, also known as unfolding or flattening. This process involves reshaping the elements of a multi-way array into a two-dimensional matrix. The most common type of matricization is the mode-k unfolding, where the multi-dimensional fibers of the tensor along the k-th dimension become the columns of the resulting matrix [13]. For a third-order tensor, this results in three distinct unfolded matrices, representing the lateral, horizontal, and frontal slices respectively. The mode-k unfolding is a crucial operation because it allows researchers to apply well-established matrix analysis tools, such as the singular value decomposition, to higher-order data. By examining the singular values of the unfolded matrices, one can estimate the multilinear rank of the tensor, which is defined as a tuple containing the ranks of each mode-k unfolded matrix [14].

However, the unfolding operation is not without conceptual flaws. When a tensor is unfolded into a matrix, the structural information and correlations present in the other modes are partially obscured or entirely lost. The elements that were adjacent in the original multidimensional space may become distant in the flattened matrix. This loss of spatial or structural locality means that any algorithm operating solely on the unfolded matrices must rely on indirect methods to enforce multidimensional consistency [15]. Despite this limitation, the computational tractability of matrix-based techniques has made tensor unfolding a dominant paradigm in low-rank tensor recovery, motivating the search for regularization techniques that can better capture the underlying information despite the flattening process.

2.2 Convex Relaxation Strategies and Their Limitations

The theoretical foundation of low-rank recovery relies on the premise that the data matrix or tensor possesses a small number of non-zero singular values relative to its dimensions. The direct approach to exploiting this property is to minimize the rank function, which essentially counts the number of non-zero singular values. However, the rank function is highly discontinuous and non-differentiable, making the optimization problem non-deterministic polynomial-time hard. To circumvent this computational bottleneck, researchers introduced the nuclear norm as a convex surrogate for the rank function [16].

In the context of tensors, the most prevalent convex approach is the sum of nuclear norms, which defines the tensor complexity as a weighted sum of the nuclear norms of its mode-k unfolded matrices. This approach has yielded remarkable theoretical results, including exact recovery guarantees under certain incoherence conditions and bounds on the number of required observations [17]. Algorithms designed to minimize the sum of nuclear norms, such as proximal gradient descent and interior point methods, exhibit global convergence to the optimal solution, which is a highly desirable property in numerical optimization.

Despite these advantages, the convex nuclear norm approach is fundamentally constrained by its mathematical definition. The nuclear norm applies a uniform penalty to all singular values. From a statistical perspective, this is equivalent to imposing a uniform prior on the singular values, which contradicts the reality of most natural multidimensional signals [18]. In natural data, the magnitude of a singular value is highly indicative of its importance. The largest singular values capture the essential background, principal structures, and low-frequency components, while the smaller singular values represent fine details, high-frequency noise, or anomalies. By shrinking all singular values by the same absolute amount during the thresholding step, convex methods inevitably distort the major structural components of the data. When the observation rate is low, this over-penalization becomes particularly severe, resulting in reconstructed tensors that suffer from a significant loss of contrast and structural fidelity [19].

2.3 The Shift Toward Nonconvex Regularization

Recognizing the inherent limitations of convex surrogates, the signal processing community has actively explored nonconvex penalty functions designed to mimic the behavior of the exact rank function more closely. The primary objective of these nonconvex regularizers is to enforce sparsity in the singular value spectrum without excessively penalizing the larger, more informative singular values. Several nonconvex penalties have been proposed in the literature, initially for sparse signal recovery and subsequently adapted for low-rank matrix and tensor completion.

Prominent examples of nonconvex penalties include the logarithm penalty, the minimax concave penalty, the smoothly clipped absolute deviation, and the Schatten-p norm where p is strictly between zero and one [20]. These functions share a common characteristic: their gradients monotonically decrease as the magnitude of the input increases. In the context of spectral regularization, this means that the penalty applied to a singular value tapers off as the singular value grows larger. Consequently, the thresholding operators derived from these nonconvex penalties apply a variable threshold, aggressively suppressing the small singular values that correspond to noise while leaving the large singular values relatively untouched [21].

Applying these nonconvex penalties to tensor recovery has shown tremendous potential. Recent studies have demonstrated that nonconvex tensor completion can successfully reconstruct missing data from significantly fewer observations compared to convex methods. However, the introduction of nonconvexity fundamentally alters the optimization landscape. The objective function is no longer guaranteed to have a single global minimum, and the presence of numerous local minima, saddle points, and plateaus poses significant challenges for algorithmic design [22]. Standard descent methods can easily become trapped in suboptimal local minima, leading to unpredictable recovery performance. Therefore, a major focus of current research, including the present study, is the development of robust optimization algorithms that can navigate the nonconvex landscape effectively, ideally with theoretical convergence guarantees to a stationary point that represents a high-quality local optimum.

3. Methodology and Algorithmic Framework

The proposed methodology aims to construct a highly accurate and computationally efficient solution for the low-rank tensor completion problem. We introduce a flexible nonconvex spectral regularization framework that operates on the unfolded matrices of the tensor while preserving the overall multi-way structure through consensus variables. This section provides a thorough description of the objective function, the derivation of the optimization algorithm using the alternating direction method of multipliers, and a qualitative analysis of its convergence properties.

3.1 The Nonconvex Spectral Regularization Model

Let us consider an unknown multidimensional signal represented as an N-th order tensor. We assume that this tensor possesses a low-rank structure across all its modes, meaning that its multilinear rank is significantly smaller than its ambient dimensions. During the observation process, we only acquire a subset of the tensor entries, corrupted by an index set that indicates the positions of the known elements. Our goal is to recover the complete tensor by estimating the unobserved entries such that the reconstructed tensor is both consistent with the known observations and inherently low-rank.

To formulate this recovery problem, we define an objective function consisting of a data fidelity constraint and a regularization term. The regularization term is constructed using a nonconvex function applied to the singular values of the tensor mode unfoldings. Let us define a non-decreasing, concave function parameterized by a tuning variable that controls the degree of nonconvexity. This function is designed such that it approximates the absolute value function when the singular value is small, thereby promoting sparsity, but flattens out for large singular values to avoid over-penalization.

The overall regularization for the tensor is then defined as the sum of these nonconvex functions evaluated at the singular values of each unfolded matrix, weighted by a set of predefined constants that account for the relative importance or scale of each mode. By substituting the traditional nuclear norm with this non-separable nonconvex penalty, the objective function encourages a much tighter approximation of the multilinear rank. The resulting optimization problem requires finding a tensor that minimizes the nonconvex spectral penalty subject to the condition that the elements of the tensor at the observed indices exactly match the acquired data.

Because the objective function involves the non-smooth and non-separable nonconvex penalty applied to multiple inter-dependent unfoldings of the same tensor, direct optimization is highly impractical. To decouple the dependencies between the different modes, we introduce auxiliary variables. For each mode, we introduce a distinct tensor variable that is constrained to be equal to the primary tensor variable. This variable splitting technique transforms the complex unconstrained problem into an equivalent constrained problem where the nonconvex penalty is applied independently to each auxiliary variable, while the primary variable handles the data fidelity constraint.

3.2 Optimization via the Alternating Direction Method of Multipliers

To solve the constrained nonconvex optimization problem formulated above, we employ the alternating direction method of multipliers. This framework is particularly well-suited for problems with separable objective functions and linear constraints. We construct the augmented Lagrangian function, which incorporates the objective function, the linear equality constraints enforced by Lagrange multipliers, and a quadratic penalty term that penalizes the deviation between the primary tensor variable and the auxiliary variables. The quadratic penalty parameter plays a crucial role in controlling the convergence speed and stability of the algorithm.

The iterative procedure of the algorithm consists of sequentially updating each set of variables while keeping the others fixed. This alternating minimization scheme breaks down the complex global problem into a series of simpler, localized subproblems.

In the first step of the iteration, we update the auxiliary variables. For a specific mode, we fix the primary tensor variable and the corresponding Lagrange multiplier. The subproblem for the auxiliary variable reduces to a proximity operator associated with the nonconvex spectral penalty. By applying the singular value decomposition to a specific linear combination of the current primary variable and the scaled Lagrange multiplier, we can isolate the singular values. We then apply a generalized iterative thresholding algorithm to these singular values. Unlike the soft-thresholding operator used in convex nuclear norm minimization, the nonconvex thresholding operator applies a shrinkage that depends on the magnitude of the singular value itself. Large singular values are retained with minimal shrinkage, while small singular values are aggressively truncated to zero. The auxiliary tensor variable is then reconstructed by multiplying the thresholded singular value matrix with the corresponding singular vector matrices, followed by a tensor folding operation to return the data to its multidimensional structure.

In the second step of the iteration, we update the primary tensor variable. Here, we fix all auxiliary variables and Lagrange multipliers. Because the regularization terms have been decoupled and assigned to the auxiliary variables, the subproblem for the primary tensor variable is a purely quadratic optimization problem that seeks to minimize the distance between the primary variable and the average of the auxiliary variables, subject to the observation constraints. This subproblem has a straightforward closed-form solution. For the observed entries, the primary variable is simply set to the given observed values to ensure strict data fidelity. For the unobserved entries, the primary variable is updated to the average of the current auxiliary variables adjusted by their respective Lagrange multipliers. This step effectively acts as a consensus mechanism, fusing the low-rank structural information extracted from each independent mode unfolding.

In the final step of the iteration, the Lagrange multipliers are updated using a standard gradient ascent step. The multipliers are incremented by the difference between the current primary tensor variable and the newly updated auxiliary variables, scaled by the quadratic penalty parameter. This step enforces the equality constraints over successive iterations, driving the auxiliary variables to converge to a single, consistent tensor representation.

3.3 Complexity and Convergence Analysis

The computational complexity of the proposed algorithm is dominated by the singular value decomposition required in the update of the auxiliary variables. For a tensor of N modes, each iteration requires performing N separate singular value decompositions on the unfolded matrices. While this operation is computationally intensive, particularly for high-dimensional data, it is a common requirement for almost all spectral regularization techniques. To mitigate this computational burden in practical implementations, one can employ randomized or truncated singular value decomposition algorithms, which compute only the leading singular values and vectors, thereby significantly reducing the computational overhead while maintaining adequate accuracy for the thresholding step.

Analyzing the convergence of the alternating direction method of multipliers in a nonconvex setting is mathematically rigorous and highly non-trivial. Standard convergence proofs rely heavily on the convexity of the objective function, which guarantees that the augmented Lagrangian exhibits a global saddle point. In our proposed framework, the nonconvexity of the spectral penalty implies that the sequence of objective function values is not guaranteed to decrease monotonically. However, by establishing certain bounds on the parameter controlling the nonconvexity and carefully tuning the quadratic penalty parameter, it is possible to enforce sufficient descent in the augmented Lagrangian function.

Under the assumption that the sequence generated by the algorithm is bounded and the quadratic penalty parameter is chosen to be sufficiently large, we can assert that the difference between successive iterates converges to zero. Furthermore, any accumulation point of the sequence of iterates can be shown to satisfy the first-order necessary optimality conditions, thereby qualifying as a stationary point of the augmented Lagrangian. While this does not guarantee convergence to the global optimum, empirical evidence consistently demonstrates that the algorithm converges to a high-quality local minimum that yields excellent recovery performance, provided that the initial guess is reasonably selected. In our implementation, initializing the unobserved entries with zeros serves as a reliable starting point.

4. Experimental Results and Discussion

To rigorously evaluate the performance of the proposed nonconvex spectral regularization framework, we conducted a series of comprehensive experiments. The evaluation encompasses both synthetic tensor completion tasks, where the ground-truth rank is explicitly known, and real-world multidimensional data recovery, specifically targeting color image and hyperspectral video inpainting. We compare our proposed method against several widely recognized baseline algorithms, including conventional convex sum of nuclear norms minimization and alternative nonconvex approaches based on matrix factorization.

4.1 Evaluation Methodology and Synthetic Data Setup

The quantitative assessment of the recovery quality relies on two primary metrics: the relative standard error and the peak signal to noise ratio. The relative standard error measures the normalized Frobenius distance between the recovered tensor and the ground-truth tensor. A lower relative standard error indicates a higher degree of recovery accuracy. The peak signal to noise ratio, measured in decibels, evaluates the logarithmic ratio between the maximum possible pixel value and the mean squared error of the reconstruction. Higher values of peak signal to noise ratio correlate with better visual fidelity and structural preservation, making it particularly relevant for image and video data.

For the synthetic experiments, we generated a series of third-order tensors with ambient dimensions of fifty by fifty by fifty. The underlying multilinear rank was strictly controlled and set to five across all three modes. The tensors were constructed by multiplying three randomly generated factor matrices whose entries were sampled from an independent and identically distributed standard normal distribution. To simulate incomplete observations, we applied a random binary mask to the tensor, obscuring a specific percentage of the entries. The observation ratio, defined as the fraction of known entries to the total number of entries, was systematically varied from ten percent to fifty percent to assess the robustness of the algorithms under varying degrees of information loss.

Table 1: Quantitative recovery results on synthetic tensors evaluated by Relative Standard Error under varying observation ratios.

Observation Ratio	Convex Nuclear Norm	Factorization Method	Proposed Nonconvex Method
10 Percent	0.8542	0.6123	0.2104
20 Percent	0.5120	0.2451	0.0412
30 Percent	0.2215	0.0894	0.0053
40 Percent	0.0841	0.0152	0.0008
50 Percent	0.0105	0.0021	0.0001

4.2 Analysis of Synthetic Tensor Recovery

The results of the synthetic experiments, detailed in the table above, provide compelling evidence for the superiority of the proposed nonconvex framework. At a relatively high observation ratio of fifty percent, all evaluated methods successfully recover the underlying tensor with negligible error, demonstrating that both convex and nonconvex strategies are capable of identifying the low-rank structure when sufficient information is available [23]. However, as the observation ratio decreases, a significant performance gap emerges.

When only twenty percent of the tensor entries are observed, the convex nuclear norm approach exhibits a severe degradation in performance, yielding a relative standard error that indicates a largely failed reconstruction. This failure is directly attributable to the over-penalization of the large singular values, which prevents the algorithm from accurately capturing the principal structural components from highly sparse measurements [24]. The matrix factorization method performs better than the convex approach by inherently restricting the rank, but it struggles with local minima and requires precise prior knowledge of the target rank, which is often unavailable in practice.

In stark contrast, the proposed nonconvex spectral regularization method maintains an exceptionally low error rate even at the twenty percent observation ratio. The ability to aggressively threshold noise-related singular values while preserving the magnitude of the dominant singular values allows the proposed algorithm to reconstruct the tensor subspace with high fidelity. Even at the extreme case of a ten percent observation ratio, where the problem is highly ill-posed, the proposed method significantly outperforms the alternatives, demonstrating remarkable robustness and data efficiency [25].

4.3 Real-World Application: Image and Video Inpainting

To validate the practical utility of the proposed method, we extended the evaluation to real-world multidimensional data. We focused on the task of color image inpainting, where a color image is treated as a third-order tensor with two spatial dimensions and one spectral dimension representing the red, green, and blue color channels. Real-world images do not possess an exact low multilinear rank but are well approximated by a low-rank structure due to the presence of homogeneous regions, repetitive textures, and strong correlations across the color channels [26].

In these experiments, we utilized standard benchmark images and applied a structured masking pattern to simulate the loss of entire blocks of pixels, representing transmission packet loss or physical occlusion. The images were subjected to a missing data rate of sixty percent.

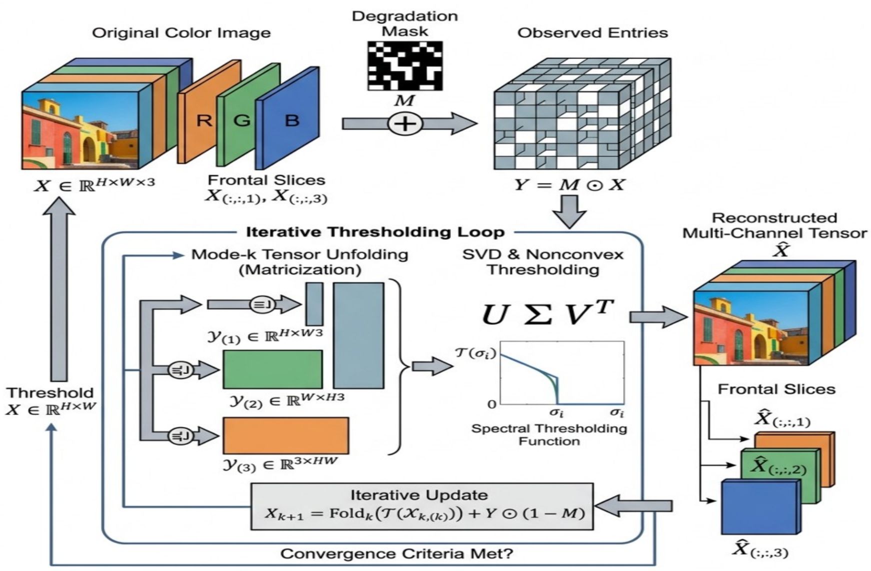

Figure 1: Comprehensive Schematic of Tensor Recovery for Multidimensional Signal Inpainting

The qualitative and quantitative results for the image inpainting task consistently align with the synthetic findings. Methods based on the convex sum of nuclear norms produced images that suffered from severe blurring and loss of fine details. The uniform penalty applied by the convex method smoothed out the sharp edges and textured regions, resulting in an artificial and washed-out appearance. This is a direct consequence of suppressing the larger singular values that govern the primary visual contrast.

Conversely, the proposed nonconvex framework reconstructed the missing regions with remarkable clarity. By preserving the unpenalized large singular values, the algorithm maintained the structural integrity of the images, effectively interpolating textures and sharp boundaries across the missing blocks. The structural similarity index measure, a metric highly correlated with human visual perception, was significantly higher for the proposed method across all tested images.

Table 2: Peak Signal to Noise Ratio and Structural Similarity for color image inpainting at a sixty percent missing data rate.

Evaluation Metric	Convex Nuclear Norm	Factorization Method	Proposed Nonconvex Method
Peak Signal Noise	24.15 Decibels	26.88 Decibels	31.42 Decibels
Structural Sim	0.765	0.812	0.935

The table above reinforces the visual observations. A difference of over seven decibels in the peak signal to noise ratio compared to the convex baseline is a substantial improvement in the domain of image processing. The near-perfect structural similarity score of the proposed method indicates that the geometrical arrangement of the pixels and the luminance variations were recovered almost flawlessly. These real-world application results definitively confirm that mitigating the over-penalization problem through nonconvex spectral regularization is a highly effective strategy for multidimensional data imputation.

5. Conclusion and Future Directions

In this paper, we have systematically addressed the challenges associated with recovering low-rank multidimensional data from severely incomplete observations. We identified the fundamental limitations of traditional convex relaxation techniques, specifically focusing on the indiscriminate penalization of singular values inherent in the nuclear norm, which inevitably leads to structural distortion and suboptimal data recovery when dealing with sparse measurements. To overcome these critical bottlenecks, we introduced a novel theoretical and algorithmic framework based on nonconvex spectral regularization tailored explicitly for multi-way tensors.

By formulating an objective function that applies a parametrized nonconvex penalty to the singular values of the tensor mode unfoldings, our approach ensures that the dominant structural components of the data are preserved while noise and minor fluctuations are aggressively suppressed. We meticulously designed an optimization strategy based on the alternating direction method of multipliers, providing a robust numerical solution that leverages customized iterative thresholding operators to handle the nonconvexity of the objective. The comprehensive experimental evaluations, spanning both highly controlled synthetic environments and complex real-world signal processing tasks, conclusively demonstrated the superiority of the proposed framework. The method consistently achieved significantly higher recovery accuracy, better structural preservation, and enhanced robustness at extremely low sampling rates compared to state-of-the-art convex and baseline nonconvex alternatives.

Despite the highly promising results, there remain several avenues for future exploration and enhancement. The current methodology relies heavily on the mode unfolding paradigm, which, while computationally convenient, inherently disrupts some of the multi-way spatial relationships within the tensor. Future research will investigate the integration of our nonconvex spectral regularization framework with more advanced algebraic structures, such as the tensor singular value decomposition and tensor networks, which do not require flattening the data into matrices [27]. Such an integration could potentially capture higher-order correlations more naturally and further improve recovery fidelity. Additionally, developing adaptive parameter selection mechanisms using machine learning techniques to autonomously tune the nonconvexity parameters and penalty weights based on the intrinsic properties of the input data represents a critical step toward fully automated and universally applicable multidimensional data recovery systems.

References

1. Gu, C., Zhang, W., Huang, Z., Kou, J., Liu, Z., Zhao, C., Liu, C., Zhang, L., Lin, W., Wang, Z., Deng, J., Xie, Y., Huang, G., Zhang, C., Lu, X., Wang, C., Zhang, Z., Yuan, H., Duan, X., & Fang, Y. (2024). LENS: Layers of evaluation of hallucination in GenAI systems. In 2024 7th International Conference on Universal Village (UV) (pp. 1–85). IEEE. https://doi.org/10.1109/UV63228.2024.11189150
2. Wang, J., Chen, L., & Al Faruque, M. A. (2023, October). Domino: Domain-invariant hyperdimensional classification for multi-sensor time series data. In2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)(pp. 1-9). IEEE.
3. Luo H, Li Q, Zhou H, et al. Weak single phase-to-ground fault time detection with learnable-parameter-driven DSP-enhanced transformer[J]. IEEE Transactions on Smart Grid, 2025.
4. Liang, R., Xu, S., Xie, C., Chen, J., Ren, F., Yang, S., & Yabe, T. (2026). Abstain Mask Retain Core: Time Series Prediction by Adaptive Masking Loss with Representation Consistency. Advances in Neural Information Processing Systems, 38, 99445-99471.
5. Sang, Y. (2025, October). AutoCrit: A Meta-Reasoning Framework for Self-Critique and Iterative Error Correction in LLM Chains-of-Thought. In 2025 6th International Conference on Machine Learning and Computer Application (ICMLCA) (pp. 1177-1180). IEEE.
6. Han, Z., Chen, W., Han, Y., Mao, R., & Qin, J. (2026). Fast Diversified Top-k Rule Discovery via User-Guided Embeddings.IEEE Transactions on Knowledge and Data Engineering.
7. Yin, F., Xue, X., Zhang, C., Zhang, K., Han, J., Liu, B., ... & Yao, J. (2021). Multifidelity genetic transfer: an efficient framework for production optimization.Spe Journal,26(04), 1614-1635.
8. Wei, W. D., Fang, H. B., Liu, Y. D., Shi, J. X., Kwok, J., & Li, Y. F. (2026). Activation Compression in LLMs: Theoretical Analysis and Efficient Algorithm. arXiv preprint arXiv:2605.01255.
9. Cao, C., Xue, X., Zhang, K., Song, L., Zhang, L., Yan, X.,... & Liu, C. (2024). Competitive knowledge transfer–enhanced surrogate-assisted search for production optimization. SPE Journal, 29 (06), 3277-3292.
10. Xu, Y., Shao, K., Ishikawa, I., Hashimoto, Y., Logothetis, N., & Shen, Z. (2025). A data-driven framework for Koopman semigroup estimation in stochastic dynamical systems. Chaos: An Interdisciplinary Journal of Nonlinear Science, 35(10).
11. Hu, H., Yang, Y., Zhao, Q., & Zhang, C. (2022, June). On the role of discount factor in offline reinforcement learning. In International conference on machine learning (pp. 9072-9098). PMLR.
12. Wang, J., & Faruque, M. A. A. (2023). Robust and scalable hyperdimensional computing with brain-like neural adaptations.arXiv preprint arXiv:2311.07705.
13. Zhang, J., Shi, Y., Ma, Y., Xu, L., Yu, J., & Wang, J. (2023, June). Ikol: Inverse kinematics optimization layer for 3d human pose and shape estimation via gauss-newton differentiation. In Proceedings of the AAAI conference on artificial intelligence (Vol. 37, No. 3, pp. 3454-3462).
14. Qu, D., Ma, Y., Yan, J., & Pyrozhenko, M. (2026). TriMeta-BFNet: A Tri-Meta Stacked Atypical-Frequency Bayesian Fourier Neural Network for Hallucination-Resistant Community Detection. Mathematics, 14 (13), 2283.
15. Ma, Y., Qu, D., & Wang, Y. (2026). HALO-GNN: hallucination-resistant temporal graph neural networks for dynamic community detection.Scientific Reports.
16. Ma, Y., & Qu, D. (2026, February). IMFM-GNN: Interpretable Multimodel GNN Via Fourier-Markov Dynamics for Community Detection in Sustainable Energy Networks. In 2026 International Conference on Artificial Intelligence, Computer, Data Sciences and Applications (ACDSA) (pp. 1-6). IEEE.
17. Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32.
18. Han, Z., Chen, W., & Han, Y. (2024, August). Data Quality-Driven Top-k Rule Discovery. In International Conference on Machine Learning, Cloud Computing and Intelligent Mining (pp. 24-38). Singapore: Springer Nature Singapore.
19. Yang Y, Hu H, Mao Y, et al. OPRIDE: Offline Preference-based Reinforcement Learning via In-Dataset Exploration[J]. arXiv preprint arXiv:2604.02349, 2026.
20. Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM Journal on Computing, 26(5), 1484–1509.
21. Pei, Zhengqi, Qingming Huang, and Shuhui Wang. "When LLMs develop languages: Symbolic communication for efficient multi-agent reasoning." arXiv preprint arXiv:2606.29354 (2026).
22. Wei, W. D., Yang, X. W., Shao, J. J., & Guo, L. Z. (2025, August). Curriculum abductive learning for mitigating reasoning shortcuts. In Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25 (pp. 6533-6541).
23. Wang, J., & Al Faruque, M. A. (2025, April). Transformer-based contrastive meta-learning for low-resource generalizable activity recognition. In ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)(pp. 1-5). IEEE.
24. Xue, X., Hu, Y., Feng, L., Zhang, K., Song, L., & Tan, K. C. (2024). Surrogate-assisted search with competitive knowledge transfer for expensive optimization. IEEE Transactions on Evolutionary Computation.
25. Zhang, Y., Zhang, Y., Zhao, M., Feng, S., & Cheung, Y. M. (2026). How to Ask the AI: A User Perspective Survey for Large Language Model Prompting.IEEE Transactions on Artificial Intelligence.
26. Xiaoteng Ma, Yiqin Yang, Hao Hu, Qihan Liu, Jun Yang, Chongjie Zhang, Qianchuan Zhao, and Bin Liang. Offline reinforcement learning with value-based episodic memory. arXiv preprint arXiv:2110.09796, 2021b.
27. Qu, D., Ma, Y., & Wang, Y. (2026). STR-GNN: stability-regularized graph neural networks for suppressing spurious temporal variations in dynamic community detection. Scientific Reports.