Deep Tensor Factorization for Modeling Non-Uniform Temporal Dynamics with Adaptive Temporal Smoothing

Hongyu Chen*, Beining Bao*, Zichun Liu*, Tianshuo Wei*, Maolin Wang†, Xiaopeng Li‡ and Hing Cheung So§
*Department of Data Science, City University of Hong Kong, Kowloon, Hong Kong
Email: {hychan2355-c, beininbao2-c, zichunliu6-c, tianshwei2-c}@my.cityu.edu.hk
†The Hong Kong Institute of AI for Science (HKAI-Sci), City University of Hong Kong, Kowloon, Hong Kong
Email: morin.wang@cityu.edu.hk
‡Department of the State Key Laboratory of Radio Frequency Heterogeneous Integration, Shenzhen University, Shenzhen, China
Email: x.p.li@szu.edu.cn
§Department of Electrical Engineering, City University of Hong Kong, Kowloon, Hong Kong
Email: hcso@ee.cityu.edu.hk

***Abstract*—Recovering tensors with *tube-wise missing* data is challenging because discrete tensor models treat the third mode as an index set rather than a physically continuous coordinate, while uniform regularization often over-smooths structural boundaries. To address this issue, we propose the Deep Transform Framework with Adaptive Temporal Smoothing (DTF-ATS). The framework combines a Continuous-Coordinate CP (CC-CP) prior with a time-aware deep transform network. CC-CP replaces the discrete third-mode factor with a coordinate-conditioned neural generator, allowing the model to interpolate missing intervals along a continuous dimension. Adaptive Temporal Smoothing (ATS) further regularizes the recovered factors through a learnable gate driven by local factor variation, so that smooth regions are stabilized while edge-like transitions are preserved. Experiments on Yale B and MRI datasets show clear gains on Yale B and competitive performance on MRI, demonstrating the value of coupling continuous low-rank priors with adaptive temporal smoothing for tube-wise missing recovery.**

***Index Terms*—Tensor completion, continuous CP decomposition, adaptive temporal smoothing, deep nonlinear transform, implicit neural representation.**

I. Introduction

High-dimensional visual data, such as facial image stacks and volumetric MRI [1], are naturally modeled as tensors [2]. Recovering them from *tube-wise missing* degradation—where entire fibers along one mode are absent—remains challenging [3]. Existing low-rank tensor completion and deep unfolding methods usually treat the third mode as a discrete index set [4]–[6]. When that mode actually follows a continuous physical coordinate, such as illumination or anatomical depth, this discretization weakens interpolation over large missing intervals and often introduces artifacts [7], [8]. Standard regularizers such as TV may further over-smooth fine structures [9].

To address this mismatch, we propose the **Deep Transform Framework with Adaptive Temporal Smoothing (DTF-ATS)**. The method couples a **Continuous-Coordinate CP (CC-CP)** prior, which replaces the discrete third-mode factor with a coordinate-conditioned neural generator, and a time-aware deep transform network that captures local non-linear dependencies beyond low-rank structure.

Complementing the continuous prior, **Adaptive Temporal Smoothing (ATS)** acts on factor variations and coordinate gaps so that smooth regions are regularized while abrupt transitions are preserved. Experiments on Yale B [10] and MRI [11] show clearer gains on Yale B and competitive results on MRI. Our contributions are: 1) a dual-stream DTF-ATS framework for tube-wise missing tensor completion; 2) a problem-specific combination of CC-CP and ATS for continuous low-rank modeling with adaptive regularization; and 3) quantitative, qualitative, and ablation evidence supporting this design.

II. Methodology

A. Problem Formulation

Let $\mathcal{X} \in \mathbb{R}^{N_1 \times N_2 \times N_3}$ be the third-order target tensor (e.g., MRI or video). We observe a partial tensor $\mathcal{Y} = \mathcal{P}_\Omega(\mathcal{X})$, where Ω denotes the observed indices.

We focus on the *tube-wise missing* problem, where fibers along the third dimension are entirely lost. Unlike traditional methods treating the third mode as a discrete index set k, we model it as a continuous physical coordinate $t_k \in [0, 1]$ (e.g., time or depth). Our goal is to recover $\mathcal{X}$ by learning a continuous neural mapping that leverages the underlying manifold continuity to interpolate missing intervals.

B. Overall Framework

The DTF-ATS framework (Fig. 1) adopts a dual-stream architecture to capture both local details and global structures. The **Main Stream** contains one forward block and one backward block. The forward block uses a 3D convolution with kernel size $(1, 3, 3)$ and LeakyReLU to expand the third-mode channels from N_3 to $4N_3$, then applies time-aware

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. This is the accepted version of the paper published in the 2026 IEEE 14th Sensor Array and Multichannel Signal Processing Workshop (SAM), DOI: 10.1109/SAM68950.2026.11687944.

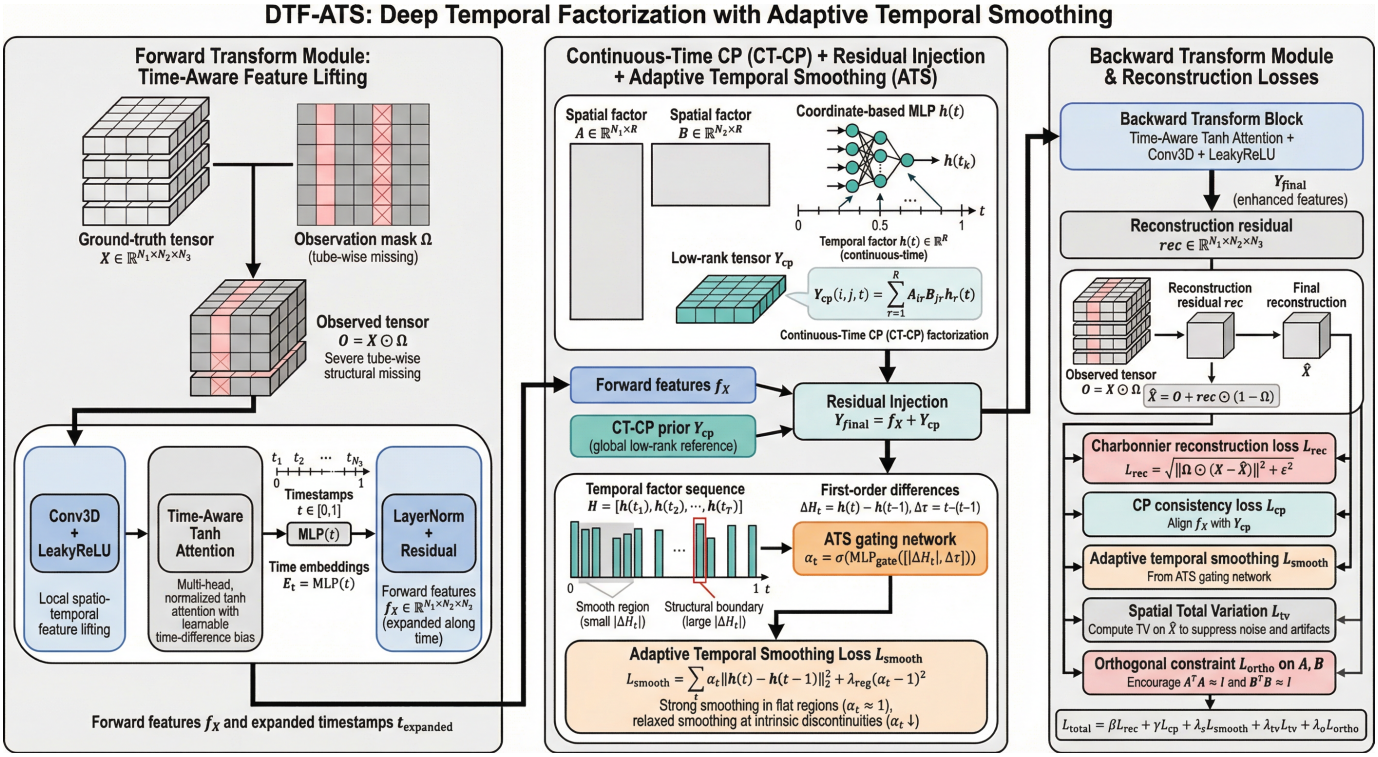

Fig. 1: **Overview of DTF-ATS.** The framework combines a time-aware Main Stream, a continuous CC-CP prior stream, and the Adaptive Temporal Smoothing (ATS) gate for tube-wise missing recovery.

tanh attention and residual LayerNorm. The backward block uses the same attention-normalization pattern before a second 3D convolution maps the features from $4N_3$ back to N_3. Thus, the Main Stream consists of two convolution layers, two attention layers, and residual normalization around attention, while the coordinate vector is linearly interpolated to match the expanded resolution.

The **Prior Stream (CC-CP)** generates a global continuous low-rank tensor $\mathcal{X}_{cp}$, which is additively injected into the Main Stream features before the backward block. The two streams are fused via residual injection:

$$\mathcal{X}_{\text{rec}} = \mathcal{B}(\mathcal{F}(\mathcal{Z}_{\text{latent}}, \mathbf{t}) + \mathcal{X}_{\text{cp}}, \mathbf{t}), \tag{1}$$

where $\mathcal{Z}_{\text{latent}}$ is the latent tensor and $\mathbf{t}$ denotes the vector of continuous coordinates.

C. Continuous-Coordinate CP (CC-CP) Module

Standard CP factors are discrete and cannot directly handle missing intervals. We therefore replace the discrete third-mode factor with a continuous generator $\mathbf{h}(t) = g_\theta(t)$, implemented by a lightweight time encoder and MLP. Given spatial factors $\mathbf{A}$ and $\mathbf{B}$, the tensor is reconstructed as:

$$\mathcal{X}_{\text{cp}}(i,j,t) = \sum_{r=1}^{R} \mathbf{A}_{ir}\mathbf{B}_{jr}\mathbf{h}_r(t). \tag{2}$$

This formulation allows querying the tensor at any coordinate $t \in \mathbf{t}$ and enables robust extrapolation over missing tubes.

D. Adaptive Temporal Smoothing (ATS)

Global temporal continuity may blur sharp structural edges. ATS employs a gating network $\mathcal{G}_\phi$ to output smoothness weights $\alpha_{t,r} = \sigma(\mathcal{G}_\phi([|\Delta h|, \Delta\tau]))$ based on factor gradients. The adaptive loss is:

$$\mathcal{L}_{\text{ats}} = \sum_{t,r} \alpha_{t,r}(\Delta h_{t,r})^2 + \lambda_{\text{reg}}(\alpha_{t,r} - 1)^2. \tag{3}$$

This mechanism allows the model to enforce strong smoothing in flat regions ($\alpha \approx 1$) while automatically "breaking" the continuity constraint at structural edges ($\alpha \to 0$), ensuring that critical details are not over-smoothed.

E. Loss Functions

We use Charbonnier loss $\mathcal{L}_{\text{fid}}$ for fidelity, combined with consistency ($\mathcal{L}_{\text{cp}} = \|\mathcal{F}_{\text{out}} - \mathcal{X}_{\text{cp}}\|^2$), orthogonality ($\mathcal{L}_{\text{ortho}}$), and spatial TV ($\mathcal{L}_{\text{tv}}$) constraints. The total objective is:

$$\mathcal{L}_{\text{total}} = \beta\mathcal{L}_{\text{fid}} + \gamma\mathcal{L}_{\text{cp}} + \lambda_s\mathcal{L}_{\text{ats}} + \lambda_{\text{tv}}\mathcal{L}_{\text{tv}} + \lambda_{\text{ortho}}\mathcal{L}_{\text{ortho}}. \tag{4}$$

The network parameters and the latent tensor $\mathcal{Z}_{\text{latent}}$ are optimized jointly during training.

III. Experiments

A. Settings & Implementation

Datasets. We evaluate DTF-ATS on two datasets whose third mode follows a physically meaningful continuous coordinate. 1) **Extended Yale B**: face images under continuously varying illumination, resized to 48×42 and stacked into a

TABLE I: Quantitative comparison on **MRI data** ($181 \times 217 \times 181$) with tube-wise missing entries. The best results are highlighted in **bold**.

Method	SR=0.1		SR=0.2		SR=0.3	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Observed	10.58	0.316	11.10	0.344	11.68	0.374
HaLRTC	16.17	0.261	19.96	0.460	23.05	0.619
TNN	18.96	0.278	22.49	0.460	25.52	0.615
TLRR	18.94	0.426	19.96	0.509	22.47	0.594
HLRTF	17.88	0.378	24.87	**0.688**	26.43	0.763
S2NTNN	18.92	**0.449**	20.00	0.530	21.53	0.591
CoNoT	18.49	0.415	19.31	0.480	23.48	0.614
TNN-G	18.95	0.425	19.96	0.507	20.74	0.553
KBR	17.76	0.240	25.53	0.611	29.73	0.781
DTF-ATS	**19.03**	0.334	25.42	0.660	**29.84**	**0.821**

tensor with 100 observations along the third mode. 2) **MRI Data** [11], [12] ($181 \times 217 \times 181$): a T1-weighted brain volume whose third mode represents anatomical depth. We simulate **tube-wise missing** by removing fibers at sampling rates (SRs) of $\{0.1, 0.3, 0.5\}$ for Yale B and $\{0.1, 0.2, 0.3\}$ for MRI.

Implementation. The model is implemented in PyTorch and trained by full-tensor optimization on a single GPU, jointly updating the network parameters and a latent tensor (effective batch size 1). We use Adam with learning rates of 2×10^{-3} and 2×10^{-2} for the network and latent tensor, weight decay 10^{-5} on the network, a cosine annealing schedule, and 3100 iterations per setting. The fidelity term uses the **Charbonnier loss**. The rank R is set to $\{18, 24, 32\}$ for Yale B at SRs $\{0.1, 0.3, 0.5\}$ and $\{56, 80, 96\}$ for MRI at SRs $\{0.1, 0.2, 0.3\}$; lower sampling rates use smaller R to impose a stronger low-rank bias, while higher sampling rates allow larger R for improved fitting capacity. A small sensitivity study on Yale B at $\mathrm{SR} = 0.3$ further showed that performance improves from small ranks and remains strong around $R = 32 \sim 40$, supporting the empirical rank tuning strategy used in our experiments. Representative parameter counts are 1.69M on Yale B at $\mathrm{SR} = 0.3$ and 2.52M on MRI at $\mathrm{SR} = 0.3$. The extra cost mainly comes from the unfolding backbone, the time-aware attention blocks, and the continuous-coordinate prior.

B. Comparisons

We compare DTF-ATS against representative tensor completion baselines, including traditional low-rank models (HaLRTC [13], TNN [14]), subspace and Kronecker-based methods (TLRR [15], KBR [16]), and advanced regularization or deep approaches (HLRTF [17], S2NTNN [18], CoNoT [19], and TNN-G [20]).

C. Quantitative Evaluation

Quantitative comparisons on **MRI data** are summarized in Table I. DTF-ATS achieves the best PSNR at $\mathrm{SR} = 0.1$ and $\mathrm{SR} = 0.3$, and remains competitive at $\mathrm{SR} = 0.2$, where KBR attains a slightly higher PSNR and HLRTF yields the highest SSIM.

TABLE II: Quantitative comparison on the **Yale B face dataset** ($48 \times 100 \times 42$) with tube-wise missing entries. The best results are highlighted in **bold**.

Method	SR=0.1		SR=0.3		SR=0.5	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Observed	7.52	0.035	8.51	0.096	10.01	0.184
HaLRTC	9.88	0.118	18.20	0.538	23.17	0.790
TNN	15.66	0.252	21.26	0.652	26.09	0.854
TLRR	18.04	0.387	20.97	0.644	23.50	0.785
HLRTF	19.59	0.510	23.27	0.746	26.12	0.852
S2NTNN	17.91	0.409	20.92	0.654	23.46	0.786
CoNoT	17.27	0.326	19.37	0.519	21.49	0.672
TNN-G	18.04	0.398	20.96	0.651	23.46	0.784
KBR	15.98	0.248	21.92	0.686	25.57	0.853
DTF-ATS	**19.79**	**0.542**	**24.58**	**0.810**	**28.99**	**0.925**

TABLE III: Ablation study of different components on the Yale B dataset ($\mathrm{SR} = 0.1$). Bold numbers indicate the best performance.

Component	**Setting**	**PSNR (dB)** ↑	**SSIM** ↑
Full Model	**CC-CP + ATS**	**19.79**	**0.5417**
CC-CP Effectiveness	w/ Discrete CP	19.75	0.5399
	w/ Static CP	19.52	0.5285
ATS Mechanism	w/o ATS	19.01	0.5022
	w/ Standard TV	19.18	0.4977
Architecture	Main Stream Only	19.69	0.5406
	Prior Stream Only	15.62	0.2843

Table II presents the results on **Extended Yale B**, where the illumination dimension more directly matches the continuous-coordinate assumption. Here the gains are clearer: DTF-ATS reaches 24.58 dB at $\mathrm{SR} = 0.3$ and 28.99 dB at $\mathrm{SR} = 0.5$, outperforming the strongest competing results by 1.31 dB and 2.87 dB, respectively.

D. Ablation Study

To validate the contribution of each core component, we conducted a comprehensive ablation study on the Yale B dataset (SR=0.1), as summarized in Table III.

We first investigate the efficacy of the proposed **CC-CP module**. By leveraging neural implicit representations, this continuous modeling strategy achieves a PSNR of 19.79 dB, surpassing both the discrete CP decomposition (19.75 dB) and the static factor matrix baseline (19.52 dB). This performance gain indicates that learning a continuous mapping from physical coordinates to latent factors provides a stronger regularization effect than traditional discrete embeddings, particularly when recovering structural information from severely undersampled data.

Next, we analyze the role of the **ATS** mechanism. The removal of the ATS module results in a substantial performance drop of 0.78 dB, confirming that regularization is indispensable for stabilizing reconstruction under extreme sparsity. Furthermore, our ATS outperforms conventional global TV regularization by 0.61 dB. Unlike standard TV, which imposes

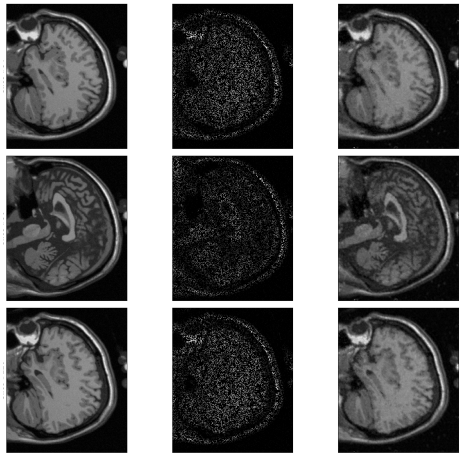

(a) MRI volumetric data at SR=0.3.

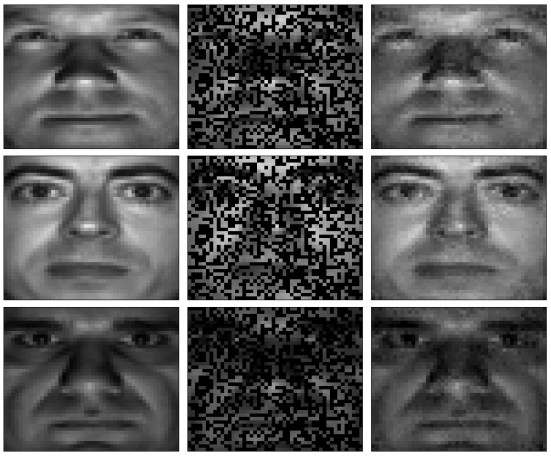

(b) Extended Yale B at SR=0.5.

Fig. 2: Qualitative results. From left to right: **Ground Truth**, **Observed**, and **Recovery**. DTF-ATS preserves anatomical structures and shadow transitions under severe tube-wise missing.

a uniform penalty that often leads to global over-smoothing, ATS incorporates structural awareness into the temporal domain. It dynamically adjusts smoothing intensity based on local structural variations, thereby effectively preserving sharp facial contours and critical shadow boundaries that are typically lost by content-agnostic regularizers.

Finally, we evaluate the synergy within the **Dual-Stream Architecture**. Relying solely on the global Prior Stream yields a poor PSNR of 15.62 dB, as the low-rank assumption alone struggles to capture complex high-frequency textures. While the non-linear Main Stream independently achieves a respectable 19.69 dB, integrating it with the Prior Stream in the full model boosts performance to 19.79 dB. This demonstrates that the residual injection of the global prior provides a robust structural scaffold, effectively complementing the deep transformation flow in capturing both global geometry and local details.

IV. Related Work

Tensor completion has been extensively studied, with traditional Low-Rank Tensor Completion (LRTC) methods, such as HaLRTC and TNN-based models, focusing on recovering data by minimizing rank [3], [4]. In recent years, deep learning approaches like unfolding networks [21] and nonlinear frameworks (HLRTF, LRTFR) have evolved to integrate convolutional or attention mechanisms for capturing complex dependencies [22], [23]. Despite these advancements, both paradigms typically rely on discrete algebraic assumptions, treating continuous dimensions as finite indices. This discretization overlooks physical continuity, which often leads to artifacts, particularly in tube-wise missing scenarios [24]. To incorporate continuous priors, neural implicit representations (neural fields) have emerged as a powerful paradigm, mapping coordinates to signals via MLPs to enable robust continuous inference [25]–[28]. Our CC-CP module adopts this philosophy to bridge tensor decomposition with continuous manifold priors [29]. Furthermore, addressing the ill-posed nature of reconstruction requires effective regularization. While standard Total Variation (TV) is widely used, it tends to produce staircasing artifacts and often over-smooths fine textures [30], [31]. Although adaptive smoothing techniques exist [32], [33], our ATS mechanism advances this direction via a learnable gating network that dynamically balances structural fidelity and noise suppression [34].

V. Conclusion

In this paper, we presented the Deep Transform Framework with Adaptive Temporal Smoothing (DTF-ATS) to address tensor completion under severe tube-wise missing patterns [35]–[37]. By integrating the Continuous-Coordinate CP (CC-CP) module, the framework models the third mode as a continuous neural function and therefore preserves physically meaningful continuity over large missing intervals [38], [39]. The Adaptive Temporal Smoothing (ATS) mechanism complements this prior by adjusting the regularization strength according to local factor variations, which helps suppress noise without erasing sharp structural boundaries [40]. Experiments on the Extended Yale B dataset and MRI volumetric data show that DTF-ATS yields the clearest gains on Yale B and competitive performance on MRI, indicating that the benefit is strongest when the third mode closely follows a smooth physical trajectory. These results support the value of coupling continuous low-rank priors with deep non-linear transforms, rather than claiming uniform dominance across all datasets. A current limitation is the additional computational cost introduced by the unfolding backbone, the time-aware attention blocks, and the continuous-coordinate prior. Future work will therefore focus on improving efficiency, profiling the computational tradeoffs more systematically, and extending the framework to dynamic video sequence completion [41].

REFERENCES

[1] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.

[2] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM review*, vol. 51, no. 3, pp. 455–500, 2009.

[3] J. Liu, P. Musialski, P. Wonka, and J. Ye, "Tensor completion for estimating missing values in visual data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 1, pp. 208–220, 2012.

[4] S. Gandy, B. Recht, and I. Yamada, "Tensor completion and low-n-rank tensor recovery via convex optimization," *Inverse problems*, vol. 27, no. 2, p. 025010, 2011.

[5] T. T. N. Mai, E. Y. Lam, and C. Lee, "Deep unrolled low-rank tensor completion for high dynamic range imaging," *IEEE Transactions on Image Processing*, vol. 31, pp. 5774–5787, 2022.

[6] Y. Liu and F. Shang, "An efficient matrix factorization method for tensor completion," *IEEE Signal Processing Letters*, vol. 20, no. 4, pp. 307–310, 2013.

[7] Y. Chen, S. Liu, and X. Wang, "Learning continuous image representation with local implicit image function," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 8628–8638.

[8] M. Sun, W. Wang, X. Zhu, and J. Liu, "Reparameterizing and dynamically quantizing image features for image generation," *Pattern Recognition*, vol. 146, p. 109962, 2024.

[9] Z. Tian, X. Jia, K. Yuan, T. Pan, and S. B. Jiang, "Low-dose CT reconstruction via edge-preserving total variation regularization," *Physics in Medicine & Biology*, vol. 56, no. 18, p. 5949, 2011.

[10] A. S. Georghiades, P. N. Belhumeur, and D. J. Kriegman, "From few to many: Illumination cone models for face recognition under variable lighting and pose," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 6, pp. 643–660, 2002.

[11] Y. Liu, Y. Shi, F. Mu, J. Cheng, C. Li, and X. Chen, "Multimodal mri volumetric data fusion with convolutional neural networks," *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–15, 2022.

[12] H.-M. Chen, H.-C. Chen, C. C.-C. Chen, Y.-C. Chang, Y.-Y. Wu, W.-H. Chen, C.-C. Sung, J.-W. Chai, and S.-K. Lee, "Comparison of Multispectral Image-Processing Methods for Brain Tissue Classification in BrainWeb Synthetic Data and Real MR Images," *BioMed Research International*, vol. 2021, no. 1, p. 9820145, 2021.

[13] L. Sun and Y. Lin, "Spectrum Completion Based on HaLRTC," in *2022 9th International Conference on Dependable Systems and Their Applications (DSA)*. IEEE, 2022, pp. 986–992.

[14] Z. Zhang and S. Aeron, "Exact tensor completion using t-SVD," *IEEE Transactions on Signal Processing*, vol. 65, no. 6, pp. 1511–1526, 2016.

[15] R. Wang, A. Kaul, and A. O. Sperry, "TLRR (lrrc67) interacts with PP1 and is associated with a cytoskeletal complex in the testis," *Biology of the Cell*, vol. 102, no. 3, pp. 173–189, 2010.

[16] Q. Xie, Q. Zhao, D. Meng, and Z. Xu, "Kronecker-basis-representation based tensor sparsity and its applications to tensor recovery," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 8, pp. 1888–1902, 2018.

[17] Y. Luo, X.-L. Zhao, D. Meng, and T.-X. Jiang, "HLRTF: Hierarchical low-rank tensor factorization for inverse problems in multi-dimensional imaging," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 19 303–19 312.

[18] Y.-S. Luo, X.-L. Zhao, T.-X. Jiang, Y. Chang, M. K. Ng, and C. Li, "Self-supervised nonlinear transform-based tensor nuclear norm for multi-dimensional image recovery," *IEEE Transactions on Image Processing*, vol. 31, pp. 3793–3808, 2022.

[19] J.-L. Wang, T.-Z. Huang, X.-L. Zhao, Y.-S. Luo, and T.-X. Jiang, "CoNoT: Coupled Nonlinear Transform-Based Low-Rank Tensor Representation for Multidimensional Image Completion," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 35, no. 7, pp. 8969–8983, 2024.

[20] B.-Z. Li, X.-L. Zhao, X. Zhang, T.-Y. Ji, X. Chen, and M. K. Ng, "A learnable group-tube transform induced tensor nuclear norm and its application for tensor completion," *SIAM Journal on Imaging Sciences*, vol. 16, no. 3, pp. 1370–1397, 2023.

[21] Y. Yang, J. Sun, H. Li, and Z. Xu, "ADMM-CSNet: A deep learning approach for image compressive sensing," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 3, pp. 521–538, 2018.

[22] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[23] Y.-B. Zheng, T.-Z. Huang, X.-L. Zhao, Q. Zhao, and T.-X. Jiang, "Fully-connected tensor network decomposition and its application to higher-order tensor completion," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 12, 2021, pp. 11 071–11 078.

[24] B. Li, F. Cheng, H. Zheng, Z. Shi, and C. Zhou, "Structural-missing tensor completion for robust DOA estimation with sensor failure," *Applied Sciences*, vol. 13, no. 23, p. 12740, 2023.

[25] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," *Advances in Neural Information Processing Systems*, vol. 33, pp. 7462–7473, 2020.

[26] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.

[27] Y. Xie, T. Takikawa, S. Saito, O. Litany, S. Yan, N. Khan, F. Tombari, J. Tompkin, V. Sitzmann, and S. Sridhar, "Neural fields in visual computing and beyond," in *Computer Graphics Forum*, vol. 41, no. 2. Wiley Online Library, 2022, pp. 641–676.

[28] M. Tancik, B. Mildenhall, T. Wang, D. Schmidt, P. P. Srinivasan, J. T. Barron, and R. Ng, "Learned initializations for optimizing coordinate-based neural representations," in *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 2021, pp. 2846–2855.

[29] C. Battaglino, G. Ballard, and T. G. Kolda, "A practical randomized CP tensor decomposition," *SIAM Journal on Matrix Analysis and Applications*, vol. 39, no. 2, pp. 876–901, 2018.

[30] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1-4, pp. 259–268, 1992.

[31] E. Y. Sidky, C.-M. Kao, and X. Pan, "Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT," *Journal of X-ray Science and Technology*, vol. 14, no. 2, pp. 119–139, 2006.

[32] D. Strong and T. Chan, "Edge-preserving and scale-dependent properties of total variation regularization," *Inverse Problems*, vol. 19, no. 6, p. S165, 2003.

[33] M. J. Black, G. Sapiro, D. H. Marimont, and D. Heeger, "Robust anisotropic diffusion," *IEEE Transactions on image processing*, vol. 7, no. 3, pp. 421–432, 1998.

[34] E. Kobler, A. Effland, K. Kunisch, and T. Pock, "Total deep variation for linear inverse problems," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 7549–7558.

[35] D. Spencer, R. Guhaniyogi, and R. Prado, "BAMITA: Bayesian multiple imputation for tensor arrays," *arXiv preprint*, vol. abs/2410.23412, 2024.

[36] L. Li, J. Zhang, Y. Wang, and B. Ran, "Missing data detection and imputation for urban ANPR system using an iterative tensor decomposition approach," *Transportation Research Part C: Emerging Technologies*, vol. 107, pp. 337–355, 2019.

[37] E. Acar, D. M. Dunlavy, T. G. Kolda, and M. Mørup, "Scalable tensor factorizations for incomplete data," *Chemometrics and Intelligent Laboratory Systems*, vol. 106, no. 1, pp. 41–56, 2011.

[38] X. Zhang, Y. Li, Y. Luo, Q. Xie, and D. Meng, "Online Functional Tensor Decomposition via Continual Learning for Streaming Data Completion," in *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

[39] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng, "Fourier features let networks learn high frequency functions in low dimensional domains," *Advances in neural information processing systems*, vol. 33, pp. 7537–7547, 2020.

[40] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 12, no. 7, pp. 629–639, 2002.

[41] S. S. Sofi and L. De Lathauwer, "Tensor train completion from fiberwise observations along a single mode," *arXiv preprint arXiv:2509.18149*, 2025.