

Recovery of Future Data via Convolution Nuclear Norm Minimization

Guangcan Liu, *Senior Member, IEEE*, Wayne Zhang

Abstract—This paper studies the problem of *time series forecasting* (TSF) from the perspective of *compressed sensing*. First of all, we convert TSF into a more inclusive problem called *tensor completion with arbitrary sampling* (TCAS), which is to restore a tensor from a subset of its entries sampled in an *arbitrary* manner. While it is known that, in the framework of Tucker low-rankness, it is theoretically impossible to identify the target tensor based on some arbitrarily selected entries, in this work we shall show that TCAS is indeed tackleable in the light of a new concept called *convolutional low-rankness*, which is a generalization of the well-known *Fourier sparsity*. Then we introduce a convex program termed Convolution Nuclear Norm Minimization (CNNM), and we prove that CNNM succeeds in solving TCAS as long as a sampling condition—which depends on the *convolution rank* of the target tensor—is obeyed. This theory provides a meaningful answer to the fundamental question of what is the minimum sampling size needed for making a given number of forecasts. Experiments on univariate time series, images and videos show encouraging results.

Index Terms—compressed sensing, sparsity and low-rankness, time series forecasting, Fourier transform, convolution.



1 INTRODUCTION

CAN we predict the future? While seems mysterious to the general public, this question is of particular interest to a wide range of scientific areas, ranging from mathematics, physics and philosophy to finance, meteorology and engineering. *Time series forecasting* (TSF) [1], which aims to predict future observations coming ahead of time based on historical data, plays an important role in the development of forecasting techniques. Formally, the problem can be described as follows:

Problem 1.1 (Time Series Forecasting). *Suppose that $\{\mathbf{M}_t\}_{t=1}^{p+h}$ is a sequence of $p+h$ order- $(n-1)$ tensors with $n \geq 1$, i.e., $\mathbf{M}_t \in \mathbb{R}^{m_1 \times \dots \times m_{n-1}}$. Given the historical part $\{\mathbf{M}_t\}_{t=1}^p$ consisting of p observed samples, the goal is to predict the next h unseen samples $\{\mathbf{M}_t\}_{t=p+1}^{p+h}$, where h is often called the ‘forecast horizon’.*

Usually, the forecast horizon h is provided by users and is determined according to the demands of planning or decision making. For example, in M4 Competition [2], h is a positive integer between 6 and 48. In general, the above problem has a wide scope that covers various settings of forecasting. When $n = 1$ or $n = 2$, Problem 1.1 corresponds to the classic *univariate* or *multivariate* TSF problems [1], respectively. In the case of $n = 3$, \mathbf{M}_t is of matrix-valued, thus Problem 1.1 embodies the *video prediction* task investigated in [3]. Interestingly, as one will see, the proposed methods can handle the general case of $n \geq 1$ in a universal way.

- G. Liu is with the School of Automation, Southeast University, Nanjing, China 210018. Email: gcliu1982@gmail.com.
- W. Zhang is with SenseTime Research, Harbour View 1 Podium, 12 Science Park East Ave, Hong Kong Science Park, Shatin, N.T., Hong Kong 999077. Email: wayne.zhang@sensetime.com.

Copyright (c) 2017 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Yet, it is worth mentioning that the most essential case is in fact $n = 1$, i.e., univariate TSF, as acknowledged by many forecasting competitions [2, 4].

Despite its brief definition, the TSF problem is extremely difficult to crack and still demands new scenarios after decades of research [5]. In particular, to our knowledge, the following fundamental question has not been previously answered:

Question 1 (Sampling Complexity). *What is the minimum sampling size p needed for predicting h future samples, or, equivalently, what is the maximum forecast horizon h when p historical samples are available?*

Seeking a valid answer to the above question is of great practical significance, especially given the increasing demand for long-term forecasting (i.e., h is large) [6]. To this end, we would like to try approaching the TSF problem from the perspective of *compressed sensing* [7–9]. Namely, we convert Problem 1.1 into a more general problem as follows. For a tensor-valued time series $\{\mathbf{M}_t\}_{t=1}^{p+h}$ with $\mathbf{M}_t \in \mathbb{R}^{m_1 \times \dots \times m_{n-1}}$, we define an order- n tensor $\mathbf{L}_0 \in \mathbb{R}^{m_1 \times \dots \times m_n}$ as

$$[\mathbf{L}_0]_{i_1, \dots, i_n} = [\mathbf{M}_{i_n}]_{i_1, \dots, i_{n-1}}, 1 \leq i_j \leq m_j, 1 \leq j \leq n, \quad (1)$$

where $m_n = p + h$ and $[\cdot]_{i_1, \dots, i_n}$ denotes the (i_1, \dots, i_n) th entry of an order- n tensor. In other words, \mathbf{L}_0 is formed by concatenating a sequence of order- $(n-1)$ tensors into an order- n one. Subsequently, we define a sampling set $\Omega \subset \{1, \dots, m_1\} \times \dots \times \{1, \dots, m_n\}$ as in the following:

$$\Omega = \{(i_1, \dots, i_n) : 1 \leq i_j \leq m_j, \forall 1 \leq j \leq n-1 \text{ and } 1 \leq i_n \leq p\}. \quad (2)$$

That is, Ω is a set consisting of the locations of the observations in $\{\mathbf{M}_t\}_{t=1}^p$, while the complement set of Ω , denoted as Ω^\perp , stores the locations of the future entries in $\{\mathbf{M}_t\}_{t=p+1}^{p+h}$.

With these notations, we turn to a more inclusive problem named *tensor completion with arbitrary sampling* (TCAS).

Problem 1.2 (Tensor Completion with Arbitrary Sampling). *Let $\mathbf{L}_0 \in \mathbb{R}^{m_1 \times \dots \times m_n}$ be the target tensor of interest. Suppose that we are given a subset of the entries in \mathbf{L}_0 and a sampling set $\Omega \subset \{1, \dots, m_1\} \times \dots \times \{1, \dots, m_n\}$ consisting of the locations of observed entries. The configuration of Ω is quite progressive, in a sense that the locations of the observed entries are allowed to be distributed **arbitrarily**. Can we restore the target tensor \mathbf{L}_0 ? If so, under which conditions?*

In this way, Problem 1.1 is incorporated into the scope of *tensor completion* [10–13], which is to fill in the missing entries of a partially observed tensor. Such a scenario has an immediate advantage; that is, it becomes straightforward to handle the difficult cases where the historical part $\{\mathbf{M}_t\}_{t=1}^p$ itself is incomplete. This is because, regardless of whether the historical part is complete or not, the task always boils down to recovering \mathbf{L}_0 from $\mathcal{P}_\Omega(\mathbf{L}_0)$, where $\mathcal{P}_\Omega(\cdot)$ is the orthogonal projection onto the subspace of tensors supported on Ω . Moreover, to obtain an accurate answer to Question 1, there is no much loss to consider instead Problem 1.2: The success condition for solving TCAS is dominated by the worst configuration of Ω , and, coincidentally, the sampling pattern of forecasting is typically the worst case, as we will confirm in Section 5.1.

While appealing, the goal of TCAS seems “unrealistic” according to the existing tensor completion theories [14–18]. Consider the case of $n = 2$, i.e., matrix completion with arbitrary sampling. In this case, for the target matrix \mathbf{L}_0 to be recoverable from a subset of its entries, it has been shown in [19, 20] that the together of \mathbf{L}_0 and Ω ought to obey an *isomeric condition*, which is obviously violated by the setup of arbitrary sampling—note that arbitrary sampling allows some columns and rows of \mathbf{L}_0 to be *wholly missing*. Yet, we would like to clarify that the isomeric condition is indispensable only when the circumstance is restricted within the scope of Low-Rank Matrix Completion (LRMC) [14], where the target \mathbf{L}_0 has the lowest Tucker rank among all the matrices that satisfy the constraint given by the observed entries. Beyond the field of Tucker low-rankness, as we will show, it is indeed possible to solve TCAS without imposing any restrictions on the sampling pattern.

Interestingly, the tool for addressing TCAS had already been established several decades ago, namely the prominent *Fourier sparsity*—the phenomenon that most of the Fourier coefficients of a signal are zero or approximately so. Concretely, TCAS can be addressed by a straightforward, but not specifically studied method termed Discrete Fourier Transform (DFT) based ℓ_1 minimization (DFT $_{\ell_1}$):

$$\min_{\mathbf{L}} \|\mathcal{F}(\mathbf{L})\|_1, \text{ s.t. } \mathcal{P}_\Omega(\mathbf{L} - \mathbf{L}_0) = 0, \quad (3)$$

where $\mathcal{F}(\cdot)$ is the DFT operator, and $\|\cdot\|_1$ denotes the ℓ_1 norm of a tensor seen as a long vector. Hereafter, denote by $\text{card}(\cdot)$ and $\text{vec}(\cdot)$ the cardinality of a set and the vectorization of a tensor, respectively. Let $\mathbf{z} = \text{vec}(\mathcal{F}(\mathbf{L}))$ in (3). Then it is easy to see that DFT $_{\ell_1}$ is a special case of the well-known ℓ_1 minimization problem:

$$\min_{\mathbf{z} \in \mathbb{C}^m} \|\mathbf{z}\|_1, \text{ s.t. } \Phi \mathbf{z} = \mathbf{y}, \quad (4)$$

where $m = \prod_{i=1}^n m_i$, $\Phi \in \mathbb{C}^{\tilde{p} \times m}$ is called the *sensing matrix*, and $\mathbf{y} \in \mathbb{C}^{\tilde{p}}$ is a vector consisting of \tilde{p} observations. The general case of (4) had been extensively studied in the literature, e.g., [7–9, 21, 22]. As for our DFT $_{\ell_1}$, it is configured that $\tilde{p} = \text{card}(\Omega)$, Φ is determined by the together of the *Fourier transform matrices* (see Section 2.2) and the sampling set Ω , and \mathbf{y} is composed of the observed entries in \mathbf{L}_0 . In this particular case, whenever $n = 1$ (i.e., \mathbf{L}_0 is a vector) and the observed entries are chosen randomly or specifically, as proven in [23, 24], the sensing matrix Φ satisfies the Restricted Isometry Property (RIP) [8] and thus program (4) can be successful, implying that DFT $_{\ell_1}$ can identify the target \mathbf{L}_0 . These results, however, are inapplicable to TCAS, in which the sampling pattern is allowed to be *arbitrary* and n may be any natural number. Moreover, as we will show in Section 3.3, the optimization problem in (3) can be efficiently solved within $\mathcal{O}(m \log m)$ time, whereas, by sharp contrast, $\mathcal{O}(m^3)$ time is usually needed for finding the optimal solution to (4).




Fig. 1. On the sampling condition for solving TCAS. (a) A 1000-dimensional signal we wish to recover. (b) 100 observations selected uniformly at random. (c) 100 observations chosen according to the setup of forecasting. While it should be doubtless that the signal in (a) can be reconstructed based on the observations in (b), it is not so real that the observations in (c) are sufficient for identifying the signal in (a).

We will prove that DFT $_{\ell_1}$ strictly succeeds in recovering the target \mathbf{L}_0 , as long as the sampling complexity, $\rho_0 = \text{card}(\Omega)/m$ ($m = \prod_{i=1}^n m_i$ and $\text{card}(\Omega)$ is the number of observed entries), satisfies

$$\rho_0 > 1 - \mathcal{O}(1/\|\mathcal{F}(\mathbf{L}_0)\|_0), \quad (5)$$

where $\|\cdot\|_0$ is the ℓ_0 norm, i.e., the number of nonzero entries of a tensor. As one may have noticed, the complexity $1 - \mathcal{O}(1/\|\mathcal{F}(\mathbf{L}_0)\|_0)$ could be very close to 1. Such a seemingly “cumbersome” sampling condition, in fact, is not caused by the limits of mathematical proofs but instead due to the nature of the setup of arbitrary sampling, as illustrated in Figure 1. Notably, the condition is useful for forecasting, as it gives a reasonable answer to Question 1. Consider the setup of Problem 1.1, where the observed samples are complete. Whenever $\mathcal{F}(\mathbf{L}_0)$ is sparse in a sense that $\|\mathcal{F}(\mathbf{L}_0)\|_0 \leq \mathcal{O}(m^\beta)$ with $0 \leq \beta < 1$, the condition $\rho_0 > 1 - \mathcal{O}(1/\|\mathcal{F}(\mathbf{L}_0)\|_0)$ requires

$$p > \mathcal{O}(\tilde{m}^{\beta/(1-\beta)} h^{1/(1-\beta)}), \quad (6)$$

where $\tilde{m} = \prod_{j=1}^{n-1} m_j$ is the dimension of the tensor \mathbf{M}_t , $1 \leq t \leq p + h$. In other words, the minimum sampling size grows polynomially as the forecast horizon h , and thus long-term forecasting is entirely possible—but would be rather challenging.

While theoretically effective and computationally efficient, DFT_{ℓ_1} suffers from a drawback that the time dimension is treated in the same way as the non-time dimensions, probably leading to imperfect results. To achieve better recovery performance, we further propose a regularizer called *convolution nuclear norm*, which is the sum of *convolution eigenvalues* [25], and which are generalization of Fourier frequencies. Specifically, the convolution nuclear norm of a tensor is the *nuclear norm* [26, 27] of its *convolution matrix* (see Section 2.3), depicting the Tucker low-rankness of the convolution matrix of the tensor—this is the so-called *convolutional low-rankness*. The derived method, Convolution Nuclear Norm Minimization (CNNM), attempts to tackle TCAS by

$$\min_{\mathbf{L}} \|\mathcal{A}_k(\mathbf{L})\|_*, \text{ s.t. } \mathcal{P}_\Omega(\mathbf{L} - \mathbf{L}_0) = 0, \quad (7)$$

where $\|\cdot\|_*$ denotes the nuclear norm of a matrix, $\mathcal{A}_k(\cdot)$ is a linear map from $\mathbb{R}^{m_1 \times \dots \times m_n}$ to $\mathbb{R}^{m \times k}$ ($m = \prod_{i=1}^n m_i$) such that $\mathcal{A}_k(\mathbf{L})$ produces the convolution matrix of \mathbf{L} , $k = \prod_{i=1}^n k_i$ is the kernel size used in defining $\mathcal{A}_k(\cdot)$, and $\{k_i : 1 \leq k_i \leq m_i\}_{i=1}^n$ is a set of n hyper-parameters (usually $k_i = \mathcal{O}(m_i)$). When $k_i = m_i, \forall i$, the convolution nuclear norm is exactly the ℓ_1 norm of the Fourier transform, and thus CNNM falls back to DFT_{ℓ_1} .

As expected, CNNM also guarantees to recover the target \mathbf{L}_0 from a subset of its entries chosen arbitrarily, as long as

$$\rho_0 > 1 - \mathcal{O}(1/r_k(\mathbf{L}_0)), \quad (8)$$

where $r_k(\cdot)$, called *convolution rank*, is the rank of the convolution matrix of a tensor. Generally speaking, the lower bound of ρ_0 , called the *sampling bound*, is functionally dependent on the kernel size k . Whenever $k_i = m_i, \forall i$ and thus $k = m$, we have $r_m(\mathbf{L}_0) = \|\mathcal{F}(\mathbf{L}_0)\|_0$ (see Section 2.4). In this case, the sampling condition also rolls back to $\rho_0 > 1 - \mathcal{O}(1/\|\mathcal{F}(\mathbf{L}_0)\|_0)$. By choosing proper kernel sizes with an attempt to minimize the sampling bound (see Section 3.4), CNNM may outperform dramatically DFT_{ℓ_1} , in terms of recovery accuracy. To summarize, the contributions of this paper mainly include:

- ◊ We explore the TSF problem from the view point of a new, inclusive problem named TCAS, and propose a novel method termed CNNM for studying TCAS, accomplishing an effective method for TSF as well. Remarkably, for the first time, our analysis provides a meaningful answer to the fundamental question of how many historical samples are required for predicting a given number of future samples.
- ◊ Even though the general ℓ_1 minimization problem in (4) was investigated in a great many papers, DFT_{ℓ_1} has not been carefully explored before. To our knowledge, we are the first to prove that DFT_{ℓ_1} strictly succeeds in solving TCAS, as long as the sampling complexity ρ_0 surpasses certain threshold.
- ◊ The proposed convolutional low-rankness provides a new tool for generalizing and penetrating the classic concept of Fourier sparsity. Whenever a signal is smooth in the sense of circular shift, it is provable that the convolution matrix of the signal is approximately low-rank (see Section 3.2). This result helps

to explain the widely observed phenomenon that smooth, bounded signals (e.g., images and videos) are often approximately sparse in the Fourier domain.

The rest of this paper is organized as follows. Section 2 summarizes the mathematical notations used throughout the paper and introduces some preliminary knowledge, explaining the technical insights behind the proposed methods as well. Section 3 is mainly consist of theoretical analysis. Section 4 shows the mathematical proofs of the presented theorems. Section 5 demonstrates empirical results and Section 6 concludes this paper.

2 NOTATIONS AND PRELIMINARIES

2.1 Summary of Main Notations

Bold capital letters such as \mathbf{X} denote order- n tensors ($n \geq 1$), and single numbers are denoted by either lowercase or Greek letters. Capital letters (e.g., X) and bold lowercase letters (e.g., \mathbf{x}) are used to represent matrices and vectors, respectively. For an order- n tensor \mathbf{X} , $[\mathbf{X}]_{i_1, \dots, i_n}$ is the (i_1, \dots, i_n) th entry of \mathbf{X} . Two types of tensor norms are used frequently throughout the paper: the Frobenius norm given by $\|\mathbf{X}\|_F = \sqrt{\sum_{i_1, \dots, i_n} |[\mathbf{X}]_{i_1, \dots, i_n}|^2}$, and the ℓ_1 norm denoted by $\|\cdot\|_1$ and given by $\|\mathbf{X}\|_1 = \sum_{i_1, \dots, i_n} |[\mathbf{X}]_{i_1, \dots, i_n}|$, where $|\cdot|$ denotes the magnitude of a real or complex number. Another two frequently used norms are the *operator norm* and *nuclear norm* [26, 27] of order-2 tensors (i.e., matrices), denoted by $\|\cdot\|$ and $\|\cdot\|_*$, respectively.

Calligraphic letters, such as \mathcal{F}, \mathcal{P} and \mathcal{A} , are used to denote linear operators. In particular, \mathcal{I} denotes the identity operator and \mathbf{I} is the identity matrix. For a linear operator $\mathcal{L} : \mathbb{H}_1 \rightarrow \mathbb{H}_2$ between Hilbert spaces, its Hermitian adjoint (or conjugate) is denoted as \mathcal{L}^* and given by

$$\langle \mathcal{L}(\mathbf{X}), \mathbf{Y} \rangle_{\mathbb{H}_2} = \langle \mathbf{X}, \mathcal{L}^*(\mathbf{Y}) \rangle_{\mathbb{H}_1}, \forall \mathbf{X} \in \mathbb{H}_1, \mathbf{Y} \in \mathbb{H}_2, \quad (9)$$

where $\langle \cdot, \cdot \rangle_{\mathbb{H}_i}$ is the inner product in the Hilbert space \mathbb{H}_i ($i = 1$ or 2). But the subscript is omitted whenever \mathbb{H}_i refers to a Euclidian space.

The symbol Ω is reserved to denote the sampling set consisting of the locations of observed entries. For $\Omega \subset \{1, \dots, m_1\} \times \dots \times \{1, \dots, m_n\}$, its *mask tensor* is denoted by Θ_Ω and given by

$$[\Theta_\Omega]_{i_1, \dots, i_n} = \begin{cases} 1, & \text{if } (i_1, \dots, i_n) \in \Omega, \\ 0, & \text{otherwise.} \end{cases} \quad (10)$$

Denote by \mathcal{P}_Ω the orthogonal projection onto Ω . Then we have the following:

$$\mathcal{P}_\Omega(\mathbf{X}) = \Theta_\Omega \circ \mathbf{X} \quad \text{and} \quad \Omega = \text{supp}(\Theta_\Omega), \quad (11)$$

where \circ denotes the Hadamard product and $\text{supp}(\cdot)$ is the support set of a tensor.

In most cases, we work with real-valued matrices (i.e., order-2 tensors). For a matrix X , $[X]_{i,:}$ is its i th row, and $[X]_{:,j}$ is its j th column. Let $\omega = \{j_1, \dots, j_l\}$ be a 1D sampling set. Then $[X]_{\omega,:}$ denotes the submatrix of X obtained by selecting the rows with indices j_1, \dots, j_l , and similarly for $[X]_{:, \omega}$. For a 2D sampling set $\bar{\Omega} \subset \{1, \dots, \bar{m}_1\} \times \{1, \dots, \bar{m}_2\}$, we imagine it as a sparse matrix

and define its ‘‘rows’’, ‘‘columns’’ and ‘‘transpose’’ accordingly. The i th row of $\bar{\Omega}$ is denoted by $\bar{\Omega}_i$ and given by $\bar{\Omega}_i = \{i_2 : (i_1, i_2) \in \bar{\Omega}, i_1 = i\}$, the j th column is defined as $\bar{\Omega}^j = \{i_1 : (i_1, i_2) \in \bar{\Omega}, i_2 = j\}$, and the transpose is given by $\bar{\Omega}^T = \{(i_2, i_1) : (i_1, i_2) \in \bar{\Omega}\}$.

Letters U and V are reserved for the left and right singular vectors of matrices, respectively. The orthogonal projection onto the column space U is denoted by \mathcal{P}_U and given by $\mathcal{P}_U(X) = UU^T X$, and similarly for the row space $\mathcal{P}_V(X) = XVV^T$. The same notation is also used to represent a subspace of matrices, e.g., we say that $X \in \mathcal{P}_U$ for any matrix X obeying $\mathcal{P}_U(X) = X$.

2.2 Multi-Dimensional DFT

Multi-dimensional DFT, also known as multi-directional DFT, is one of the most widely-used tool for signal processing. Its definition is commonly available in some public documents. Here, we shall present a definition that would be easy for engineers to understand. First consider the case of $n = 1$, i.e., the DFT of a vector $\mathbf{x} \in \mathbb{R}^m$. In this particular case, $\mathcal{F}(\mathbf{x})$ can be simply expressed as

$$\mathcal{F}(\mathbf{x}) = U\mathbf{x}, \quad (12)$$

with $U \in \mathbb{C}^{m \times m}$ being a complex-valued, symmetric matrix that satisfies $U^H U = U U^H = m\mathbf{I}$, where $(\cdot)^H$ is the conjugate transpose of a complex-valued matrix. The matrix U , called the *Fourier transform matrix*, is consist of data-independent numbers. More precisely, the real and imaginary components of $[U]_{i,j}$ are $\cos(2\pi(i-1)(j-1)/m)$ and $-\sin(2\pi(i-1)(j-1)/m)$, respectively.

Similarly, when $n = 2$, the DFT of a matrix $X \in \mathbb{R}^{m_1 \times m_2}$ is given by

$$\mathcal{F}(X) = U_1 X U_2, \quad (13)$$

where $U_1 \in \mathbb{C}^{m_1 \times m_1}$ and $U_2 \in \mathbb{C}^{m_2 \times m_2}$ are the Fourier transform matrices for the columns and rows of X , respectively. For a general order- n ($n \geq 1$) tensor $\mathbf{X} \in \mathbb{R}^{m_1 \times \dots \times m_n}$, its DFT is given by

$$\mathcal{F}(\mathbf{X}) = \mathbf{X} \times_1 U_1 \cdots \times_n U_n, \quad (14)$$

where $U_i \in \mathbb{C}^{m_i \times m_i}$ is the Fourier transform matrix for the i th dimension, and \times_j ($1 \leq j \leq n$) is the *mode- j product* [28] between tensors and matrices.

2.3 Convolution Matrix

Discrete convolution, which is probably the most fundamental concept in signal processing, plays an important role in this paper. Its definition—though mostly unique—has multiple variants, depending on which *boundary condition* is used. What we consider in this paper is the *circular convolution*, i.e., the convolution equipped with *circular boundary condition* [29]. Let’s begin with the simple case of $n = 1$, i.e., the circular convolution procedure of converting $\mathbf{x} \in \mathbb{R}^m$ and $\mathbf{k} \in \mathbb{R}^k$ ($k \leq m$) into $\mathbf{x} \star \mathbf{k} \in \mathbb{R}^m$:

$$[\mathbf{x} \star \mathbf{k}]_i = \sum_{s=1}^k [\mathbf{x}]_{i-s} [\mathbf{k}]_s, i = 1, \dots, m, \quad (15)$$

where \star denotes the convolution operator, and it is assumed that $[\mathbf{x}]_{i-s} = [\mathbf{x}]_{i-s+m}$ for $i \leq s$; this is the so-called

circulant boundary condition. Throughout this paper, we assume $k \leq m$ and refer to the smaller tensor \mathbf{k} as the *kernel*. In general, the convolution operator is linear and can be converted into matrix multiplication:

$$\mathbf{x} \star \mathbf{k} = \mathcal{A}_k(\mathbf{x})\mathbf{k}, \forall \mathbf{x} \in \mathbb{R}^m, \mathbf{k} \in \mathbb{R}^k, \quad (16)$$

where $\mathcal{A}_k(\cdot)$ is the *convolution matrix* of a vector, and the subscript k is put to remind the readers that the convolution matrix is always associated with a certain kernel size k . In the light of circular convolution, the convolution matrix of a vector $\mathbf{x} = [x_1, \dots, x_m]^T$ (i.e., $[\mathbf{x}]_i = x_i$) is a truncated circular matrix of size $m \times k$:

$$\mathcal{A}_k(\mathbf{x}) = \begin{bmatrix} x_1 & x_m & \cdots & x_{m-k+2} \\ x_2 & x_1 & \cdots & x_{m-k+3} \\ \vdots & \vdots & \ddots & \vdots \\ x_m & x_{m-1} & \cdots & x_{m-k+1} \end{bmatrix}. \quad (17)$$

In other words, the j th column of $\mathcal{A}_k(\mathbf{x})$ is given by $\mathcal{S}^{j-1}(\mathbf{x})$, where \mathcal{S} is an operator that circularly shifts the elements of a vector by one position; namely,

$$\mathcal{S}(\mathbf{x}) = [x_m, x_1, x_2, \dots, x_{m-1}]^T. \quad (18)$$

This shift operator is implemented by the Matlab function ‘‘circshift’’. In the special case of $k = m$, the convolution matrix $\mathcal{A}_m(\mathbf{x})$ is exactly an $m \times m$ circular matrix.

Now we turn to the general case of any $n \geq 1$. Let $\mathbf{X} \in \mathbb{R}^{m_1 \times \dots \times m_n}$ and $\mathbf{K} \in \mathbb{R}^{k_1 \times \dots \times k_n}$ be two order- n tensors. Again, it is configured that $k_j \leq m_j, \forall 1 \leq j \leq n$, and \mathbf{K} is called the kernel. Then the procedure of circularly convoluting \mathbf{X} and \mathbf{K} into $\mathbf{X} \star \mathbf{K} \in \mathbb{R}^{m_1 \times \dots \times m_n}$ is given by

$$[\mathbf{X} \star \mathbf{K}]_{i_1, \dots, i_n} = \sum_{s_1, \dots, s_n} [\mathbf{X}]_{i_1-s_1, \dots, i_n-s_n} [\mathbf{K}]_{s_1, \dots, s_n}. \quad (19)$$

The above convolution procedure can be also converted into matrix multiplication. Let $\text{vec}(\cdot)$ be the vectorization of a tensor, then we have

$$\text{vec}(\mathbf{X} \star \mathbf{K}) = \mathcal{A}_k(\mathbf{X})\text{vec}(\mathbf{K}), \forall \mathbf{X}, \mathbf{K}, \quad (20)$$

where $\mathcal{A}_k(\mathbf{X}) \in \mathbb{R}^{m \times k}$, with $m = \prod_{j=1}^n m_j$ and $k = \prod_{j=1}^n k_j$, is the convolution matrix of the tensor \mathbf{X} . To compute the convolution matrix of an order- n tensor, one just needs to replace the one-directional circular shift operator given in (18) with a multi-directional one, so as to stay in step with the structure of high-order tensors. To be more precise, let $\mathcal{S}(\mathbf{X}, u, v)$ be the operator that circularly shifts the elements in the tensor \mathbf{X} by u positions along the v th dimension (or direction); this operator is also implemented by the Matlab function ‘‘circshift’’. For any index (t_1, \dots, t_n) with $1 \leq t_q \leq k_q$ ($1 \leq q \leq n$), define an invertible operator $\mathcal{T}_{(t_1, \dots, t_n)} : \mathbb{R}^{m_1 \times \dots \times m_n} \rightarrow \mathbb{R}^m$ as

$$\mathcal{T}_{(t_1, \dots, t_n)}(\mathbf{X}) = \text{vec}(\mathbf{X}_n), \quad (21)$$

where \mathbf{X}_n is determined by the following recursion rule:

$$\mathbf{X}_q = \mathcal{S}(\mathbf{X}_{q-1}, t_q - 1, q), 1 \leq q \leq n, \mathbf{X}_0 = \mathbf{X}. \quad (22)$$

Assume conveniently that $k_0 = 1$, and let $j = 1 + \sum_{q=1}^n (t_q - 1) \prod_{b=0}^{q-1} k_b$. Then it may be easily seen that j ranges from 1 to

k while t_q increases from 1 to k_q , $q = 1, \dots, n$, and the j th column of $\mathcal{A}_k(\mathbf{X})$ is given by

$$[\mathcal{A}_k(\mathbf{X})]_{:,j} = \mathcal{T}_{(t_1, \dots, t_n)}(\mathbf{X}). \quad (23)$$

As a result, when $k_i = m_i, \forall i$, the convolution matrix of an order- n tensor is a *block-circular matrix* [30] taking a form as in the following (up to some permutation):

$$\begin{bmatrix} C_1 & C_{\bar{m}} & \cdots & C_2 \\ C_2 & C_1 & \cdots & C_3 \\ \vdots & \vdots & \ddots & \vdots \\ C_{\bar{m}-1} & C_{\bar{m}-2} & \cdots & C_{\bar{m}} \\ C_{\bar{m}} & C_{\bar{m}-1} & \cdots & C_1 \end{bmatrix}, \quad (24)$$

where $C_i \in \mathbb{R}^{m_1 \times m_1}$ ($1 \leq i \leq \bar{m}$) is also a circular matrix and $\bar{m} = m/m_1$. The general case of $k_i \leq m_i$ is similar. Namely, $\mathcal{A}_k(\cdot)$ is a truncated block-circular matrix consisting of m/m_1 row partitions and k/k_1 column partitions, with each block being a truncated circular matrix of size $m_1 \times k_1$.

2.4 Connections Between DFT and Convolution

Whenever the kernel \mathbf{K} has the same size as the tensor \mathbf{X} , i.e., $k_j = m_j, \forall 1 \leq j \leq n$, the resultant convolution matrix, $\mathcal{A}_m(\mathbf{X})$, is block-circular and therefore can be diagonalized by DFT. The cases of $n = 1$ and $n = 2$ are well-known and have been widely used in the literature, e.g., [31, 32]. In effect, the conclusion holds for any $n \geq 1$, as pointed out by Proposition 2.5 of [30]. More precisely, let the DFT of \mathbf{X} be $\mathcal{F}(\mathbf{X}) = \mathbf{X} \times_1 U_1 \cdots \times_n U_n$, and denote $U = U_1 \otimes \cdots \otimes U_n$ with \otimes being the Kronecker product. Then $U\mathcal{A}_m(\mathbf{X})U^H$ is a diagonal matrix, namely $U\mathcal{A}_m(\mathbf{X})U^H = m\Sigma$ with $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_m)$. Since the first column of U is a vector of all ones, it is easy to see that

$$\begin{aligned} \text{vec}(\mathcal{F}(\mathbf{X})) &= U \text{vec}(\mathbf{X}) = U[\mathcal{A}_m(\mathbf{X})]_{:,1} \\ &= [U\mathcal{A}_m(\mathbf{X})]_{:,1} = [\Sigma U]_{:,1} = [\sigma_1, \dots, \sigma_m]^T. \end{aligned} \quad (25)$$

That is, the eigenvalues of the convolution matrix $\mathcal{A}_m(\mathbf{X})$ are exactly the Fourier frequencies given by $\mathcal{F}(\mathbf{X})$. Hence, for any $\mathbf{X} \in \mathbb{R}^{m_1 \times \cdots \times m_n}$, we have

$$\|\mathcal{F}(\mathbf{X})\|_0 = \text{rank}(\mathcal{A}_m(\mathbf{X})), \quad \|\mathcal{F}(\mathbf{X})\|_1 = \|\mathcal{A}_m(\mathbf{X})\|_*, \quad (26)$$

where $\|\cdot\|_0$ is the ℓ_0 norm of a tensor, and $\|\cdot\|_*$ is the nuclear norm of a matrix. As a consequence, the DFT_{ℓ_1} program (3) is equivalent to the following real-valued convex optimization problem:

$$\min_{\mathbf{L}} \|\mathcal{A}_m(\mathbf{L})\|_*, \quad \text{s.t. } \mathcal{P}_\Omega(\mathbf{L} - \mathbf{L}_0) = 0. \quad (27)$$

Hence, DFT_{ℓ_1} is a special case of CNNM. Although real-valued and convex, the above problem is hard to be solved efficiently, thereby the formulation is used only for the purpose of theoretical analysis.

2.5 Convolution Eigenvalues

The concept of convolution eigenvalues is first proposed and investigated by [25], under the context of image deblurring. Though made specific to order-2 tensors (i.e., matrices),




Fig. 2. Visualizing the convolution eigenvectors and filtered signals, using a 200×200 Boats image as the experimental data and setting $k_1 = k_2 = 13$. For the sake of visual effects, the convolution eigenvector $\kappa_i(\mathbf{X})$ is processed such that its values are in the range of 0 and 1, and the filtered signal $\mathbf{X} \star \kappa_i(\mathbf{X})$ is normalized to have a maximum of 1.

the definitions given in [25] can be easily generalized to order- n tensors with any $n \geq 1$.

Definition 2.1 (Convolution Eigenvalues and Eigenvectors [25]). For a tensor $\mathbf{X} \in \mathbb{R}^{m_1 \times \cdots \times m_n}$ associated with a certain kernel size $k_1 \times \cdots \times k_n$ ($k_j \leq m_j, \forall j$), its first convolution eigenvalue is denoted as $\sigma_1(\mathbf{X})$ and given by

$$\sigma_1(\mathbf{X}) = \max_{\mathbf{K} \in \mathbb{R}^{k_1 \times \cdots \times k_n}} \|\mathbf{X} \star \mathbf{K}\|_F, \quad \text{s.t. } \|\mathbf{K}\|_F = 1. \quad (28)$$

The maximizer to above problem is called the first convolution eigenvector, denoted as $\kappa_1(\mathbf{X}) \in \mathbb{R}^{k_1 \times \cdots \times k_n}$. Similarly, the i th ($i = 2, \dots, k$, $k = \prod_{j=1}^n k_j$) convolution eigenvalue, denoted as $\sigma_i(\mathbf{X})$, is defined as

$$\begin{aligned} \sigma_i(\mathbf{X}) &= \max_{\mathbf{K} \in \mathbb{R}^{k_1 \times \cdots \times k_n}} \|\mathbf{X} \star \mathbf{K}\|_F, \\ \text{s.t. } &\|\mathbf{K}\|_F = 1, \langle \mathbf{K}, \kappa_j(\mathbf{X}) \rangle = 0, \forall j < i. \end{aligned} \quad (29)$$

The maximizer to above problem is the i th convolution eigenvector, denoted as $\kappa_i(\mathbf{X}) \in \mathbb{R}^{k_1 \times \cdots \times k_n}$.

Figure 2 shows the first four convolution eigenvectors of a natural image. It can be seen that the convolution eigenvectors are essentially a series of ordered filters, in which the later ones have higher cut-off frequencies. Note that, unlike the standard singular values, the convolution eigenvalues of a fixed tensor are not fixed and may vary along with the kernel size.

Due to the relationship given in (20), the convolution eigenvalues are nothing more than the singular values of the convolution matrix. Thus, the number of nonzero convolution eigenvalues of a tensor, called the *convolution rank*, is simply the rank of the convolution matrix of the tensor. The so-called *convolution nuclear norm*, defined as the sum of convolution eigenvalues, strictly equals to the nuclear norm of the convolution matrix.

2.6 Convolution Coherence

The same as in most studies about matrix completion, we also need to access the concept of *coherence* [14, 33]. For a rank- r matrix $Y \in \mathbb{R}^{a \times b}$, let its skinny SVD be $Y = U\Sigma V^T$ with $\Sigma \in \mathbb{R}^{r \times r}$. Then there are two coherence parameters, μ_u and μ_v , for characterizing some properties of Y :

$$\mu_u(Y) = \frac{a}{r} \max_{1 \leq i \leq a} \| [U]_{i,:} \|_F^2, \quad \mu_v(Y) = \frac{b}{r} \max_{1 \leq i \leq b} \| [V]_{i,:} \|_F^2, \quad (30)$$

where $[\cdot]_{i,:}$ is the i th row of a matrix. Accordingly, the *convolution coherence* of a tensor $\mathbf{X} \in \mathbb{R}^{m_1 \times \dots \times m_n}$, denoted as $\mu_k(\mathbf{X})$, is defined as the coherence of the convolution matrix $\mathcal{A}_k(\mathbf{X})$:

$$\mu_k(\mathbf{X}) = \max(\mu_u(\mathcal{A}_k(\mathbf{X})), \mu_v(\mathcal{A}_k(\mathbf{X}))). \quad (31)$$

Let $m = \prod_{i=1}^n m_i$ and $k = \prod_{i=1}^n k_i$. By definitions, $1 \leq \mu_u(\mathcal{A}_k(\mathbf{X})) \leq m$ and $1 \leq \mu_v(\mathcal{A}_k(\mathbf{X})) \leq k$, which simply lead to $1 \leq \mu_k(\mathbf{X}) \leq m$. Yet, due to the special property of the convolution matrix, the upper bound m is hardly attainable and a more accurate range is given by

$$1 \leq \mu_k(\mathbf{X}) \leq \frac{m}{k} \nu^2, \quad (32)$$

where ν is the *condition number* of $\mathcal{A}_k(\mathbf{X})$, i.e., ν is the ratio of the largest singular value of $\mathcal{A}_k(\mathbf{X})$ to its smallest nonzero singular value. It can be verified that the above lower bound is achieved if \mathbf{X} is a constant tensor whose elements are all the same, and the upper bound is attained if \mathbf{X} is a standard basis of $\mathbb{R}^{m_1 \times \dots \times m_n}$, $\forall 1 \leq k \leq m$.

Proof. Let the skinny SVD of $\mathcal{A}_k(\mathbf{X})$ be $\mathcal{A}_k(\mathbf{X}) = U\Sigma V^T$, where $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_r)$, and $\sigma_1 \geq \dots \geq \sigma_r > 0$ are the nonzero singular values of $\mathcal{A}_k(\mathbf{X})$. Considering the ℓ_2 norm of the j th column of $\mathcal{A}_k(\mathbf{X})$, we have

$$\| \mathbf{X} \|_F^2 = \| [U\Sigma V^T]_{:,j} \|_F^2 = \| \Sigma [V^T]_{:,j} \|_F^2 \geq \sigma_r^2 \| [V]_{j,:} \|_F^2, \quad (33)$$

which gives that

$$\| [V]_{j,:} \|_F^2 \leq \frac{\| \mathbf{X} \|_F^2}{\sigma_r^2} = \frac{\sum_{i=1}^r \sigma_i^2}{k \sigma_r^2} \leq \frac{r}{k} \nu^2, \quad (34)$$

from which it follows that $\mu_v(\mathcal{A}_k(\mathbf{X})) \leq \nu^2$. Furthermore, since the ℓ_2 norm of any row of $\mathcal{A}_k(\mathbf{X})$ is at most $\| \mathbf{X} \|_F$, we have $\mu_u(\mathcal{A}_k(\mathbf{X})) \leq m \nu^2 / k$. \square

2.7 Convolution Sampling Set

According to the theories in [19, 20], it seems ‘‘impossible’’ to restore the target \mathbf{L}_0 when some of its slices are wholly missing. Nevertheless, as aforementioned, this assertion is made based on the premise that the target \mathbf{L}_0 has the lowest Tucker rank among all possible completions. In the light of convolutional low-rankness, recovery with arbitrary sampling pattern is indeed feasible.

First of all, we would like to introduce a concept called *convolution sampling set*:

Definition 2.2 (Convolution Sampling Set). For $\Omega \subset \{1, \dots, m_1\} \times \dots \times \{1, \dots, m_n\}$ with $m = \prod_{j=1}^n m_j$, its *convolution sampling set* associated with kernel size $k_1 \times \dots \times k_n$ is denoted by $\Omega_{\mathcal{A}}$ and given by

$$\Theta_{\Omega_{\mathcal{A}}} = \mathcal{A}_k(\Theta_{\Omega}) \quad \text{and} \quad \Omega_{\mathcal{A}} = \text{supp}(\Theta_{\Omega_{\mathcal{A}}}), \quad (35)$$

where $\Theta_{\Omega} \in \mathbb{R}^{m_1 \times \dots \times m_n}$ and $\Theta_{\Omega_{\mathcal{A}}} \in \mathbb{R}^{m \times k}$ are the mask tensor and the mask matrix of Ω and $\Omega_{\mathcal{A}}$, respectively. Note here that the subscript k is omitted from $\Omega_{\mathcal{A}}$ for the sake of simplicity.




Fig. 3. Illustrating the effects of convolution. Left: the mask Θ_{Ω} of a 2D sampling set Ω with $m_1 = m_2 = 10$, where the last 5 columns are wholly missing. Right: the mask $\Theta_{\Omega_{\mathcal{A}}}$ of the convolution sampling set $\Omega_{\mathcal{A}}$ with $k_1 = k_2 = 10$.

In general, $\Omega_{\mathcal{A}}$ is a convolution counterpart of Ω , and the corresponding orthogonal projection onto $\Omega_{\mathcal{A}}$ is given by $\mathcal{P}_{\Omega_{\mathcal{A}}}(Y) = \Theta_{\Omega_{\mathcal{A}}} \circ Y, \forall Y \in \mathbb{R}^{m \times k}$. No matter how the observed entries are selected, the convolution sampling set $\Omega_{\mathcal{A}}$ always exhibits a well-posed pattern. Namely, each column of the mask matrix $\Theta_{\Omega_{\mathcal{A}}}$ has exactly $\rho_0 m$ ones and $(1 - \rho_0)m$ zeros, and each row of $\Theta_{\Omega_{\mathcal{A}}}$ has at most $(1 - \rho_0)m$ zeros. Whenever $k_j = m_j, \forall j$, each row of $\Theta_{\Omega_{\mathcal{A}}}$ has also exactly $\rho_0 m$ ones and $(1 - \rho_0)m$ zeros, as shown in Figure 3.

The following lemma shows some algebraic properties about $\Omega_{\mathcal{A}}$, playing a key role in the proofs.

Lemma 2.1. Let $\Omega \subset \{1, \dots, m_1\} \times \dots \times \{1, \dots, m_n\}$, and let the kernel size used to define \mathcal{A}_k be $k_1 \times \dots \times k_n, 1 \leq k_j \leq m_j, \forall 1 \leq j \leq n$. Denote $m = \prod_{j=1}^n m_j$ and $k = \prod_{j=1}^n k_j$. Let $\Omega_{\mathcal{A}} \subset \{1, \dots, m\} \times \{1, \dots, k\}$ be the 2D convolution sampling set defined in (35). For any $Y \in \mathbb{R}^{m \times k}$ and $\mathbf{X} \in \mathbb{R}^{m_1 \times \dots \times m_n}$, we have the following:

$$\begin{aligned} \mathcal{A}_k^* \mathcal{A}_k(\mathbf{X}) &= k \mathbf{X}, \\ \mathcal{A}_k \mathcal{P}_{\Omega}(\mathbf{X}) &= \mathcal{P}_{\Omega_{\mathcal{A}}} \mathcal{A}_k(\mathbf{X}), \\ \mathcal{A}_k^* \mathcal{P}_{\Omega_{\mathcal{A}}}(Y) &= \mathcal{P}_{\Omega} \mathcal{A}_k^*(Y), \end{aligned} \quad (36)$$

where \mathcal{A}_k^* is the Hermitian adjoint of \mathcal{A}_k .

2.8 Isomerism and Relative Well-Conditionedness

By the definition of convolution matrix, recovering \mathbf{L}_0 can ensure to recover its convolution matrix $\mathcal{A}_k(\mathbf{L}_0)$. On the other hand, Lemma 2.1 implies that obtaining $\mathcal{A}_k(\mathbf{L}_0)$ also suffices to identify \mathbf{L}_0 . So, Problem 1.2 can be equivalently converted into a standard matrix completion problem:

Problem 2.1 (Dual Problem). Use the same notations as in Problem 1.2. Denote by $\Omega_{\mathcal{A}}$ the convolution sampling set of Ω . Given $\mathcal{P}_{\Omega_{\mathcal{A}}}(\mathcal{A}_k(\mathbf{L}_0))$, the goal is to recover $\mathcal{A}_k(\mathbf{L}_0)$.

As aforementioned, the pattern of $\Omega_{\mathcal{A}}$ is always well-posed, in a sense that some observations are available at every column and row of the matrix.¹ Hence, provided that \mathbf{L}_0 is convolutionally low-rank, i.e., $\mathcal{A}_k(\mathbf{L}_0)$ is low-rank, Problem 2.1 is exactly the LRMC problem widely studied

1. To meet this, the kernel size needs to be chosen properly. For example, under the setup of Problem 1.1, k_n should be greater than h .

in the literature [14, 15, 33–35]. However, unlike the setting of random sampling adopted by most studies, the sampling regime here is *deterministic* rather than random, thereby we have to count on the techniques established by [19, 20]. For the completeness of presentation, we would briefly introduce the concepts of *isomeric condition* (or *isomerism*) [19] and *relative well-conditionedness* [20].

Definition 2.3 ($\bar{\Omega}/\bar{\Omega}^T$ -Isomeric [19]). *Let $X \in \mathbb{R}^{a \times b}$ be a matrix and $\bar{\Omega} \subset \{1, \dots, a\} \times \{1, \dots, b\}$ be a 2D sampling set. Suppose that $\bar{\Omega}_i \neq \emptyset$ (empty set) and $\bar{\Omega}^j \neq \emptyset, \forall i, j$. Then X is $\bar{\Omega}$ -isomeric iff*

$$\text{rank}([X]_{\bar{\Omega}^j, :}) = \text{rank}(X), \forall j = 1, \dots, b. \quad (37)$$

Furthermore, the matrix X is called $\bar{\Omega}/\bar{\Omega}^T$ -isomeric iff X is $\bar{\Omega}$ -isomeric and X^T is $\bar{\Omega}^T$ -isomeric.

Definition 2.4 ($\bar{\Omega}/\bar{\Omega}^T$ -Relative Condition Number [20]). *Use the same notations as in Definition 2.3. Suppose that $[X]_{\bar{\Omega}^j, :} \neq 0$ and $[X]_{:, \bar{\Omega}_i} \neq 0, \forall i, j$. Then the $\bar{\Omega}$ -relative condition number of X is denoted by $\gamma_{\bar{\Omega}}(X)$ and given by*

$$\gamma_{\bar{\Omega}}(X) = \min_{1 \leq j \leq b} 1/\|X([X]_{\bar{\Omega}^j, :})^+\|^2, \quad (38)$$

where $(\cdot)^+$ is the Moore-Penrose pseudo-inverse of a matrix. Furthermore, the $\bar{\Omega}/\bar{\Omega}^T$ -relative condition number of X is denoted by $\gamma_{\bar{\Omega}, \bar{\Omega}^T}(X)$ and given by $\gamma_{\bar{\Omega}, \bar{\Omega}^T}(X) = \min(\gamma_{\bar{\Omega}}(X), \gamma_{\bar{\Omega}^T}(X^T))$.

In order to show that CNNM succeeds in recovering \mathbf{L}_0 even when the observed entries are arbitrarily placed, we just need to prove that $\mathcal{A}_k(\mathbf{L}_0)$ is $\Omega_{\mathcal{A}}/\Omega_{\mathcal{A}}^T$ -isomeric and $\gamma_{\Omega_{\mathcal{A}}, \Omega_{\mathcal{A}}^T}(\mathcal{A}_k(\mathbf{L}_0))$ is reasonably large as well. To do this, the following lemma is useful.

Lemma 2.2 ([20]). *Use the same notations as in Definition 2.3. Let $\mu_X = \max(\mu_u(X), \mu_v(X))$ be the coherence of the matrix X , and let r_X be the rank of X . Define a quantity ρ as*

$$\rho = \min(\min_{1 \leq i \leq a} \text{card}(\bar{\Omega}_i)/b, \min_{1 \leq j \leq b} \text{card}(\bar{\Omega}^j)/a). \quad (39)$$

For any $0 \leq \alpha < 1$, if $\rho > 1 - (1 - \alpha)/(\mu_X r_X)$ then X is $\bar{\Omega}/\bar{\Omega}^T$ -isomeric and $\gamma_{\bar{\Omega}, \bar{\Omega}^T}(X) > \alpha$.

3 ANALYSIS AND ALGORITHMS

In this section, we will provide theoretical analysis to validate the recovery ability of CNNM (and DFT_{ℓ_1}), uncovering the mystery on why CNNM can work with arbitrary sampling patterns.

3.1 Main Results

First consider the ideal case where the observed data is precise and noiseless. In this case, the CNNM program (7) guarantees to exactly recover the target \mathbf{L}_0 under a certain sampling condition, as shown in the following theorem.

Theorem 3.1 (Noiseless). *Let $\mathbf{L}_0 \in \mathbb{R}^{m_1 \times \dots \times m_n}$ and $\Omega \subset \{1, \dots, m_1\} \times \dots \times \{1, \dots, m_n\}$. Let the adopted kernel size be $k_1 \times \dots \times k_n$ with $k_j \leq m_j, \forall 1 \leq j \leq n$. Denote $k = \prod_{j=1}^n k_j$, $m = \prod_{j=1}^n m_j$ and $\rho_0 = \text{card}(\Omega)/m$. Denote by $r_k(\mathbf{L}_0)$ and $\mu_k(\mathbf{L}_0)$ the convolution rank and convolution coherence of the*

target \mathbf{L}_0 , respectively. Then $\mathbf{L} = \mathbf{L}_0$ is the unique minimizer to the CNNM program (7), as long as

$$\rho_0 > 1 - \frac{0.25k}{\mu_k(\mathbf{L}_0)r_k(\mathbf{L}_0)m}. \quad (40)$$

The above theorem illustrates that, to maximize the recovery ability of CNNM, the kernel size k should be chosen to minimize the sampling bound $1 - 0.25k/(\mu_k(\mathbf{L}_0)r_k(\mathbf{L}_0)m)$, which suggests to minimize $r_k(\mathbf{L}_0)/k$. Note here that $r_k(\mathbf{L}_0) \leq k$ and $r_k(\mathbf{L}_0)$ is a non-decreasing function of k , thereby $r_k(\mathbf{L}_0)/k$ could be a U-shaped function of k . This is a useful clue for determining the key parameters $\{k_1, \dots, k_n\}$ in realistic environments, as we will elaborate in Section 3.4. By setting the kernel to have the same size with the target \mathbf{L}_0 , CNNM falls back to DFT_{ℓ_1} . Thus, the following is an immediate consequence of Theorem 3.1.

Corollary 3.1 (Noiseless). *Use the same notations as in Theorem 3.1, and set $k_j = m_j, \forall 1 \leq j \leq n$. Then $\mathbf{L} = \mathbf{L}_0$ is the unique minimizer to the DFT_{ℓ_1} program (3), as long as*

$$\rho_0 > 1 - \frac{0.25}{\mu_m(\mathbf{L}_0)\|\mathcal{F}(\mathbf{L}_0)\|_0}. \quad (41)$$




Fig. 4. Investigating the difference between future data and randomly chosen missing entries. (a) The sine sequence used for experiments: $\{\mathbf{M}_t\}_{t=1}^m$ with $\mathbf{M}_t = \sin(2t\pi/m)$, thereby \mathbf{L}_0 is an m -dimensional vector with $r_k(\mathbf{L}_0) = 2$ and $\mu_k(\mathbf{L}_0) = 1, \forall m \geq k > 2$. (b) The sampling bound under the setup of forecasting, where $\{\mathbf{M}_t\}_{t=\rho_0 m+1}^m$ is the missing data. (c) The sampling bound under the context of random sampling. In these experiments, the sampling bound is calculated as the smallest fraction of observed entries for the methods to succeed in recovering \mathbf{L}_0 , in a sense that the recovery accuracy measured by Peak Signal-to-Noise Ratio (PSNR) is greater than 50.

As aforementioned, the kernel size should be made positively proportional to the tensor size, i.e., $k = \mathcal{O}(m)$. Thus, the above theorems suggest that the success of CNNM (and DFT_{ℓ_1}) requires $\rho_0 > 1 - \mathcal{O}(1/r_k(\mathbf{L}_0))$, which says that the sampling bound, the lower bound of ρ_0 , has no direct link to the tensor dimension m . This is quite unlike the random sampling based matrix completion theories (e.g., [14, 36]). To be more precise, consider the case of $k = m$ for simplicity. According to [36], the sampling complexity required for recovering the convolution matrix, $\mathcal{A}_m(\mathbf{L}_0) \in \mathbb{R}^{m \times m}$, can be as low as $\mathcal{O}(r_m(\mathbf{L}_0)(\log m)^2/m)$, which gives that the sampling bound should tend to decrease as m grows. In fact, there is no conflict because the sampling regime under forecasting is deterministic rather than random. Figure 4 illustrates that the results derived from random sampling cannot apply to forecasting, confirming the certainty of our result. Even more, the sampling bound $1 - \mathcal{O}(1/r_k(\mathbf{L}_0))$ is

pretty tight under the setup of forecasting, as we will show in Section 5.1.

In practice, the observed data is often a noisy version of $\mathcal{P}_\Omega(\mathbf{L}_0)$ —or the target \mathbf{L}_0 is not strictly convolutionally low-rank as equal. In this case, one should relax the equality constraint and consider instead the following:

$$\min_{\mathbf{L}} \|\mathcal{A}_k(\mathbf{L})\|_*, \text{ s.t. } \|\mathcal{P}_\Omega(\mathbf{L} - \mathbf{M})\|_F \leq \epsilon, \quad (42)$$

where $\mathcal{P}_\Omega(\mathbf{M})$ denotes an observation of $\mathcal{P}_\Omega(\mathbf{L}_0)$, and $\epsilon \geq 0$ is a parameter. The following theorem guarantees the recovery accuracy of (42).

Theorem 3.2 (Noisy). *Use the same notations as in Theorem 3.1. Suppose that $\|\mathcal{P}_\Omega(\mathbf{M} - \mathbf{L}_0)\|_F \leq \epsilon$. If*

$$\rho_0 > 1 - \frac{0.22k}{\mu_k(\mathbf{L}_0)r_k(\mathbf{L}_0)m}, \quad (43)$$

then any optimal solution \mathbf{L}_o to the CNNM program (42) gives a near recovery to the target tensor \mathbf{L}_0 , in a sense that

$$\|\mathbf{L}_o - \mathbf{L}_0\|_F \leq (1 + \sqrt{2})(38\sqrt{k} + 2)\epsilon. \quad (44)$$

Due to the error bound in (44), it seems that the recovery error produced by CNNM may increase with \sqrt{k} . Similar phenomena appear in many papers on matrix completion, e.g., [37]. This, however, is unlikely to be optimal and is indeed an unpleasant effect of the proof techniques commonly used in the community. As a result, it would be incorrect to speculate, based on the error bound in (44), that smaller kernel size k always results in less recovery error. To boost the recovery performance of CNNM, in fact, k should be chosen to minimize the sampling bound, as we have pointed out below Theorem 3.1. Once again, the relationship between CNNM and DFT $_{\ell_1}$ leads to the following result:

Corollary 3.2 (Noisy). *Use the same notations as in Theorem 3.1, and set $k_j = m_j, \forall j$. Suppose that $\|\mathcal{P}_\Omega(\mathbf{M} - \mathbf{L}_0)\|_F \leq \epsilon$ and \mathbf{L}_o is an optimal solution to the following DFT $_{\ell_1}$ program:*

$$\min_{\mathbf{L}} \|\mathcal{F}(\mathbf{L})\|_1, \text{ s.t. } \|\mathcal{P}_\Omega(\mathbf{L} - \mathbf{M})\|_F \leq \epsilon. \quad (45)$$

If $\rho_0 > 1 - 0.22/(\mu_m(\mathbf{L}_0)\|\mathcal{F}(\mathbf{L}_0)\|_0)$, then \mathbf{L}_o gives a near recovery to \mathbf{L}_0 , in a sense that

$$\|\mathbf{L}_o - \mathbf{L}_0\|_F \leq (1 + \sqrt{2})(38\sqrt{m} + 2)\epsilon. \quad (46)$$

For any data tensor \mathbf{M} , one can always decompose it into $\mathbf{M} = \mathbf{L}_0 + \mathbf{N}$ with $\mathcal{A}_k(\mathbf{L}_0)$ being strictly low-rank and $\|\mathbf{N}\|_F \leq \epsilon$. So, the recovery error of CNNM is consistently bounded from above regardless the structure of data. Yet, this does not mean that CNNM can work well on all kinds of data: Whenever $\mathcal{A}_k(\mathbf{M})$ is far from being low-rank, the residual ϵ could be large and the obtained recovery is unnecessarily accurate.

3.2 On Convolutional Low-Rankness

The analyses presented in the above subsection illustrate that, for CNNM to work well, it is important that the convolution matrix of the data tensor \mathbf{M} is strictly or close to be low-rank, i.e., \mathbf{M} itself is convolutionally low-rank or approximately so. This condition cannot be met by all kinds of data, but there do exist many examples of com- pliable. For

example, suppose that $\mathbf{M} \in \mathbb{R}^{m_1 \times \dots \times m_n}$ is a periodic tensor with period (π_1, \dots, π_n) ($1 \leq \pi_q \ll m_q, \forall q$); namely,

$$\mathbf{M} = \mathcal{S}(\mathbf{M}, \pi_q, q), \forall q = 1, \dots, n, \quad (47)$$

where $\mathcal{S}(\cdot, \cdot, \cdot)$ is the multi-directional shift operator used in (22). Then it may be easily seen that

$$r_k(\mathbf{M}) = \text{rank}(\mathcal{A}_k(\mathbf{M})) \leq \prod_{q=1}^n \pi_q, \forall k. \quad (48)$$

That is, periodicity can lead to convolutional low-rankness in a strict manner. This result also implies that, while coping with periodic tensors, the kernel should be made as large as possible, i.e., $k_q = m_q, \forall q$.

Besides periodicity, the smoothness of a signal can also result in convolutional low-rankness in an approximate fashion. To see why, we shall begin with the case of $n = 1$, i.e., \mathbf{M} is a vector. In this case, the j th column of the convolution matrix $\mathcal{A}_k(\mathbf{M})$ is simply the vector obtained by circularly shifting the entries in \mathbf{M} by $j - 1$ positions. Intuitively, when \mathbf{M} possesses substantial smoothness and the shift degree is relatively small, the signals before and after circular shift are mostly the same and therefore $\mathcal{A}_k(\mathbf{M})$ may possess low-rankness. To reach a rigorous conclusion, we consider the *rank- r approximation error* of matrices, which is denoted by $\epsilon_r(\cdot)$ and defined as

$$\epsilon_r(Y) = \min_X \|X - Y\|_F, \text{ s.t. } \text{rank}(X) \leq r, \quad (49)$$

where $Y \in \mathbb{R}^{a \times b}$ is a matrix and r is an integer between 1 and $\min(a, b)$. In the context of circular convolution, the smoothness of a vector \mathbf{x} can be measured by a quantity, denoted as $\delta(\cdot)$, that is similar to the well-known *total variation*:

$$\delta(\mathbf{x}) = \|\mathbf{x} - \mathcal{S}(\mathbf{x})\|_2, \quad (50)$$

where $\|\cdot\|_2$ is the ℓ_2 norm of a vector and $\mathcal{S}(\cdot)$ is the circular shift operator used in (18). With these notations, it is provable that the rank- r approximation error of the convolution matrix of \mathbf{M} satisfies

$$\epsilon_r(\mathcal{A}_k(\mathbf{M})) \leq \frac{\lceil \frac{k}{r} \rceil (k - r)}{2} \delta(\mathbf{M}), \quad (51)$$

where $\lceil \frac{k}{r} \rceil$ stands for the smallest integer greater than or equal to k/r . Hence, the smoothness, which has been adapted to the circulant boundary condition, provably leads to convolutional low-rankness in an approximate sense.

Proof. When \mathbf{M} is an m -dimensional vector, we have $\mathcal{A}_k(\mathbf{M}) = [\mathbf{M}, \mathcal{S}(\mathbf{M}), \dots, \mathcal{S}^{k-1}(\mathbf{M})]$. Decompose $\mathcal{A}_k(\mathbf{M})$ into the concatenation of r submatrices, namely $\mathcal{A}_k(\mathbf{M}) = [A_1, A_2, \dots, A_r]$, such that A_i has b_i columns with $1 \leq b_i \leq \lceil \frac{k}{r} \rceil$ and $\sum_{i=1}^r b_i = k$. For $A_i \in \mathbb{R}^{m \times b_i}$, construct a rank-1 matrix $\hat{A}_i \in \mathbb{R}^{m \times b_i}$ by repeating the first column of A_i for b_i times. Then we have

$$\begin{aligned} \epsilon_1(A_i) &\leq \|A_i - \hat{A}_i\|_F \leq \sum_{c=0}^{b_i-1} c\delta(\mathbf{M}) \\ &= \frac{b_i(b_i - 1)}{2} \delta(\mathbf{M}) \leq \frac{\lceil \frac{k}{r} \rceil (b_i - 1)}{2} \delta(\mathbf{M}), \end{aligned} \quad (52)$$

which gives that

$$\varepsilon_r(\mathcal{A}_k(\mathbf{M})) \leq \sum_{i=1}^r \varepsilon_1(A_i) \leq \frac{\lceil \frac{k}{r} \rceil (k-r)}{2} \delta(\mathbf{M}). \quad (53)$$

□

The above arguments can be easily extended to the general case of $n \geq 1$. To prove that the conclusion in (51) holds for any tensors of order $n \geq 1$, one just needs to generalize the definition in (50) to the following:

$$\delta(\mathbf{X}) = \max_{1 \leq q \leq n} \|\mathbf{X} - \mathcal{S}(\mathbf{X}, 1, q)\|_F, \quad (54)$$

where \mathbf{X} is an order- n tensor with $n \geq 1$ and $\mathcal{S}(\cdot, \cdot, \cdot)$ is the multi-directional shift operator.

Since $\|\mathcal{F}(\mathbf{M})\|_0 = r_m(\mathbf{M})$, $\forall \mathbf{M}$, the result in (51) is also helpful for understanding the phenomenon of Fourier sparsity, which appears frequently in many domains, ranging from images [38] and videos [39] to Boolean functions [40] and wideband channels [41]. More precisely, regarding the smooth signals with bounded values, e.g., images and videos, the smoothness quantity defined in (54) is often small, thereby convolutional low-rankness is very likely to occur, and so for Fourier sparsity.

3.3 Optimization Algorithms

Algorithm for CNNM: For the ease of implementation, we shall not try to solve problem (42) directly, but instead consider its equivalent version as in the following:

$$\min_{\mathbf{L}} \|\mathcal{A}_k(\mathbf{L})\|_* + \frac{\lambda k}{2} \|\mathcal{P}_\Omega(\mathbf{L} - \mathbf{M})\|_F^2, \quad (55)$$

where we amplify the parameter λ by a factor of $k = \prod_{j=1}^n k_j$ for the purpose of normalizing the two objectives to a similar scale. This problem is convex and can be solved by Alternating Direction Method of Multipliers (ADMM) [42, 43]. We first convert it to the following equivalent problem:

$$\min_{\mathbf{L}, Z} \|\mathbf{L}\|_* + \frac{\lambda k}{2} \|\mathcal{P}_\Omega(\mathbf{L} - \mathbf{M})\|_F^2, \text{ s.t. } \mathcal{A}_k(\mathbf{L}) = Z. \quad (56)$$

Then the ADMM algorithm minimizes the augmented Lagrangian function,

$$\|\mathbf{L}\|_* + \frac{\lambda k}{2} \|\mathcal{P}_\Omega(\mathbf{L} - \mathbf{M})\|_F^2 + \langle \mathcal{A}_k(\mathbf{L}) - Z, Y \rangle + \frac{\tau}{2} \|\mathcal{A}_k(\mathbf{L}) - Z\|_F^2, \quad (57)$$

with respect to \mathbf{L} and Z , respectively, by fixing the other variables and then updating the Lagrange multiplier Y and the penalty parameter τ . Namely, while fixing the other variables, the variable Z is updated by

$$Z = \arg \min_Z \frac{1}{\tau} \|Z\|_* + \frac{1}{2} \left\| Z - \left(\mathcal{A}_k(\mathbf{L}) + \frac{Y}{\tau} \right) \right\|_F^2, \quad (58)$$

which is solved via Singular Value Thresholding (SVT) [44]. While fixing the others, the variable \mathbf{L} is updated via

$$\mathbf{L} = (\lambda \mathcal{P}_\Omega + \tau \mathcal{I})^{-1} \left(\frac{\mathcal{A}_k^*(\tau Z - Y)}{k} + \lambda \mathcal{P}_\Omega(\mathbf{M}) \right), \quad (59)$$

where the inverse operator is simply the entry-wise tensor division. The convergence of ADMM with two or fewer

blocks has been well understood, and researchers had even developed advanced techniques to improve its convergence speed, see [42, 45]. While solving the CNNM problem, the computation of each ADMM iteration is dominated by the SVT step, which has a complexity of $\mathcal{O}(mk^2)$. Usually, depending on the increase rate of the penalty parameter τ , the number of iterations for ADMM to get converged may range from tens to hundreds.

Algorithm for DFT $_{\ell_1}$: We shall consider the following problem that is equivalent to the DFT $_{\ell_1}$ program (45):

$$\min_{\mathbf{L}} \|\mathcal{F}(\mathbf{L})\|_1 + \frac{\lambda m}{2} \|\mathcal{P}_\Omega(\mathbf{L} - \mathbf{M})\|_F^2, \quad (60)$$

where $\lambda > 0$ is a parameter. Due to the connection given in (26), the solution to (60) can be determined by finding a solution to (55) with $k_j = m_j, \forall j$. Yet, for the sake of efficiency, it is better to implement a specialized algorithm that can utilize the advantage of DFT.

Generally, the problem in (60) is solved in a similar way to the CNNM problem (55). The main difference happens in updating \mathbf{Z} and \mathbf{L} . While fixing the other variables and updating \mathbf{Z} , one needs to solve the following convex problem:

$$\mathbf{Z} = \arg \min_{\mathbf{Z}} \frac{1}{\tau} \|\mathbf{Z}\|_1 + \frac{1}{2} \left\| \mathbf{Z} - \left(\mathcal{F}(\mathbf{L}) + \frac{\mathbf{Y}}{\tau} \right) \right\|_F^2. \quad (61)$$

Note here that the variable \mathbf{Z} is of complex-valued, and thus one needs to invoke Lemma 4.1 of [46] to obtain a closed-form solution; namely,

$$\mathbf{Z} = f_{1/\tau} \left(\mathcal{F}(\mathbf{L}) + \frac{\mathbf{Y}}{\tau} \right), \quad (62)$$

where $f_\alpha(\cdot)$, a mapping parameterized by $\alpha > 0$, is an entry-wise shrinkage operator given by

$$f_\alpha(z) = \begin{cases} \frac{|z|-\alpha}{|z|} z, & \text{if } |z| > \alpha, \\ 0, & \text{otherwise,} \end{cases} \quad \forall z \in \mathbb{C}. \quad (63)$$

While fixing the others, the variable \mathbf{L} is updated via

$$\mathbf{L} = (\lambda \mathcal{P}_\Omega + \tau \mathcal{I})^{-1} \left(\frac{\mathcal{F}^*(\tau \mathbf{Z} - \mathbf{Y})}{m} + \lambda \mathcal{P}_\Omega(\mathbf{M}) \right), \quad (64)$$

where \mathcal{F}^* denotes the Hermitian adjoint of DFT and is given by $m\mathcal{F}^{-1}$. As can be seen, the computational load is dominated by the calculations of the DFT operator as well as its inverse. Due to the strengths of the Fast Fourier Transform (FFT) algorithm [47], calculating DFT or inverse DFT for an $m_1 \times \dots \times m_n$ tensor needs a computational complexity of only $\mathcal{O}(m \log m)$ with $m = \prod_{j=1}^n m_j$.

3.4 Discussions

On Influences of Parameters: The hyper-parameters in CNNM mainly include the kernel size $k_1 \times \dots \times k_n$ and the regularization parameter λ .

According to the results in Figure 4, it seems beneficial to use large kernels. But this is not the case with most real-world data. As shown in Figure 5(a), the recovery accuracy of CNNM increases as the enlargement of the adopted kernel size at first, but then drops eventually as the kernel size continues to grow. In fact, both phenomena are consistent




Fig. 5. Exploring the influences of the parameters in CNNM, using a 50×50 image patch as the experimental data. (a) Plotting the recovery accuracy as a function of the kernel size. (b) Plotting the averaged coding length as a function of the kernel size. (c) Plotting the recovery accuracy as a function of the parameter λ . For the experiments in (c), the missing rate is set as 35% and the observed entries are contaminated by iid Gaussian noise with mean 0 and standard deviation σ . Note that in this paper the PSNR measure is evaluated only on the missing entries.

with the our theories, which say that the sampling bound is positively related to $r_k(\mathbf{L}_0)/k$ with $k = \prod_{j=1}^n k_j$. For the particular example in Figure 4, $r_k(\mathbf{L}_0) \equiv 2, \forall k \geq 2$, and thus large k produces better recovery. However, on realistic data, the convolution rank may increase as the kernel size grows. To show the consistence in this case, we would like to investigate empirically the *coding length* [48] of the convolution matrix of $\mathbf{L}_0 \in \mathbb{R}^{m_1 \times \dots \times m_n}$:

$$\begin{aligned} \text{CL}_\theta(\mathcal{A}_k(\mathbf{L}_0)) & \\ &= \frac{1}{2}(m+k) \log \det \left(\mathbf{I} + \frac{m}{k\theta^2} \mathcal{A}_k(\mathbf{L}_0)(\mathcal{A}_k(\mathbf{L}_0))^T \right), \end{aligned} \quad (65)$$

where $m = \prod_{j=1}^n m_j$, $\det(\cdot)$ is the determinant of a matrix, and $\theta > 0$ is a parameter. In general, $\text{CL}_\theta(\mathcal{A}_k(\mathbf{L}_0))$ is no more than a computationally-friendly approximate to the convolution rank $r_k(\mathbf{L}_0)$, thereby a reasonable approximate to $r_k(\mathbf{L}_0)/k$ is given by

$$\begin{aligned} \text{ACL}_\theta(\mathcal{A}_k(\mathbf{L}_0)) &= \frac{\text{CL}_\theta(\mathcal{A}_k(\mathbf{L}_0))}{k} \\ &= \frac{1}{2} \left(\frac{m}{k} + 1 \right) \log \det \left(\mathbf{I} + \frac{m}{k\theta^2} \mathcal{A}_k(\mathbf{L}_0)(\mathcal{A}_k(\mathbf{L}_0))^T \right), \end{aligned} \quad (66)$$

where $\text{ACL}_\theta(\cdot)$ is the *averaged coding length* of a matrix.

As we can see from Figure 5(b), the averaged coding length of $\mathcal{A}_k(\mathbf{L}_0)$ is minimized at some value between $k = 1$ and $k = m$, and, interestingly, the minimizer can coincide with the point that maximizes the recovery accuracy provided that the parameter θ is chosen properly. So, to gain the “best” performance, the kernel size in CNNM should be set properly according to the structure of the target \mathbf{L}_0 . Since \mathbf{L}_0 is unknown, estimating the parameters $\{k_1, \dots, k_n\}$ is essentially a challenging model selection problem, and there is no magic to ascertain the best choice. In fact, the uncertainty in determining $\{k_1, \dots, k_n\}$ is related to *interval forecasting* (see [49]). Regarding the setup of *point forecasting* adopted in this paper, we would suggest some empirical rules. Let $k_j = \alpha_j m_j$, then the hyper-parameters $\{\alpha_1, \dots, \alpha_n\}$ could be set as follows:

- The parameter α_n , which is associated with the time-dimension, should satisfy $h/m_n < \alpha_n \leq 1$, where h is the forecast horizon. In most cases, $\alpha_n = 0.5$ is a suitable choice.
- The parameters $\{\alpha_1, \dots, \alpha_{n-1}\}$, which correspond to the non-time dimensions, just need to obey $0 <$

$\alpha_j \leq 1, \forall 1 \leq j \leq n-1$. Empirically, while handling natural images and videos, near optimal recovery performance is often attained at $\alpha_j = 0.25, \forall j$.

Figure 5(c) shows the influence of the parameter λ . As we can see, there seems no need to tune λ carefully. The reason is probably because the nuclear norm is already good at handling heavy-tailed data. So, no matter whether the observations are contaminated by noise or not, we would suggest setting $\lambda = 1000$ for CNNM and DFT_{ℓ_1} .

Is Convolution Rank Minimization NP-Hard? Consider the original form of the CNNM problem (7):

$$\min_{\mathbf{L}} \text{rank}(\mathcal{A}_k(\mathbf{L})), \text{ s.t. } \mathcal{P}_\Omega(\mathbf{L} - \mathbf{L}_0) = 0. \quad (67)$$

It seems not easy to figure out whether the above problem is NP-hard, as it is seemingly difficult to be reduced from some existing NP-hard problem. So we shall examine instead the special case of $k_i = m_i, 1 \leq i \leq n$, i.e., the original form of the DFT_{ℓ_1} problem in (3):

$$\min_{\mathbf{L}} \|\mathcal{F}(\mathbf{L})\|_0, \text{ s.t. } \mathcal{P}_\Omega(\mathbf{L} - \mathbf{L}_0) = 0. \quad (68)$$

According to the derivations around (4), the above ℓ_0 minimization problem is equivalent to the classic problem of finding the sparsest vector in an affine subspace, which is known to be NP-hard [50]. As a consequence, in general cases, the problem in (68) is NP-hard, and so for (67).

It is worth mentioning that the convolution nuclear norm might not be the tightest convex approximation of the convolution rank. This is because, on the set of convolution matrices, the matrix nuclear norm is probably no longer the convex envelope of the rank function. That said, our theorems prove that the convolution nuclear norm minimization can produce exact solutions, as long as certain sampling complexity conditions are met.

4 MATHEMATICAL PROOFS

This section presents in detail the proofs to the proposed lemmas and theorems.

4.1 Proof to Lemma 2.1

Proof. We shall revisit the notations defined in Section 2.3. Denote by \mathcal{S} the “circshift” operator in Matlab; namely, $\mathcal{S}(\mathbf{G}, u, v)$ circularly shifts the elements in tensor \mathbf{G} by u positions along the v th direction. For any $(i_1, \dots, i_n) \in \{1, \dots, k_1\} \times \dots \times \{1, \dots, k_n\}$, we define an invertible operator $\mathcal{T}_{(i_1, \dots, i_n)}: \mathbb{R}^{m_1 \times \dots \times m_n} \rightarrow \mathbb{R}^m$ as

$$\mathcal{T}_{(i_1, \dots, i_n)}(\mathbf{G}) = \text{vec}(\mathbf{G}_n), \forall \mathbf{G} \in \mathbb{R}^{m_1 \times \dots \times m_n}, \quad (69)$$

where $\text{vec}(\cdot)$ is the vectorization operator and \mathbf{G}_n is determined by the following recursion rule:

$$\mathbf{G}_0 = \mathbf{G}, \mathbf{G}_q = \mathcal{S}(\mathbf{G}_{q-1}, i_q - 1, q), 1 \leq q \leq n. \quad (70)$$

Suppose that $j = 1 + \sum_{a=1}^n (i_a - 1) \prod_{b=0}^{a-1} k_b$, where it is assumed conveniently that $k_0 = 1$. Then we have

$$[\mathcal{A}_k(\mathbf{X})]_{:,j} = \mathcal{T}_{(i_1, \dots, i_n)}(\mathbf{X}). \quad (71)$$

According to the definition of the Hermitian adjoint operator given in (9), we have

$$\mathcal{A}_k^*(Z) = \sum_{i_1, \dots, i_n} \mathcal{T}_{(i_1, \dots, i_n)}^{-1}([Z]_{:,j}), \forall Z \in \mathbb{R}^{m \times k}, \quad (72)$$

where it is worth noting that the number j functionally depends on the index (i_1, \dots, i_n) . By (71) and (72),

$$\mathcal{A}_k^* \mathcal{A}_k(\mathbf{X}) = \sum_{i_1, \dots, i_n} \mathcal{T}_{(i_1, \dots, i_n)}^{-1} \mathcal{T}_{(i_1, \dots, i_n)}(\mathbf{X}) = k\mathbf{X}. \quad (73)$$

The second claim is easy to prove. By (11), (35) and (71),

$$\begin{aligned} [\mathcal{A}_k \mathcal{P}_\Omega(\mathbf{X})]_{:,j} &= \mathcal{T}_{(i_1, \dots, i_n)}(\Theta_\Omega \circ \mathbf{X}) = \\ \mathcal{T}_{(i_1, \dots, i_n)}(\Theta_\Omega) \circ \mathcal{T}_{(i_1, \dots, i_n)}(\mathbf{X}) &= [\Theta_{\Omega_A}]_{:,j} \circ [\mathcal{A}_k(\mathbf{X})]_{:,j} \\ &= [\mathcal{P}_{\Omega_A} \mathcal{A}_k(\mathbf{X})]_{:,j}. \end{aligned} \quad (74)$$

It remains to prove the third claim. By (11) and (35),

$$\mathcal{A}_k^* \mathcal{P}_{\Omega_A}(Y) = \mathcal{A}_k^*(\Theta_{\Omega_A} \circ Y) = \mathcal{A}_k^*(\mathcal{A}_k(\Theta_\Omega) \circ Y), \quad (75)$$

which, together with (71) and (72), gives that

$$\begin{aligned} \mathcal{A}_k^* \mathcal{P}_{\Omega_A}(Y) &= \sum_{i_1, \dots, i_n} \mathcal{T}_{(i_1, \dots, i_n)}^{-1}([\mathcal{A}_k(\Theta_\Omega) \circ Y]_{:,j}) \\ &= \sum_{i_1, \dots, i_n} \mathcal{T}_{(i_1, \dots, i_n)}^{-1}([\mathcal{A}_k(\Theta_\Omega)]_{:,j}) \circ \mathcal{T}_{(i_1, \dots, i_n)}^{-1}([Y]_{:,j}) \\ &= \sum_{i_1, \dots, i_n} \Theta_\Omega \circ \mathcal{T}_{(i_1, \dots, i_n)}^{-1}([Y]_{:,j}) = \mathcal{P}_\Omega \mathcal{A}_k^*(Y). \end{aligned} \quad (76)$$

□

4.2 Proof to Theorem 3.1

The proof process is quite standard. We shall first prove the following lemma that establishes the conditions under which the solution to (7) is unique and exact.

Lemma 4.1. *Suppose the skinny SVD of the convolution matrix of \mathbf{L}_0 is given by $\mathcal{A}_k(\mathbf{L}_0) = U_0 \Sigma_0 V_0^T$. Denote by $\mathcal{P}_{T_0}(\cdot) = U_0 U_0^T(\cdot) + (\cdot) V_0 V_0^T - U_0 U_0^T(\cdot) V_0 V_0^T$ the orthogonal projection onto the sum of U_0 and V_0 . Then \mathbf{L}_0 is the unique minimizer to the problem in (7) provided that:*

1. $\mathcal{P}_{\Omega_A}^\perp \cap \mathcal{P}_{T_0} = \{0\}$.
2. There exists $Y \in \mathbb{R}^{m \times k}$ such that $\mathcal{P}_{T_0} \mathcal{P}_{\Omega_A}(Y) = U_0 V_0^T$ and $\|\mathcal{P}_{T_0}^\perp \mathcal{P}_{\Omega_A}(Y)\| < 1$.

Proof. Take $W = \mathcal{P}_{T_0}^\perp \mathcal{P}_{\Omega_A}(Y)$. Then $\mathcal{A}_k^*(U_0 V_0^T + W) = \mathcal{A}_k^* \mathcal{P}_{\Omega_A}(Y)$. By Lemma 2.1,

$$\mathcal{A}_k^* \mathcal{P}_{\Omega_A}(Y) = \mathcal{P}_\Omega \mathcal{A}_k^*(Y) \in \mathcal{P}_\Omega. \quad (77)$$

By the standard convexity arguments shown in [51], \mathbf{L}_0 is an optimal solution to the convex optimization problem in (7).

It remains to prove that \mathbf{L}_0 is the unique minimizer. To do this, we consider a feasible solution $\mathbf{L}_0 + \Delta$ with $\mathcal{P}_\Omega(\Delta) = 0$, and we shall show that the objective value strictly increases unless $\Delta = 0$. Due to the convexity of the convolution nuclear norm, we have

$$\begin{aligned} \|\mathcal{A}_k(\mathbf{L}_0 + \Delta)\|_* - \|\mathcal{A}_k(\mathbf{L}_0)\|_* &\geq \langle \mathcal{A}_k^*(U_0 V_0^T + H), \Delta \rangle \\ &= \langle U_0 V_0^T + H, \mathcal{A}_k(\Delta) \rangle, \end{aligned} \quad (78)$$

where $H \in \mathcal{P}_{T_0}^\perp$ and $\|H\| \leq 1$. The first inequality above follows from a basic property of convexity. Namely, if $f(\cdot)$

is a convex function, we have $f(\mathbf{L}_0 + \Delta) - f(\mathbf{L}_0) \geq \langle \partial_{\mathbf{L}_0} f, \Delta \rangle$, where $\partial_{\mathbf{L}_0} f$ is the subgradient of f at \mathbf{L}_0 . For $f(\mathbf{L}) = \|\mathcal{A}_k(\mathbf{L})\|_*$, we have $\partial_{\mathbf{L}_0} f = \mathcal{A}_k^*(U_0 V_0^T + H)$, where $U_0 V_0^T + H$ is the subgradient of the nuclear norm at $\mathcal{A}_k(\mathbf{L}_0)$.

By the duality between the operator and nuclear norms, we can always choose an H such that

$$\langle H, \mathcal{A}_k(\Delta) \rangle = \|\mathcal{P}_{T_0}^\perp \mathcal{A}_k(\Delta)\|_*. \quad (79)$$

In addition, it follows from Lemma 2.1 that

$$\begin{aligned} \langle \mathcal{P}_{\Omega_A}(Y), \mathcal{A}_k(\Delta) \rangle &= \langle Y, \mathcal{P}_{\Omega_A} \mathcal{A}_k(\Delta) \rangle \\ &= \langle Y, \mathcal{A}_k \mathcal{P}_\Omega(\Delta) \rangle = 0. \end{aligned} \quad (80)$$

Hence, we have

$$\begin{aligned} \langle U_0 V_0^T + H, \mathcal{A}_k(\Delta) \rangle &= \langle \mathcal{P}_{\Omega_A}(Y) + H - W, \mathcal{A}_k(\Delta) \rangle \\ &= \langle H - W, \mathcal{A}_k(\Delta) \rangle \geq (1 - \|W\|) \|\mathcal{P}_{T_0}^\perp \mathcal{A}_k(\Delta)\|_*. \end{aligned} \quad (81)$$

Since $\|W\| < 1$, $\|\mathcal{A}_k(\mathbf{L}_0 + \Delta)\|_*$ is greater than $\|\mathcal{A}_k(\mathbf{L}_0)\|_*$ unless $\mathcal{A}_k(\Delta) \in \mathcal{P}_{T_0}$. Note that $\mathcal{P}_{\Omega_A} \mathcal{A}_k(\Delta) = \mathcal{A}_k \mathcal{P}_\Omega(\Delta) = 0$, i.e., $\mathcal{A}_k(\Delta) \in \mathcal{P}_{\Omega_A}^\perp$. Since $\mathcal{P}_{\Omega_A}^\perp \cap \mathcal{P}_{T_0} = \{0\}$, it follows that $\mathcal{A}_k(\Delta) = 0$, which immediately leads to $\Delta = 0$. □

In the rest of the proof, we shall show how we will prove the dual conditions listed in Lemma 4.1. Notice that, even if the locations of the missing entries are arbitrarily distributed, each column of Ω_A has exactly a cardinality of $\rho_0 m$, and each row of Ω_A contains at least $k - (1 - \rho_0)m$ elements. Denote by ρ the smallest fraction of observed entries in each row and column of $\mathcal{A}_k(\mathbf{L}_0)$. Provided that $\rho_0 > 1 - 0.25k/(\mu_k(\mathbf{L}_0)r_k(\mathbf{L}_0)m)$, we have

$$\rho \geq \frac{k - (1 - \rho_0)m}{k} > 1 - \frac{0.25}{\mu_k(\mathbf{L}_0)r_k(\mathbf{L}_0)}. \quad (82)$$

Then it follows from Lemma 2.2 that $\mathcal{A}_k(\mathbf{L}_0)$ is Ω_A/Ω_A^T -isomeric and $\gamma_{\Omega_A, \Omega_A^T}(\mathcal{A}_k(\mathbf{L}_0)) > 0.75$. Thus, according to Lemma 5.11 of [20], we have

$$\|\mathcal{P}_{T_0} \mathcal{P}_{\Omega_A}^\perp \mathcal{P}_{T_0}\| \leq 2(1 - \gamma_{\Omega_A, \Omega_A^T}(\mathcal{A}_k(\mathbf{L}_0))) < 0.5 < 1, \quad (83)$$

which, together with Lemma 5.6 of [20], results in $\mathcal{P}_{\Omega_A}^\perp \cap \mathcal{P}_{T_0} = \{0\}$. As a consequence, we could define Y as

$$Y = \mathcal{P}_{\Omega_A} \mathcal{P}_{T_0} (\mathcal{P}_{T_0} \mathcal{P}_{\Omega_A} \mathcal{P}_{T_0})^{-1} (U_0 V_0^T). \quad (84)$$

It can be verified that $\mathcal{P}_{T_0} \mathcal{P}_{\Omega_A}(Y) = U_0 V_0^T$. Moreover, it follows from Lemma 5.12 of [20] that

$$\begin{aligned} \|\mathcal{P}_{T_0}^\perp \mathcal{P}_{\Omega_A}(Y)\| &\leq \|\mathcal{P}_{T_0}^\perp \mathcal{P}_{\Omega_A} \mathcal{P}_{T_0} (\mathcal{P}_{T_0} \mathcal{P}_{\Omega_A} \mathcal{P}_{T_0})^{-1}\| \|U_0 V_0^T\| \\ &= \sqrt{\frac{1}{1 - \|\mathcal{P}_{T_0} \mathcal{P}_{\Omega_A}^\perp \mathcal{P}_{T_0}\|} - 1} < 1, \end{aligned} \quad (85)$$

which finishes to construct the dual certificate.

4.3 Proof to Theorem 3.2

Proof. Let $\mathbf{N} = \mathbf{L}_o - \mathbf{L}_0$ and denote $N_A = \mathcal{A}_k(\mathbf{N})$. Notice that $\|\mathcal{P}_\Omega(\mathbf{L}_o - \mathbf{M})\|_F \leq \epsilon$ and $\|\mathcal{P}_\Omega(\mathbf{M} - \mathbf{L}_0)\|_F \leq \epsilon$. By triangle inequality, $\|\mathcal{P}_\Omega(\mathbf{N})\|_F \leq 2\epsilon$. Thus,

$$\|\mathcal{P}_{\Omega_A}(N_A)\|_F^2 = \|\mathcal{A}_k \mathcal{P}_\Omega(\mathbf{N})\|_F^2 = k \|\mathcal{P}_\Omega(\mathbf{N})\|_F^2 \leq 4k\epsilon^2. \quad (86)$$

To bound $\|\mathbf{N}\|_F$, it is sufficient to bound $\|N_{\mathcal{A}}\|_F$. So, it remains to bound $\|\mathcal{P}_{\Omega_{\mathcal{A}}}^{\perp}(N_{\mathcal{A}})\|_F$. To do this, we define Y and W in the same way as in the proof to Theorem 3.1. Since $\mathbf{L}_o = \mathbf{L}_0 + \mathbf{N}$ is an optimal solution to (42), we have the following:

$$\begin{aligned} 0 &\geq \|\mathcal{A}_k(\mathbf{L}_0 + \mathbf{N})\|_* - \|\mathcal{A}_k(\mathbf{L}_0)\|_* \\ &\geq (1 - \|W\|)\|\mathcal{P}_{T_0}^{\perp}(N_{\mathcal{A}})\|_* + \langle \mathcal{P}_{\Omega_{\mathcal{A}}}(Y), N_{\mathcal{A}} \rangle. \end{aligned} \quad (87)$$

Provided that $\rho_0 > 1 - 0.22/(\mu(\mathbf{L}_0)r(\mathbf{L}_0))$, we can prove that $\|W\| = \|\mathcal{P}_{T_0}^{\perp}\mathcal{P}_{\Omega_{\mathcal{A}}}(Y)\| < 0.9$. As a consequence, we have the following:

$$\begin{aligned} \|\mathcal{P}_{T_0}^{\perp}(N_{\mathcal{A}})\|_* &\leq -10\langle \mathcal{P}_{\Omega_{\mathcal{A}}}(Y), \mathcal{P}_{\Omega_{\mathcal{A}}}(N_{\mathcal{A}}) \rangle \\ &\leq 10\|\mathcal{P}_{\Omega_{\mathcal{A}}}(Y)\|\|\mathcal{P}_{\Omega_{\mathcal{A}}}(N_{\mathcal{A}})\|_* \leq 19\|\mathcal{P}_{\Omega_{\mathcal{A}}}(N_{\mathcal{A}})\|_* \\ &\leq 19\sqrt{k}\|\mathcal{P}_{\Omega_{\mathcal{A}}}(N_{\mathcal{A}})\|_F \leq 38k\epsilon, \end{aligned} \quad (88)$$

from which it follows that $\|\mathcal{P}_{T_0}^{\perp}(N_{\mathcal{A}})\|_F \leq \|\mathcal{P}_{T_0}^{\perp}(N_{\mathcal{A}})\|_* \leq 38k\epsilon$, and which simply leads to

$$\|\mathcal{P}_{T_0}^{\perp}\mathcal{P}_{\Omega_{\mathcal{A}}}^{\perp}(N_{\mathcal{A}})\|_F \leq (38k + 2\sqrt{k})\epsilon. \quad (89)$$

We also have

$$\begin{aligned} &\|\mathcal{P}_{\Omega_{\mathcal{A}}}\mathcal{P}_{T_0}\mathcal{P}_{\Omega_{\mathcal{A}}}^{\perp}(N_{\mathcal{A}})\|_F^2 \\ &= \langle \mathcal{P}_{T_0}\mathcal{P}_{\Omega_{\mathcal{A}}}\mathcal{P}_{T_0}\mathcal{P}_{\Omega_{\mathcal{A}}}^{\perp}(N_{\mathcal{A}}), \mathcal{P}_{T_0}\mathcal{P}_{\Omega_{\mathcal{A}}}^{\perp}(N_{\mathcal{A}}) \rangle \\ &\geq (1 - \|\mathcal{P}_{T_0}\mathcal{P}_{\Omega_{\mathcal{A}}}\mathcal{P}_{T_0}\|)\|\mathcal{P}_{T_0}\mathcal{P}_{\Omega_{\mathcal{A}}}^{\perp}(N_{\mathcal{A}})\|_F^2 \\ &\geq \frac{1}{2}\|\mathcal{P}_{T_0}\mathcal{P}_{\Omega_{\mathcal{A}}}^{\perp}(N_{\mathcal{A}})\|_F^2, \end{aligned} \quad (90)$$

which gives that

$$\begin{aligned} \|\mathcal{P}_{T_0}\mathcal{P}_{\Omega_{\mathcal{A}}}^{\perp}(N_{\mathcal{A}})\|_F^2 &\leq 2\|\mathcal{P}_{\Omega_{\mathcal{A}}}\mathcal{P}_{T_0}\mathcal{P}_{\Omega_{\mathcal{A}}}^{\perp}(N_{\mathcal{A}})\|_F^2 \\ &= 2\|\mathcal{P}_{\Omega_{\mathcal{A}}}\mathcal{P}_{T_0}^{\perp}\mathcal{P}_{\Omega_{\mathcal{A}}}^{\perp}(N_{\mathcal{A}})\|_F^2 \leq 2(38k + 2\sqrt{k})^2\epsilon^2. \end{aligned} \quad (91)$$

Combining the above justifications, we have

$$\begin{aligned} \|N_{\mathcal{A}}\|_F &\leq \|\mathcal{P}_{T_0}\mathcal{P}_{\Omega_{\mathcal{A}}}^{\perp}(N_{\mathcal{A}})\|_F + \|\mathcal{P}_{T_0}\mathcal{P}_{\Omega_{\mathcal{A}}}(N_{\mathcal{A}})\|_F \\ &+ \|\mathcal{P}_{T_0}^{\perp}(N_{\mathcal{A}})\|_F \leq (\sqrt{2} + 1)(38k + 2\sqrt{k})\epsilon. \end{aligned} \quad (92)$$

Finally, the fact $\|\mathbf{N}\|_F = \|N_{\mathcal{A}}\|_F/\sqrt{k}$ finishes the proof. \square




Fig. 6. Investigating the recovery performance of DFT_{ℓ_1} under the setup of random sampling and forecasting (i.e., deterministic sampling). The white and black areas mean “success” and “failure”, respectively, where the recovery is regarded as being successful iff $\text{PSNR} > 50$.

5 EXPERIMENTS

All experiments are conducted on the Matlab 2019a platform. The source codes are available at <https://github.com/gcliu1982/CNNM>.

5.1 Simulations

We shall experiment with synthetic data to verify the theorems proven in this paper. We generate \mathbf{M} and \mathbf{L}_0 according to a model as follows: $[\mathbf{M}]_t = [\mathbf{L}_0]_t = \sum_{i=1}^a \sin(2ti\pi/m)$, where $m = 1000$, $t = 1, \dots, m$ and $a = 1, \dots, 19$. So the target \mathbf{L}_0 is a univariate time series of dimension 1000, with $\|\mathcal{F}(\mathbf{L}_0)\|_0 = 2a$ and $\mu_m(\mathbf{L}_0) = 1$. The values in \mathbf{L}_0 are further normalized to have a maximum of 1. Regarding the locations of observed entries, we consider two settings: One is to randomly select a subset of entries as in Figure 4(c), referred to as “random sampling”, the other is the forecasting setup used in Figure 4(b), referred to as “deterministic sampling”. The observation fraction is set as $\rho_0 = 0.05, 0.1, \dots, 0.95$. For the setup of random sampling, we run 20 trials, and thus there are $19 \times 19 \times 21 = 7581$ simulations in total.

The results are shown in Figure 6, the left part of which illustrates that the sampling complexity proven in Corollary 3.1 is indeed tight: The boundary between success and failure is quite consistent with the curve given by $1 - \mathcal{O}(1/\|\mathcal{F}(\mathbf{L}_0)\|_0)$. Also, as we can see, the recovery of future observations is much more challenging than restoring the missing entries chosen uniformly at random. In fact, the sampling pattern of forecasting is almost the *worst case* of arbitrary sampling. In these experiments, the convolution rank $r_k(\mathbf{L}_0)$ does not grow with the kernel size, so there is no benefit to try different kernel sizes in CNNM.

5.2 Univariate Time Series Forecasting

Now we consider two real-world time series downloaded from Time Series Data Library (TSDL): One for the annual Wolfer sunspot numbers from 1770 to 1869 with $m = 100$, and the other for the highest mean monthly levels of Lake Michigan from 1860 to 1955 with $m = 96$. We consider for comparison the well-known methods of Auto-Regressive Moving Average (ARMA) and Long Short-Term Memory (LSTM). ARMA contains many hyper-parameters, which are manually tuned to maximize its recovery accuracy (in terms of PSNR) on the first sequence. The LSTM architecture used for experiment is consist of four layers, including an input layer with 1 unit, a hidden LSTM layer with 200 units, a fully connected layer and a regression layer.

The results are shown in Figure 7. Via manually choosing the best parameters, ARMA can achieve the best performance on the first series. But its results are unsatisfactory while applying the same parametric setting to the second one. By contrast, DFT_{ℓ_1} produces reasonable forecasts on both series that differs greatly in the evolution rules. This is not incredible, because the method never assumes explicitly how the future entries are related to the previously observed samples, and it is indeed the Fourier sparsity of the series itself that enables the recovery of the unseen future data. Via choosing proper kernel sizes, CNNM can future outperform DFT_{ℓ_1} —though the improvements here are mild. A more comprehensive evaluation about the forecasting performance of CNNM can be found in [49].

5.3 Image Completion

The proposed CNNM is indeed a general method for completing partial tensors rather than specific forecasting mod-




Fig. 7. Results for univariate time series forecasting. The left figure shows the entire series used for experiments, where the observed and future entries are plotted with black and green markers, respectively. The sparsity degree of the Fourier transform of a series is measured by Gini [52].




Fig. 8. Evaluating the recovery performance of various methods, using the 200×200 Boats image as the experimental data. The missing entries are chosen uniformly at random, and the numbers plotted above are averaged from 20 random trials.

TABLE 1

PSNR and running time on the 200×200 Boats image. In these experiments, 60% of the image pixels are randomly chosen to be missing. The numbers shown below are collected from 20 runs.

Methods	PSNR	Time (seconds)
LRMC	21.43 ± 0.11	15.9 ± 0.3
DFT_{ℓ_1}	23.25 ± 0.10	0.49 ± 0.06
CNNM(13×13)	24.92 ± 0.11	33.2 ± 0.4
CNNM(23×23)	25.15 ± 0.12	160 ± 1
CNNM(33×33)	25.23 ± 0.12	429 ± 2
CNNM(43×43)	25.27 ± 0.12	968 ± 10
CNNM(53×53)	25.28 ± 0.12	2255 ± 19
CNNM(63×63)	25.26 ± 0.12	4925 ± 42

els. To validate their completion performance, we consider the task of restoring the 200×200 Boats image from its incomplete versions. We also include the LRMC method established by [14] into comparison.

Figure 8 evaluates the recovery performance of various methods. It is clear that LRMC is distinctly outperformed by DFT_{ℓ_1} , which is further outperformed largely by CNNM. Figure 9 shows that the images restored by CNNM is visually better than DFT_{ℓ_1} , whose results contain many artifacts

but are still better than LRMC. In particular, the second row of Figure 9 illustrates that CNNM can well handle the “unidentifiable” cases where some rows and columns of an image are wholly missing. Table 1 shows some detailed evaluation results. As we can see, CNNM works almost equally well under a wide range of kernel sizes. However, since the computational complexity of CNNM is $\mathcal{O}(mk^2)$, the running time grows fast as the enlargement of the kernel size. So, when the computational efficiency is a high priority, it is desirable to use smaller kernels, e.g., $k_1 = k_2 = 13$.

5.4 Video Completion and Prediction

We create a $50 \times 50 \times 62$ video consisting of a sequence of 50×50 images patches quoted from the CD-net 2014 database [53], as demonstrated in Figure 10. We first consider a completion task of restoring the video from some randomly chosen entries. To show the advantages of the proposed methods, we include for comparison three tensor completion methods, including Tensor as a Matrix (TaM) [14], Tensor Nuclear Norm Minimization (TNNM) [16] and Mixture of Low-Rank Tensors (MLRT) [54]. As shown in Figure 11, DFT_{ℓ_1} dramatically outperforms all competing methods that utilize (Tucker) low-rankness to approximate the structure of videos. But DFT_{ℓ_1} is further outperformed distinctly by CNNM ($13 \times 13 \times 13$); this confirms the benefits of controlling the kernel size.

We now consider a forecasting task of recovering the last 6 frames from the former 56 frames. Similar to the experimental setup of Section 5.2, the LSTM network used here also contains 4 layers. The number of input units and LSTM hidden units are 2500 and 500, respectively. Table 2 compares the performance of various methods, in terms of PSNR. It can be seen that TaM, TNNM and MLRT produce very poor forecasts. In fact, all the methods built upon Tucker low-rankness may use zero to predict the unseen future entries. As shown in Figure 12, DFT_{ℓ_1} owns certain ability to forecast the future data, but the obtained images




Fig. 9. Examples of image completion. From left to right: the input data, the results by LRM, the results by DFT_{ℓ_1} , and the results by CNNM. The Gini index of the Fourier transform of the original Boats image is 0.63. Again, note that the PSNR values reported in this paper are computed only on the missing entries.



Fig. 10. All 62 frames of the $50 \times 50 \times 62$ video used for experiments. This sequence records the entire process that a bus passes through a certain location on a highway. The Gini index of the Fourier transform of this dataset is 0.73.

TABLE 2

Evaluation results (PSNR) of video prediction. The task is to forecast the last 6 frames given the former 56 frames. For CNNM, the first two quantities of the kernel size are fixed as $k_1 = k_2 = 13$.

Methods	1st	2st	3st	4st	5st	6st
TaM	5.22	5.49	5.75	5.96	6.15	6.23
TNNM	5.22	5.49	5.75	5.96	6.15	6.23
MLRT	5.22	5.49	5.75	5.96	6.15	6.23
LSTM	20.17	18.80	17.01	16.49	14.73	11.83
DFT_{ℓ_1}	22.56	19.53	18.72	19.43	21.21	26.20
CNNM($k_3 = 13$)	23.48	20.64	20.02	20.80	23.07	28.18
CNNM($k_3 = 31$)	24.27	21.65	21.17	22.10	24.49	29.72
CNNM($k_3 = 62$)	24.57	22.01	21.56	22.56	24.98	30.32

are full of artifacts. By contrast, the quality of the images predicted by CNNM is much higher.




Fig. 11. Evaluation results of restoring the $50 \times 50 \times 62$ video from randomly selected entries. The numbers plotted above are averaged from 20 random trials.

6 CONCLUSION AND FUTURE WORK

In this work, we studied the problem of predicting the future values of a time series based on some previously observed data, i.e., the well-known *time series forecasting* (TSF). We first reformulated the TSF problem to an inclusive compressed sensing problem termed *tensor completion with arbitrary sampling* (TCAS), which targets at restoring a tensor from a subset of its entries sampled in an arbitrary fashion. We then showed that the TCAS problem, under certain

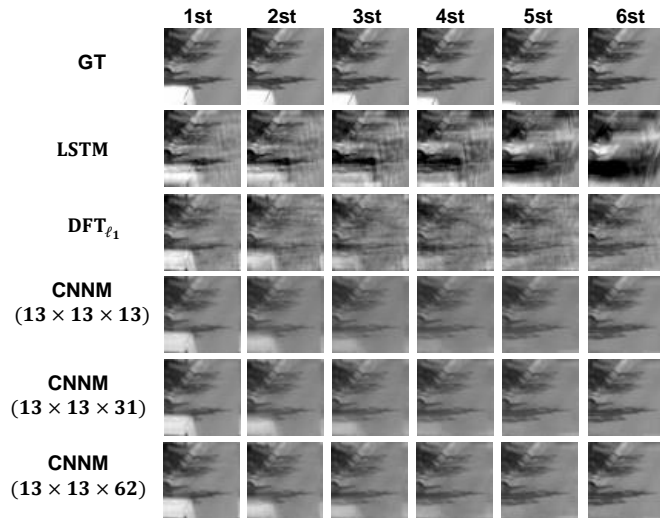


Fig. 12. Visual results of video prediction. The goal of the task is to forecast the last 6 frames given the former 56 frames. GT standards for “ground truth”.

situations, can be solved by a novel method termed CNNM. Namely, we proved that, whenever the sampling complexity exceeds certain threshold that depends on the convolution rank of the target tensor, CNNM ensures the success of recovery absolutely—or at least to some extent. Experiments on some realistic datasets demonstrated that CNNM is a promising method for data completion and forecasting.

While effective, as one may have noticed, CNNM contributes only to a certain case of the TCAS problem; namely, the data tensor should own particular structures such that its convolution matrix is low-rank or approximately so. There are many kinds of data that may violate such a convolutional low-rankness condition, for which the idea of *representation learning* is quite helpful, as shown in [49].

ACKNOWLEDGEMENT

This work is supported in part by New Generation AI Major Project of Ministry of Science and Technology of China under Grant 2018AAA0102501, in part by national Natural Science Foundation of China (NSFC) under Grant U21B2027.

REFERENCES

- [1] J. G. Gooijera and R. J. Hyndman, “25 years of time series forecasting,” *International Journal of Forecasting*, vol. 22, no. 3, pp. 443–473, 2006.
- [2] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, “The m4 competition: 100,000 time series and 61 forecasting methods,” *International Journal of Forecasting*, vol. 36, no. 1, pp. 54–74, 2020.
- [3] M. Mathieu, C. Couprie, and Y. LeCun, “Deep multi-scale video prediction beyond mean square error,” in *International Conference on Learning Representations*, 2016, pp. 1–14.
- [4] R. J. Hyndman, “A brief history of forecasting competitions,” *International Journal of Forecasting*, vol. 36, no. 1, pp. 7–14, 2020.
- [5] P. Montero-Manso and R. J. Hyndman, “Principles and algorithms for forecasting groups of time series: Locality and globality,” *International Journal of Forecasting*, vol. 37, no. 4, pp. 1632–1653, 2021.
- [6] W. Wang, W. Pedrycz, and X. Liu, “Time series long-term forecasting model based on information granules and fuzzy clustering,” *Engineering Applications of Artificial Intelligence*, vol. 41, pp. 17–24, 2015.
- [7] E. Candès and M. Wakin, “An introduction to compressive sampling,” *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, 2008.
- [8] E. Candès and T. Tao, “Decoding by linear programming,” *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.
- [9] M. A. Davenport, M. F. Duarte, Y. C. Eldar, and G. Kutyniok, “Introduction to compressed sensing,” in *Compressed Sensing: Theory and Applications*, 2012, pp. 1–68.
- [10] A. Krishnamurthy and A. Singh, “Low-rank matrix and tensor completion via adaptive sampling,” in *Neural Information Processing Systems*, 2013, pp. 836–844.
- [11] M. Yuan and C.-H. Zhang, “On tensor completion via nuclear norm minimization,” *Foundations of Computational Mathematics*, vol. 16, pp. 1031–1068, 2016.
- [12] X.-Y. Liu, S. Aeron, V. Aggarwal, and X. Wang, “Low-tubal-rank tensor completion using alternating minimization,” *IEEE Transactions on Information Theory*, vol. 66, no. 3, pp. 1714–1737, 2020.
- [13] J.-L. Wang, T.-Z. Huang, X.-L. Zhao, T.-X. Jiang, and M. K. Ng, “Multi-dimensional visual data completion via low-rank tensor representation under coupled transform,” *IEEE Transactions on Image Processing*, vol. 30, pp. 3581–3596, 2021.
- [14] E. Candès and B. Recht, “Exact matrix completion via convex optimization,” *Foundations of Computational Mathematics*, vol. 9, no. 6, pp. 717–772, 2009.
- [15] D. Gross, “Recovering low-rank matrices from few coefficients in any basis,” *IEEE Transactions on Information Theory*, vol. 57, no. 3, pp. 1548–1566, 2011.
- [16] 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, Jan 2013.
- [17] P. Zhou, C. Lu, Z. Lin, and C. Zhang, “Tensor factorization for low-rank tensor completion,” *IEEE Transactions on Image Processing*, vol. 27, no. 3, pp. 1152–1163, March 2018.
- [18] C. Lu, J. Feng, Y. Chen, W. Liu, Z. Lin, and S. Yan, “Tensor robust principal component analysis with a new tensor nuclear norm,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 1, no. 1, pp. 1–14, 2019.
- [19] G. Liu, Q. Liu, and X.-T. Yuan, “A new theory for matrix completion,” in *Neural Information Processing Systems*, 2017, pp. 785–794.
- [20] G. Liu, Q. Liu, X.-T. Yuan, and M. Wang, “Matrix completion with deterministic sampling: Theories and methods,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 2, pp. 549–566, 2021.
- [21] D. Donoho and X. Huo, “Uncertainty principles and ideal atomic decomposition,” *IEEE Transactions on Information Theory*, vol. 47, no. 7, pp. 2845–2862, 2001.

- [22] G. Xu and Z. Xu, "Compressed sensing matrices from fourier matrices," *IEEE Transactions on Information Theory*, vol. 61, no. 1, pp. 469–478, 2015.
- [23] M. Rudelson and R. Vershynin, "Sparse reconstruction by convex relaxation: Fourier and gaussian measurements," in *Annual Conference on Information Sciences and Systems*, 2006, pp. 207–212.
- [24] J. Haupt, L. Applebaum, and R. Nowak, "On the restricted isometry of deterministically subsampled fourier matrices," in *Annual Conference on Information Sciences and Systems*, 2010, pp. 1–6.
- [25] G. Liu, S. Chang, and Y. Ma, "Blind image deblurring using spectral properties of convolution operators," *IEEE Transactions on Image Processing*, vol. 23, no. 12, pp. 5047–5056, 2014.
- [26] M. Fazel, "Matrix rank minimization with applications," Ph.D. dissertation, Electrical Engineering Department, Stanford University, 450 Serra Mall, Stanford, CA 94305, USA, 2002.
- [27] B. Recht, M. Fazel, and P. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Review*, vol. 52, no. 3, pp. 471–501, 2010.
- [28] L. D. Lathauwer, B. D. Moor, and J. Vandewalle, "A multilinear singular value decomposition," *SIAM Journal on Matrix Analysis and Applications*, vol. 21, no. 4, pp. 1253–1278, 2000.
- [29] M. F. Fahmy, G. M. A. Raheem, U. S. Mohamed, and O. F. Fahmy, "A new fast iterative blind deconvolution algorithm," *Journal of Signal and Information Processing*, vol. 3, no. 1, pp. 98–108, 2012.
- [30] M. Combescure, "Block-circulant matrices with circulant blocks, weil sums, and mutually unbiased bases. ii. the prime power case," *Journal of Mathematical Physics*, vol. 50, no. 3, pp. 1–14, 2009.
- [31] M. K. Ng, R. H. Chan, and W.-C. Tang, "A fast algorithm for deblurring models with neumann boundary conditions," *SIAM J. Sci. Comput.*, vol. 21, no. 3, pp. 851–866, 1999.
- [32] Y. Wang, J. Yang, W. yin, and Y. Zhang, "A new alternating minimization algorithm for total variation image reconstruction," *SIAM Journal on Imaging Sciences*, vol. 1, no. 3, pp. 248–272, 2008.
- [33] G. Liu and P. Li, "Low-rank matrix completion in the presence of high coherence," *IEEE Transactions on Signal Processing*, vol. 64, no. 21, pp. 5623–5633, 2016.
- [34] R. Ge, J. D. Lee, and T. Ma, "Matrix completion has no spurious local minimum," in *Neural Information Processing Systems*, 2016, pp. 2973–2981.
- [35] J. Yoon, J. Jordon, and M. van der Schaar, "GAIN: Missing data imputation using generative adversarial nets," in *International Conference on Machine Learning*, 2018, pp. 5689–5698.
- [36] Y. Chen, "Incoherence-optimal matrix completion," *IEEE Transactions on Information Theory*, vol. 61, no. 5, pp. 2909–2923, 2015.
- [37] E. Candès and Y. Plan, "Matrix completion with noise," in *IEEE Proceeding*, vol. 98, 2010, pp. 925–936.
- [38] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [39] A. Chandrakasan, V. Gutnik, and T. Xanthopoulos, "Data driven signal processing: an approach for energy efficient computing," in *International Symposium on Low Power Electronics and Design*, 1996, pp. 347–352.
- [40] R. O'Donnell, "Some topics in analysis of boolean functions," in *Annual ACM Symposium on Theory of Computing*, 2008, pp. 569–578.
- [41] A. P. V. Mengda Lin, "A low complexity high resolution cooperative spectrum-sensing scheme for cognitive radios," *Circuits, Systems, and Signal Processing*, vol. 31, no. 3, pp. 1127–1145, 2012.
- [42] D. Gabay and B. Mercier, "A dual algorithm for the solution of nonlinear variational problems via finite element approximation," *Computers and Mathematics with Applications*, vol. 2, no. 1, pp. 17–40, 1976.
- [43] Z. Lin, M. Chen, and Y. Ma, "The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices," *UIUC Technical Report UILU-ENG-09-2215*, 2009.
- [44] J. Cai, E. Candès, and Z. Shen, "A singular value thresholding algorithm for matrix completion," *SIAM J. on Optimization*, vol. 20, no. 4, pp. 1956–1982, 2010.
- [45] C. Fang, F. Cheng, and Z. Lin, "Faster and non-ergodic $o(1/k)$ stochastic alternating direction method of multipliers," in *Advances in Neural Information Processing Systems*, 2017, pp. 4476–4485.
- [46] G. Liu, Z. Lin, S. Yan, J. Sun, Y. Yu, and Y. Ma, "Robust recovery of subspace structures by low-rank representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 1, pp. 171–184, 2013.
- [47] M. Frigo and S. G. Johnson, "Fftw: an adaptive software architecture for the fft," in *International Conference on Acoustics, Speech and Signal Processing*, vol. 3, 1998.
- [48] Y. Ma, H. Derksen, W. Hong, and J. Wright, "Segmentation of multivariate mixed data via lossy data coding and compression," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 9, pp. 1546–1562, 2007.
- [49] G. Liu, "Time series forecasting via learning convolutionally low-rank models," *IEEE Transactions on Information Theory*, vol. 68, no. 5, pp. 3362–3380, 2022.
- [50] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM Journal on Computing*, vol. 24, no. 2, pp. 227–234, 1995.
- [51] R. T. Rockafellar, *Convex Analysis*. Princeton, NJ, USA: Princeton University Press, 1970.
- [52] N. Hurley and S. Rickard, "Comparing measures of sparsity," *IEEE Transactions on Information Theory*, vol. 55, no. 10, pp. 4723–4741, 2009.
- [53] Y. Wang, P.-M. Jodoin, F. Porikli, J. Konrad, Y. Benezeth, and P. Ishwar, "Cdnet 2014: An expanded change detection benchmark dataset," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2014, pp. 393–400.
- [54] R. Tomioka, K. Hayashi, and H. Kashima, "Estimation of low-rank tensors via convex optimization," *eprint arXiv:1010.0789*, vol. 1, no. 1, pp. 1–19, 2010.

PLACE
PHOTO
HERE

Guangcan Liu (Senior Member, IEEE) received the bachelor's degree in mathematics and the Ph.D. degree in computer science and engineering from Shanghai Jiao Tong University, Shanghai, China, in 2004 and 2010, respectively. He was a Post-Doctoral Researcher with the National University of Singapore, Singapore, from 2011 to 2012; the University of Illinois at UrbanaCChampaign, Champaign, IL, USA, from 2012 to 2013; Cornell University, Ithaca, NY, USA, from 2013 to 2014; and Rutgers University,

Piscataway, NJ, USA, in 2014. He was a Professor with the School of Automation, Nanjing University of Information Science and Technology, Nanjing, from 2014 to 2021. He is currently a Professor with the School of Automation, Southeast University, Nanjing, China. His research interests include the areas of machine learning, computer vision, and signal processing.

PLACE
PHOTO
HERE

Wayne Zhang received the B.Eng. degree in electronic engineering from the Tsinghua University, Beijing, China, in 2007, the M.Phil. degree in 2009, and Ph.D. degree in 2012, both in information engineering from The Chinese University of Hong Kong. He is currently a Senior Research Director in SenseTime Group Limited. He serves as an EXCO member of AI Specialist Group of Hong Kong Computer Society. His research interests include deep learning and computer vision.