

MTC: Multiresolution Tensor Completion from Partial and Coarse Observations

Chaoqi Yang¹, Navjot Singh¹, Cao Xiao², Cheng Qian³, Edgar Solomonik¹, Jimeng Sun¹

¹Department of Computer Science, University of Illinois Urbana-Champaign, Urbana, IL 61801

²Amplitude, San Francisco, CA 94105 ³IQVIA, Cambridge, MA 02142

ABSTRACT

Existing tensor completion formulation mostly relies on partial observations from a single tensor. However, tensors extracted from real-world data are often more complex due to: (i) **Partial observation**: Only a small subset (e.g., 5%) of tensor elements are available. (ii) **Coarse observation**: Some tensor modes only present coarse and aggregated patterns (e.g., monthly summary instead of daily reports). In this paper, we are given a subset of the tensor and some aggregated/coarse observations (along one or more modes) and seek to recover the original fine-granular tensor with low-rank factorization. We formulate a coupled tensor completion problem and propose an efficient Multi-resolution Tensor Completion model (MTC) to solve the problem. Our MTC model explores tensor mode properties and leverages the hierarchy of resolutions to recursively initialize an optimization setup, and optimizes on the coupled system using alternating least squares. MTC ensures low computational and space complexity. We evaluate our model on two COVID-19 related spatio-temporal tensors. The experiments show that MTC could provide 65.20% and 75.79% percentage of fitness (PoF) in tensor completion with only 5% fine granular observations, which is 27.96% relative improvement over the best baseline. To evaluate the learned low-rank factors, we also design a tensor prediction task for daily and cumulative disease case predictions, where MTC achieves 50% in PoF and 30% relative improvements over the best baseline.

CCS CONCEPTS

• **Computing methodologies** → *Factorization methods*; • **Information systems** → *Data mining*.

KEYWORDS

Tensor factorization, Tensor completion, Spatio-temporal analysis

ACM Reference Format:

Chaoqi Yang¹, Navjot Singh¹, Cao Xiao², Cheng Qian³, Edgar Solomonik¹, Jimeng Sun¹. 2021. MTC: Multiresolution Tensor Completion from Partial and Coarse Observations. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21), August 14–18, 2021, Virtual Event, Singapore*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3447548.3467261>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '21, August 14–18, 2021, Virtual Event, Singapore.

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8332-5/21/08...\$15.00

<https://doi.org/10.1145/3447548.3467261>

1 INTRODUCTION

Tensor completion is about estimating missing elements of multi-dimensional higher-order data (the dimensions of tensors are usually called “mode”s). The CANDECOMP/PARAFAC (CP) and Tucker factorization approaches are commonly used for this data denoising and completion task, with applications ranging from image restoration [21, 32] to healthcare data completion [1, 31], recommendation systems [28], link prediction [19], image fusion [10] and spatio-temporal prediction [7, 16].

In many real-world tensor data, missing values and observed data have unique characteristics:

Partial observations: Real-world tensor data has many missing values (only a small subset of the elements are available). This is a common assumption in most tensor completion works [1, 21]. Many existing works estimate the missing elements purely based on data statistics [14] by low-rank CP or Tucker model, while some also consider auxiliary domain knowledge [19, 31].

Coarse observations: In many applications, tensor data is available at multiple granularity. For example, a location mode can be at zipcode, county, or state level, while a time mode can be at daily, weekly, or monthly level. Coarse-granular tensors are often fully available, while fine-granular tensors are usually incomplete with many missing elements. Few works [5, 6] leverage the existing coarse level information to enhance tensor completion.

Motivating example: Given (i) **partial observations**: a small subset (e.g., 5%) of a fine-granular *location by disease by time* tensor, where each tensor element is the disease counts (by ICD-10 code¹) in a geographical location (by anonymous location identifier) at a particular time (by date); (ii) **coarse observations**: two coarse-granular tensors at *state by ICD-10 by date* level and *location by CCS by date* level, respectively. The problem is how to recover the fine-granular *location identifier by ICD-10 by date* tensor based on this partial and coarse information, where the mode aggregation mechanism is sometimes unknown, especially in healthcare [24].

To capture the data characteristics, we identify the following technical challenges:

- **Challenge in partial and coarse data fusion**. While heterogeneous information could compensate with each other and serve as potential supplements, it is challenging to combine different sources in a compact tensor factorization objective.
- **Challenge in computational efficiency and accuracy**. With the increasing data volume and multidimensional structures, it becomes challenging to choose an accurate model initialization and a fast optimization kernel.

¹A disease can be represented with different coding standards: For example, International Classification of Disease version-10 (ICD-10) is a fine-granular code with over 100K dimensions. And Clinical Classification Software for diagnosis codes (CCS) is a coarse-granular code with O(100) dimensions.

To address these challenges, we propose a Multiresolution Tensor Completion (MTC) method with the following contributions.

- MTC fuses both partial and coarse data into Frobenius norm based formulation, which is computationally efficient to deal with as they present a generic coupled-ALS form.
- MTC enables an effective initialization for model optimization to improve model efficiency and accuracy. We pack multiple linear equations into one joint normal equation during optimization and propose an ALS-based solver.

We evaluate MTC on a spatio-temporal disease dataset and a public COVID keyword searching dataset. Extensive experiments demonstrate that for tensors with only 5% observed elements, MTC could provide 65.20% and 75.79% percentage of fitness (PoF) in the completion task, which shows relative improvement over the best baseline by 27.96%. We also design a tensor prediction task to evaluate the learned factors, where our model reaches nearly 50% PoF in disease prediction and outperforms the baselines by 30%.

2 RELATED WORK ON TENSORS

2.1 Low-rank Tensor Completion

Tensor completion (or tensor imputation) benefits various applications, which is usually considered a byproduct when dealing with missing data during the decomposition [27]. The low-rank hypothesis [20] is often introduced to address the under-determined problem. Existing tensor completion works are mostly based on its partial [1, 21] or coarsely [5, 6] aggregated data. Many works have employed augmented loss functions based on statistical models, such as Wasserstein distance [3], missing values [15], noise statistics (e.g., distribution, mean or max value) [9] or the aggregation patterns [33]. Others have utilized other auxiliary domain knowledges, such as within-mode regularization [16, 31], p -norms [19], pairwise constraints [31]. Among which, a line of coupled matrix/tensor factorization [5, 7, 8, 10, 15] is related to our setting. One of the most noticeable works is coupled matrix and tensor factorization (CMTF) proposed by Acar. et al. [2], where a common factor is shared between a tensor and an auxiliary matrix. Another close work [6] recovered the tensor from only aggregated information.

However, few works combine both the partial and coarse data by a coupled formulation. This paper fills in the gap and leverages both the partial and coarse information.

2.2 Multi-scale and Randomized Methods

With the rapid growth in data volume, multi-scale and randomized tensor methods become increasingly important for higher-order structures to boost efficiency and scalability [20, 22].

Multi-scale methods can interpolate from low-resolution to high-resolution [23, 25, 29] or operate first on one part of the tensor and then progressively generalize to the whole tensor [28]. For example, Schifanella et al. [25] exploits extra domain knowledge and develops a multiresolution method to improve CP and Tucker decomposition. Song et al. [28] imputes the tensor from a small corner and gradually increases on all modes by multi-aspect streaming.

The randomized methods are largely based on sampling, which accelerate the computation of over-determined least square problem [20] in ALS for dense [4] and sparse [12] tensors by effective

strategies, such as Fast Johnson-Lindenstrauss Transform (FJLT) [4], leverage-based sampling [12] and tensor sketching.

Unlike previous methods, our model does not require a domain-specific mode hierarchy. Our multiresolution factorization relies on a heuristic mode sampling technique (mode is continuous or categorical), which is proposed to initialize the optimization phase such that fewer iterations would be needed in high resolution. Also, we tackle a more challenging coupled factorization problem.

3 PRELIMINARY

3.1 Notations

We use plain letters for scalar, such as x or X , boldface uppercase letters for matrices, e.g., \mathbf{X} , and Euler script letter for tensors and sets, e.g., \mathcal{X} . Tensors are multidimensional arrays indexed by three or more indices (modes). For example, an N -mode tensor \mathcal{X} is an N -dimensional array of size $I_1 \times \dots \times I_N$, where $\mathcal{X}(i_1, \dots, i_N)$ is the element at the (i_1, \dots, i_N) -th position. For matrices, $\mathbf{X}(:, r)$ and $\mathbf{X}(r, :)$ are the r -th column and row, and $\mathbf{X}(i, j)$ is for the (i, j) -th element. Indices typically start from 1, e.g., $\mathbf{X}(1, :)$ is the first row of the matrix. In this paper, we elaborate on three-mode tensors, while our method could also be applied to higher-mode tensors.

3.1.1 CANDECOMP/PARAFAC (CP) Decomposition. One of the common compression methods for tensors is CP decomposition [11, 13], which approximates a tensor by multiple rank-one components. For example, let $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$ be an arbitrary 3-mode tensor of CP rank R , it can be expressed exactly by three factor matrices $\mathbf{U} \in \mathbb{R}^{I_1 \times R}$, $\mathbf{V} \in \mathbb{R}^{I_2 \times R}$, $\mathbf{W} \in \mathbb{R}^{I_3 \times R}$ as

$$\mathcal{X} = \sum_{r=1}^R \mathbf{U}(:, r) \circ \mathbf{V}(:, r) \circ \mathbf{W}(:, r) = \llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket, \quad (1)$$

where \circ is the vector outer product, and the double bracket notation $\llbracket \cdot, \cdot, \cdot \rrbracket$ is the Kruskal operator.

3.1.2 Mode Product. For the same tensor \mathcal{X} , a coarse observation is given by conducting a tensor mode product with an aggregation matrix (for example, on the first mode), $\mathbf{P}_1 \in \mathbb{R}^{J_1 \times I_1}$, where $J_1 < I_1$,

$$\begin{aligned} \mathcal{C}^{(1)} &= \mathcal{X} \times_1 \mathbf{P}_1 = \sum_{r=1}^R \mathbf{P}_1 \mathbf{U}(:, r) \circ \mathbf{V}(:, r) \circ \mathbf{W}(:, r) \\ &= \llbracket \mathbf{P}_1 \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket \in \mathbb{R}^{J_1 \times I_2 \times I_3}, \end{aligned} \quad (2)$$

where we use superscripts, e.g., $\mathcal{C}^{(1)}$, to denote the aggregated tensor. Other tensor background could be found in Appendix. We organize the symbols in Table 1.

3.2 Problem Definition

Problem 1 (Tensor Completion with Partial and Coarse Observations). For an unknown tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$, we are given

- a partial observation $\mathcal{M} * \mathcal{X}$, parameterized by mask $\mathcal{M} \in \{0, 1\}^{I_1 \times I_2 \times I_3}$;
- a known set of aggregation matrices $\mathcal{P}^* \subseteq \mathcal{P} = \{\mathbf{P}_s : \mathbf{P}_s \in \{0, 1\}^{J_s \times I_s}, J_s < I_s, s = 1, 2, 3\}$, each column of \mathbf{P}_s is one-hot;
- a set of coarse tensors $\mathcal{J} = \{\mathcal{C}^{\mathcal{H}} : \mathcal{C}^{\mathcal{H}} = \mathcal{X} \times_{h \in \mathcal{H}} \mathbf{P}_h, \mathcal{H} \in \text{Powerset}\{1, 2, 3\} \setminus \emptyset, \mathbf{P}_h \in \mathcal{P}\}$. For example, a coarse tensor aggregated on both the first and the second mode can be written as $\mathcal{C}^{(1,2)} = \mathcal{X} \times_{h \in \{1,2\}} \mathbf{P}_h = \mathcal{X} \times_1 \mathbf{P}_1 \times_2 \mathbf{P}_2$.

Table 1: Notations used in MTC

Symbols	Descriptions
R	tensor CP rank
$\odot, *$	tensor Khatri-Rao/Hadamard product
$\llbracket \cdot, \cdot, \cdot \rrbracket$	tensor Kruskal operator
\mathcal{T}	the set of coarse tensors
$\mathcal{P}^*, \tilde{\mathcal{P}}, \mathcal{P}$	known/unknown/full aggregation matrix set
$\mathbf{P}_1 \in \mathbb{R}^{I_1 \times I_1}, \mathbf{P}_2 \in \mathbb{R}^{I_2 \times I_2}$	aggregation matrices
$\mathcal{X}, \tilde{\mathcal{X}}, \mathcal{R} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$	the original/interim/reconstructed tensor
$\mathcal{C}^{(1)} \in \mathbb{R}^{I_1 \times I_2 \times I_3}, \mathcal{C}^{(2)} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$	two coarse observation tensors
$\tilde{\mathcal{X}}_i, \mathcal{C}_i^{(1)}, \mathcal{C}_i^{(2)}, \mathbf{M}_i, \mathbf{R}_i, i = 1, 2, 3$	tensor unfoldings/matrixizations
$\mathcal{M} \in \{0, 1\}^{I_1 \times I_2 \times I_3}$	tensor mask
$\mathbf{U} \in \mathbb{R}^{I_1 \times R}, \mathbf{V} \in \mathbb{R}^{I_2 \times R}, \mathbf{W} \in \mathbb{R}^{I_3 \times R}$	factor matrices
$\mathbf{Q}_1 \in \mathbb{R}^{I_1 \times R}, \mathbf{Q}_2 \in \mathbb{R}^{I_2 \times R}$	auxiliary factor matrices
$(\cdot)_{(d_t)}, \mathcal{S}_{(d_t, \cdot)}$	entity or index set at d_t resolution
$\tilde{\mathbf{U}}, \tilde{\mathbf{V}}, \tilde{\mathbf{W}}, \tilde{\mathbf{Q}}_1, \tilde{\mathbf{Q}}_2$	initializations for the factors

The problem is to find low-rank CP factors $\mathbf{U}, \mathbf{V}, \mathbf{W}$, such that the following (Frobenius norm based) loss is minimized,

$$\mathcal{L}(\mathcal{M} * \mathcal{X}, \mathcal{T}, \mathcal{P}^*; \mathbf{U}, \mathbf{V}, \mathbf{W}, \tilde{\mathcal{P}}) = \|\mathcal{M} * (\mathcal{X} - \llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket)\|_F^2 + \sum_{\mathcal{H}} \lambda_{\mathcal{H}} \|\mathcal{C}^{\mathcal{H}} - \llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket \times_{h \in \mathcal{H}} \mathbf{P}_h\|_F^2, \quad (3)$$

where $\tilde{\mathcal{P}} = \mathcal{P} \setminus \mathcal{P}^*$ is the set of unknown aggregation matrices, \mathcal{H} enumerates the index of \mathcal{T} , and $\lambda_{\mathcal{H}}$ being the weights. We separate the knowns and unknowns (to estimate) by semicolon “;”. Note that, the same mode cannot be aggregated in all coarse tensors.

The problem reduces to the conventional CP tensor completion problem if the set \mathcal{T} of coarse tensors is empty, which has been explored extensively in the previous literature [1, 17, 21].

3.3 A Motivating Example

As an example, we use the aforementioned COVID disease tensor to motivate our application. For an unknown *location identifier by ICD-10 by date* tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$, we are given

- a partial observation $\mathcal{M} * \mathcal{X}$, parametrized by mask \mathcal{M} ;
- a coarse *state by ICD-10 by date* tensor, $\mathcal{C}^{(1)} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$, where $I_1 < I_1$. It satisfies $\mathcal{C}^{(1)} = \mathcal{X} \times_1 \mathbf{P}_1$ with an unknown $\mathbf{P}_1 \in \mathbb{R}^{I_1 \times I_1}$, since we have no knowledge on the mapping from anonymized location identifier to state (for privacy preservation);
- a coarse *location identifier by CCS by date* tensor, $\mathcal{C}^{(2)} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$, where $I_2 < I_2$ and it satisfies $\mathcal{C}^{(2)} = \mathcal{X} \times_2 \mathbf{P}_2$ with known *ICD-10 to CCS* mapping $\mathbf{P}_2 \in \mathbb{R}^{I_2 \times I_2}$.

We seek to find low-rank CP factors $\mathbf{U}, \mathbf{V}, \mathbf{W}$, such that the loss is minimized over parameters $(\mathbf{U}, \mathbf{V}, \mathbf{W}, \mathbf{P}_1)$,

$$\mathcal{L}(\mathcal{M} * \mathcal{X}, \mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \mathbf{P}_2; \mathbf{U}, \mathbf{V}, \mathbf{W}, \mathbf{P}_1) = \|\mathcal{M} * (\mathcal{X} - \llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket)\|_F^2 + \lambda_1 \|\mathcal{C}^{(1)} - \llbracket \mathbf{P}_1 \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket\|_F^2 + \lambda_2 \|\mathcal{C}^{(2)} - \llbracket \mathbf{U}, \mathbf{P}_2 \mathbf{V}, \mathbf{W} \rrbracket\|_F^2, \quad (4)$$

w.l.o.g., we use this example to introduce our methodology.

3.4 Solution Outline

3.4.1 Coupled Tensor Decomposition. To address the above objective, we consider to successively approximate the low-rank completion problem by coupled tensor decomposition. We first introduce

two transformations and reformulate the objective into a coupled decomposition form:

- **Parameter replacement.** We replace $\mathbf{Q}_1 = \mathbf{P}_1 \mathbf{U}$ and $\mathbf{Q}_2 = \mathbf{P}_2 \mathbf{V}$ as two new auxiliary factor matrices and add extra constraints later since \mathbf{Q}_2 and \mathbf{V} are connected by a known \mathbf{P}_2 .
- **Interim Tensor.** For the first term in Equation (4), we construct an interim tensor $\tilde{\mathcal{X}}$ from an expectation-maximization (EM) [1] approach,

$$\begin{aligned} & \mathcal{M} * (\mathcal{X} - \llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket) \\ &= \mathcal{M} * \mathcal{X} + (1 - \mathcal{M}) * \llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket - \llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket \\ &\approx \underbrace{\left[\mathcal{M} * \mathcal{X} + (1 - \mathcal{M}) * \llbracket \mathbf{U}^k, \mathbf{V}^k, \mathbf{W}^k \rrbracket \right]}_{\text{the interim tensor } \tilde{\mathcal{X}} \text{ for next iteration}} - \llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket \end{aligned} \quad (5)$$

the interim tensor $\tilde{\mathcal{X}}$ for next iteration

where $\mathbf{U}^k, \mathbf{V}^k, \mathbf{W}^k$ are the learned factors at the k -th iteration. The reconstruction $\mathcal{R} = \llbracket \mathbf{U}^k, \mathbf{V}^k, \mathbf{W}^k \rrbracket$ is practically expensive for large tensors. However, we do not need to compute it explicitly. In Section 5, we introduce a way to use the implicit form.

Afterwards, the objective in Equation (4) is approximated by a coupled tensor decomposition form (at the $(k+1)$ -th iteration),

$$\begin{aligned} \mathcal{L}(\tilde{\mathcal{X}}, \mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \mathbf{P}_2; \mathbf{U}, \mathbf{V}, \mathbf{W}, \mathbf{Q}_1, \mathbf{Q}_2) &= \|\tilde{\mathcal{X}} - \llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket\|_F^2 \\ &+ \lambda_1 \|\mathcal{C}^{(1)} - \llbracket \mathbf{Q}_1, \mathbf{V}, \mathbf{W} \rrbracket\|_F^2 \\ &+ \lambda_2 \|\mathcal{C}^{(2)} - \llbracket \mathbf{U}, \mathbf{Q}_2, \mathbf{W} \rrbracket\|_F^2 \end{aligned} \quad (6)$$

with constraint $\mathbf{Q}_2 = \mathbf{P}_2 \mathbf{V}$ and the interim tensor $\tilde{\mathcal{X}}$,

$$\tilde{\mathcal{X}} = \mathcal{M} * \mathcal{X} + (1 - \mathcal{M}) * \llbracket \mathbf{U}^k, \mathbf{V}^k, \mathbf{W}^k \rrbracket. \quad (7)$$

Now, our goal turns into optimizing a coupled decomposition problem with five shared factors, among which $\mathbf{U}, \mathbf{V}, \mathbf{W}$ are our focus.

3.4.2 Optimization Idea. As a solution, we adopt a multiresolution strategy, as in Figure 1, which leverages a hierarchy of resolution to initialize the coupled optimization effectively. The **multiresolution factorization** part (in Section 4) serves as the outer loop of our solution. At each resolution, we derive a joint normal equation for each factor and develop an efficient ALS-based **Solver** (in Section 5) to handle multiple linear systems simultaneously.

4 MULTIREOLUTION FACTORIZATION

Let us denote the resolution hierarchy as $\{d_{\max}, \dots, d_{t+1}, d_t, \dots\}$, where d_{\max} is the highest resolution (i.e., original tensors). The multiple resolutions are computed via recursive subsampling along three tensor modes simultaneously. The underlying ideas of the multiresolution factorization are: (i) to subsample high-resolution information into low resolution and solve a low-resolution problem; (ii) to then interpolate the low-resolution solution factors for initializing the high-resolution problem, such that fewer iterations are needed in high resolution.

4.1 High (d_{t+1}) to Low (d_t) Resolution

Assume that d_{t+1} is a generic notation for high resolution and d_t is for the neighboring low resolution, we use *subscript-parenthesis* notation, e.g., (d_t) to denote the entities at each resolution.

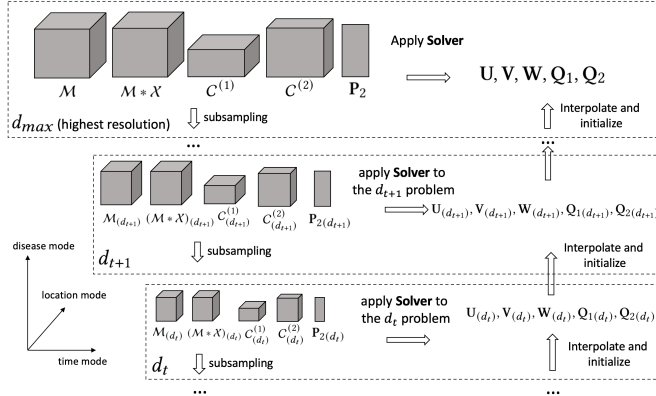


Figure 1: MTC Framework follows a multiresolution recursive flow. First, we apply subsampling on all accessible information (i.e., the tensors and the known aggregation matrices) into the lowest resolution. Then, we solve the low resolution problem by applying the optimization solver. Next, we interpolate the solution into the higher resolution to initialize the high resolution factors. We repeat this process and find a good initialization for the original fine-granular problem. The mode **subsampling and interpolation** are discussed in Section 4; the **Solver** is developed in Section 5.

Objective at d_{t+1} . The goal is to find optimal factor matrices: $U_{(d_{t+1})} \in \mathbb{R}^{I_1(d_{t+1}) \times R}$, $V_{(d_{t+1})} \in \mathbb{R}^{I_2(d_{t+1}) \times R}$, $W_{(d_{t+1})} \in \mathbb{R}^{I_3(d_{t+1}) \times R}$, $Q_1(d_{t+1}) \in \mathbb{R}^{I_1(d_{t+1}) \times R}$, $Q_2(d_{t+1}) \in \mathbb{R}^{I_2(d_{t+1}) \times R}$, such that they minimize the loss at d_{t+1} (the form could refer to Equation (6)),

$$\mathcal{L}_{(d_{t+1})} = \mathcal{L} \left(\tilde{X}_{(d_{t+1})}, C_{(d_{t+1})}^{(1)}, C_{(d_{t+1})}^{(2)}, P_{2(d_{t+1})}; U_{(d_{t+1})}, V_{(d_{t+1})}, W_{(d_{t+1})}, Q_1(d_{t+1}), Q_2(d_{t+1}) \right).$$

In the following, we conceptualize the solution to the objective:

- (1) **Subsampling:** We first subsample all accessible information (tensors and the known aggregation matrices) from d_{t+1} to d_t at once,

$$\begin{aligned} M_{(d_{t+1})}, (M * X)_{(d_{t+1})}, C_{(d_{t+1})}^{(1)}, C_{(d_{t+1})}^{(2)}, P_{2(d_{t+1})} &\xrightarrow{\text{subsample}} \\ M_{(d_t)}, (M * X)_{(d_t)}, C_{(d_t)}^{(1)}, C_{(d_t)}^{(2)}, P_{2(d_t)} &\quad (8) \end{aligned}$$

- (2) **Solve the objective at d_t :** Next, we solve the d_t -problem recursively and obtain smaller factor matrices ($U_{(d_t)} \in \mathbb{R}^{I_1(d_t) \times R}$, $V_{(d_t)} \in \mathbb{R}^{I_2(d_t) \times R}$, $W_{(d_t)} \in \mathbb{R}^{I_3(d_t) \times R}$, $Q_1(d_t) \in \mathbb{R}^{I_1(d_t) \times R}$, $Q_2(d_t) \in \mathbb{R}^{I_2(d_t) \times R}$), where the d_t -objective can be similarly written as,

$$\begin{aligned} \mathcal{L}_{(d_t)} = & \left\| \tilde{X}_{(d_t)} - \llbracket U_{(d_t)}, V_{(d_t)}, W_{(d_t)} \rrbracket \right\|_F^2 \\ & + \lambda_1 \left\| C_{(d_t)}^{(1)} - \llbracket Q_1(d_t), V_{(d_t)}, W_{(d_t)} \rrbracket \right\|_F^2 \\ & + \lambda_2 \left\| C_{(d_t)}^{(2)} - \llbracket U_{(d_t)}, Q_2(d_t), W_{(d_t)} \rrbracket \right\|_F^2, \quad (9) \end{aligned}$$

the interim tensor $\tilde{X}_{(d_t)}$ is calculated by

$$\tilde{X}_{(d_t)} = (M * X)_{(d_t)} + (1 - M_{(d_t)}) * \llbracket U_{(d_t)}^k, V_{(d_t)}^k, W_{(d_t)}^k \rrbracket, \quad (10)$$

and the d_t resolution constraint is $Q_2(d_t) = P_{2(d_t)} V_{(d_t)}$.

- (3) **Interpolation:** Then, we interpolate the d_t solutions to initialize the d_{t+1} problem. More precisely, four factor matrices are obtained by interpolation:

$$U_{(d_t)}, V_{(d_t)}, W_{(d_t)}, Q_1(d_t) \xrightarrow{\text{interpolate}} \quad (11)$$

$$\tilde{U}_{(d_{t+1})}, \tilde{V}_{(d_{t+1})}, \tilde{W}_{(d_{t+1})}, \tilde{Q}_1(d_{t+1}). \quad (12)$$

And one is calculated by the known aggregation matrix according to the d_{t+1} constraint:

$$\tilde{Q}_2(d_{t+1}) = P_{2(d_{t+1})} \tilde{V}_{(d_{t+1})}. \quad (13)$$

- (4) **Optimization at d_{t+1} :** With the initialized factor matrices, we finally can apply our optimization algorithm recursively on the d_{t+1} objective to improve the factors.

These four steps are the key to the paper. The multiresolution factorization is conducted by recursively applying four steps within Step (2). In this section, Step (1) and (3) are of our focus, while we will specify Step (4) in Section 5. The subsampling and interpolation are closely related, and we elucidate them below.

4.2 Resolution Transition from d_{t+1} to d_t

In our example, three modes are presented with d_{t+1} -specific mode sizes: (i) location mode: anonymized location identifier $I_1(d_{t+1})$, state $J_1(d_{t+1})$; (ii) disease mode: ICD-10 $I_2(d_{t+1})$, CCS $J_2(d_{t+1})$; (iii) time mode: date $I_3(d_{t+1})$. To make it clear, we use the term **aspect** to indicate a tensor mode or its projection/aggregation at a coarse granularity.

In total, we have three modes in our problem and five different aspects, corresponding to five factors. We call a mode with two aspects an *aggregated mode* and a mode with one aspect as a *single mode*. For example, the location and disease modes are aggregated modes, and the time mode is a single mode. Because location mode can be at zipcode or state aspect; disease mode can be at ICD-10 aspect (fine-granular) or CCS aspect (coarse-granular). But time mode is only represented at the date aspect.

Mode subsampling is actually conducted on each aspect. Tensors and aggregation matrices are subsampled altogether in the same way when they share the same aspect. For example, on subsampling ICD-10 aspect, M (the second mode), $M * X$ (the second mode), $C^{(1)}$ (the second mode) and P_2 (the second dimension) are targeted. We consider different subsampling scenarios:

- **Continuous or categorical.** For single modes, we provide two strategies based on mode property: (i) mode with continuous or smooth information (called continuous mode), e.g., date; (ii) mode with categorical information (called categorical mode).
- **Known or unknown aggregation.** For aggregated mode, we provide different treatments depending on whether the associated aggregation matrix P is (i) known or (ii) unknown.

In fact, the strategies for aggregated modes are based on the single mode (continuous or categorical) strategy. We first introduce the strategies on a single mode.

4.3 Single Continuous Mode

The transition for this mode type is under the smooth assumption: information along the mode contains some spatial or temporal consistency. For a continuous mode, we subsample the d_{t+1} index

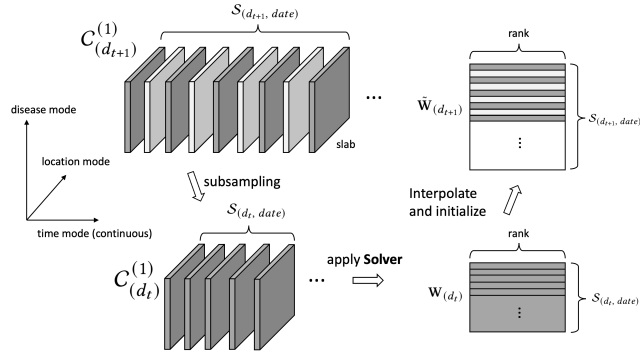


Figure 2: Continuous Mode. We use the time mode (corresponding to factor \mathbf{W}) and one of the involved tensors, $C^{(1)}$, as an example. Here, we fix other modes and only discuss the transition along the time mode (in real practice, we do subsampling along three modes simultaneously).

set (of the aspect) to d_t and later use neighborhood smoothing to interpolate the d_t factors back to initialize d_{t+1} factors.

For example, the time mode assumes smoothness, since within a time period, the disease statistics over the same geographical regions will not change abruptly. See the diagram in Figure 2.

4.3.1 Subsampling. We define

$$\mathcal{S}_{(d_{t+1}, \text{date})} = \{1, 2, \dots, I_3(d_{t+1})\} \quad (14)$$

to be the index set of the date aspect at d_{t+1} -resolution. We sample at a regular interval as

$$\mathcal{S}_{(d_t, \text{date})} = \{1, 3, 5, \dots\} \quad (15)$$

yielding a d_t -resolution index set of size $\lceil \frac{I_3(d_{t+1})}{2} \rceil$.

For example, $C^{(1)}_{(d_{t+1})}$ includes the date aspect, thus it will be subsampled using the index set $\mathcal{S}_{(d_t, \text{date})}$.

4.3.2 Interpolation. Factor $\mathbf{W}_{(d_t)} \in \mathbb{R}^{I_3(d_t) \times R}$ corresponds to the date aspect at d_t . The initialization of d_{t+1} factor $\tilde{\mathbf{W}}_{(d_{t+1})} \in \mathbb{R}^{I_3(d_{t+1}) \times R}$ is given by

$$\tilde{\mathbf{W}}_{(d_{t+1})}(2j, :) = \frac{1}{2} \left(\mathbf{W}_{(d_t)}(j, :) + \mathbf{W}_{(d_t)}(j+1, :) \right), \quad (16)$$

$$\tilde{\mathbf{W}}_{(d_{t+1})}(2j+1, :) = \mathbf{W}_{(d_t)}(j, :), \quad 0 \leq j \leq \left\lfloor \frac{I_3(d_{t+1}) - 1}{2} \right\rfloor, \quad (17)$$

if $I_3(d_{t+1})$ is an even number, then

$$\tilde{\mathbf{W}}_{(d_{t+1})}(I_3(d_{t+1}), :) = \mathbf{W}_{(d_t)} \left(\frac{I_3(d_{t+1})}{2}, : \right). \quad (18)$$

4.3.3 Intuition. As shown in Figure 2, from $C^{(1)}_{(d_{t+1})}$ to $C^{(1)}_{(d_t)}$, only the “dark” slabs are selected for the d_t -resolution. Tensor slabs of $C^{(1)}_{(d_t)}$ along the date aspect are controlled by $\mathbf{W}_{(d_t)}$, where one row corresponds to one slab. After solving the d_t problem, the solution factor $\mathbf{W}_{(d_t)}$ provides an estimate for those “dark” slabs. To also estimate the “light” slabs, we assume that it could be approximated by the neighboring “dark” slabs. The actual interpolation follows Equation (16), (17), (18), where we average neighboring rows.

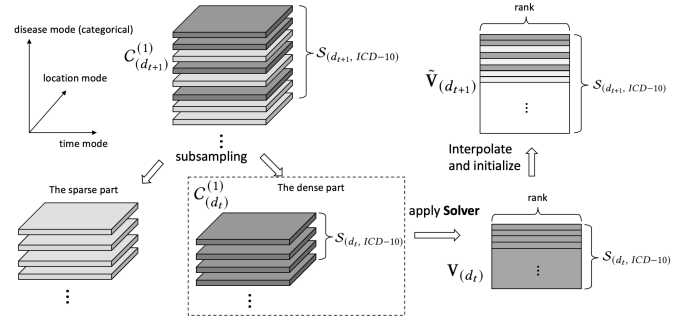


Figure 3: Categorical Mode. We use the ICD-10 aspect at disease mode (corresponding to factor \mathbf{V}) and one of the involved tensors, $C^{(1)}$, as an example. Similarly, only one aspect is discussed here.

4.4 Single Categorical Mode

For a categorical mode, we also want to subsample the d_{t+1} index set to d_t . However, we now focus on the density of the corresponding tensor slabs. Later on, we interpolate the d_t factors back to d_{t+1} .

We use the ICD-10 (fine-granular disease code) aspect as an example. This aspect corresponds to factor $\mathbf{V}_{(d_{t+1})}$ at d_{t+1} , for which we provide an illustration in Figure 3.

The ICD-10 information could be obtained from one coarse tensor, e.g., $C^{(1)}_{(d_{t+1})}$, at d_{t+1} . Along this aspect, we count the non-zero elements over each slab and obtain a count vector; each element stores the counts of one slab, corresponding to the ICD-10 code,

$$\text{Cnt}_{(d_{t+1}, \text{ICD-10})} = [\text{Cnt}_1, \text{Cnt}_2, \dots, \text{Cnt}_{I_2(d_{t+1})}]. \quad (19)$$

4.4.1 Subsampling. We define

$$\mathcal{S}_{(d_{t+1}, \text{ICD-10})} = \{1, 2, 3, \dots, I_2(d_{t+1})\} \quad (20)$$

to be the index set of ICD-10 aspect at d_{t+1} -resolution.

We order the count vector and find the indices corresponding to the largest half (composing the dense part of the tensor, while the unselected half is the sparse part), denoted as $\mathcal{S}_{(d_t, \text{ICD-10})}$, as the index set at d_t , so that

$$\text{Cnt}_i \geq \text{Cnt}_j, \quad \forall i \in \mathcal{S}_{(d_t, \text{ICD-10})}, j \in \mathcal{S}_{(d_{t+1}, \text{ICD-10})} \setminus \mathcal{S}_{(d_t, \text{ICD-10})}. \quad (21)$$

The tensors and aggregation matrices that include ICD-10 aspect, e.g., $C^{(1)}_{(d_{t+1})}$, are subsampled along this mode.

4.4.2 Interpolation. $\mathbf{V}_{(d_t)} \in \mathbb{R}^{I_2(d_t) \times R}$ is the corresponding solution factor for ICD-10 at d_t . We get the d_{t+1} initialization $\tilde{\mathbf{V}}_{(d_{t+1})} \in \mathbb{R}^{I_2(d_{t+1}) \times R}$ by: the rows of $\tilde{\mathbf{V}}_{(d_{t+1})}$ are direct copies from the rows of $\mathbf{V}_{(d_t)}$ if the corresponding indices are selected in $\mathcal{S}_{(d_t, \text{ICD-10})}$, otherwise, we fill the rows with random entries (in the paper, we use i.i.d. samples from $[-1, 1]$, and users might customize their configuration based on the applications).

4.4.3 Intuition. As for the categorical mode (in Figure 3), the rationality of our strategy comes from two underlying heuristics: (i) the dense half of the indices is likely to account for a large volume of the tensor, such that after the optimization at d_t , we already recover a dominant part at d_{t+1} ; (ii) the selected indices (slabs) preserve

the dense half of the tensor, while from randomized algorithm perspective, a slab is barely a training sample, composited by the other factor matrices. Preserving those dense slabs would be beneficial for the estimation of other factors.

4.5 Aggregated Mode

We are now ready to introduce our treatments of aggregated modes, for which the aggregation matrix \mathbf{P} may or may not be known.

Each aggregated mode associates with two aspects and thus two factor matrices. *For mode subsampling*: two aspects would be subsampled independently based on the aforementioned categorical or continuous strategy. If they are connected with a known aggregation matrix \mathbf{P} , then \mathbf{P} is subsampled accordingly along two dimensions; *For factor interpolation*: we follow the aforementioned strategies to interpolate two factors independently if the \mathbf{P} is unknown; otherwise, we use the strategy to interpolate only the fine-granular factor and manually compute the other one. For example, at d_t resolution, the aggregation matrix from $\mathbf{V}_{(d_t)}$ to $\mathbf{Q}_{2(d_t)}$, i.e., $\mathbf{P}_{2(d_t)}$, is known. In this case, we would first interpolate a high resolution $\mathbf{V}_{(d_{t+1})}$ and then compute a high resolution $\mathbf{Q}_{2(d_{t+1})}$ based on $\mathbf{Q}_{2(d_{t+1})} = \mathbf{P}_{2(d_{t+1})} \mathbf{V}_{(d_{t+1})}$, instead of interpolating up from $\mathbf{Q}_{2(d_t)}$.

5 OPTIMIZATION ALGORITHMS

This section presents an effective alternating optimization algorithm, which is proposed for the coupled decomposition problem at each resolution. We first present the ALS algorithm and then formulate our quadratic subproblems.

5.1 Coupled-ALS Algorithm

5.1.1 Alternating Least Squares (ALS). To solve a low-rank CP approximation of a given tensor, ALS is one of the most popular algorithms [26]. It has been applied for both decomposition [20, 26] and completion [17, 27]. However, our approach differs by considering a reconstructed (interim) tensor at each step in a coupled decomposition setting. Below is the standard ALS algorithm.

Given a tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$, to find a rank- R CP decomposition with factors $\mathbf{U} \in \mathbb{R}^{I_1 \times R}$, $\mathbf{V} \in \mathbb{R}^{I_2 \times R}$, $\mathbf{W} \in \mathbb{R}^{I_3 \times R}$, the following loss function is minimized,

$$\mathcal{L}(\mathcal{X}; \mathbf{U}, \mathbf{V}, \mathbf{W}) = \|\mathcal{X} - \llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket\|_F^2. \quad (22)$$

ALS algorithm optimally solves for one of the factor matrices while keeping the others fixed and then cycles through all of the factor matrices. A sub-iteration of ALS for \mathbf{U} is executed by solving

$$(\mathbf{V} \odot \mathbf{W}) \mathbf{U}^\top = \mathbf{X}_1^\top, \quad (23)$$

where $\mathbf{V} \odot \mathbf{W} \in \mathbb{R}^{I_2 \times I_3 \times R}$ is the Khatri-Rao product and $\mathbf{X}_1 \in \mathbb{R}^{I_1 \times I_2 \times I_3}$ is the unfolded tensor along the first mode.

5.1.2 Coupled-ALS. In our case, at each resolution d_t , we have a coupled tensor decomposition problem as defined in equation (9). Similar to standard ALS, if we solve for one factor matrix while keeping all the other unknowns fixed, the problem reduces to a quadratic problem which can be solved optimally by coupling normal equations with respect to the factor at each granularity.

To reduce clutter, we elaborate on the full resolution form of Equation (6). Specifically, we solve a least square problem with the following equations.

- For \mathbf{U} , equations are derived using $\tilde{\mathbf{X}}_1$ and $\mathbf{C}_1^{(2)}$, which are the unfoldings of $\tilde{\mathcal{X}}$ and $\mathcal{C}^{(2)}$ along the first mode. Thus, we have the following two equations,

$$(\mathbf{V} \odot \mathbf{W}) \mathbf{U}^\top = \tilde{\mathbf{X}}_1^\top, \quad (24)$$

$$\lambda_2(\mathbf{Q}_2 \odot \mathbf{W}) \mathbf{U}^\top = \lambda_2 \mathbf{C}_1^{(2)\top}. \quad (25)$$

- Similarly, for \mathbf{V} , equations are derived using $\tilde{\mathbf{X}}_2$ and $\mathbf{C}_2^{(1)}$.
- For \mathbf{W} , equations are derived using $\tilde{\mathbf{X}}_3$, $\mathbf{C}_3^{(1)}$ and $\mathbf{C}_3^{(2)}$.
- For \mathbf{Q}_1 , equations are derived using $\mathbf{C}_1^{(1)}$.
- For \mathbf{Q}_2 , equations are derived using $\mathbf{C}_2^{(2)}$ only. However, \mathbf{Q}_2 and \mathbf{V} satisfy the constraint $\mathbf{Q}_2 = \mathbf{P}_2 \mathbf{V}$.

5.2 Multi-resolution ALS

5.2.1 Alternating Optimization with Constraint. Considering the constraint issue, e.g., between \mathbf{V} and \mathbf{Q}_2 , we use the equations to update one of them (the one corresponds to fine granular aspect) and later update the other one. For our example, \mathbf{V} is updated by the linear equations, while \mathbf{Q}_2 is updated by the constraint $\mathbf{Q}_2 = \mathbf{P}_2 \mathbf{V}$ and new \mathbf{V} , instead of using linear equations².

The optimization flow follows the ALS routine, where during each iteration, we sequentially update \mathbf{W} , \mathbf{Q}_1 , \mathbf{U} , \mathbf{V} by the associated linear equations, and then manually update a new \mathbf{Q}_2 .

5.2.2 Joint Normal Equations. \mathbf{W} , \mathbf{Q}_1 , \mathbf{U} , \mathbf{V} are solved by using joint normal equations. For each of them, we consider stacking multiple linear systems and formulating joint normal equations with respect to each factor for the objective defined in equation (6). Using \mathbf{U} as the example, the normal equations are given as

$$\begin{bmatrix} \mathbf{V} \odot \mathbf{W} \\ (\mathbf{Q}_2 \odot \mathbf{W}) \end{bmatrix}^\top \begin{bmatrix} \mathbf{V} \odot \mathbf{W} \\ \lambda_2(\mathbf{Q}_2 \odot \mathbf{W}) \end{bmatrix} \mathbf{U}^\top = \begin{bmatrix} \mathbf{V} \odot \mathbf{W} \\ (\mathbf{Q}_2 \odot \mathbf{W}) \end{bmatrix}^\top \begin{bmatrix} \tilde{\mathbf{X}}_1^\top \\ \lambda_2 \mathbf{C}_1^{(2)\top} \end{bmatrix}, \quad (26)$$

where the components are weighted accordingly. The formulation could be further simplified as

$$\begin{aligned} & ((\mathbf{V}^\top \mathbf{V}) * (\mathbf{W}^\top \mathbf{W}) + \lambda_2(\mathbf{Q}_2^\top \mathbf{Q}_2) * (\mathbf{W}^\top \mathbf{W})) \mathbf{U}^\top \\ &= (\mathbf{V} \odot \mathbf{W})^\top \tilde{\mathbf{X}}_1^\top + \lambda_2(\mathbf{Q}_2 \odot \mathbf{W})^\top \mathbf{C}_1^{(2)\top}. \end{aligned} \quad (27)$$

This formulation is amenable to large scale problems as the right hand sides can exploit sparsity and the use of matricized tensor times Khatri-Rao product (MTTKRP) kernel with each tensor. To avoid dense MTTKRP with $\tilde{\mathcal{X}}$, the implicit form given in Equation (7) can be used to efficiently calculate the MTTKRP by

$$\begin{aligned} (\mathbf{V} \odot \mathbf{W})^\top \tilde{\mathbf{X}}_1^\top &= (\mathbf{V} \odot \mathbf{W})^\top (\mathbf{M}_1 * \mathbf{X}_1)^\top + (\mathbf{V}^\top \mathbf{V}^k * \mathbf{W}^\top \mathbf{W}^k) \mathbf{U}^{k\top} \\ &\quad - (\mathbf{V} \odot \mathbf{W})^\top (\mathbf{M}_1 * \mathbf{R}_1)^\top, \end{aligned} \quad (28)$$

where $\mathcal{R} = \llbracket \mathbf{U}^k, \mathbf{V}^k, \mathbf{W}^k \rrbracket$. The above simplification requires sparse MTTKRP and a sparse reconstruction tensor $\mathcal{M} * \mathcal{R}$, for which efficient programming abstractions are available [34].

²Note that, since these two factors are coupled, we could have transformed the \mathbf{Q}_2 equations to supplement \mathbf{V} equations by pseudo inverse (i.e. inverse of $\mathbf{P}_2^\top \mathbf{P}_2$), such that \mathbf{V} would associate with three linear equations. However, the highly non-uniqueness of $\mathbf{P}_2^\top \mathbf{P}_2$ makes the optimization unstable.

5.2.3 Solver Implementation. To solve the joint normal equations, we employ a two-stage approach,

- **Stage 1:** we start from the initialization given by the low resolution and apply the weighted Jacobi algorithm for N (e.g., 5) iterations, which smooths out the error efficiently.
- **Stage 2:** from the $N + 1$ iteration, we apply Cholesky decomposition to the left-hand side (e.g., Equation (27)) and decompose the joint normal equations into solving two triangular linear systems, which could generate more stable solutions. Finally, the triangular systems are solved exactly.

In practice, we find that the combination of Stage 1 and Stage 2 works better than directly applying Stage 2 in light of both convergence speed and final converged results, especially when the partial observation is less. We perform ablation studies for this part.

The **time and space complexity** of MTC is asymptotically equal to that of applying standard CP-ALS on original tensor \mathcal{X} , which is $O(I_1 I_2 I_3 R)$ and $O(I_1 I_2 I_3)$, respectively.

6 EXPERIMENTS

The experiments provide a comprehensive study on MTC. As an outline, the following evaluations are presented:

- (1) Ablation study on the different amount of partially observed data, compared to SOTA tensor completion and decomposition methods.
- (2) Performance comparison on tensor completion and downstream tensor prediction task, compared with baselines.
- (3) Ablation study on model initialization, compared with other initialization baselines.

6.1 Data Preparation

6.1.1 Two COVID-19 Related Databases. Two COVID-related databases are considered in our evaluation:

- The *health tensor (HT)* is a proprietary real-world dataset, constructed from IQVIA medical claim database, including a *location by disease by time* tensor, counting the COVID related disease cases in the US from Mar. to Aug. 2020 during the pandemic.
- (ii) The *Google COVID-19 Symptoms Search (GCSS)*³ is a public dataset, which shows aggregated trends in Google searches for disease-related health symptoms, signs, and conditions, in format *location identifier by keyword by time*.

6.1.2 Data Processing. For *HT* data, we use the most popular 1,200 zipcodes and 1,000 ICD-10 codes as the first two modes and use the first five months as the third mode (while the last month data is used for the prediction task), which creates \mathcal{X} . We aggregate the first mode by county and the second mode by CCS code separately to obtain two aggregated tensors, $C^{(1)}, C^{(2)}$. We randomly sample 5% elements from the tensor as the partial observation, $\mathcal{X} * \mathcal{M}$. For the *GCSS* data, we use all 2,727 location identifiers and 422 search keywords as the first two modes and a complete 2020 year span as the third time mode. To create two coarse viewed tensors, we aggregate on the state level in the first mode and week level for the third mode, separately. Assume 5% of the elements are observed in *GCSS*. For both datasets, the first two modes are assumed categorical, while the last time mode is continuous.

³https://pair-code.github.io/covid19_symptom_dataset/?country=GB

In this experiment, we assume all the aggregation matrices are unknown. In Appendix, we use synthetic data and GCSS to show that our model could achieve **OracleCPD** level accuracy if some aggregation \mathbf{P} is given. Basic data statistics are listed in Table 2.

6.2 Experimental Setup

6.2.1 Baselines. We include the following comparison models from different perspectives: (i) reference models: **CPC-ALS** [17]: state of the art tensor completion model, only using partial observed data, **OracleCPD** (CP decomposition of the original complete tensor with ALS); (ii) related baselines (they are gradient based): **Block gradient descent (BGD)** [32], **B-PREMA** [6], **CMTF-OPT** [2]. They use both partial and coarse information; (iii) initialization models: **MRTL** [23], **TendiB** [5], **Higher-order SVD (HOSVD)**; (iv) our variants: **MTC_{multi-}** removes multiresolution module; **MTC_{stage1-}** removes stage 1 (Jacobi) in the solver; **MTC_{both-}** removes both multiresolution part and stage 1. More detail in Appendix.

6.2.2 Tasks and Metrics. Evaluation metrics include: (i) *Percentage of Fitness (PoF)* [27], suppose the target tensor is \mathcal{X} and the low-rank reconstruction is $\mathcal{R} = \llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket$, the relative standard error (RSE) and PoF are defined by

$$RSE = \frac{\|\mathcal{X} - \mathcal{R}\|_F}{\|\mathcal{X}\|_F}, \quad PoF = 1 - RSE; \quad (29)$$

(ii) *CPU Time*; (iii) *Peak Memory*.

We also consider **(2) future tensor prediction**, which belongs to a broader spatio-temporal prediction domain [30]. We considered the tensor based baselines while a thorough comparison with other models is beyond our scope. Specifically, we use Gaussian Process (GP) with radial basis function and white noise kernel along the date aspect to estimate a date factor matrix for the next month

$$\mathbf{W}_{future} = GP(\mathbf{W}). \quad (30)$$

Then, the future disease tensor could be estimated by $\mathcal{R}_{future} = \llbracket \mathbf{U}, \mathbf{V}, \mathbf{W}_{future} \rrbracket$. We use the next-month real disease tensor to evaluate the prediction results, and again, we use PoF as the metric.

6.2.3 Hyperparameters. For *HT* dataset, we use CP rank $R = 50$, 20 iterations at each low resolution, and $N = 5$ in solver stage 1, within which 5 Jacobi rounds are performed per iteration. For *GCSS*, we use $R = 20$, 10 iterations at each low resolution, $N = 5$, and 10 Jacobi rounds per iteration. By default, we set 200 iterations at the finest resolution for both datasets, which ensures the convergence. The parameter λ_i is set to be $\lambda = e^{-\frac{i}{20}}$ at the i -th iteration, which varies the focus from coarse information to fine-granular information gradually. We publish our dataset and codes in this repository⁴.

6.3 Comparison with Reference Models

We first compare our MTC with tensor decomposition model OracleCPD and a state of the art tensor completion model CPC-ALS on *GCSS*. The experiment shows the importance of different data sources and the model sensitivity w.r.t. the amount of partial data. Three models use the same CP rank $R = 20$. The OracleCPD is implemented on the original tensor, which provides an ideal low-rank approximation. The CPC-ALS runs on the partially observed data only, while our MTC uses both the partial and coarse tensors.

⁴<https://github.com/ycq091044/MTC>

Table 2: Dataset Statistics

Name	1 _{st} -mode	2 _{nd} -mode	3 _{rd} -mode	Sparsity	Agg. 1 _{st} -mode	Agg. 2 _{nd} -mode	Agg. 3 _{rd} -mode	Partial Obs.
HT	zipcode (1,200)	ICD-10 (1,000)	date (128)	75.96%	county (220)	CCS (189)	--	5%
GCSS	identifier (2,727)	keyword (422)	date (362)	87.73%	state (50)	--	week (52)	5%

* sparsity x% means x% of the entries are zeros in the original tensor.

Table 3: Tensor Completion Results. Our MTC beats the baselines significantly on PoF and CPU Time with considerable space complexity.

Model	Health Tensor (HT)			Google COVID-19 Symptoms Search (GCSS)		
	PoF	CPU Time	Peak Memory	PoF	CPU Time	Peak Memory
BGD	0.4480 ± 0.0161 (7.155e-05)	11,708.02 ± 98.39s	10.27 GB	0.5552 ± 0.0098 (5.733e-06)	26,074.53 ± 207.00s	19.32 GB
B-PREMA	0.4678 ± 0.0176 (1.329e-04)	11,614.75 ± 113.34s	11.04 GB	0.5923 ± 0.0045 (1.791e-06)	26,201.12 ± 273.30s	19.40 GB
CMTF-OPT	0.4815 ± 0.0245 (4.490e-04)	10,159.50 ± 88.37s	10.53 GB	0.5720 ± 0.0198 (9.394e-05)	25,681.41 ± 141.30s	19.25 GB
MTC _{both} -	0.6109 ± 0.0234 (5.713e-02)	1,437.20 ± 67.57s	10.32 GB	0.7508 ± 0.0063 (2.022e-01)	1,821.84 ± 93.39s	19.75 GB
MTC _{multi} -	0.6432 ± 0.0144 (4.821e-01)	1,375.32 ± 52.66s	10.31 GB	0.7560 ± 0.0059 (6.880e-01)	1,874.18 ± 74.37s	20.09 GB
MTC _{stage1} -	0.6282 ± 0.0223 (1.880e-01)	1,499.26 ± 59.66s	10.55 GB	0.7524 ± 0.0044 (2.274e-01)	1,932.02 ± 25.25s	20.33 GB
MTC	0.6520 ± 0.0113	1,501.43 ± 58.36s	10.73 GB	0.7579 ± 0.0049	1,929.49 ± 40.27s	20.15 GB

* table format: mean ± standard deviation (*p*-value)

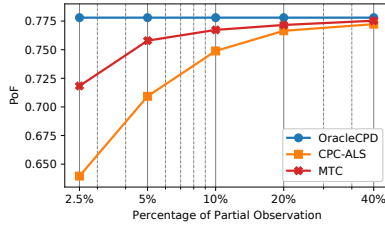


Figure 4: Ablation Study with Different Amount of Partially Observed Data on GCSS. Coarse information could improve tensor completion, especially when the partial data is not enough.

We show the results in Figure 4. When less partially observed data (e.g., 2.5%) are presented, MTC can achieve more (e.g., 7.5%) improvement on PoF over CPC-ALS, with the help of coarse-granular data. When we have more partial observations (e.g., 20%), the gaps between three models become smaller (within 5%).

6.4 Comparison on Tensor Completion

We compare our MTC and the variants with related baselines: BGD, B-PREMA, CMTF-OPT on the tensor completion task. The experiments are conducted 3 times with different random seeds. The results, standard deviations, and statistical *p*-values are shown in Table 3. We can conclude that the baselines are inferior to our proposed methods by around 15% fitness measure on both datasets.

The table also shows the memory requirements and executed time of each method. MTC reduces the time consumption significantly to about $\frac{1}{8}$ of that of the baselines, since: (i) our iterative method incurs fewer floating-point arithmetics than computing the exact gradients from Equation (6); (ii) gradient baselines require to perform expensive line search with multiple function evaluations.

We also show the optimization landscape of all models and some ablation study results in Figure 5. We find that: (i) the multiresolution factorization and solver stage 1 are better used together; (ii) our MTC is more useful when having more missing values.

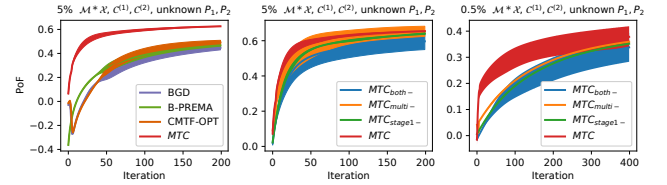


Figure 5: Optimization Landscape and Ablation Study on HT. Compared to gradient based baselines, our MTC provides a faster convergence rate and better converged results. With less partial observed information, MTC becomes more advantageous.

6.5 Comparison on Tensor Prediction

After the optimization on HT dataset, we load the CP factors of each method and conduct the *tensor prediction task* as described in the experimental setup. Specifically, we perform: cumulative and daily disease case prediction, where for one particular ICD-10 code at one particular zipcode location, we predict how many more disease cases will happen within the future *x* days.

The results are shown in Figure 6, where the dark bars are the standard deviations. The prediction results confirmed MTC and its variants could outperform the baselines by about 20% on PoF.

6.6 Ablation Study on Initializations

To evaluate the multiresolution factorization part in MTC, we compare with other tensor initialization methods. Specifically, we implement the initialization from (i) other baselines: MRTL, TendiB; (ii) typical tensor initialization method: HOSVD; (iii) no initialization, MTC_{multi}-. For a fair comparison, we use the same optimization method (in Section 5) after different initializations.

Different amount of partial observations (i.e., different *M*) are considered: 5%, 2% and 1% on HT dataset. We report the first 30 iterations and the initialization time in Figure 7. Essentially, our MTC outperforms all other models after 30 iterations. One interesting finding is that TendiB and HOSVD start from the same PoF

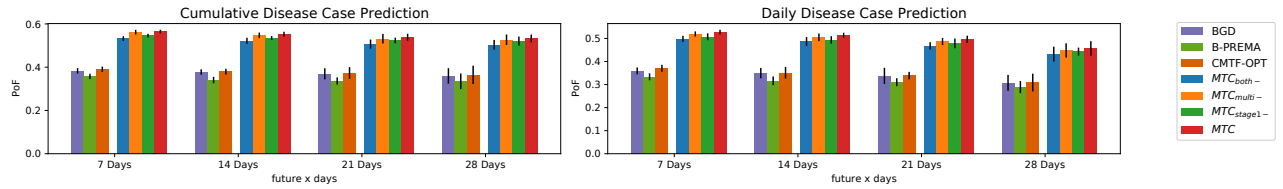


Figure 6: Predicting Future Disease Cases with the Learned Factors on HT. The results of tensor prediction accord with the tensor completion results, where our MTC and the variants outperform the baselines by a great margin. In addition, we conjecture that there might be some underlying weekly patterns in the data, which explains why model performance decreases slowly with longer prediction windows.

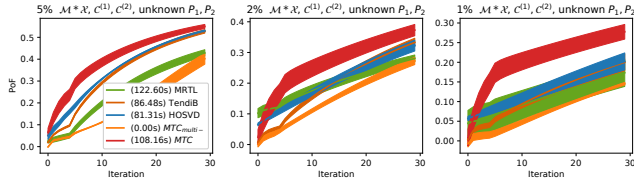


Figure 7: Initialization Comparison on HT. The multiresolution factorization in MTC outperforms other initialization methods. With the multiresolution factorization module, the 5 Jacobi iterations in stage 1 become very powerful, and then the stage 2 Cholesky solver continues to provide a stable and accurate estimation.

level regardless of the amount of partial observation since their initialization barely depends on the coarse views, i.e., $C^{(1)}$ and $C^{(2)}$. However, with only a few iterations, MTC can surpass all baselines.

7 CONCLUSION

This paper identifies and solves a new problem: *tensor completion from its partial and coarse observations*. We formulate the problem into a generic couple-ALS form and propose an efficient completion model, MTC. The module combines multiresolution factorization and an effective optimization method as the treatment. We evaluate our MTC on two COVID-19 databases and show noticeable performance gain over the baseline methods.

ACKNOWLEDGEMENTS

This work was in part supported by the National Science Foundation award SCH-2014438, PPOSS 2028839, IIS-1838042, the National Institute of Health award NIH R01 1R01NS107291-01 and OSF Healthcare.

REFERENCES

- [1] Evrim Acar, Daniel M Dunlavy, Tamara G Kolda, and Morten Mørup. 2011. Scalable tensor factorizations for incomplete data. *Chemometrics and Intelligent Laboratory Systems* (2011).
- [2] Evrim Acar, Tamara G Kolda, and Daniel M Dunlavy. 2011. All-at-once optimization for coupled matrix and tensor factorizations. *arXiv:1105.3422* (2011).
- [3] Ardavan Afshar, Kejing Yin, Sherry Yan, Cheng Qian, Joyce C Ho, Haesun Park, and Jimeng Sun. 2021. SWIFT: Scalable Wasserstein Factorization for Sparse Nonnegative Tensors. In *AAAI*.
- [4] Nir Ailon and Bernard Chazelle. 2006. Approximate nearest neighbors and the fast Johnson-Lindenstrauss transform. In *STOC*. 557–563.
- [5] Faisal M Almutairi, Charilaos I Kanatsoulis, and Nicholas D Sidiropoulos. 2020. Tendi: Tensor disaggregation from multiple coarse views. In *PAKDD*. Springer.
- [6] Faisal M Almutairi, Charilaos I Kanatsoulis, and Nicholas D Sidiropoulos. 2021. PREMA: principled tensor data recovery from multiple aggregated views. *IEEE Journal of Selected Topics in Signal Processing* (2021).
- [7] Miguel Ramos Araujo, Pedro Manuel Pinto Ribeiro, and Christos Faloutsos. 2017. TensorCast: Forecasting with context using coupled tensors. In *ICDM*.
- [8] Sanaz Bahargam and E. Papalexakis. 2018. Constrained coupled matrix-tensor factorization and its application in pattern and topic detection. In *ASONAM*.
- [9] Boaz Barak and Ankur Moitra. 2016. Noisy tensor completion via the sum-of-squares hierarchy. In *COLT*. PMLR, 417–445.
- [10] Ricardo Augusto Borsoi, Clémence Prévost, Konstantin Usevich, David Brie, José Carlos Moreira Bermudez, and C. Richard. 2020. Coupled Tensor Decomposition for Hyperspectral and Multispectral Image Fusion with Variability. *arXiv* (2020).
- [11] J Douglas Carroll and Jih-Jie Chang. 1970. Analysis of individual differences in multidimensional scaling via an N-way generalization of “Eckart-Young” decomposition. *Psychometrika* 35, 3 (1970), 283–319.
- [12] A. Eshragh, F. Roosta, A. Nazari, and M. Mahoney. 2019. LSAR: Efficient Leverage Score Sampling Algorithm for the Analysis of Big Time Series Data. *arXiv* (2019).
- [13] Frank L Hitchcock. 1927. The expression of a tensor or a polyadic as a sum of products. *Journal of Mathematics and Physics* 6, 1-4 (1927), 164–189.
- [14] David Hong, Tamara G Kolda, and Jed A Duersch. 2020. Generalized canonical polyadic tensor decomposition. *SIAM Rev.* 62, 1 (2020), 133–163.
- [15] Huyan Huang, Yipeng Liu, and Ce Zhu. 2020. A Unified Framework for Coupled Tensor Completion. *arXiv preprint arXiv:2001.02810* (2020).
- [16] Nikos Kargas, Cheng Qian, Nicholas D Sidiropoulos, Cao Xiao, Lucas M Glass, and Jimeng Sun. 2021. STELAR: Spatio-temporal Tensor Factorization with Latent Epidemiological Regularization. In *AAAI*.
- [17] Lars Karlsson, Daniel Kressner, and André Uschmajew. 2016. Parallel Algorithms for Tensor Completion in the CP Format. *Parallel Comput.* 57, C (2016).
- [18] Tamara G Kolda and Brett W Bader. 2009. Tensor decompositions and applications. *SIAM review* 51, 3 (2009), 455–500.
- [19] Timothee Lacroix, Nicolas Usunier, and Guillaume Obozinski. 2018. Canonical Tensor Decomposition for Knowledge Base Completion. In *ICML*.
- [20] Brett W Larsen and Tamara G Kolda. 2020. Practical leverage-based sampling for low-rank tensor decomposition. *arXiv preprint arXiv:2006.16438* (2020).
- [21] Ji Liu, Przemyslaw Musialski, Peter Wonka, and Jieping Ye. 2012. Tensor completion for estimating missing values in visual data. *IEEE TPAMI* 35, 1 (2012).
- [22] Alp Özdemir, Mark A Iwen, and Selin Ayiyente. 2016. Multiscale tensor decomposition. In *ACSCS*. IEEE, 625–629.
- [23] J. Park, K. Carr, S. Zheng, Y. Yue, and R. Yu. 2020. Multiresolution Tensor Learning for Efficient and Interpretable Spatial Analysis. In *ICML*. PMLR, 7499–7509.
- [24] Yubin Park and Joydeep Ghosh. 2014. Ludia: An aggregate-constrained low-rank reconstruction algorithm to leverage publicly released health data. In *KDD*.
- [25] Claudio Schifanella, K Selçuk Candan, and Maria Luisa Sapino. 2014. Multiresolution tensor decompositions with mode hierarchies. *TKDD* 8, 2 (2014), 1–38.
- [26] Nicholas D Sidiropoulos, Lieven De Lathauwer, Xiao Fu, Kejun Huang, Evangelos E Papalexakis, and Christos Faloutsos. 2017. Tensor decomposition for signal processing and machine learning. *IEEE TSP* 65, 13 (2017), 3551–3582.
- [27] Qingquan Song, Hancheng Ge, James Caverlee, and Xia Hu. 2019. Tensor completion algorithms in big data analytics. *ACM TKDD* 13, 1 (2019), 1–48.
- [28] Qingquan Song, Xiao Huang, Hancheng Ge, James Caverlee, and Xia Hu. 2017. Multi-aspect streaming tensor completion. In *KDD*. 435–443.
- [29] Hans De Sterck and Killian Miller. 2013. An adaptive algebraic multigrid algorithm for low-rank canonical tensor decomposition. *SIAM SISC* 35, 1 (2013), B1–B24.
- [30] Senzhang Wang, Jiannong Cao, and Philip Yu. 2020. Deep learning for spatio-temporal data mining: A survey. *TKDE* (2020).
- [31] Yichen Wang, Robert Chen, Joydeep Ghosh, Joshua C Denny, Abel Kho, You Chen, Bradley A Malin, and Jimeng Sun. 2015. Rubik: Knowledge guided tensor factorization and completion for health data analytics. In *KDD*.
- [32] Yangyang Xu and Wotao Yin. 2013. A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion. *SIAM Journal on imaging sciences* 6, 3 (2013).
- [33] Kejing Yin, William K Cheung, Yang Liu, Benjamin CM Fung, and Jonathan Poon. 2019. Joint Learning of Phenotypes and Diagnosis-Medication Correspondence via Hidden Interaction Tensor Factorization. In *IJCAI*.
- [34] Zecheng Zhang, Xiaoxiao Wu, Naijing Zhang, Siyuan Zhang, and Edgar Solomonik. 2019. Enabling distributed-memory tensor completion in Python using new sparse tensor kernels. *arXiv preprint arXiv:1910.02371* (2019).

A BASICS OF TENSOR COMPUTATION

Kronecker Product. One important product for matrices is Kronecker product. For $\mathbf{A} \in \mathbb{R}^{I \times J}$ and $\mathbf{B} \in \mathbb{R}^{K \times L}$, their Kronecker product is defined by (each block is a scalar times matrix)

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} \mathbf{A}(1,1) \cdot \mathbf{B} & \mathbf{A}(1,2) \cdot \mathbf{B} & \cdots & \mathbf{A}(1,J) \cdot \mathbf{B} \\ \mathbf{A}(2,1) \cdot \mathbf{B} & \mathbf{A}(2,2) \cdot \mathbf{B} & \cdots & \mathbf{A}(2,J) \cdot \mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{A}(I,1) \cdot \mathbf{B} & \mathbf{A}(I,2) \cdot \mathbf{B} & \cdots & \mathbf{A}(I,J) \cdot \mathbf{B} \end{bmatrix} \in \mathbb{R}^{IK \times JL}.$$

Khatri-Rao Product. Khatri-Rao product is another important product for matrices, specifically, for matrices with same number of columns. The Khatri-Rao product of $\mathbf{A} \in \mathbb{R}^{I \times J}$ and $\mathbf{B} \in \mathbb{R}^{K \times J}$ can be viewed as column-wise Kronecker product,

$$\mathbf{A} \odot \mathbf{B} = [\mathbf{A}(:,1) \otimes \mathbf{B}(:,1), \mathbf{A}(:,2) \otimes \mathbf{B}(:,2), \dots, \mathbf{A}(:,L) \otimes \mathbf{B}(:,L)],$$

where $\mathbf{A} \odot \mathbf{B} \in \mathbb{R}^{IK \times L}$.

Tensor Unfolding. This operation is to matricize a tensor along one mode. For tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$, we could unfold it along the first mode into a matrix $\mathbf{X}_1 \in \mathbb{R}^{I_1 \times I_2 I_3}$ (we use subscript notation). Specifically, each row of \mathbf{X}_1 is a vectorization of a slab in the original tensor; we have

$$\mathbf{X}_1(i, j \times I_3 + k) = \mathcal{X}(i, j, k).$$

Similarly, for the unfolding operation along the second or third mode, we have

$$\mathbf{X}_2(j, i \times I_3 + k) = \mathcal{X}(i, j, k) \in \mathbb{R}^{I_2 \times I_1 I_3},$$

$$\mathbf{X}_3(k, i \times I_2 + j) = \mathcal{X}(i, j, k) \in \mathbb{R}^{I_3 \times I_1 I_2}.$$

Hadamard Product. The Hadamard product is the element-wise product for tensors of the same size. For example, the Hadamard product of two 3-mode tensors $\mathcal{X}, \mathcal{Y} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$ is

$$\mathcal{Z} = \mathcal{X} * \mathcal{Y} \in \mathbb{R}^{I_1 \times I_2 \times I_3}.$$

B SOLVER IMPLEMENTATION

We show the implementation of two stages in our solver.

Stage 1. Suppose that we want to solve for the factor matrix \mathbf{U} , given Equation (27). To simplify the derivation, we assume

$$\mathbf{A} = (\mathbf{V}^\top \mathbf{V}) * (\mathbf{W}^\top \mathbf{W}) + \lambda_2 (\mathbf{Q}_2^\top \mathbf{Q}_2) * (\mathbf{W}^\top \mathbf{W}), \quad (31)$$

$$\mathbf{B} = (\mathbf{V} \odot \mathbf{W})^\top \tilde{\mathbf{X}}_1^\top + \lambda_2 (\mathbf{Q}_2 \odot \mathbf{W})^\top \mathbf{C}_1^{(2)\top}. \quad (32)$$

The Jacobi method first decomposes \mathbf{A} into a diagonal matrix \mathbf{D} and an off-diagonal matrix \mathbf{R} , such that $\mathbf{A} = \mathbf{D} + \mathbf{R}$. After moving the off-diagonal part to the right, Equation (27) becomes,

$$(\mathbf{D} + \mathbf{R})\mathbf{U}^\top = \mathbf{B} \Rightarrow \mathbf{D}\mathbf{U}^\top = \mathbf{B} - \mathbf{R}\mathbf{U}^\top. \quad (33)$$

It is cheap to take the inverse of the diagonal \mathbf{D} (before the inverse, we add a small $\epsilon = 1e-5$ to the diagonal values to improve numerical stability) on both sides. The Jacobi iteration is given by the following recursive equation,

$$\mathbf{U}^\top = \mathbf{D}^{-1}\mathbf{B} - \mathbf{D}^{-1}\mathbf{R}\mathbf{U}^\top \Rightarrow \mathbf{U}^{\top(k+1)} = \mathbf{D}^{-1}\mathbf{B} - \mathbf{D}^{-1}\mathbf{R}\mathbf{U}^{\top(k)}. \quad (34)$$

The convergence property is ensured by the largest absolute value of eigenvalues (i.e., spectral radius) of $\mathbf{D}^{-1}\mathbf{R}$. If the eigenvalues of $\mathbf{D}^{-1}\mathbf{R}$ are all in $(-1, 1)$, then the iteration will converge. The iteration is enhanced by adding the first order-momentum.

$$\begin{aligned} \mathbf{U}^{\top(k+1)} &= (1 - w)\mathbf{U}^{\top(k)} + w \underbrace{(\mathbf{D}^{-1}\mathbf{B} - \mathbf{D}^{-1}\mathbf{R}\mathbf{U}^{\top(k)})}_{\text{the Jacobi term}} \\ &= w\mathbf{D}^{-1}\mathbf{B} + (1 - w\mathbf{D}^{-1}\mathbf{A})\mathbf{U}^{\top(k)}. \end{aligned} \quad (35)$$

Now, the new spectral radius of $1 - w\mathbf{D}^{-1}\mathbf{A}$ is controlled by w .

Stage 2. This stage uses Cholesky decomposition to solve the joint normal equation. We keep the assignments in Equation (31) (32). First, a Cholesky decomposition is applied on the symmetric \mathbf{A} ,

$$\mathbf{A} \Rightarrow \mathbf{L}\mathbf{L}^\top, \quad (36)$$

where \mathbf{L} is a lower triangular matrix. Now, the problem reduces to solving two triangular systems in sequence.

$$\mathbf{L}\mathbf{Y} = \mathbf{B}, \quad (37)$$

$$\mathbf{L}^\top \mathbf{U}^\top = \mathbf{Y}. \quad (38)$$

We then use `scipy.linalg.solve_triangular` to get the exact solution.

C DETAILS IN EXPERIMENTS

We publish the dataset and codebase in this anonymous repo⁵. We list the GCSS data, processing files, the synthetic data, codes of our models, and all other baselines in the repo.

C.1 Baseline Implementation

We provide more details about baselines used in the experiments:

- **CPC-ALS** is a SOTA tensor completion method. It could be integrated with the sparse structure on parallel machines. The implementation could refer to [17]. We provide our implementation in the link.
- **OracleCPD** is implemented by standard CP decomposition on the original tensor, which provides the target low-rank factors. We implement it with our proposed solver. After convergence, the PoF results and objective functions are similar to the results provided by python Package `tensorly.decomposition.parafac`.
- **Block Gradient Descent (BGD)** uses $\mathbf{Q}_1, \mathbf{Q}_2$ to replace $\mathbf{P}_1\mathbf{U}, \mathbf{P}_2\mathbf{V}$ in Equation (4), and then apply the block gradient descent algorithm. The learning rate is chosen by a binary line search algorithm with depth 3.
- **B-PREMA** [6]: this work is similar to BGD but with extra loss terms,

$$\beta_1 \|\mathbf{1}^\top \mathbf{Q}_1 - \mathbf{1}^\top \mathbf{U}\|^2 + \beta_2 \|\mathbf{1}^\top \mathbf{Q}_2 - \mathbf{1}^\top \mathbf{V}\|^2, \quad (39)$$

ensuring that after the aggregation, the respective column sums of this two factors should be equal. Then, it also uses line search to find the learning rate.

- **CMTF-OPT** [2] is a popular model for coupled matrix and tensor factorization. We adopt it for coupled tensor factorization. Instead of alternatively optimizing factor matrices, it updates all factors at once. We use gradient descent to implement it, and the learning rate is also selected by line search.
- **MRTL** [23] designs a multiresolution tensor algorithm to build good initialization for accelerating their classification task. We adopt their tensor initialization method and compare it to our proposed multiresolution factorization.

⁵<https://github.com/ycq091044/MTC>

- **TendiB** [5] provides a heuristic initialization trick for the coupled factorization task. We only adopt its initialization based on our problem setting. In our case, we implement it by: we apply CPD on $C^{(1)}$ to estimate Q_1 , V and W first, and then we freeze W and use it to estimate U and Q_2 by applying another CPD on $C^{(2)}$. We have also tried to interpolate U and Q_2 from the estimated Q_1 and V . However, it performs poorly. We finally adopted the first implementation.
- **Higher-order singular value decomposition (HOSVD)** [18] is commonly used for initializing tensor CP-ALS decomposition. It first alternatively estimates orthogonal factors by higher-order orthogonal iteration (HOOI) and then applies ALS on the core tensor. The final initialization is the compositions of the orthogonal factor and core ALS factor. We adopt the HOSVD method as another initialization baseline.

C.2 Tasks and Metrics

We provide more details and rationality on the metrics.

- **Percentage of Fitness (PoF).** Several performance evaluation metrics have been introduced in a survey paper [27], and PoF is the most common one. For this metric, we also performed two-tailed Student T-test and showed the p -value in Table 3.
- **CPU Time.** We make the methods execute for the same number of iterations and then compare their running time. The running time includes data loading, initialization (in our methods, it also includes the expense on low resolutions) and excludes the metric computation time.
- **Peak Memory.** To evaluate space complexity, we record the memory usage during the optimization process and report the peak memory load.

C.3 Implementation and Hyperparameters

All the experiments are implemented by *Python 3.7.8*, *scipy 1.5.2* and *numpy 1.19.1* on Intel Cascade Lake Linux platform with 64 GB memory and 80 vCPU cores. By default, we conduct each experiment three times with different random seeds.

Since the CPD accepts a scaling invariance property, we apply the following trick after each iteration. For example, a set of factors $\{U, V, W\}$ would be identical to another factor set $\{2U, 0.5V, W\}$ in terms of both the completion task or the prediction task. During the implementation of all models, we have adopted the following rescaling strategy to minimize the float round-off error:

$$U(:, i) \leftarrow f_i \cdot \frac{U(:, i)}{\|U(:, i)\|} \quad (40)$$

$$V(:, i) \leftarrow f_i \cdot \frac{V(:, i)}{\|V(:, i)\|} \quad (41)$$

$$W(:, i) \leftarrow f_i \cdot \frac{W(:, i)}{\|W(:, i)\|} \quad (42)$$

where

$$f_i = (\|U(:, i)\| \|V(:, i)\| \|W(:, i)\|)^{\frac{1}{3}}, \forall i = [1..R] \quad (43)$$

which essentially equilibrates the norms of the factors of each component (also do it for Q_2 and Q_3).

C.4 Experiments on Synthetic Data

The experiments on synthetic data verify our model in the ideal low-rank case.

Synthetic Tensor. The data is generated by three uniformly random rank-10 factors with 125 as the mode size consistently, $U, V, W \in \mathbb{R}^{125 \times 10}$. We sort each column of V, W , to make sure the mode smoothness. Thus, a tensor $X = \llbracket U, V, W \rrbracket \in \mathbb{R}^{125 \times 125 \times 125}$ is constructed, assuming that the first mode is categorical and the second/third modes are continuous. We also generate two aggregation matrix $P_1, P_2 \in \{0, 1\}^{12 \times 125}$ for the first and second modes, separately, to obtain coarse tensors, $C^1 = X \times_1 P_1$, $C^2 = X \times_2 P_2$.

Verification on Synthetic Data. On the synthetic data, our MTC is compared with three variants. We find that with only 2% ~ 3% partially observed elements, all models can achieve near 100% PoF, which is an exact completion of the low-rank synthetic X . We thus further reduce the partial observation rate and consider three scenarios: (i) only 1% observed data; (ii) 1% observation with C^1, C^2 , where P_1 is unknown and P_2 is known; (iii) 1% observation with C^1, C^2 , where P_1, P_2 are both unknown.

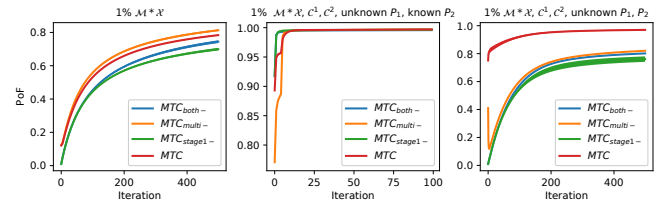


Figure 8: Comparison on Synthetic Data

Results are shown in Figure 8. Basically, we conclude that: (i) our model is more advantageous with coarse level information; (ii) the multiresolution factorization and stage 1 are better used together.

C.5 Additional Experiments on GCSS with Known Aggregation

The main paper experiments do not use any aggregation information since, in those scenarios, the differences between model variants are clearer. This appendix supplies one aggregation matrix along the third mode and shows that all our variants can achieve a similar tensor completion performance as OracleCPD.

Specifically, on GCSS dataset, we keep other settings unchanged and supply the aggregation matrix (from date to week) along the third mode. All variants converge quickly, and thus we only run the experiments for 50 iterations. The results are shown below.

Table 4: Results with Known Aggregation on GCSS

Model	PoF
OracleCPD	0.7780 ± 0.0012
MTC _{both-}	0.7706 ± 0.0058
MTC _{multi-}	0.7694 ± 0.0025
MTC _{stage1-}	0.7704 ± 0.0037
MTC	0.7708 ± 0.0026