Linear Algebra and its Applications 481 (2015) 131–173

Contents lists available at ScienceDirect

Linear Algebra and its Applications

www.elsevier.com/locate/laa

Optimization on the Hierarchical Tucker manifold – Applications to tensor completion

Curt Da Silva [a,*], Felix J. Herrmann [b]

[a] Department of Mathematics, University of British Columbia, Canada
[b] Department of Earth and Ocean Sciences, University of British Columbia, Canada

ARTICLE INFO

Article history:
Received 15 April 2014
Accepted 10 April 2015
Available online 14 May 2015
Submitted by H. Rauhut

MSC:
15A69
57N35
65K10
65Y20
53B21
90C90

Keywords:
Hierarchical Tucker tensors
Tensor completion
Riemannian manifold optimization
Gauss–Newton
Differential geometry
Low-rank tensor

ABSTRACT

In this work, we develop an optimization framework for problems whose solutions are well-approximated by *Hierarchical Tucker* (HT) tensors, an efficient structured tensor format based on recursive subspace factorizations. By exploiting the smooth manifold structure of these tensors, we construct standard optimization algorithms such as Steepest Descent and Conjugate Gradient for completing tensors from missing entries. Our algorithmic framework is fast and scalable to large problem sizes as we do not require SVDs on the ambient tensor space, as required by other methods. Moreover, we exploit the structure of the Gramian matrices associated with the HT format to regularize our problem, reducing overfitting for high subsampling ratios. We also find that the organization of the tensor can have a major impact on completion from realistic seismic acquisition geometries. These samplings are far from idealized randomized samplings that are usually considered in the literature but are realizable in practical scenarios. Using these algorithms, we successfully interpolate large-scale seismic data sets and demonstrate the competitive computational scaling of our algorithms as the problem sizes grow.

* Corresponding author.
E-mail address: curtd@math.ubc.ca (C. Da Silva).

http://dx.doi.org/10.1016/j.laa.2015.04.015

1. Introduction

The matrix completion problem is concerned with interpolating an $m \times n$ matrix from a subset of its entries. The amount of recent successes in developing solution techniques to this problem is a result of assuming a *low-rank* model on the 2-D signal of interest and a uniform random sampling scheme [9,8,10]. The original signal is recovered by promoting low-rank structures subject to data constraints.

Using a similar approach, we consider the problem of interpolating a d-dimensional tensor from samples of its entries. That is, we aim to solve,

$$\min_{\mathbf{X} \in \mathcal{H}} \frac{1}{2} \|P_\Omega \mathbf{X} - b\|_2^2, \tag{1}$$

where P_Ω is a linear operator $P_\Omega : \mathbb{R}^{n_1 \times n_2 \times \ldots \times n_d} \to \mathbb{R}^m$, $b \in \mathbb{R}^m$ is our subsampled data satisfying $b = P_\Omega \mathbf{X}^*$ for some "solution" tensor $\mathbf{X}^*$ and $\mathcal{H}$ is a *specific* class of low-rank tensors to be specified later. Under the assumption that $\mathbf{X}^*$ is well approximated by an element in $\mathcal{H}$, our goal is to recover $\mathbf{X}^*$ by solving (1). For concreteness, we concern ourselves with the case when P_Ω is a restriction operator, i.e.,

$$P_\Omega \mathbf{X} = \mathbf{X}_{i_1, i_2, \ldots, i_d} \quad \text{if } (i_1, i_2, \ldots, i_d) \in \Omega,$$

and $\Omega \subset [n_1] \times [n_2] \times \cdots \times [n_d]$ is the so-called *sampling set*, where $[n] = \{1, \ldots, n\}$. In the above equation, we suppose that $|\Omega| = m \ll n_1 n_2 \ldots n_d$, so that P_Ω is a subsampling operator.

Unlike the matrix case, there is no unique notion of rank for tensors, as we shall see in Section 1.1, and there are multiple tensor formats that generalize a particular notion of *separability* from the matrix case—i.e., there is no unique extension of the SVD to tensors. Although each tensor format can lead to compressible representations of their respective class of low-rank signals, the truncation of a general signal to one of these formats requires access to the *fully* sampled tensor $\mathbf{X}$ (or at the very least *query*-based access to the tensor [4]) in order to achieve reasonable accuracy, owing to the use of truncated SVDs acting on various *matricizations* of the tensor. As in matrix completion, randomized missing entries change the behavior of the singular values and vectors of these matricizations and hence of the final approximation. Moreover, when the tensor of interest is a discretized continuous signal, there can be a number of constraints, physical or otherwise, that limit our ability to ideally sample it. For instance, in the seismic case, the tensor of interest is a multi-dimensional wavefield in the earth's subsurface sampled at an array of receivers located at the surface. In real-world seismic experiments, budgetary constraints or environmental obstructions can limit both the total amount of time available for data acquisition as well as the number and placement of active sources and receivers. Many scientific areas, including seismic processing, rely on having fully sampled data for drawing accurate inferences and, as a result, tensor completion is an important technique for acquiring multidimensional data.