Journées du GDR AMORE, Dec. 2016

Tensor numerical methods for high-dimensional problems

—

Part 2

Approximation in low-rank formats

We present algorithms for computing low-rank approximations of the solution of variational problems

$$\min_{v \in V} \mathcal{J}(v),$$

where V is a tensor space.

Different contexts

- For the approximation of a given tensor u with respect to a certain norm,

$$\mathcal{J}(v) = \|u - v\|.$$

Here, the aim is the compression of u or the extraction of information from u (data analysis).

- For the solution of an equation $Au = b$, the functional $\mathcal{J}(v)$ will measure some distance between u and the approximation v, e.g.

$$\mathcal{J}(v) = \|Av - b\|.$$

The aim is here to obtain an approximation of the solution u with a low computational complexity.

Different contexts

- In tensor completion,
$$\mathcal{J}(v) = \sum_{i \in \Omega} |u(i) - v(i)|^2,$$
where $\Omega \subset I$ is a set of known entries of the tensor. The aim is here to recover (or complete) a tensor from partial information, by exploiting low-rank structures of the tensor.
- For inverse problems, where we want to identify a tensor u from indirect and partial observations, the functional $\mathcal{J}(v)$ measures some distance between observations y and a prediction Av, where A is an observation map:
$$\mathcal{J}(v) = d(y, Av).$$
Exploiting low-rank structures in u allows to reduce the number of parameters to estimate and possibly makes the problem well-posed.

Different contexts

- For least-squares approximation of a function $u(X)$ from samples $\{u(x^k)\}_{k=1}^n$,

$$\mathcal{J}(v) = \frac{1}{n}\sum_{k=1}^{n}(u(x^k) - v(x^k))^2$$

- Other problems in statistics and machine learning
  - Supervised learning of the relation between a random variable Y and another random variable X from samples $\{(x_k, y_k)\}_{k=1}^n$: minimization of a risk functional

$$\mathcal{J}(v) = \frac{1}{n}\sum_{k=1}^{n}\ell(y^k, v(x^k))$$

  - Estimation of the density of a random variable X from samples $\{x_k\}_{k=1}^n$: minimizing the log-likelihood function

$$\mathcal{J}(v) = -\sum_{k=1}^{n}\log(v(x^k))$$

  - ...

Outline

1. Higher-order singular value decomposition and tensor truncation
2. Direct optimization in subsets of low-rank tensors
3. Iterative solvers with tensor truncation
4. Greedy algorithms

We consider a tensor u in a Hilbert tensor space $V = V^1 \otimes \ldots \otimes V^d$ and we assume that u is given as a full tensor or in a certain low-rank format.

We present truncation schemes for finding a low-rank approximation of u with reduced complexity, relying on the standard singular value decomposition of order-two tensors.

We denote by $\|\cdot\|$ the canonical tensor norm on V.

Truncated singular value decomposition for order-two tensors

Let u be an order-two tensor in the Hilbert space $V \otimes W$, where V and W are Hilbert spaces, and let $\|\cdot\|$ denote the canonical norm on $V \otimes W$ (the Frobenius norm for u a matrix).

Let consider a tensor u in $V \otimes W$ with singular value decomposition

$$u = \sum_{k=1}^{N} \sigma_k v_k \otimes w_k,$$

where the singular values are sorted by decreasing order.

An element of best approximation of u in the set of tensors with rank bounded by r is provided by the truncated singular value decomposition

$$u_r = \sum_{k=1}^{r} \sigma_k v_k \otimes w_k,$$

such that

$$\|u - u_r\|^2 = \min_{\text{rank}(v) \leq r} \|u - v\|^2 = \sum_{k=r+1}^{N} \sigma_k^2.$$

Truncated singular value decomposition for order-two tensors

An approximation u_r with relative precision ϵ, such that

$$\|u - u_r\| \leq \epsilon \|u\|,$$

can be obtained by choosing a rank r such that

$$\sum_{k=r+1}^{N} \sigma_k^2 \leq \epsilon \sum_{k=1}^{N} \sigma_k^2.$$

Remark.

The complexity of computing the singular value decomposition of a tensor u is $O(n^3)$ if $\dim(V) = \dim(W) = O(n)$. If u is given in low-rank format $u = \sum_{k=1}^{R} a_k \otimes b_k$, with a rank $R < n$, the complexity breaks down to $O(R^3 + 2Rn^2)$.

Higher-order tensors as order-two tensors...

For a non-empty subset α in $D = \{1, \dots, d\}$, a tensor $u \in V^1 \otimes \dots \otimes V^d$ can be identified with its matricisation

$$\mathcal{M}_\alpha(u) \in V^\alpha \otimes V^{\alpha^c},$$

an order-two tensor which admits a singular value decomposition

$$\mathcal{M}_\alpha(u) = \sum_{k \geq 1} \sigma_k^\alpha v_k^\alpha \otimes w_k^{\alpha^c}.$$

The set $\sigma^\alpha(u) := \{\sigma_k^\alpha\}_{k \geq 1}$ is called the set of α-singular values of u. The α-rank of u is the number of non-zero α-singular values

$$\text{rank}_\alpha(u) = \|\sigma^\alpha(u)\|_0.$$

Higher-order tensors as order-two tensors...

By sorting the α-singular values by decreasing order, an approximation u_r with α-rank r can be obtained by retaining the r largest singular values, i.e.

$$u_r \quad \text{such that} \quad \mathcal{M}_\alpha(u_r) = \sum_{k=1}^{r} \sigma_k^\alpha v_k^\alpha \otimes w_k^{\alpha^c},$$

which satisfies

$$\|u - u_r\|^2 = \min_{\text{rank}_\alpha(v) \leq r} \|u - v\|^2 = \sum_{k>r} (\sigma_k^\alpha)^2.$$

There are 2^{d-1} different binary partitions $\alpha \cup \alpha^c$ of D, to each of which corresponds a singular value decomposition and a way to truncate a higher-order tensor !

Truncation scheme for the approximation in Tucker format

For each $\nu \in \{1, \dots, d\}$, we consider the singular value decomposition of the matricisation $\mathcal{M}_\nu(u)$ of a tensor u

$$\mathcal{M}_\nu(u) = \sum_{k\geq 1} \sigma_k^\nu v_k^\nu \otimes w_k^\nu.$$

Let $U_{r_\nu}^\nu = \text{span}\{v_k^\nu\}_{k=1}^{r_\nu}$ be the subspace of V^ν generated by the r_ν dominant left singular vectors of $\mathcal{M}_\nu(u)$, and by $P_{U_{r_\nu}}$ the orthogonal projection from V^ν to $U_{r_\nu}^\nu$.

The tensor

$$u_r = (P_{U_{r_1}^1} \otimes \dots \otimes P_{U_{r_d}^d})u$$

is a projection of u onto the reduced tensor space

$$U_{r_1}^1 \otimes \dots \otimes U_{r_d}^d$$

and therefore,

$$u_r \in \mathcal{T}_r = \{v \in U^1 \otimes \dots \otimes U^d : U^\nu \subset V^\nu, \dim(U^\nu) = r_\nu, 1 \leq \nu \leq d\}.$$

Higher-order singular value decomposition for Tucker format

The operator

$$\mathcal{P}^{\nu}_{r_\nu} = \mathcal{M}_\nu^{-1} P_{U^\nu_{r_\nu}} \mathcal{M}_\nu = I \otimes \ldots \otimes P_{U^\nu_{r_\nu}} \otimes \ldots \otimes I$$

is the orthogonal projection from V onto

$$V^1 \otimes \ldots \otimes U^\nu_{r_\nu} \otimes \ldots \otimes V^d,$$

which is such that

$$\|u - \mathcal{P}^{\nu}_{r_\nu} u\| = \min_{\text{rank}_\nu(v) \leq r_\nu} \|u - v\| = \sum_{k \geq r_\nu + 1} (\sigma_k^\nu)^2.$$

The approximation u_r can then be written

$$u_r = \mathcal{P}^{1}_{r_1} \ldots \mathcal{P}^{d}_{r_d} u,$$

and satisfies

$$\|u - u_r\|^2 = \sum_{\nu=1}^{d} \|u - \mathcal{P}^{\nu}_{r_\nu} u\|^2 = \sum_{\nu=1}^{d} \min_{\text{rank}_\nu(v) \leq r_\nu} \|u - v\|^2,$$

from which we deduce the quasi-optimality property

$$\|u - u_r\| \leq \sqrt{d} \min_{v \in \mathcal{T}_r} \|u - v\|.$$

Truncation scheme for the approximation in Tucker format

Also, from

$$\|u - u_r\|^2 = \sum_{\nu=1}^{d} \|u - \mathcal{P}_{r_\nu}^{\nu} u\|^2 = \sum_{\nu=1}^{d} \sum_{k_\nu > r_\nu} (\sigma_{k_\nu}^{\nu})^2,$$

we deduce that if we select the ranks $(r_1, \ldots, r_d)$ such that for each ν

$$\sum_{k_\nu > r_\nu} (\sigma_{k_\nu}^{\nu})^2 \leq \frac{\epsilon^2}{d} \sum_{k_\nu \geq 1} (\sigma_{k_\nu}^{\nu})^2 = \frac{\epsilon^2}{d} \|u\|^2,$$

then the truncated singular value decomposition $\mathcal{P}_{r_\nu}^{\nu} u$ has a relative precision $\epsilon/\sqrt{d}$ and we finally obtain an approximation u_r with relative precision ϵ,

$$\|u - u_r\| \leq \epsilon \|u\|.$$

Note that the definition of u_r is independent on the order of the projections $\mathcal{P}_{r_\nu}^{\nu}$.

Truncation scheme for tree-based tensor formats

For tree-based (hierarchical) low-rank tensor formats

$$\mathcal{T}_r^T = \{v : \text{rank}_\alpha(v) \leq r_\alpha, \alpha \in T\},$$

where T is a subset of a dimension partition tree over $D = \{1, \dots, d\}$, a higher order singular value decomposition (also called hierarchical singular value decomposition) can also be defined from singular value decompositions of matricisations $\mathcal{M}_\alpha(u)$ of a tensor u.

Truncation scheme for tree-based tensor formats

Letting $U^{\alpha}_{r_\alpha}$ be the subspace generated by the r_α dominant left singular vectors of $\mathcal{M}_\alpha(u)$, and letting $P_{U^{\alpha}_{r_\alpha}}$ be the orthogonal projector from V^α to $U^{\alpha}_{r_\alpha}$, we define the orthogonal projection

$$\mathcal{P}^{\alpha}_{r_\alpha} = \mathcal{M}_\alpha^{-1} P_{U^{\alpha}_{r_\alpha}} \mathcal{M}_\alpha .$$

Then, an approximation with tree-based rank $r = (r_\alpha)_{\alpha \in T}$ can be defined by

$$u_r = \mathcal{P}_r^{T,(L)} \mathcal{P}_r^{T,(L-1)} \dots \mathcal{P}_r^{T,(1)} u \quad \text{with} \quad \mathcal{P}^{T,(\ell)} = \prod_{\substack{\alpha \in T \\ \text{level}(\alpha)=\ell}} \mathcal{P}^{\alpha}_{r_\alpha}$$

where we apply to u a sequence of projections ordered by increasing level in the tree (from the root to the leaves). Here $L = \max_{\alpha \in T} \text{level}(\alpha)$.

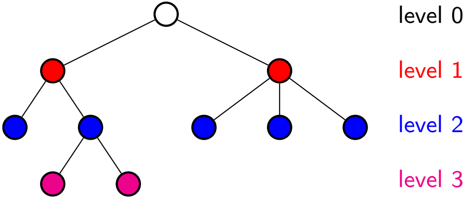

Truncation scheme for tree-based tensor formats

The obtained approximation u_r is such that

$$\|u - u_r\| \leq \sqrt{\#T} \min_{v \in \mathcal{T}_r^T} \|u - v\|,$$

where $\#T \leq 2d - 2$.

Also, if we select the ranks $(r_\alpha)_{\alpha \in T \setminus D}$ such that for all α

$$\sum_{k_\alpha > r_\alpha} (\sigma^\alpha_{k_\alpha})^2 \leq \frac{\epsilon^2}{\#T} \|u\|^2,$$

we obtain an approximation u_r with relative precision ϵ,

$$\|u - u_r\| \leq \epsilon \|u\|.$$

Outline

1. Higher-order singular value decomposition and tensor truncation
2. Direct optimization in subsets of low-rank tensors
3. Iterative solvers with tensor truncation
4. Greedy algorithms

Direct optimization in subsets of low-rank tensors

Let $\mathcal{M}_r$ be a subset of tensors in a certain low-rank format $\mathcal{M}_r$ with a multilinear parametrization of the form

$$v(i_1, \ldots, i_d) = \sum_{k_1=1}^{r_1} \ldots \sum_{k_L=1}^{r_L} \prod_{\nu=1}^{d} p^{(\nu)}(i_\nu, (k_i)_{i \in S_\nu}) \prod_{\nu=d+1}^{M} p^{(\nu)}((k_i)_{i \in S_\nu})$$

and let

$$\mathcal{M}_r = \{v = \Psi(p^{(1)}, \ldots, p^{(M)}) : p^{(\nu)} \in P^{(\nu)}, 1 \leq \nu \leq M\},$$

where Ψ is a multilinear map.

The problem

$$\min_{v \in \mathcal{M}_r} \mathcal{J}(v)$$

can be written as an optimization problem over the parameters

$$\min_{p^{(1)}} \ldots \min_{p^{(M)}} \mathcal{J}(\Psi(p^{(1)}, \ldots, p^{(M)})).$$

Alternating minimization algorithm

The alternating minimization algorithm consists in solving successively minimization problems

$$\min_{p^{(\nu)} \in P^{(\nu)}} \mathcal{J}(\Psi(p^{(1)}, \dots, p^{(\nu)}, \dots, p^{(M)})) := \min_{p^{(\nu)} \in P^{(\nu)}} \mathcal{J}_\nu(p^{(\nu)}) \quad (1)$$

over the parameter $p^{(\nu)}$, letting the other parameters $p^{(\eta)}$, $\eta \neq \nu$, fixed.

When $P^{(\nu)}$ is a linear vector space, problem (1) is a linear approximation problem.

If $\mathcal{J}$ is a convex (resp. differentiable) functional, then $\mathcal{J}_\nu$ is a convex (resp. differentiable) functional.

Modified alternating minimization algorithm

Modified alternating minimization algorithm[1] is a modification of the alternating minimization algorithm which allows for an automatic rank adaptation.

It can be used for optimization in tree-based tensor formats or more general tensor networks.

At each step of the algorithm, we consider two nodes ν and η connected by an edge e and we update simultaneously the associated parameters $p^{(\nu)}$ and $p^{(\eta)}$.

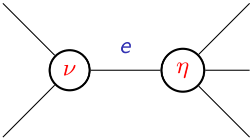

[1] known as DMRG algorithm (for Density Matrix Renormalization Group) for tensor networks.

Modified alternating minimization algorithm

In the expression of a tensor $v = \Psi(p^{(1)}, \dots, p^{(M)})$, the two tensors p^{ν} and p^{η} connected by the edge e appear as

$$\sum_{k_e=1}^{r_e} p^{(\nu)}(k_e, ...)p^{(\eta)}(k_e, ...) := p^{(e)}(...)$$

where $p^{(e)}$ is a tensor of order

$$\text{order}(p^{(e)}) = \text{order}(p^{(\nu)}) + \text{order}(p^{(\eta)}) - 2.$$

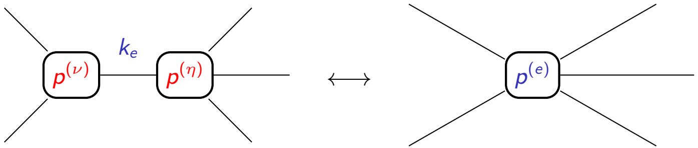

This corresponds to a new tensor networks where the nodes ν and η and edge e are replaced by a single node e, and a new parametrization

$$v = \Psi^e(\dots, p^{(e)}, \dots).$$

Modified alternating minimization algorithm

We first solve an optimization problem

$$\min_{p^{(e)}} \mathcal{J}(\Psi^e(\dots, p^{(e)}, \dots))$$

for obtaining an new value of the tensor $p^{(e)}$.

Then, we compute a low-rank approximation of the tensor $p^{(e)}$

$$p^{(e)}(...) \approx \sum_{k_e=1}^{r_e} p^{(\nu)}(k_e, ...)p^{(\eta)}(k_e, ...)$$

where the rank r_e in general differs from the initial rank.

In practice, the approximation is obtained using truncated singular value decomposition.

Direct optimization in subsets of low-rank tensors

Other optimization algorithms (e.g. gradient descent, Newton) can be used, possibly exploiting the geometry of low-rank tensor manifolds $\mathcal{M}_r$.

Under rather standard assumptions, some results have been obtained for the convergence of algorithms: local convergence to a global optimizer, or global convergence to stationary points.

Up to now, there is no available algorithm for obtaining a global optimizer of a general (even convex) functional in a subset of low-rank tensors.

1. Higher-order singular value decomposition and tensor truncation
2. Direct optimization in subsets of low-rank tensors
3. Iterative solvers with tensor truncation
4. Greedy algorithms

Iterative solvers with tensor truncation

Another strategy for solving an operator equation

$$Au = b$$

or a more general optimization problem

$$\min_{v \in V} \mathcal{J}(v)$$

is to rely on classical iterative solvers by interpreting all standard algebraic operations on vector spaces as algebraic operations in tensor spaces.

Iterative solvers with tensor truncation

As a motivating example, consider a simple Richardson algorithm

$$u^n = u^{n-1} - \omega(Au^{n-1} - b).$$

For A and b given in low-rank formats, computing u^n involves standard algebraic operations.

However, the representation rank of the iterates dramatically increases since

$$\mathrm{rank}(u^n) \approx \mathrm{rank}(A)\,\mathrm{rank}(u^{n-1}) + \mathrm{rank}(u^{n-1}) + \mathrm{rank}(b).$$

This requires additional truncation steps for reducing the ranks of the iterates, such as

$$u^n = T(u^{n-1} - \omega(Au^{n-1} - b)),$$

where $T(v)$ provides a low-rank approximation of v.

We now analyze the behavior of these algorithms depending on the properties of the truncation operator T.

Fixed point iterations algorithm

Let us consider a problem which can be written as a fixed point problem

$$F(u) = u,$$

where $F : V \to V$ is a contractive map, such that for all $u, v \in V$,

$$\|F(u) - F(v)\| \leq \rho \|u - v\|,$$

with $0 \leq \rho < 1$.

Then, consider the fixed point iterations algorithm

$$u^{n+1} = F(u^n)$$

which provides a sequence $(u^n)_{n \geq 1}$ which converges to u, such that

$$\|u - u^n\| \leq \rho^n \|u - u^0\|.$$

Example 1

For a problem $Au = b$, consider $F(u) = u - \omega(Au - b)$, with ω such that $\|I - \omega A\| < 1$. Fixed point iterations $u^{n+1} = u^n - \omega(Au^n - b)$ correspond to Richardson iterations.

Perturbed fixed point iterations algorithm

Now consider the perturbed fixed point iterations

$$v^{n+1} = F(u^n), \quad u^{n+1} = T(v^{n+1})$$

where T is a mapping which for a tensor v provides an approximation (called truncation) $T(v)$ in a certain low-rank format $\mathcal{M}_r$.

Truncations with controlled relative precision

Suppose that the mapping T provides an approximation with relative precision ϵ, i.e.

$$\|T(v) - v\| \leq \epsilon \|v\|.$$

This is made possible by using an adaptation of the ranks.

Then the sequence $(u^n)_{n\geq 1}$ is such that

$$\|u - u^n\| \leq \gamma^n \|u - u^0\| + \frac{\epsilon}{1-\gamma}\|u\|,$$

with $\gamma = \rho(1+\epsilon)$. Therefore, if $\gamma < 1$

$$\limsup_{n\to\infty} \|u - u^n\| \leq \frac{\epsilon}{1-\gamma}\|u\|$$

which means that the sequence tends to enter a neighborhood of u with radius $\frac{\epsilon}{1-\gamma}\|u\|$.

The drawback of this algorithm is that the ranks of the iterates are not controlled and may become very high during the iterations.

Truncations in fixed subsets

Now consider that the mapping T provides an approximation in a fixed subset of tensors $\mathcal{M}_r$ with rank bounded by r.

Let us assume that for all v, $T(v)$ provides a quasi-optimal approximation of v such that

$$\|T(v) - v\| \leq C \min_{w \in \mathcal{M}_r} \|v - w\|. \tag{2}$$

A practical realization of a mapping T verifying (2) is provided by truncated higher-order singular value decompositions, where

$$C = O(\sqrt{d}).$$

Truncations in fixed subsets

Let u_r be an element of best approximation of u, with

$$\|u - u_r\| = \min_{v \in \mathcal{M}_r} \|u - v\|.$$

The sequence $(u^n)_{n \geq 1}$ is such that

$$\|u - u^n\| \leq \gamma^n \|u - u^0\| + \frac{C}{1-\gamma} \|u - u_r\|,$$

with $\gamma = \rho(1 + C)$. If $\gamma < 1$ (which may be quite restrictive on ρ), we obtain

$$\limsup_{n \to \infty} \|u - u^n\| \leq \frac{C}{1-\gamma} \min_{v \in \mathcal{M}_r} \|u - v\|,$$

which means that the sequence tends to enter a neighborhood of u with radius $\frac{C}{1-\gamma}\sigma_r$, where σ_r is the best approximation error of u by elements of $\mathcal{M}_r$.

An advantage of this approach is that the ranks of the iterates are controlled. A drawback is that the condition $\gamma < 1$ imposes to rely on an iterative solver with small contractivity constant $\rho < (1 + C)^{-1}$, which may be quite restrictive (requires good preconditioners).

Truncations with non-expansive maps

Now we assume that the mapping T providing an approximation in low-rank format is non-expansive, i.e.

$$\|T(v) - T(w)\| \leq \|v - w\| \quad (3)$$

The sequence u^n is defined by

$$u^{n+1} = G(u^n),$$

where $G = T \circ F$ is a contractive mapping with the same contractivity constant ρ as F. Therefore, the sequence u^n converges to the unique fixed point $u^\star$ of G such that

$$G(u^\star) = u^\star,$$

with

$$\|u^\star - u^n\| \leq \rho^n \|u^\star - u^0\|.$$

The obtained approximation $u^\star$ is such that

$$(1+\rho)^{-1}\|u - T(u)\| \leq \|u - u^\star\| \leq (1-\rho)^{-1}\|u - T(u)\|.$$

A practical realization of a mapping T verifying (2) is provided by the soft singular values thresholding operator. The ranks of the iterates are not controlled. However, it is observed in practice that the ranks of iterates are usually lower than with truncations with controlled relative precision.

Convex relaxation

Consider an order two tensor u in a Hilbert tensor space $V \otimes W$ equipped with the canonical norm.

The constrained optimization problem

$$\min_{\mathrm{rank}(v) \leq r} \mathcal{J}(v)$$

is equivalent to

$$\min_{v} \mathcal{J}(v) + \tau \, \mathrm{rank}(v)$$

for some value of τ.

A convex optimization problem is obtained by replacing $\mathrm{rank}(v) = \|\sigma(v)\|_0$ by the function $\|\sigma(v)\|_1 = \|v\|_*$ (the trace norm of v)

$$\min_{v} \mathcal{J}(v) + \tau \|v\|_*$$

Soft thresholding of singular values

Let $u = \sum_{k\geq 1} \sigma_k v_k \otimes w_k$ be the singular value decomposition of u.

The solution of

$$\min_v \frac{1}{2}\|u - v\|^2 + \tau\|v\|_*$$

is given by

$$T_\tau(u) = \sum_{k\geq 1}(\sigma_k - \tau)_+ v_k \otimes w_k$$

which corresponds to a soft-thresholding of singular values.

In convex analysis, $T_\tau(u)$ is known as the proximal operator of the convex function $v \mapsto \tau\|v\|_*$.

The function $v \mapsto T_\tau(v)$ is non expansive:

$$\|T_\tau(u) - T_\tau(v)\| \leq \|u - v\|.$$

Proximal algorithms

Consider the problem

$$\min_v \mathcal{J}(v) + \tau\|v\|_*$$

A proximal algorithm constructs a sequence $(u^n)_{n\geq 1}$ as follows.

At iteration n, we linearize the function $\mathcal{J}$ around u^n and define u^{n+1} as the solution of

$$\min_v \mathcal{J}(u^n) + (\nabla\mathcal{J}(u^n), v - u^n) + \frac{\beta}{2}\|u - u^n\|^2 + \tau\|v\|_*$$

where β is a parameter.

This is equivalent to solving

$$\min_v \frac{1}{2}\|v - (u^n - \beta^{-1}\nabla\mathcal{J}(u^n))\|^2 + \frac{\tau}{\beta}\|v\|_*$$

whose solution is provided by

$$u^{n+1} = T_{\tau/\beta}(u^n - \beta^{-1}\nabla\mathcal{J}(u^n))$$

where $T_{\tau/\beta}$ is the proximal operator of $v \mapsto \frac{\tau}{\beta}\|v\|_*$.

1 Higher-order singular value decomposition and tensor truncation

2 Direct optimization in subsets of low-rank tensors

3 Iterative solvers with tensor truncation

4 Greedy algorithms
- Greedy algorithms for canonical format
- Greedy algorithms for Tucker format
- Partially greedy algorithms for Tucker format

1. Higher-order singular value decomposition and tensor truncation
2. Direct optimization in subsets of low-rank tensors
3. Iterative solvers with tensor truncation
4. Greedy algorithms
   - Greedy algorithms for canonical format
   - Greedy algorithms for Tucker format
   - Partially greedy algorithms for Tucker format

Greedy algorithms for canonical format

A tensor $v \in \mathcal{R}_r$ with canonical rank r can be written as a sum of r rank-one tensors

$$v = \sum_{k=1}^{r} c_k w_k, \quad w_k \in \mathcal{R}_1.$$

Therefore, v can be interpreted as a n-sparse element with respect to dictionary of rank-one tensors $\mathcal{R}_1$.

Greedy algorithms for canonical format

Standard greedy algorithms can be used to construct a sequence of approximations u^n with increasing canonical rank

$$u^n = \sum_{k=1}^{n} c_k^n w_k, \quad c_k^n \in \mathbb{R},$$

where

$$w_n = w_n^{(1)} \otimes \ldots \otimes w_n^{(d)} \in \mathcal{R}_1$$

is such that

$$w_n \in \arg\min_{w \in \mathcal{R}_1} \mathcal{J}(u^{n-1} + w), \tag{4}$$

and where the coefficients c_k^n can be either taken as $c_k^n = 1$ (for a pure greedy algorithm), or as the solution of

$$\min_{c_1,\ldots,c_n} \mathcal{J}(\sum_{k=1}^{n} c_k w_k). \tag{5}$$

Each step requires to solve an optimization problem in $\mathcal{R}_1$, for which we can rely on an alternating minimization algorithm or other optimization algorithms.

Greedy algorithms with dictionary of low-rank tensors

These algorithms are essentially used for the approximation in canonical format but $\mathcal{R}_1$ could be replaced by another subset of low-rank tensors $\mathcal{M}$ containing $\mathcal{R}_1$.

Convergence is guaranteed under quite general assumptions on $\mathcal{J}$ (strongly convex, differentiable with Lipschitz differential) and the set $\mathcal{M}$ ($\mathcal{M}$ closed, $\operatorname{span} \mathcal{M} = V$).

Greedy algorithms with a dictionary $\mathcal{R}_1$ of rank-one tensors often present a slow convergence compared to the ideal performance of n-term approximations

$$\inf_{v \in \mathcal{R}_n} \mathcal{J}(v).$$

Also, these algorithms do not really exploit the structure of tensors.

1. Higher-order singular value decomposition and tensor truncation
2. Direct optimization in subsets of low-rank tensors
3. Iterative solvers with tensor truncation
4. Greedy algorithms
   - Greedy algorithms for canonical format
   - Greedy algorithms for Tucker format
   - Partially greedy algorithms for Tucker format

Approximation in Tucker format: a subspace point of view

The set $\mathcal{T}_r$ of tensors with Tucker rank bounded by $r = (r_1, \dots, r_d)$ is defined by

$$\mathcal{T}_r = \left\{ v = \sum_{1 \le k_1 \le r_1} \dots \sum_{1 \le k_d \le r_d} C_{k_1,\dots,k_d} v^{(1)}_{k_1} \otimes \dots \otimes v^{(d)}_{k_d} : C \in \mathbb{R}^{r_1 \times \dots \times r_d}, v^{(\nu)}_{k_\nu} \in V_\nu \right\}.$$

It can be equivalenlty parametrized by subspaces

$$\mathcal{T}_r = \{ v : v \in U_1 \otimes \dots \otimes U_d \text{ with } U_\nu \subset V_\nu, \dim(U_\nu) = r_\nu \}.$$

Then, an optimization problem on $\mathcal{T}_r$ can be interpreted as a problem of finding optimal low-dimensional spaces:

$$\min_{v \in \mathcal{T}_r} \mathcal{J}(v) = \min_{\dim(U_1) = r_1} \dots \min_{\dim(U_d) = r_d} \min_{v \in U_1 \otimes \dots \otimes U_d} \mathcal{J}(v).$$

This is a multilinear version of projection-based model order-reduction methods, where an approximation is searched in a tensor product $U_1^{r_1} \otimes \dots \otimes U_d^{r_d}$ of optimal subspaces $U_\nu^{r_\nu}$ of dimension r_ν.

Greedy algorithms for approximation in Tucker format

Greedy algorithms with a subspace point of view, which are similar to greedy algorithms for reduced basis methods, can be introduced for the construction of approximations u^n in an increasing sequence of tensor subspaces

$$U_1^n \otimes \ldots \otimes U_d^n, \quad n \geq 1,$$

with

$$U_\nu^1 \subset \ldots \subset U_\nu^n \subset \ldots, \quad 1 \leq \nu \leq d.$$

Greedy algorithms for approximation in Tucker format

At step n of these algorithms, we have an approximation u^{n-1} and associated subspaces U_ν^{n-1} of dimension r_ν^{n-1}, $1 \leq \nu \leq d$.

Assume that we have selected a set of dimensions $D_n \subset \{1, \dots, d\}$ to be enriched ($D_n = \{1, \dots, d\}$ for an isotropic enrichment).

For $\nu \notin D_n$, we let $U_\nu^n = U_\nu^{n-1}$, and for $\nu \in D_n$ we construct new spaces U_ν^n with dimension $r_\nu^n = r_\nu^{n-1} + 1$ and such that $U_\nu^n \supset U_\nu^{n-1}$.

An optimal greedy algorithm would consist in solving

$$\mathcal{J}(u^n) = \min_{\substack{\dim(U_\nu^n) = r_\nu^n \\ U_\nu^n \supset U_\nu^{n-1} \\ \nu \in D_n}} \min_{v \in U_1^n \otimes \dots \otimes U_d^n} \mathcal{J}(v)$$

Greedy algorithms for approximation in Tucker format

A practical greedy algorithm consists in computing an optimal rank-one correction of u^{n-1}

$$\mathcal{J}(u^{n-1} + w_n^{(1)} \otimes \ldots \otimes w_n^{(d)}) = \min_{w \in \mathcal{R}_1} \mathcal{J}(u^{n-1} + w),$$

in enriching the spaces according to

$$U_\nu^n = U_\nu^{n-1} + \text{span}(w_n^{(\nu)}), \quad \nu \in D_n,$$

and finally in computing the best approximation u^n in the tensor space $U_1^n \otimes \ldots \otimes U_d^n$ by solving

$$\mathcal{J}(u^n) = \min_{v \in U_1^n \otimes \ldots \otimes U_d^n} \mathcal{J}(v)$$

or

$$\min_{C \in \mathbb{R}^{r_1^n \times \ldots \times r_d^n}} \mathcal{J}\Big(\sum_{1 \le k_1 \le r_1^n} \ldots \sum_{1 \le k_d \le r_d^n} C_k v_{k_1}^{(1)} \otimes \ldots \otimes v_{k_d}^{(d)} \Big) \tag{6}$$

where $\{v_i^{(\nu)}\}_{i=1}^{r_\nu^n}$ is a basis of U_ν^n.

For high-dimensional problems, the practical solution of (6) requires a structured approximation of the tensor C, e.g. using sparse or low-rank formats. Note that if we add the constraint of having a super-diagonal tensor C, we recover a standard greedy algorithm for approximation in canonical format.

1. Higher-order singular value decomposition and tensor truncation
2. Direct optimization in subsets of low-rank tensors
3. Iterative solvers with tensor truncation
4. Greedy algorithms
   - Greedy algorithms for canonical format
   - Greedy algorithms for Tucker format
   - Partially greedy algorithms for Tucker format

Partially greedy algorithms for Tucker format

For order-two tensors in $V_1 \otimes V_2$, greedy algorithms for Tucker format construct a sequence of spaces

$$U^n = U_1^n \otimes U_2^n,$$

with a greedy enrichment of both left and right spaces, and a corresponding sequence of rank-n approximations u^n with

$$\mathcal{J}(u^n) = \min_{v \in U_1^n \otimes U_2^n} \mathcal{J}(v) = \min_{C \in \mathbb{R}^{n \times n}} \mathcal{J}(\sum_{i,j=1}^{n} v_i^{(1)} \otimes v_j^{(2)} C_{i,j})$$

A partially greedy strategy consists in constructing a sequence of spaces

$$U^n = U_1^n \otimes V_2,$$

where only the left spaces are constructed in a greedy fashion.

Partially greedy algorithms for Tucker format

At step n, a suboptimal algorithm consists in computing a rank-one correction of u^{n-1}

$$\mathcal{J}(u^{n-1} + w_n^{(1)} \otimes w_n^{(2)}) = \min_{w^{(1)}, w^{(2)}} \mathcal{J}(u^{n-1} + w^{(1)} \otimes w^{(2)}),$$

in enriching the left subspace according to

$$U_1^n = U_1^{n-1} + \mathrm{span}(w_n^{(1)}),$$

and then in computing an approximation u^n in $U_1^n \otimes V_2$ by solving

$$\mathcal{J}(u^n) = \min_{v \in U_1^n \otimes V_2} \mathcal{J}(v) = \min_{v_1^{(2)}, \ldots, v_n^{(2)}} \mathcal{J}\Big(\sum_{i=1}^{n} v_i^{(1)} \otimes v_i^{(2)}\Big)$$

where $\{v_i^{(1)}\}_{i=1}^n$ is a basis of U_1^n.

Other topics

- Approximation power of low-rank formats
- Interpolation methods for low-rank approximation
- Geometry of low-rank formats and its consequences in model order reduction of dynamical systems and optimization.
- Selection of a tensor format
- Exploiting sparsity in tensor representations
- Higher-order tensor methods for low-dimensional problems : quantization
- ...