

ZERO ANNOUNCEMENT

**3rd International Conference
on Structured Matrices and Tensors
Hong Kong, HKBU and PolyU
19-22 January 2010**

CONTACT E-MAIL: mng@math.hkbu.edu.hk (Michael Ng)

SUBLINEAR ALGEBRA FOR TENSOR APPROXIMATIONS OF VECTORS AND MATRICES

Eugene Tyrtysnikov

`tee@inm.ras.ru`

Institute of Numerical Mathematics
Russian Academy of Sciences

(joint work with I. Oseledets and D. Savostyanov)



OVERVIEW OF THE TALK

- Huge-scale data calls for *sublinear* complexity
- Tensor background
 - Tucker decomposition
 - Canonical decomposition
- General rank reduction methods
 - Tensor SVD (Higher Order SVD)
 - Alternating Least Squares (ALS)
- Recompression methods in matrix-by-matrix multiplication
 - Pre-recompression stage
 - Tucker-Tucker recompression algorithms
- Matrix inversion with sublinear complexity
 - Newton–Schultz and its modification
 - General theory for approximate iterations
 - Examples: 3D Laplacian and Newton potential
 - Inversion of a two-level Toeplitz matrix
 - Tensor eigensolver perspectives
- Concluding remarks and future work

WHAT IS A HUGE-SCALE PROBLEM?

Solve

$$\int_D \frac{1}{|x - y|} \varphi(y) dy = f(x), \quad x, y \in D = [0, 1]^d$$

Let $d = 3$. Subdivide the cube D into subcubes D_{ijk} :

$$D_{ijk} = [a_{i-1}, a_i] \times [a_{j-1}, a_j] \times [a_{k-1}, a_k], \quad 0 = a_0 < a_1 < \dots < a_n = 1$$

$$\varphi(y) \approx u_{ijk} = \text{const} \quad \text{на } D_{ijk} \quad (\text{collocation}) \quad \Rightarrow \quad Au = f$$

Vectors = discrete functions u_{ijk}, f_{ijk} on a grid with n^3 nodes.

Matrix A contains n^6 nonzero entries, none of them can be neglected, no visible structure in the case of nonuniform grids.

$n = 64 \quad \Rightarrow \quad \text{STORAGE FOR } A = 512\text{Gb} \text{ — not few!}$

$n = 256 \quad \Rightarrow \quad \text{STORAGE FOR } A = 2\text{Pb} \text{ (1 Pb} = 2^{50} \text{ byte)}.$

A big problem is already with storage for matrix coefficients!

CAN WE STILL SOLVE IT?

The only idea is to find a sufficiently close problem with a flaunting low-parametric structure.

Store only **few parameters** of this structure.

Apply **gain-of-the-structure** methods.

USE SMOOTHNESS IN DATA!

SMOOTHNESS = RANK STRUCTURES =
TENSOR STRUCTURES

WHAT IS A TENSOR PROBLEM?

It is one where all data on input and output are given exactly or approximately in tensor formats defined by a small number of parameters compared to the total amount of data.

For such problems we propose to seek for algorithms that work with data exclusively in tensor formats, the price we pay is a contamination of data through recompression (approximation) at each operation.

WHAT ARE TENSORS?

TENSOR = MULTI-INDEX ARRAY = MULTI-WAY ARRAY =
MULTI-DIMENSIONAL MATRIX:

$$\mathbf{A} = [a_{ij\dots k}]$$

$$i \in I, \quad j \in J, \quad \dots, \quad k \in K$$

Number of different indices is *dimension*.

Indices are called also *modes*.

Cardinalities of index ranges $\mathbf{I}, \mathbf{J}, \dots, \mathbf{K}$ are *mode sizes*.

In case of dimension \mathbf{d} and mode sizes $\mathbf{n}_1, \mathbf{n}_2, \dots, \mathbf{n}_d$,
 \mathbf{A} is a *tensor of size* $\mathbf{n}_1 \times \mathbf{n}_2 \times \dots \times \mathbf{n}_d$.

Talking of tensors, tacitly assume that $\mathbf{d} \geq 3$.

TENSORS AND MATRICES

Let $\mathbf{A} = [a_{ijklm}]$.

Consider pairs of complementary *long indices*

(ij) and (klm)

(kl) and (ijm)

.....

Then \mathbf{A} gives rise to several matrices:

$$\mathbf{B}_1 = [b_{(ij),(klm)}],$$

$$\mathbf{B}_2 = [b_{(kl),(ijm)}]$$

.....

with

$$b_{(ij),(klm)} = b_{(kl),(ijm)} = \dots = a_{ijklm}$$

MODE UNFOLDING MATRICES

$$\mathbf{A}_1 = [\mathbf{a}_{i,(jklm)}]$$

$$\mathbf{A}_2 = [\mathbf{a}_{j,(iklm)}]$$

$$\mathbf{A}_3 = [\mathbf{a}_{k,(ijlm)}]$$

$$\mathbf{A}_4 = [\mathbf{a}_{l,(ijkm)}]$$

$$\mathbf{A}_5 = [\mathbf{a}_{m,(ijkl)}]$$

Columns of unfolding matrices are called *mode vectors*.

If $\mathbf{d} = \mathbf{3}$, typical names are *columns, rows, fibers*.

Ranks of unfolding matrices are called *mode ranks* or *Tucker ranks*.

L. R. Tucker, Some mathematical notes on three-mode factor analysis,
Psychometrika, V. 31, P. 279–311 (1966).

TENSOR-BY-MATRIX MULTIPLICATIONS

Also called *mode contractions*.

Given a tensor $\mathbf{A} = [\mathbf{a}_{ijk}]$ and matrices

$$\mathbf{U} = [\mathbf{u}_{i'i}], \quad \mathbf{V} = [\mathbf{v}_{j'j}], \quad \mathbf{W} = [\mathbf{w}_{k'k}],$$

define new tensors

$$\begin{aligned} \mathbf{A}^U &= \mathbf{A} \times_1 \mathbf{U} = [\mathbf{a}_{i'jk}^U] \\ \mathbf{A}^V &= \mathbf{A} \times_2 \mathbf{V} = [\mathbf{a}_{ij'k}^V] \\ \mathbf{A}^W &= \mathbf{A} \times_3 \mathbf{W} = [\mathbf{a}_{ijk'}^W] \end{aligned}$$

as follows:

$$\begin{aligned} \mathbf{a}_{i'jk}^U &= \sum_i \mathbf{u}_{i'i} \mathbf{a}_{ijk} & \Leftrightarrow & \mathbf{A}_1^U = \mathbf{U} \mathbf{A}_1 \\ \mathbf{a}_{ij'k}^V &= \sum_j \mathbf{v}_{j'j} \mathbf{a}_{ijk} & \Leftrightarrow & \mathbf{A}_2^V = \mathbf{V} \mathbf{A}_2 \\ \mathbf{a}_{ijk'}^W &= \sum_k \mathbf{w}_{k'k} \mathbf{a}_{ijk} & \Leftrightarrow & \mathbf{A}_3^W = \mathbf{W} \mathbf{A}_3 \end{aligned}$$

WHY CONTRACTIONS?

Let $\mathbf{A} = [\mathbf{a}_{ijk}]$ be $n \times n \times n$ and mode ranks be equal to $r \ll n$.
Consider \mathbf{QR} decompositions of unfolding matrices

$$\mathbf{A}_1 = \mathbf{Q}_1 \mathbf{R}_1, \quad \mathbf{A}_2 = \mathbf{Q}_2 \mathbf{R}_2, \quad \mathbf{A}_3 = \mathbf{Q}_3 \mathbf{R}_2$$

$\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3$ are orthogonal $n \times r$ matrices.

Define the *Tucker core* tensor $\mathbf{G} = [\mathbf{g}_{\alpha\beta\gamma}]$
of *contracted* size $r \times r \times r$:

$$\mathbf{G} = \mathbf{A} \times_1 \mathbf{Q}_1^\top \times_2 \mathbf{Q}_2^\top \times_3 \mathbf{Q}_3^\top \quad \text{i.e.} \quad \mathbf{g}_{\alpha\beta\gamma} = \sum_{i,j,k} \mathbf{a}_{ijk} \mathbf{q}_{i\alpha}^1 \mathbf{q}_{j\beta}^2 \mathbf{q}_{k\gamma}^3$$

THEOREM

$$\mathbf{A} = \mathbf{G} \times_1 \mathbf{Q}_1 \times_2 \mathbf{Q}_2 \times_3 \mathbf{Q}_3 \quad \text{i.e.} \quad \mathbf{a}_{ijk} = \sum_{\alpha,\beta,\gamma} \mathbf{g}_{\alpha\beta\gamma} \mathbf{q}_{i\alpha}^1 \mathbf{q}_{j\beta}^2 \mathbf{q}_{k\gamma}^3$$

IMPORTANT: \mathbf{A} is now represented in a *contracted form*
with only $3nr + r^3 \ll n^3$ parameters.

TUCKER DECOMPOSITION

Regarded as *Tensor SVD* or *Higher Order SVD*:

$$\mathbf{A} = \mathbf{G} \times_1 \mathbf{Q}_1 \times_2 \mathbf{Q}_2 \times_3 \mathbf{Q}_3 \quad \text{i.e.} \quad a_{ijk} = \sum_{\alpha, \beta, \gamma} g_{\alpha\beta\gamma} q_{i\alpha}^1 q_{j\beta}^2 q_{k\gamma}^3$$

Orthogonal matrices $\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3$ are *Tucker factors* or *frame matrices*.

THEOREM

Rows in each of unfolding matrices for the Tucker core can be made *orthogonal* and arranged in *length-decreasing order*.

Row lengths of unfoldings for \mathbf{G} = singular values of unfoldings for \mathbf{A} .

PROOF is easy via SVD of unfolding matrices:

if $\mathbf{A}_1 = \mathbf{Q}_1 \mathbf{\Sigma}_1 \mathbf{V}_1$ then $(\mathbf{A} \times_1 \mathbf{Q}_1^\top)_1 = \mathbf{\Sigma}_1 \mathbf{V}_1$.

Same for other modes.

TUCKER APPROXIMATIONS

$$a_{ijk} \approx \sum_{\alpha, \beta, \gamma} g_{\alpha\beta\gamma} q_{i\alpha}^1 q_{j\beta}^2 q_{k\gamma}^3$$

APPLICATIONS:

- Multi-way Principal Component Analysis
(senior frame matrices are most informative).
- Tensor data compression
(ignore small and get to reduced Tucker ranks \ll mode sizes).
- New generation of numerical algorithms
with [all data in the Tucker format](#).
Enjoy linear and even sublinear complexity in total size of data
(could be petabytes).

[I. Oseledets, D. Savostyanov, E. Tyrtyshnikov,](#)
Linear algebra for tensor problems, submitted to *Computing* (2008).

G. Beylkin, M. Mohlenkamp, Algorithms for numerical analysis in high dimensions, *SIAM J. Sci. Comput.*, 26 (6), pp. 2133-2159 (2005).

CANONICAL DECOMPOSITION

$$a_{ij\dots k} = \sum_{t=1}^{\rho} u_{it} v_{jt\dots} w_{kt}$$

Minimal $\rho = \mathbf{tRank}$ is called *canonical rank* or *tensor rank* of \mathbf{A} .

THEOREM

Let mode ranks be equal to r . Then

$$r \leq \mathbf{tRank}(\mathbf{A}) \leq r^2.$$

CANONICAL APPROXIMATIONS

$$a_{ij\dots k} \approx \sum_{t=1}^{\rho} u_{it} v_{jt\dots} w_{kt}$$

play same compression role as Tucker.

Could be better but not necessarily!

TENSOR RANKS IN COMPLEXITY THEORY

In the “row-by-column” rule for multiplication of $n \times n$ matrices we have n^2 multiplications. Can we reduce this number?

$$\begin{bmatrix} c_1 & c_2 \\ c_3 & c_4 \end{bmatrix} = \begin{bmatrix} a_1 & a_2 \\ a_3 & a_4 \end{bmatrix} = \begin{bmatrix} b_1 & b_2 \\ b_3 & b_4 \end{bmatrix}$$

$$c_k = \sum_{i=1}^n \sum_{j=1}^n h_{ijk} a_i b_j$$

Let ρ = tensor rank of h_{ijk} and canonical decomposition read

$$h_{ijk} = \sum_{t=1}^{\rho} u_{it} v_{jt} w_{kt} \Rightarrow$$

$$c_k = \sum_{t=1}^{\rho} w_{kt} \left(\sum_{i=1}^4 u_{it} a_i \right) \left(\sum_{j=1}^n v_{jt} b_j \right)$$

Now we have ρ multiplications!

If $n = 2$ then $\rho = 7$ (Strassen, 1965).

By recursion \Rightarrow only $O(n^{\log_2 7})$ multiplications for arbitrary n .

TUCKER VS CANONICAL FOR MATRICES

$$a_{ij} = \sum_{\alpha=1}^r \sum_{\beta=1}^r g_{\alpha\beta} q_{i\alpha}^1 q_{j\beta}^2 \Leftrightarrow A = Q_1 G Q_2^\top$$

Tucker = a *pseudo-skeleton* decomposition of A .

$$a_{ij} = \sum_{t=1}^{\rho} u_{it} v_{jt} \Leftrightarrow A = UV^\top$$

Canonical = a *skeleton* or *dyadic* decomposition of A .

Tensor (canonical) rank seems to be a true generalization of the matrix rank concept.

However, tensor rank for dimension ≥ 3 and matrix rank have *noticably different properties*.

KRONECKER PRODUCT REPRESENTATION

Tucker decomposition:

$$\mathbf{A} = \sum_{\alpha, \beta, \gamma} g_{\alpha\beta\gamma} \mathbf{u}_{\alpha} \otimes \mathbf{v}_{\beta} \otimes \mathbf{w}_{\gamma}$$

Canonical decomposition:

$$\mathbf{A} = \sum_t \mathbf{u}_t \otimes \mathbf{v}_t \otimes \mathbf{w}_t$$

BASIC OPERATION

Given two matrices \mathbf{A} and \mathbf{X} , compute a tensor-structured approximation $\tilde{\mathbf{Y}}$ to their product

$$\tilde{\mathbf{Y}} \approx \mathbf{Y} = \mathbf{A}\mathbf{X}$$

so that \mathbf{Y} never appears as a full matrix.

Assume that \mathbf{A} is $\mathbf{N} \times \mathbf{N}$.

Important cases for \mathbf{X} :

- (a) \mathbf{X} is a vector (a rectangular matrix of size $\mathbf{N} \times \mathbf{1}$);
- (b) \mathbf{X} is a square matrix of the same order \mathbf{N} .

Case (a) is a basic operation in all iterative solvers for linear systems or eigenvalue problems with the coefficient matrix \mathbf{A} .

Case (b) is a workhorse operation in computation of matrix functions of \mathbf{A} , in particular \mathbf{A}^{-1} .

CANONICAL FORMAT

$$A = \sum_{t=1}^{\rho} A_t \otimes B_t \otimes C_t$$
$$X = \sum_{\tau=1}^{\rho} U_{\tau} \otimes V_{\tau} \otimes W_{\tau}$$

TUCKER FORMAT

$$A = \sum_{\sigma=1}^{r_1} \sum_{\delta=1}^{r_2} \sum_{\tau=1}^{r_3} f_{\sigma\delta\tau} A_{\sigma} \otimes B_{\delta} \otimes C_{\tau}$$
$$X = \sum_{\alpha=1}^{r_1} \sum_{\beta=1}^{r_2} \sum_{\gamma=1}^{r_3} g_{\alpha\beta\gamma} U_{\alpha} \otimes V_{\beta} \otimes W_{\gamma}$$

Thus, $N = n_1 n_2 \dots n_d$.

Proceed as if $d = 3$, $n = n_1 = \dots = n_d$, $r = r_1 = \dots = r_d$.

COMBINATIONS OF FORMATS

(CC) \mathbf{A} and \mathbf{X} are both in the canonical.

(CT) \mathbf{A} is in the canonical, \mathbf{X} is in the Tucker.

(TT) \mathbf{A} and \mathbf{X} are both in the Tucker.

The result $\tilde{\mathbf{Y}} \approx \mathbf{Y} = \mathbf{A}\mathbf{X}$ is assumed to keep the format of \mathbf{X} .

We focus only on \mathbf{CT} and \mathbf{TT} .

PRE-RECOMPRESSION STAGE

$$A = \sum_{\sigma, \delta, \tau} f_{\sigma\delta\tau} A_{\sigma} \otimes B_{\delta} \otimes C_{\tau}$$

$$X = \sum_{\alpha, \beta, \gamma} g_{\alpha\beta\gamma} U_{\alpha} \otimes V_{\beta} \otimes W_{\gamma}$$

$$\Rightarrow Y = \sum_{\sigma, \delta, \tau} \sum_{\alpha, \beta, \gamma} f_{\sigma\delta\tau} g_{\alpha\beta\gamma} A_{\sigma} U_{\alpha} \otimes B_{\delta} V_{\beta} \otimes C_{\tau} W_{\gamma}$$

Pre-recompression stage consists in the computation of matrices

$$A'_{\sigma\alpha} = A_{\sigma} U_{\alpha}, \quad B'_{\delta\beta} = B_{\delta} V_{\beta}, \quad C'_{\tau\gamma} = C_{\tau} W_{\gamma}.$$

Consequently,

$$Y = \sum_{\sigma, \delta, \tau} \sum_{\alpha, \beta, \gamma} f_{\sigma\delta\tau} g_{\alpha\beta\gamma} A'_{\sigma\alpha} \otimes B'_{\delta\beta} \otimes C'_{\tau\gamma}.$$

GENERAL RANK REDUCTION METHODS

Consider a general tensor $\mathbf{Z} = [z_{ijk}]$ with the mode sizes \mathbf{n} .

Tensor SVD (Higher Order SVD):

- Construct the unfolding matrices $\mathbf{Z}_1 = [z_{i,(jk)}]$, $\mathbf{Z}_2 = [z_{j,(ik)}]$, $\mathbf{Z}_3 = [z_{k,(ij)}]$.
- Perform rank revealing decompositions (via SVD)

$$\mathbf{Z}_1 = \mathbf{Q}_1 \mathbf{R}_1 + \mathbf{E}_1, \quad \mathbf{Z}_2 = \mathbf{Q}_2 \mathbf{R}_2 + \mathbf{E}_2, \quad \mathbf{Z}_3 = \mathbf{Q}_3 \mathbf{R}_3 + \mathbf{E}_3,$$

where

$$\mathbf{Q}_1 = [q_{i\alpha}^1], \quad \mathbf{Q}_2 = [q_{j\beta}^2], \quad \mathbf{Q}_3 = [q_{k\gamma}^3]$$

are orthogonal $\mathbf{n} \times \mathbf{r}$ matrices.

- Compute the Tucker core tensor

$$h_{\alpha\beta\gamma} = \sum_{i,j,k} q_{i\alpha}^1 q_{j\beta}^2 q_{k\gamma}^3 z_{ijk}.$$

- Finish with the Tucker approximation in the form

$$\mathbf{y}_{ijk} \approx \tilde{\mathbf{y}}_{ijk} = \sum_{\alpha\beta\gamma} h_{\alpha\beta\gamma} q_{i\alpha}^1 q_{j\beta}^2 q_{k\gamma}^3.$$

$$\text{COMPLEXITY} = O(n^4)$$

GENERAL RANK REDUCTION METHODS

Alternating Least Squares (ALS):

- Freeze Q_2, Q_3 and compute

$$h_{i\beta\gamma}^1 = \sum_{j,k} q_{j\beta}^2 q_{k\gamma}^3 z_{ijk}.$$

Consider a matrix $H_1 = [h_{i\beta\gamma}^1]$ of size $n \times r^2$ (here β, γ form a “long index” for columns) and find Q_1 from a rank revealing decomposition

$$H_1 = Q_1 R_1 + F_1$$

with minimal $\|F_1\|_F$. Note that Q_1 is a maximizer for $\|Q^\top H_1\|_F$ over all matrices Q with r orthonormal columns.

- Freeze Q_1, Q_3 and find the best fit Q_2 .
- Freeze Q_1, Q_2 and find the best fit Q_3 .
- Repeat until convergence, then compute the Tucker core for the obtained mode frame matrices.

$$\text{COMPLEXITY} = O(n^3 r + n^2 r^2 + n r^3)$$

GENERAL RECOMPRESSION ALGORITHM

Given $\mathbf{z}_{ijk} = \sum_{\alpha,\beta,\gamma} \mathbf{h}_{\alpha\beta\gamma} \mathbf{a}_{i\alpha} \mathbf{b}_{j\beta} \mathbf{c}_{k\gamma}$ with the mode sizes \mathbf{n} and mode ranks \mathbf{r}_0 , find its approximation of lesser rank with $\mathbf{r} < \mathbf{r}_0$.

- Compute orthogonal $\mathbf{n} \times \mathbf{r}_0$ matrices $\mathbf{Q}_1 = [\mathbf{q}_{i\alpha'}^1]$, $\mathbf{Q}_2 = [\mathbf{q}_{j\beta'}^2]$, $\mathbf{Q}_3 = [\mathbf{q}_{k\gamma'}^3]$ s.t.
 $[\mathbf{a}_{i\alpha}] = \mathbf{Q}_1 \mathbf{R}_1$, $[\mathbf{b}_{j\beta}] = \mathbf{Q}_2 \mathbf{R}_2$, $[\mathbf{c}_{k\gamma}] = \mathbf{Q}_3 \mathbf{R}_3$.
- Find an auxiliary $\mathbf{r}_0 \times \mathbf{r}_0 \times \mathbf{r}_0$ tensor $\mathbf{h}'_{\alpha'\beta'\gamma'} = \sum_{\alpha,\beta,\gamma} \mathbf{r}_{\alpha'\alpha}^1 \mathbf{r}_{\beta'\beta}^2 \mathbf{r}_{\gamma'\gamma}^3 \mathbf{h}_{\alpha\beta\gamma}$ via 3 steps:

$$\mathbf{h}_{\alpha'\beta\gamma}^* = \sum_{\alpha} \mathbf{r}_{\alpha'\alpha}^1 \mathbf{h}_{\alpha\beta\gamma}, \quad \mathbf{h}_{\alpha'\beta'\gamma}^{**} = \sum_{\beta} \mathbf{r}_{\beta'\beta}^2 \mathbf{h}_{\alpha'\beta\gamma}^*, \quad \mathbf{h}'_{\alpha'\beta'\gamma'} = \sum_{\gamma} \mathbf{r}_{\gamma'\gamma}^3 \mathbf{h}_{\alpha'\beta'\gamma}^{**}.$$

- Reduce its mode ranks: $\mathbf{h}'_{\alpha'\beta'\gamma'} \approx \sum_{\sigma=1}^r \sum_{\delta=1}^r \sum_{\tau=1}^r \mathbf{p}_{\alpha'\sigma}^1 \mathbf{p}_{\beta'\delta}^2 \mathbf{p}_{\gamma'\tau}^3 \tilde{\mathbf{h}}_{\sigma\delta\tau}$.

- Finally, $\mathbf{z}_{ijk} \approx \sum_{\sigma=1}^r \sum_{\delta=1}^r \sum_{\tau=1}^r \tilde{\mathbf{q}}_{i\sigma}^1 \tilde{\mathbf{q}}_{j\delta}^2 \tilde{\mathbf{q}}_{k\tau}^3 \tilde{\mathbf{h}}_{\sigma\delta\tau}$ with the Tucker factors

$$\tilde{\mathbf{q}}_{i\sigma}^1 = \sum_{\alpha'=1}^{\mathbf{r}_0} \mathbf{q}_{i\alpha'}^1 \mathbf{p}_{\alpha'\sigma}^1, \quad \tilde{\mathbf{q}}_{j\delta}^2 = \sum_{\beta'=1}^{\mathbf{r}_0} \mathbf{q}_{j\beta'}^2 \mathbf{p}_{\beta'\delta}^2, \quad \tilde{\mathbf{q}}_{k\tau}^3 = \sum_{\gamma'=1}^{\mathbf{r}_0} \mathbf{q}_{k\gamma'}^3 \mathbf{p}_{\gamma'\tau}^3.$$

COMPLEXITY = $O(\mathbf{n} \mathbf{r}_0^2 + \mathbf{n} \mathbf{r}_0 \mathbf{r} + \mathbf{r}_0^4)$ (possibly $\mathbf{r}_0^3 \mathbf{r}$ instead of \mathbf{r}_0^4)

TUCKER-TUCKER RECOMPRESSION ALGORITHM

- Find orthogonal $n \times r^2$ matrices $Q_1 = [q_{i,(\sigma'\alpha')}^1]$, $Q_2 = [q_{j,(\delta'\beta')}^2]$, $Q_3 = [q_{k,(\tau'\gamma')}^3]$ s.t.

$$[a_{i,(\sigma\alpha)}] = Q_1 R_1, \quad [b_{j,(\delta\beta)}] = Q_2 R_2, \quad [c_{k,(\tau\gamma)}] = Q_3 R_3.$$

- Define the auxiliary core tensor by

$$h_{\sigma'\alpha'\delta'\beta'\tau'\gamma'} = \sum_{\sigma,\delta,\tau} \sum_{\alpha,\beta,\gamma} r_{\sigma'\alpha'\sigma\alpha}^1 r_{\delta'\beta'\delta\beta}^2 r_{\tau'\gamma'\tau\gamma}^3 f_{\sigma\delta\tau} g_{\alpha\beta\gamma}$$

and compute it through the following prescriptions:

$$\begin{aligned} h'_{\sigma'\alpha'\delta\beta\tau\gamma} &= \sum_{\sigma,\alpha} r_{\sigma'\alpha'\sigma\alpha}^2 f_{\sigma\delta\tau} g_{\alpha\beta\gamma}, \\ h''_{\sigma'\alpha'\delta'\beta'\tau\gamma} &= \sum_{\delta,\beta} r_{\delta'\beta'\delta\beta}^2 h'_{\sigma'\alpha'\delta\beta\tau\gamma}, \\ h_{\sigma'\alpha'\delta'\beta'\tau'\gamma'} &= \sum_{\tau,\gamma} r_{\tau'\gamma'\tau\gamma}^3 h''_{\sigma'\alpha'\delta'\beta'\tau\gamma}. \end{aligned}$$

- Apply a mode rank reduction algorithm to the auxiliary tensor.
- Recompute the Tucker factors in the final Tucker approximation with mode ranks r .

$$\text{COMPLEXITY} = O(nr^4 + r^8) \quad \text{MEMORY} = O(nr^2 + r^6)$$

TUCKER-TUCKER ALS RECOMPRESSION ALGORITHM

- Freeze \mathbf{Q}_2 , \mathbf{Q}_3 and find the best fit for \mathbf{Q}_1 : compute

$$\mathbf{v}_{\beta'\delta\beta} = \sum_j \mathbf{q}_{j\beta'}^2 \mathbf{b}_{j\delta\beta}, \quad \mathbf{w}_{\gamma'\tau\gamma} = \sum_k \mathbf{q}_{k\gamma'}^3 \mathbf{c}_{k\tau\gamma},$$

then acquire

$$\mathbf{h}_{i\beta'\gamma'} = \sum_{\sigma,\delta,\tau} \sum_{\alpha,\beta,\gamma} \mathbf{f}_{\sigma\delta\tau} \mathbf{g}_{\alpha\beta\gamma} \mathbf{v}_{\beta'\delta\beta} \mathbf{w}_{\gamma'\tau\gamma} \mathbf{a}_{i\sigma\alpha}$$

via the following prescriptions:

$$\begin{aligned} \mathbf{u}'_{\alpha\beta\gamma'\tau} &= \sum_{\gamma} \mathbf{g}_{\alpha\beta\gamma} \mathbf{w}_{\gamma'\tau\gamma}, & \mathbf{u}''_{\alpha\beta\gamma'\sigma\delta} &= \sum_{\tau} \mathbf{u}'_{\alpha\beta\gamma'\tau} \mathbf{f}_{\sigma\delta\tau}, \\ \mathbf{u}_{\alpha\gamma'\sigma\beta'} &= \sum_{\delta,\beta} \mathbf{u}''_{\alpha\beta\gamma'\sigma\delta} \mathbf{v}_{\beta'\delta\beta}, & \mathbf{h}_{i\beta'\gamma'} &= \sum_{\sigma,\alpha} \mathbf{u}_{\alpha\gamma'\sigma\beta'} \mathbf{a}_{i\sigma\alpha}. \end{aligned}$$

Obtain \mathbf{Q}_1 from a rank revealing decomposition of the matrix $\mathbf{H}_1 = [\mathbf{h}_{i,(\beta'\gamma')}]$ of size $n \times r^2$.

- Similarly, freeze \mathbf{Q}_1 , \mathbf{Q}_3 and find the best fit for \mathbf{Q}_2 , then freeze \mathbf{Q}_1 , \mathbf{Q}_2 and find the best fit for \mathbf{Q}_3 .
- Repeat until convergence.

$$\text{COMPLEXITY} = O(nr^4 + r^6), \quad \text{MEMORY} = O(nr^2 + r^5).$$

BEWARE OF A POSSIBLE PITFALL

One may want to circumvent the computation of a full $\mathbf{r}^2 \times \mathbf{r}^2 \times \mathbf{r}^2$ core and try to compress first the factor matrices via SVD (instead of QR).

This may not work!

EXAMPLE:

$$\mathbf{A} = \mathbf{U}\mathbf{V}^\top = (\mathbf{a} \ \varepsilon^2 \mathbf{b}) \begin{pmatrix} \mathbf{c}^\top \\ \varepsilon^{-1} \mathbf{d}^\top \end{pmatrix} = \mathbf{a}\mathbf{c}^\top + \varepsilon \mathbf{b}\mathbf{d}^\top$$

$\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}$ are unit-length vectors of size \mathbf{n} and $(\mathbf{a}, \mathbf{b}) = (\mathbf{c}, \mathbf{d}) = 0$

If we “compress” \mathbf{U} and \mathbf{V} separately, then the senior singular vectors of \mathbf{U} and \mathbf{V} would be \mathbf{a} and \mathbf{d} . Hence,

$$\mathbf{A} \approx \gamma \mathbf{a}\mathbf{d}^\top,$$

which leaves us with no hope.

Thus, we have to compute the full core and treat all the factors *simultaneously*, not separately.

MATRIX INVERSION ALGORITHMS

Newton–Schultz method

$$\mathbf{X}_{k+1} = 2\mathbf{X}_k - \mathbf{X}_k \mathbf{A} \mathbf{X}_k, \quad k = 0, 1, \dots$$

converges quadratically (in exact arithmetics) if $||\mathbf{I} - \mathbf{A}\mathbf{X}_0|| < 1$.

Let \mathbf{A} be given in the Tucker format and the same structure be maintained for all iterates \mathbf{X}_k :

$$\mathbf{Z}_{k+1} = 2\mathbf{X}_k - \mathbf{X}_k \mathbf{A} \mathbf{X}_k, \quad \mathbf{X}_{k+1} = \mathcal{P}(\mathbf{Z}_{k+1}).$$

\mathcal{P} is a (nonlinear) projector onto the manifold of matrices in the Tucker format with mode ranks $\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3$ selected and fixed before the method starts.

Truncated iteration converge still quadratically if

$$\mathcal{P}(\mathbf{A}^{-1}) = \mathbf{A}^{-1} + \mathbf{E}, \quad ||\mathbf{E}|| \leq \epsilon$$

until $||\mathbf{X}_k - \mathbf{A}^{-1}|| > c\epsilon$ for some $c > 1$.

GENERAL THEORY OF APPROXIMATE ITERATIONS

V a normed space

$B \in V$ the target of computaion

ITERATIVE PROCESS: $X_k = \Phi_k(X_{k-1})$

LEMMA.

Assume $\exists \alpha > 1, \varepsilon_\Phi, c_\Phi$ s.t.

$$\|X - B\| \leq \varepsilon_\Phi \Rightarrow$$

$$\|\Phi_k(X) - B\| \leq c_\Phi \|X - B\|^\alpha.$$

Then

$$\|X_0 - B\| < \varepsilon \Rightarrow$$

$$\|X_k - B\| \leq c^{-1} (c \|X_0 - B\|)^{\alpha^k}, \quad k = 1, 2 \dots$$

$$\varepsilon = \min(\varepsilon_\Phi, c^{-1}), \quad c = c_\Phi^{\frac{1}{\alpha-1}}$$

$\mathcal{S} \subset \mathcal{V}$ a subset of “structured” elements (e.g. structured matrices)

$R : \mathcal{V} \rightarrow \mathcal{S}$ a truncation operator.

$$X \in \mathcal{S} \quad \Rightarrow \quad R(X) = X.$$

TRUNCATED ITERATIVE PROCESS:

$$Y_0 = R(X_0)$$

$$Y_k = R(\Phi_k(Y_{k-1}))$$

THEOREM 1.

Assume that

- (1) Premises of Lemma are fulfilled.
- (2) $R(B) = B$.
- (3) $\|X - B\| \leq \varepsilon_\Phi \Rightarrow \|X - R(X)\| \leq c_R \|X - B\|$.

Then $\exists \delta > 0$ s.t.

$$Y_0 = R(Y_0), \quad \|Y_0 - B\| < \delta \Rightarrow$$

$$\|Y_k - B\| \leq c_{R\Phi} \|Y_{k-1} - B\|^\alpha, \quad k = 1, 2, \dots$$

$$c_{R\Phi} = (c_R + 1)c_\Phi$$

W. Hackbusch, B.N. Khoromskij, E. Tyrtysnikov. *Approximate iteration for structured matrices*. Preprint no. 112, Max-Planck-Institut für Mathematik in den Naturwissenschaften, Leipzig 2005. Numer. Math., DOI 10.1007/s00211-008-0143-0, 2008.

THEOREM 2.

Assume $\exists \quad \varepsilon_\Phi, c_R, \varepsilon_{RB}$ s.t

$$\begin{aligned} ||X - B|| &\leq \varepsilon_\Phi \Rightarrow \\ ||X - R(X)|| &\leq c_R ||X - B|| + \varepsilon_{RB}. \end{aligned}$$

Let m be the minimal k s.t.

$$e_{k-1}^\alpha \leq \frac{\varepsilon_{RB}}{c_{R\Phi}}, \quad c_{R\Phi} = (c_R + 1)c_\Phi.$$

Then errors $e_k = ||Y_k - B||$ of trun. iter. decrease superlinearly until $k \leq m$:

$$k \leq m - 1 \Rightarrow e_k \leq 2c_{R\Phi} e_{k-1}^\alpha,$$

$$k \geq m \Rightarrow e_m \leq 2\varepsilon_{RB}.$$

PROOF. $Z_k := \Phi_k(Y_{k-1})$

$$||Y_k - B|| \leq ||Y_k - Z_k|| + ||Z_k - B|| \leq (c_R + 1)||Z_k - B|| + \varepsilon_{RB} \Rightarrow$$

$$e_k \leq c_{R\Phi} e_{k-1}^\alpha + \varepsilon_{RB} \leq 2c_{R\Phi} e_{k-1}^\alpha.$$

MODIFIED NEWTON-SCHULTZ

$$\begin{aligned} X_{k+1} &= X_k(2I - AX_k) = X_k(2I - Y_k) \text{ with } Y_k = AX_k \Rightarrow \\ AX_{k+1} &= AX_k(2I - Y_k) \Rightarrow \end{aligned}$$

$$Y_{k+1} = Y_k(2I - Y_k), \quad X_{k+1} = X_k(2I - Y_k)$$

X_0 is an initial guess for A^{-1} , $Y_0 = AX_0$.

TRUNCATED VERSION:

$$H_k = \mathcal{P}(2I - Y_k), \quad Y_{k+1} = \mathcal{P}_1(Y_k H_k), \quad X_{k+1} = \mathcal{P}_2(X_k H_k)$$

Projectors $\mathcal{P}_1, \mathcal{P}_2$ ought to maintain the required accuracy and \mathcal{P} can be a way less accurate.

If X_0 is close enough to A^{-1} then Y_k and H_k are close to I , a “perfect structure” matrix. Experiments confirm that the modified Newton method is faster and more accurate than the standard Newton method.

INVERSION OF A 3D LAPLACIAN

$$\mathbf{A} = \Delta \otimes \mathbf{I} \otimes \mathbf{I} + \mathbf{I} \otimes \Delta \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{I} \otimes \Delta$$

$$\Delta = \text{tridiag}(-1, 2, -1)$$

$N = n^3$	32^3	64^3	128^3	256^3
Time	9 sec	11 sec	24 sec	227 sec
$\ \mathbf{A}\mathbf{X} - \mathbf{I}\ _F / \ \mathbf{I}\ _F$	$3 \cdot 10^{-6}$	$4 \cdot 10^{-6}$	$2 \cdot 10^{-5}$	10^{-4}

For the Laplacian, projectors are chosen as follows:

$$\mathcal{P} = \mathcal{P}_{(2,2,2)}, \quad \mathcal{P}_1 = \mathcal{P}_2 = \mathcal{P}_{(12,12,12)}.$$

$\mathcal{P}_{(r_1 r_2 r_3)}$ is a projector onto tensors with mode ranks $\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3$.

INVERSION OF A NEWTON-POTENTIAL MATRIX

Mode ranks for the Newton potential matrix, $\varepsilon = 10^{-5}$

$N = n^3$	32^3	64^3	128^3	256^3
r	10	12	13	15

Timings for the Newton potential matrix

$N = n^3$	32^3	64^3	128^3	256^3
Time	60 sec	107 sec	360 sec	1574 sec
$ \mathbf{A}\mathbf{X} - \mathbf{I} _F / \mathbf{I} _F$	10^{-2}	$9 \cdot 10^{-3}$	$5 \cdot 10^{-2}$	$4 \cdot 10^{-2}$

For the Newton potential matrix, to cause the method to converge we had to select

$$\mathcal{P} = \mathcal{P}_1 = \mathcal{P}_2 = \mathcal{P}_{(10,10,10)}.$$

SUBLINEAR COMPLEXITY FOR 2D TOEPLITZ INVERSION

Consider a 5-point Laplacian of order $N = n^2$ (2-level Toeplitz matrix).
The inverse matrix is approximated by the Newton–Schultz method

$$X_{k+1} = \text{APPROXIMATION}(2X_k - X_k A X_k)$$

with a rank-structured approximation of all computed matrices: by matrices of limited tensor rank and limited displacement rank of each block.

n	64^2	128^2	256^2	512^2
Tensor rank A^{-1}	9	10	11	12
Averaged displacement rank of A^{-1}	13.5	13.5	16.8	18.6

Inversion of the 5-point Laplacian

Time behaves as $\mathcal{O}(\sqrt{N}r_{\text{mean}}^2)$,
 $r_{\text{mean}} = \text{averaged displacement rank}$.

V.Olshevsky, I.Oseledets, E.Tyrtyshnikov,

Superfast inversion of two-level Toeplitz matrices using Newton iteration and tensor-displacement structure, *Operator Theory Advances and Applications*, vol. 179, pp. 229–240 (2007).

Low-parametric representations of inverse matrices contain $\mathcal{O}(N)$ parameters.

Hence, all difficulties are relegated to the representation of vectors, not matrices!

TENSOR STRUCTURE IN VECTORS

$$x = \begin{bmatrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \\ 9 \end{bmatrix} \Leftrightarrow X = \begin{bmatrix} 1 & 4 & 7 \\ 2 & 5 & 8 \\ 3 & 6 & 9 \end{bmatrix}$$

$$X = \text{MATRIX}(x) \quad \Leftrightarrow \quad x = \text{VECTOR}(X)$$

$$X = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \begin{bmatrix} 0 & 1 & 2 \end{bmatrix} = u_1 v_1^\top + u_2 v_2^\top$$

$$x = v_1 \otimes u_1 + v_2 \otimes u_2$$

EIGENVECTOR STRUCTURE

$$M = A \otimes B + C \otimes D$$

If A , C and B , D are two pairs of commuting matrices, then any eigenvector has a tensor rank-1 structure.

DISCRETE LAPLACIAN CASE

$$Mu = \lambda u, \quad M = A \otimes I + I \otimes A$$

$$A = \begin{bmatrix} 2 & -1 & & & \\ -1 & 2 & -1 & & \\ & \dots & \dots & \dots & \\ & & -1 & 2 & -1 \\ & & & -1 & 2 \end{bmatrix}$$

$$x_{kl} = u^k \otimes v^l \quad u_s^k = \sin \frac{\pi ks}{n+1}, \quad v_t^l = \sin \frac{\pi lt}{n+1}$$

$$\lambda_{kl} = 4 \sin^2 \frac{\pi k}{2(n+1)} + 4 \sin^2 \frac{\pi l}{2(n+1)}, \quad 1 \leq k, l \leq n$$

USE TENSOR VECTORS IN EIGENSOLVERS

LANCZOS:

- Choose an initial vector \mathbf{p}_1 with $||\mathbf{p}_1|| = 1$ and set $\mathbf{p}_0 = \mathbf{0}$, $b_0 = 0$.
- For $k = 1, 2, \dots$ compute

$$\mathbf{z}_k = M\mathbf{p}_k$$

$$a_k = (\mathbf{z}_k, \mathbf{p}_k)$$

$$\mathbf{q}_k = \mathbf{z}_k - a_k\mathbf{p}_k - b_{k-1}\mathbf{p}_{k-1}$$

$$b_k = ||\mathbf{q}_k||$$

$$\mathbf{p}_{k+1} = \mathbf{q}_k/b_k$$

- Compute the Rietz values as the eigenvalues of a projection $k \times k$ matrix (a symmetric tridiagonal matrix consisting of the values a_k, b_k)

$$M_k = P_k^\top M P_k, \quad P_k = [\mathbf{p}_1, \dots, \mathbf{p}_k].$$

TENSOR LANCZOS

- Choose an initial vector \mathbf{p}_1 with $||\mathbf{p}_1|| = 1$ and set $\mathbf{p}_0 = \mathbf{0}$, $b_0 = 0$.
- For $k = 1, 2, \dots$ compute

$$\mathbf{z}_k = M\mathbf{p}_k$$

$$a_k = (\mathbf{z}_k, \mathbf{p}_k)$$

$$\mathbf{q}_k = T_\varepsilon(\mathbf{z}_k - a_k\mathbf{p}_k - b_{k-1}\mathbf{p}_{k-1})$$

$$b_k = ||\mathbf{q}_k||$$

$$\mathbf{p}_{k+1} = \mathbf{q}_k/b_k$$

- Compute the Rietz values using $k \times k$ projection matrices.

STANDARD VERSUS TENSOR (50 iterations)

n	1000	2000	4000	6000
Lanczos time (sec)	2.8	12.1	76.7	224.9
Tensor Lanczos time (sec)	0.4	0.7	1.5	2.2

For $n = 6000$ we observe a **100** times acceleration.

MORE EXAMPLES

$$\mathbf{M}_r = \mathbf{M} - \sum_{t=1}^{\rho} \mathbf{D}_t \otimes \mathbf{D}_t$$

$\mathbf{M} = -\text{Laplacian}$. \mathbf{D}_t are diagonal matrices with positive entries.

Compare maximal eigenvalues on the 50th iteration.

$n = 300^2$, truncation rank = 10, $\varepsilon = 10^{-2}$.

ρ	1	3	5	7	9
Standard Lanczos	7.989	7.957	7.925	7.900	7.893
Tensor Lanczos	7.977	7.940	7.917	7.893	7.906

Entries of \mathbf{D}_t are uniform grid values of $(\mathbf{1} + \mathbf{T}_t(\mathbf{x}))/10$,

\mathbf{T}_t is the Chebyshev polynomial of degree t .

ρ	1	3	5	7	9
Standard Lanczos	7.862	7.615	7.302	6.800	6.460
Tensor Lanczos	7.852	7.608	7.292	6.789	6.452

Entries of \mathbf{D}_t come from random vectors with uniform distribution on $[0, 1]$.

SUBLINEAR COMPLEXITY

For 2-dimensional problems time grows
as SQUARE ROOT
of total number of nodes!

For d -dimensional problems time should grow
as ROOT OF DEGREE d
of total number of nodes!

FUTURE RESEARCH

- Additional structure for matrix factors (wavelet sparsification, Toeplitz-like structure, circulant-plus-low rank structure).
- Iterative methods for matrix functions (sign functions, square root, matrix exponential).
- Iterative methods for linear systems with structured vectors (PCG, BiCGStab) and preconditioners computed by the Newton method.
- Eigenvalue solvers with tensor-structured eigenvectors.

FUTURE RESEARCH

A challenge topic for the linear algebra and matrix analysis community:

Which properties of matrices do account for a Tucker approximation with low mode ranks for the inverses, matrix exponentials and other matrix functions?

All experiments are in favour of the following hypothesis: for all practical operators of mathematical physics (differential operators, integral operators with smooth, singular and hypersingular kernels) on tensor grids, standard matrix functions can be approximated in the Tucker format with ranks of order

$$r \sim \log^\alpha n \log^\beta \varepsilon^{-1},$$

with some constants α, β . The known theorems on tensor structure of the matrix functions rely on analytical considerations. Matrix-grounds attempts are very few:

I.V. Oseledets, E.E. Tyrtyshnikov, and N.L. Zamarashkin, Matrix inversion cases with size-independent tensor rank estimates, *Linear Algebra Appl.*, submitted, 2008.

E.E. Tyrtyshnikov, Tensor Ranks for Inversion of Tensor-Product Binomials, submitted, 2008.