The contraction variety of a tensor

Alexander Taveira Blomenhofer[1], Benjamin Lovitz[2]

[1]University of Copenhagen [2]Northeastern University, Boston

The contraction variety of a tensor

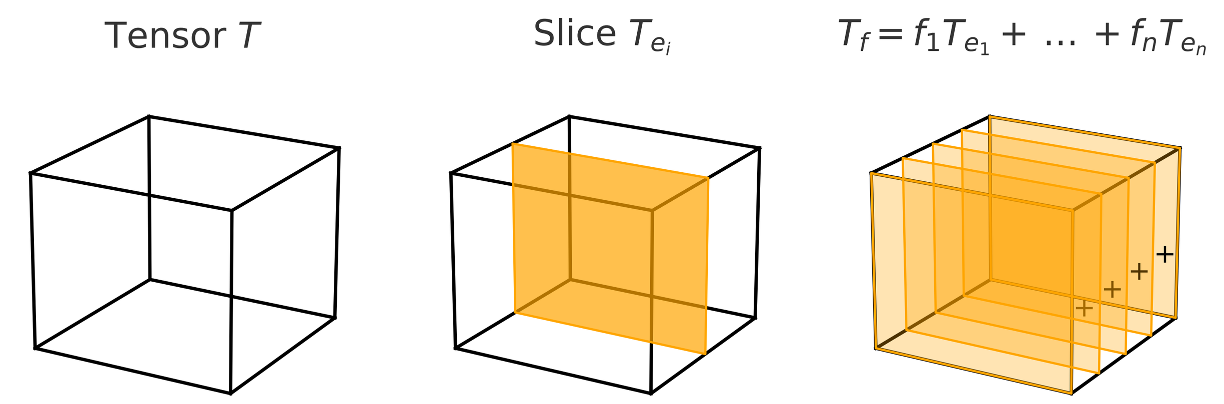

Figure 1. Slicing a 3-tensor gives a space $\mathcal{L}_T$ of matrices T_f, which is spanned by the slices T_{e_i}. We search for the singular matrices in this space. Those form the contraction variety X_T.

Given a concise 3-tensor $T \in \mathbb{C}^{n\times n\times n}$, we consider its contractions $T_f \in \mathbb{C}^{n\times n}$ along the first index via a vector $f \in \mathbb{C}^n$. The space $\mathcal{L}_T$ spanned by all contractions T_f is the span of its slices.

The contraction variety X_T of T is given by

$$X_T = \{f \in \mathbb{C}^n \mid \det(T_f) = 0\} \tag{1}$$

Sometimes, the contraction variety can be explicitly computed. See the following overview.[a]

Overview

(A) **Tensor Rank** If $T = \sum_{i=1}^n a_i^{\otimes 3}$ and $a_1, \ldots, a_n$ are linearly independent, then

$$X_T = \bigcup_{i=1}^n \langle a_i\rangle^\perp$$

(B) **Skew Rank** For $n = 3r$, if $T = \sum_{i=1}^r a_{i1} \wedge a_{i2} \wedge a_{i3}$ and $\{a_{ij}\}_{i=1,\ldots,r,j=1,2,3}$ is linearly independent, then

$$X_T = \bigcup_{i=1}^r \langle a_{i1}, a_{i2}, a_{i3}\rangle^\perp.$$

(C) **Chow Rank** For $n = 3r$, if $T = \sum_{i=1}^r a_{i1} \otimes_{\mathrm{sym}} a_{i2} \otimes_{\mathrm{sym}} a_{i3}$ and $\{a_{ij}\}_{i=1,\ldots,r,j=1,2,3}$ is linearly independent, then

$$X_T = \bigcup_{i=1}^r \langle a_{i1}\rangle^\perp \cup \langle a_{i2}\rangle^\perp \cup \langle a_{i3}\rangle^\perp.$$

(D) **W-States** For $n = 2r$, if $T = \sum_{i=1}^r a_{i1}^{\otimes 2} \otimes_{\mathrm{sym}} a_{i2}$ and $\{a_{ij}\}_{i=1,\ldots,r,j=1,2}$ is linearly independent, then

$$X_T = \bigcup_{i=1}^r \langle a_{i1}\rangle^\perp$$

.

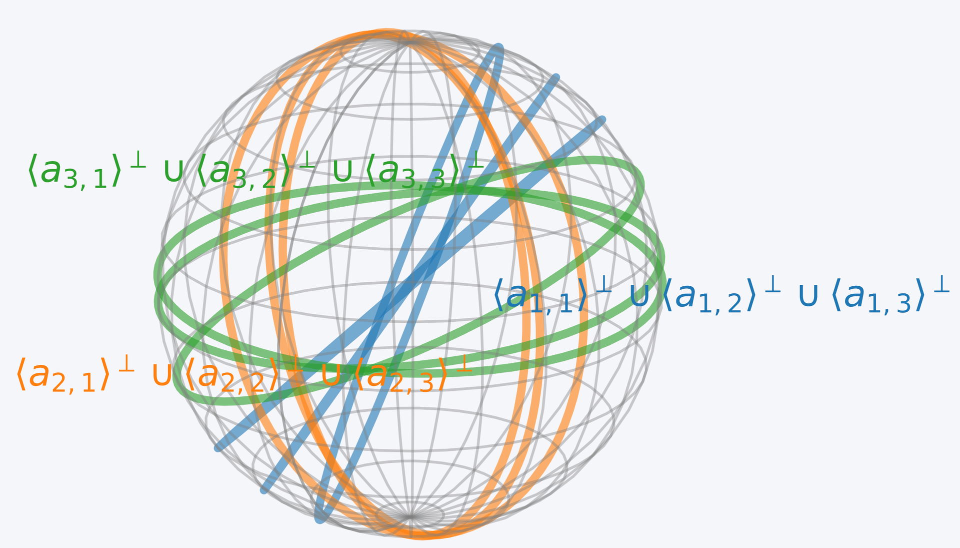

Figure 2. Visualization of the Chow contraction variety in $\mathbb{P}^2 \subseteq \mathbb{P}^{n-1}$ as a union of hyperplanes $\langle a_{ij}\rangle^\perp$.

[a] NB: If all matrices T_f are singular, then the contraction variety T_f is defined as the set of all $f \in \mathbb{C}^n$ such that the rank of T_f is smaller than the rank of a generic contraction. This is only important for the skew rank, since there, f always lies in the kernel of T_f.

Key Idea for Chow decomposition

If X_T is a union of n hyperplanes, then a generic line $f - \lambda g$ intersects X_T in n simple, distinct points. These points are the solutions of $\det(T_f - \lambda T_g) = 0$, which is a generalized eigenequation. We turn this observation into a linear-time algorithm for Chow rank decomposition.

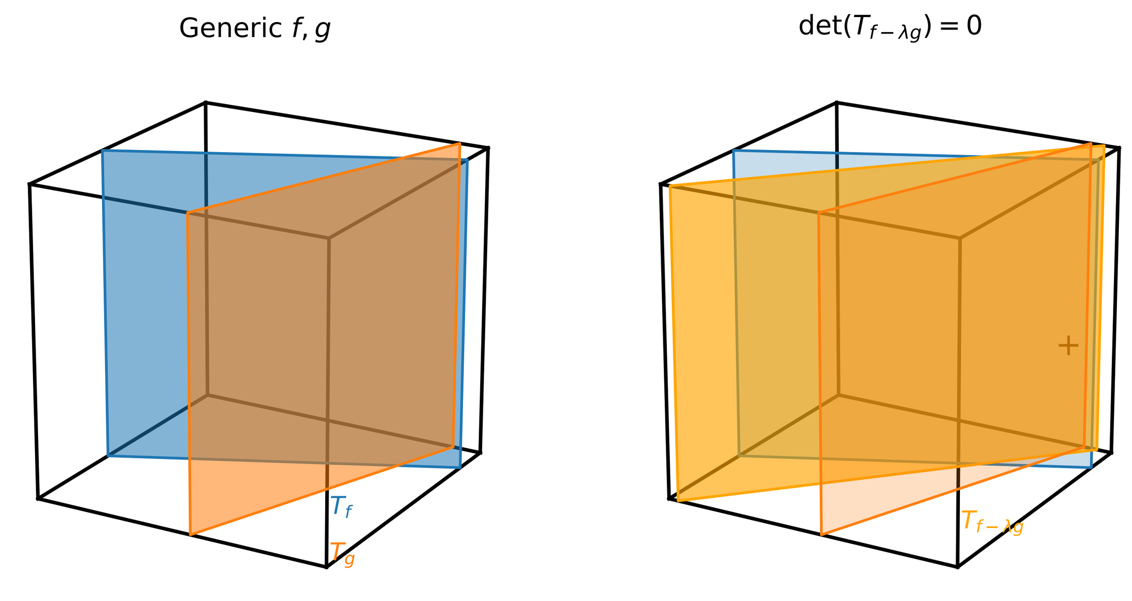

Figure 3. Visualization of a general matrix pencil obtained from contracting T and its n singular points $T_f - \lambda_i T_g$.

Linear time algorithm for Chow rank decomposition

Input: A 3-tensor $T \in \mathbb{C}^{n\times n\times n}$.
Assumption: T has a Chow decomposition $T = \sum_{i=1}^r a_{i1} \otimes_{\mathrm{sym}} a_{i2} \otimes_{\mathrm{sym}} a_{i3}$ with linearly independent $\{a_{ij}\}_{i=1,\ldots,r,j=1,2,3}$.
Output: The unique Chow rank decomposition $\{a_{ij}\}_{ij}$ of T.
Runtime: $\mathcal{O}(n^3)$, thus linear time in the size of T.

Procedure:

1. Pick three generic $f, g, h \in \mathbb{C}^n$ and compute the contractions T_f, T_g. (Time: $\mathcal{O}(n^3)$).
2. Compute the n (pairwise distinct) generalized eigenvalues of the pair T_f, T_g. That is, compute the n values of $\lambda \in \mathbb{C}$, so that
$$\det(T_f - \lambda T_g) = 0.$$
Denote by x_λ the corresponding eigenvalues. (Time: $\mathcal{O}(n^3)$)
3. Match the vectors b_λ into groups of three: Start with an eigenvalue λ and find the two other eigenvalues μ such that $\langle x_\mu, T_h x_\lambda\rangle \neq 0$. (Time: $\mathcal{O}(n^2)$ to compute $T_h x_\lambda$, time n^2 to compute the inner products of $T_h x_\lambda$ with all x_μ).
4. After matching, write the vectors as b_{ij} where $i = 1, \ldots, r$ and $j = 1, 2, 3$. Denote $U_i = \langle b_{i1}, b_{i2}, b_{i3}\rangle$.
5. Compute nonzero vectors $a'_{ik} \in \langle x_{ik}\rangle^\perp \cap \langle h_{ik}\rangle^\perp \cap U_i$ for $k = 1, 2, 3$ and all i. These vectors will be multiples of the a_{ik}. (Time: $\mathcal{O}(n^2)$)
6. Find the missing scalars $\rho_1, \ldots, \rho_r$ such that $T = \sum_{i=1}^r \rho_i b_{i1} b_{i2} b_{i3}$. To this end, solve the $n \times n$ linear system
$$T_{ff} = \sum_{i=1}^r \rho_i (a'_{i1} a'_{i2} a'_{i3})_{ff}.$$
This takes time $\mathcal{O}(n^3)$.
7. Output $\rho_1 \cdot a'_{11} \otimes_{\mathrm{sym}} a'_{12} \otimes_{\mathrm{sym}} a'_{13}, \ldots, \rho_r \cdot a'_{r1} \otimes_{\mathrm{sym}} a'_{r2} \otimes_{\mathrm{sym}} b_{r3}$.

Extensions

1. Similar algorithm for W-states, but not linear time $\mathcal{O}(n^4)$.
2. In general algorithms possible when $U_i = \langle a_{i1}, a_{i2}, a_{i3}\rangle$ are in direct sum.
3. For higher odd orders $T \in (\mathbb{C}^n)^{\otimes 2d+1}$, we can decompose generic Chow tensors of rank $r \approx n^d$ via flattening $T \mapsto \hat{T} \in \mathbb{C}^{n\times n^d\times n^d}$. This writes $\hat{T}_f = \sum_{i=1}^r M_{if}$ as a sum of matrices M_{if} of rank $\binom{2d+1}{d}$. Images of M_{if} are still skew by [?].

Correctness of the Algorithm (for brave & caffeinated readers)

Write $T_f = \sum_{i=1}^r M_{if}$, where M_{if} is the contraction of $a_{i1} \otimes_{\mathrm{sym}} a_{i2} \otimes_{\mathrm{sym}} a_{i3}$.

Step 0: First, we verify that $\langle a_{i1}\rangle^\perp \cup \langle a_{i2}\rangle^\perp \cup \langle a_{i3}\rangle^\perp$. Indeed, the matrices M_{if} have their images contained in the spaces $\langle a_{i1}, a_{i2}, a_{i3}\rangle$. Due to linear independence, these spaces are in direct sum. Therefore, the matrix T_f with image $\operatorname{im} T_f = \bigoplus_{i=1}^r \operatorname{im} M_{if}$ drops rank if and only if one of the M_{if} drops rank. Let us look at $i = 1$. To see when, say, M_{1f} drops rank, we can w.l.o.g. look apply a change of coordinates and assume that $a_{11}, a_{12}, a_{13} = e_1, e_2, e_3$. Writing $\alpha = \langle f, a_{11}\rangle, \beta = \langle f, a_{12}\rangle$ and $\gamma = \langle f, a_{13}\rangle$, we see that $M_{1f} = \alpha e_2 e_3 + \beta e_1 e_3 + \gamma e_1 e_2$, i.e., the upper 3×3 block of M_{1f} equals

$$\begin{pmatrix} 0 & \alpha & \beta \\ \alpha & 0 & \gamma \\ \beta & \gamma & 0 \end{pmatrix},$$

which has determinant $2\alpha\beta\gamma$. So M_{if} is rank deficient if and only if $f \in \langle a_{i1}\rangle^\perp \cup \langle a_{i2}\rangle^\perp \cup \langle a_{i3}\rangle^\perp$.

Step 3: Since X_T is a union of codimension 1 subspaces, we obtain a point in each component when intersecting it with a general line $\lambda \mapsto f - \lambda g$. Thus, for each solution λ of $\det(T_f - \lambda T_g) = 0$, the vector $h_\lambda = f - \lambda g$ is orthogonal to precisely one a_{ij} (assuming f, g generic). W.l.o.g., let $h_\lambda \perp a_{11}$ and remember that $a_{11}, a_{12}, a_{13} = e_1, e_2, e_3$. Then, we find nonzero β', γ' such that $M_{1h_\lambda} = e_1(\beta' e_3 + \gamma' e_2)$. For the eigenvector x_λ, we have $0 = T_{h_\lambda} x_\lambda = \bigoplus_{i=1}^r M_{ih_\lambda} x_\lambda$. Due to the direct sum property, this means that $x_\lambda \perp \langle a_{j1}, a_{j2}, a_{j2}\rangle$ for each $j \geq 2$. Since $\operatorname{im} M_{1h_\lambda}$ contains $e_1 = a_{11}$, we also have $x_\lambda \perp a_{11}$.

In particular, $x_\mu \perp T_h x_\lambda$, if they correspond to different spaces U_i. It can also be verified by direct calculation that $x_\mu \perp T_h x_\lambda \neq 0$ if both $T_h x_\mu$ and $T_h x_\lambda$ lie in the same U_i.

Step 5: After triple-pairing, write the eigenvalues, eigenvectors etc as $\lambda_{ik}, x_{ik}, h_{ik}$. By the above steps, a_{ik} is a vector in U_i which is orthogonal to both h_{ik} and x_{ik}. This uniquely determines a_{ik} up to scalar multiple, since $\dim U_i = 3$. Therefore, $\langle a_{ik}\rangle = \langle a'_{ik}\rangle$.

Step 6: After grouping, solving for the missing scalars can be done by a straightforward linear system $T = \sum_{i=1}^r \rho_i a'_{i1} a'_{i2} a'_{i3}$. However, this $n^3 \times n^3$ system is too large to be solved in linear time. Therefore, we double contract by f first to get a system of size $n \times n$, which can be solved in linear time.

Consequence

If a tensor T has a decomposition as in (A), (B) or (C), then it is the **unique** minimum tensor rank (A), skew rank (B) or Chow rank (B) decomposition.

For instance, the tensor $f = e_1 \otimes_{\mathrm{sym}} e_2 \otimes_{\mathrm{sym}} e_3 + e_4 \otimes_{\mathrm{sym}} e_5 \otimes_{\mathrm{sym}} e_6 \in \mathbb{C}^{6\times 6\times 6}$ has a Chow rank 2 and its minimum Chow rank representation is unique.

Proof.

For any other decomposition $T = \sum_{i=1}^r b_i^{\otimes 3}$ with $r \leq n$ or $T = \sum_{i=1}^r b_{i1} \wedge b_{i2} \wedge b_{i3}$ $T = \sum_{i=1}^r b_{i1} \otimes_{\mathrm{sym}} b_{i2} \otimes_{\mathrm{sym}} b_{i3}$ with $3r \leq n$, the b-vectors must be linearly independent, since the $(1, 2, 3)$-flattening of T has rank n. Therefore, we get two decompositions of the variety X_T into irreducible components. E.g., for Chow rank,

$$\bigcup_{i=1}^r \langle a_{i1}\rangle^\perp \cup \langle a_{i2}\rangle^\perp \cup \langle a_{i3}\rangle^\perp = X_T = \bigcup_{i=1}^r \langle b_{i1}\rangle^\perp \cup \langle b_{i2}\rangle^\perp \cup \langle b_{i3}\rangle^\perp.$$

By uniqueness of the irreducible decomposition, the decompositions are equal. □

References

[1] Alexander Taveira Blomenhofer and Benjamin Lovitz. On chow decompositions and the contraction variety, 2025. In preparation.