

Tensor decomposition beyond uniqueness, with an application to the minrank problem

Pascal Koiran* Rafael Oliveira†

October 31, 2025

Abstract

We prove a generalization to Jennrich's uniqueness theorem for tensor decompositions in the undercomplete setting. Our uniqueness theorem is based on an alternative definition of the standard tensor decomposition, which we call matrix-vector decomposition. Moreover, in the same settings in which our uniqueness theorem applies, we also design and analyze an efficient randomized algorithm to compute the unique minimum matrix-vector decomposition (and thus a tensor rank decomposition of minimum rank). As an application of our uniqueness theorem and our efficient algorithm, we show how to compute all matrices of minimum rank (up to scalar multiples) in certain generic vector spaces of matrices.

1 Introduction

A tensor can be viewed as a multidimensional array with entries in some field $\mathbb{K}$. In this paper, we will only consider tensors of order 3, i.e., elements of $\mathbb{K}^{m \times n \times p} = \mathbb{K}^m \otimes \mathbb{K}^n \otimes \mathbb{K}^p$. Given 3 vectors $u \in \mathbb{K}^m, v \in \mathbb{K}^n, w \in \mathbb{K}^p$ we recall that their tensor product $u \otimes v \otimes w$ is the tensor $T \in \mathbb{K}^{m \times n \times p}$ with entries: $T_{ijk} = u_i v_j w_k$. By definition, a tensor of this form with $u, v, w \neq 0$ is said to be of *rank one*. The rank of an arbitrary tensor T is defined as the smallest integer r such that T can be written as a sum of r tensors of rank one (and the rank of $T = 0$ is 0). The decomposition

$$T = \sum_{i=1}^r u_i \otimes v_i \otimes w_i \tag{1}$$

is said to be unique (or sometimes, "essentially unique") if up to a permutation, the rank-1 terms $u_i \otimes v_i \otimes w_i$ are the same in all decompositions of T as a sum of r tensors of rank one.

*Univ Lyon, EnsL, UCBL, CNRS, LIP, F-69342, LYON Cedex 07, France. Email: pascal.koiran@ens-lyon.fr.

†University of Waterloo. Email: rafael@uwaterloo.ca

The starting point of this paper is a uniqueness theorem for decomposition of order 3 tensors and an associated decomposition algorithm:

Theorem 1.1 (Jennrich’s uniqueness theorem). *Let $T = \sum_{i=1}^r u_i \otimes v_i \otimes w_i$ be a tensor in $\mathbb{K}^{m \times n \times p}$ such that:*

- (i) *The vectors u_i are linearly independent.*
- (ii) *The vectors v_i are linearly independent.*
- (iii) *The vectors w_i are pairwise independent.*

Then $\text{rank}(T) = r$, and the decomposition of T as a sum of r rank one tensors is essentially unique.

This result is a special case of Kruskal’s uniqueness theorem [Kru77]. In contrast to Kruskal’s theorem, Theorem 1.1 has an efficient algorithmic proof. The resulting algorithm is known as the “simultaneous diagonalization” or “Jennrich’s algorithm.”¹ The polynomial running time of this algorithm is a remarkable feature since tensor decomposition (or just computing the rank of a tensor) is in general NP-hard [Hås90]. This property has made the simultaneous diagonalization algorithm a cornerstone of further algorithmic work on tensors. For instance, this algorithm was used as a subroutine in [MSS16] to obtain an algorithm for overcomplete decomposition of random tensors. The term *overcomplete* refers to the situation when the tensor rank is larger than the dimensions of the tensor; by contrast, in the *undercomplete* setting of Theorem 1.1 we must have $r \leq \min(m, n)$. Some of the ideas behind Jennrich’s algorithm have also inspired the spectral algorithm in [HSS19]. More recently, the simultaneous diagonalization algorithm was used in [Koi25] to obtain the first efficient algorithm for overcomplete decomposition of generic tensors of order 3. Jennrich’s algorithm was also used in [JLV23] to find low rank matrices in matrix subspaces. This result of [JLV23] was in turn used in [KMW25] to give another efficient algorithm for overcomplete decomposition of generic tensors of order 3. In light of all these results, it is quite natural to look for more applications and generalizations of Jennrich’s uniqueness theorem and of the corresponding decomposition algorithm. This is the main focus of this paper.

1.1 Our results

In this paper we work in the undercomplete setting and prove three main results, which we now outline. Our first result is a generalization of Jennrich’s

¹The algorithm was not actually discovered by Robert Jennrich, so this name should be viewed as a tribute to his contributions to tensor decomposition (see [Har70, Har72]) rather than as a historically accurate attribution. The first published version of the algorithm seems to be from [LRA93].

uniqueness theorem ([Theorem 1.1](#)) where condition (iii) is removed. Our second result is on the algorithmic side, where we also generalize Jennrich’s decomposition algorithm to the setting of our uniqueness theorem. As a consequence of these results, our third result is to obtain an efficient algorithm that finds all matrices of minimum rank in certain generically chosen subspaces of matrices.

To properly describe our results, we introduce some basic notation that will be used throughout the paper. We denote by $M_{m,n}(\mathbb{K})$ the set of matrices with m rows, n columns and entries in $\mathbb{K}$. We denote by $M_n(\mathbb{K})$ the set of square matrices of size n , by $GL_n(\mathbb{K})$ the group of invertible matrices of size n , and by I_n the identity matrix of size n .

Before we formally state our results, we motivate the conceptual aspect behind our uniqueness theorem, which we call *matrix-vector decompositions*.

It is well known that without condition (iii) in [Theorem 1.1](#), the tensor decomposition of minimal rank is no longer unique. One way in which uniqueness fails already happens in the decompositions of matrices: it is easy to see that the decomposition of a rank 2 matrix as the sum of two matrices of rank 1 is never unique; see [Proposition 2.4](#) in [Section 2.2](#) for a proof of this fact. This failure of uniqueness may at first seem problematic, as most of the efficient tensor decomposition algorithms apply in a setting where the decomposition of smallest rank is known to be unique. Informally speaking, uniqueness (and the ingredients in a uniqueness proof) help a decomposition algorithm “zero in” on the correct decomposition.

The way we deal with the above obstacle to uniqueness is by considering a more relaxed decomposition of a tensor as a sum of tensor products of the form $M \otimes w$ where M is a matrix and w is a vector. More precisely, we have the following definition:

Definition 1.2 (Matrix-vector decompositions). *Let $T \in \mathbb{K}^{m \times n \times p}$ be a tensor. A matrix-vector decomposition of T is a decomposition of the form*

$$T = \sum_{\ell=1}^q M_{\ell} \otimes w_{\ell} \tag{2}$$

where $M_{\ell} \in M_{m,n}(\mathbb{K}) \setminus \{0\}$, $w_{\ell} \in \mathbb{K}^p \setminus \{0\}$ for every $\ell = 1, \dots, q$, and no two vectors w_{ℓ} in this list are colinear.

The rank of this decomposition is defined as $\sum_{\ell=1}^q \text{rank}(M_{\ell})$.

In [Proposition 2.3](#) we show that the smallest rank of a matrix-vector decomposition equals the tensor rank of T . Thus, the above decomposition generalizes the traditional tensor decomposition (the latter also requires that each M_{ℓ} be of rank 1), and it has the advantage that it avoids the aforementioned issue for matrices. In contrast to condition (iii) in [Theorem 1.1](#), non-colinearity of the vectors w_{ℓ} can be assumed without loss of generality: if two vectors $w_{\ell}, w_{\ell'}$ are scalar multiples of each other, we can add up the

corresponding (properly scaled) matrices $M_\ell, M_{\ell'}$ (and the number of terms in the decomposition goes from q to $q - 1$).

A related notion of decomposition is studied in [JLV23].² One key difference is that they want to minimize the number of terms in a decomposition, whereas we want to minimize the sum of the ranks of the M_ℓ .

So far, it looks like all we've achieved with Definition 1.2 is to redefine tensor rank in a slightly unusual way. But we will now see that this definition leads to new uniqueness results and decomposition algorithms. This is the content of our main uniqueness theorem.

Theorem 1.3 (Uniqueness theorem). *Let $T = \sum_{i=1}^r u_i \otimes v_i \otimes w_i$ be a tensor in $\mathbb{K}^{m \times n \times p}$ such that:*

- (i) *The vectors u_i are linearly independent.*
- (ii) *The vectors v_i are linearly independent.*
- (iii) *Every vector w_i is nonzero.*

Then $\text{rank}(T) = r$, and T has a unique rank r matrix-vector decomposition.

We now give an equivalent version of the above theorem which is stated solely in terms of matrix-vector decompositions. This version will be more convenient to use in certain parts of the paper.

Theorem 1.4 (Uniqueness theorem, equivalent formulation). *Suppose that a tensor $T \in \mathbb{K}^{m \times n \times p}$ has a matrix-vector decomposition of the form:*

$$T = \sum_{\ell=1}^q M_\ell \otimes w_\ell \tag{3}$$

where the linear spaces $\text{Im}(M_\ell)$ are in direct sum and where the linear spaces $\text{Im}(M_\ell^T)$ are also in direct sum. Then $\text{rank}(T) = \sum_{\ell=1}^q \text{rank}(M_\ell)$, and (3) is the unique matrix-vector decomposition of T of minimum rank.

We prove the equivalence of these two theorems in Section 3.1, and we prove the above theorems in Section 3.2.

With the above uniqueness theorems at hand, we are now ready to state our algorithmic contribution: under the uniqueness conditions, we can compute the minimum matrix-vector decomposition.

²An improved analysis of one of their results can be found in [DWW25].

Theorem 1.5 (Matrix-vector decomposition algorithm). *Suppose that a tensor $T \in \mathbb{K}^{m \times n \times p}$ has a matrix-vector decomposition of the form:*

$$T = \sum_{\ell=1}^q M_{\ell} \otimes w_{\ell}$$

where the linear spaces $\text{Im}(M_{\ell})$ are in direct sum and where the linear spaces $\text{Im}(M_{\ell}^T)$ are also in direct sum.

There is a randomized, polynomial-time algorithm (Algorithm 3) such that, on input T as above, it outputs the above matrix-vector decomposition (as usual, the M_{ℓ} and w_{ℓ} are determined only up to scaling and permutation).

Note that our decomposition algorithm applies under the same conditions as in Theorem 1.3 (or equivalently Theorem 1.4).

When $\mathbb{K} = \mathbb{Q}$, our algorithm can be implemented in the Turing machine model of computation, and it runs in time polynomial in the bit size of the input tensor T (Remark 4.8).

In the case of general fields $\mathbb{K}$, we assume that we have access to an algorithm for the computation of the roots of a univariate polynomial with coefficients in $\mathbb{K}$ (for step 4 of Algorithm 1). This is a fairly standard assumption in the study of tensor decomposition algorithms. For a more thorough discussion of this issue, see [Koi25, Section 1.4].

With the above algorithmic result at hand, we can state our third result: an efficient algorithm which finds all the matrices of minimum rank in subspaces of matrices with a basis satisfying certain special properties.

Theorem 1.6 (Minrank algorithm). *Suppose that $V \subset M_{m,n}(\mathbb{K})$ is a subspace spanned by a basis $M_1, \dots, M_p$ (the “hidden basis”) where the linear spaces $\text{Im}(M_1), \dots, \text{Im}(M_p)$ are in direct sum, and where the linear spaces $\text{Im}(M_1^T), \dots, \text{Im}(M_p^T)$ are also in direct sum.*

There is a randomized, polynomial-time algorithm (Algorithm 4) such that, when given as input any basis $Z_1, \dots, Z_p$ of V , it correctly finds the hidden basis and it outputs $\rho = \min_{M \in V, M \neq 0} \text{rank } M$, as well as matrices $A_1, \dots, A_s$ which are, up to scalar multiplication, the only matrices of rank ρ in V . Moreover, the matrices A_i are a subset of the hidden basis.

The above theorem follows as a corollary of our decomposition algorithm. Given as input a basis $Z_1, \dots, Z_p$ of V , we obtain a matrix-vector decomposition of the tensor $T \in \mathbb{K}^{m \times n \times p}$ obtained by “stacking up” the matrices $Z_1, \dots, Z_p$ (i.e., these matrices are the 3-slices of T). From this matrix-vector decomposition, we show that the direct sum conditions imply that the matrices of minimum rank in this decomposition are the desired minimum rank matrices. A proof of the above theorem and of the aforementioned claims can be found in Section 5.2.

Genericity of hypotheses. It is worth noting that the hypotheses in all of the theorems stated above are *generic properties*, in the algebraic geometric sense (which we discuss in more detail in [Section 2](#)). In particular, they will work for randomly chosen tensors of rank r or vector spaces of matrices spanned by matrices $M_1, \dots, M_p$ randomly chosen such that

$$\sum_{i=1}^p \text{rank } M_i \leq \min(m, n).$$

1.2 Previous & related works

Comparison with Jennrich’s algorithm Let $T_1, \dots, T_p$ be the 3-slices of the input tensor $T \in \mathbb{K}^{m \times n \times p}$. Each slice is an $m \times n$ matrix. Assume first for simplicity that one of the slices is invertible, for instance T_1 . This implies that $m = n = r$ in [Theorem 1.1](#). In this case, the matrices $T_k T_1^{-1}$ for $k = 2, \dots, p$ turn out to be the simultaneously diagonalizable and the u_i are the eigenvectors. This can be exploited algorithmically as follows: we compute two random linear combinations of the 3-slices. With high probability all the eigenvalues of $T_a T_b^{-1}$ are distinct, and the u_i are the (unique) eigenvectors. The v_i can be computed with a similar procedure, and each u_i can be paired with the corresponding v_i by comparing the respective eigenvalues. Finally, the w_i can be obtained by solving a linear system. When $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{C}$ and we no longer assume that $m = n = r$, T_b^{-1} can be replaced by the Moore-Penrose pseudoinverse $T_b^\dagger$. For more details on the resulting “Jennrich” or “simultaneous diagonalization” algorithm refer to [Section A](#), where we have followed the presentation in [[Moi18](#), Section 3.3].

As explained above, for our decomposition algorithm we drop condition (iii) in [Theorem 1.1](#). The resulting algorithm can be thought of as a version of Jennrich’s for multiple eigenvalues (the relevant matrices can have multiplicity greater than 1). Also, we no longer have access to the Moore-Penrose inverse since we aim for a uniqueness theorem and a corresponding decomposition algorithm that apply to arbitrary fields. For this reason, instead of diagonalizing $T_a T_b^\dagger$ we solve the generalized eigenvalue problem $T_a x = \lambda T_b x$. For the field of real and complex numbers, we also provide a version of the algorithm which uses the Moore-Penrose inverse (compare [Sections C](#) and [4.1](#)). Even when there are no multiple eigenvalues, as explained in [Remark 4.6](#) this version of the algorithm slightly differs from the standard Jennrich algorithm as presented in [[Moi18](#), Section 3.3].

When the matrices $T_k T_1^{-1}$ are simultaneously diagonalizable, they must in particular commute. In this paper we generalize Jennrich’s uniqueness theorem and the corresponding decomposition algorithm from the case where there are no multiple eigenvalues to the case where the multiplicities can be larger than 1. In a forthcoming paper we plan to generalize this one step

further, from the diagonalizable case to the case where the relevant matrices may no longer be diagonalizable, but still commute.

Another natural question to pursue is the computation of matrix-vector decompositions in the overcomplete setting. Two algorithms for the decomposition of generic tensors in that setting were recently given in [Koi25, KMW25]. In particular, the algorithm in [Koi25] uses Jennrich’s algorithm to decompose an auxiliary tensor T' computed from the input tensor T . If we can instead decompose T' with the algorithm from the present paper, this would likely increase the range of applicability of the algorithm in [Koi25] (and of the corresponding uniqueness theorem).

Comparison with [JLV23]. The work of [JLV23] studies the computational problem of determining the intersection of an algebraic variety with a generic linear subspace of appropriate dimension, such that the intersection is zero-dimensional. Their motivation to study this problem is due to the fact that special cases of this problem have applications in quantum information theory and tensor decompositions. Two of the main applications of the technical results of [JLV23], which are also related to our work, are [JLV23, Corollary 4] on the min-rank problem and [JLV23, Corollary 8] on the tensor decomposition problem.

In Corollary 4, they give an algorithm that takes as inputs $r \in \mathbb{N}$, dimensions $n_1, n_2 > r$, and a basis for a linear subspace $\mathcal{U} \subset \mathbb{K}^{n_1 \times n_2}$ of dimension

$$R \leq \frac{\binom{n_1}{r+1} \cdot \binom{n_2}{r+1}}{(r+1)! \cdot \binom{n_1 n_2 + r}{r+1}} \cdot (n_1 n_2 + r) \leq \frac{n_1 n_2}{((r+1)!)^2}$$

where $\mathcal{U}$ satisfies $\mathcal{U} := \text{span}_{\mathbb{K}}\{A_1, \dots, A_s\} + \text{span}_{\mathbb{K}}\{B_1, \dots, B_{R-s}\}$. Here, the A_i are *generic* matrices of rank at most r and the B_i are *generic* matrices in $\mathbb{K}^{n_1 \times n_2}$ (both A_i, B_j are not known in advance). The authors prove that $\mathcal{U}$ intersects the space of matrices of rank at most r at exactly s matrices (up to scalar multiples) and [JLV23, Algorithm 1] returns these elements (up to scalar multiples) in $(n_1 n_2)^{O(r)}$ time.

While their result yields polynomial time algorithms for any constant r , and in this regime for spaces of dimension $O(n_1 n_2)$, as soon as r is not constant the running time of their algorithm rapidly deteriorates, and as soon as $r \sim \frac{1}{2} \cdot \log(n_1 n_2)$ their algorithm stops working due to the rank condition never being satisfied. In comparison, our Theorem 1.6 (and Remark 5.4) works for generic instances with input $\mathcal{U} := \text{span}_{\mathbb{K}}\{A_1, \dots, A_R\}$ where the matrices A_i (which can be thought of as the “hidden basis”) satisfy the following inequality:

$$\sum_{i=1}^R \text{rank } A_i \leq \min\{n_1, n_2\}.$$

In the above case, our algorithm fully recovers the hidden basis in polynomial time. In particular, our algorithm improves [JLV23, Corollary 4] in the “pure low rank” version of their problem (i.e., when $\mathcal{U}$ is only generated by low-rank matrices – the A_i ’s) whenever $r \gtrsim \log(n_1 n_2)$.

In [JLV23, Corollary 8], the authors show that their main algorithm ([JLV23, Algorithm 1]) can compute in randomized polynomial time the tensor rank along with the tensor rank decomposition of generically chosen tensors $T \in \mathbb{K}^{n_1} \otimes \mathbb{K}^{n_2} \otimes \mathbb{K}^{n_3}$ of rank upper bounded by $\min \left\{ \frac{(n_1 - 1)(n_2 - 1)}{4}, n_3 \right\}$. Their main algorithm is remarkably simple in nature, as it simply consists of applying Jennrich’s algorithm to a suitable 3-tensor (in spaces of larger dimensions) constructed from T . The generic conditions in their paper ([JLV23, Proposition 25]) are distinct from the ones we consider in this paper, and it is easy to see that no generic condition is a subset of the other whenever the parameters allow for both to occur. However, it is worth noting that their genericity assumption holds for random tensors of rank upper bounded by $\min \left\{ \frac{(n_1 - 1)(n_2 - 1)}{4}, n_3 \right\}$, which makes their elegant algorithm useful beyond what our approach can handle.

2 Preliminaries

In this section we establish the technical preliminaries that we will need in the rest of the paper. We begin with a discussion and precise definition of generic properties.

Genericity. Throughout the paper, we use the term “generic” in its standard algebro-geometric sense: a set is generic if it contains a Zariski-open set (i.e., if it contains the complement of an algebraic set). More precisely, if $\mathbb{K}$ is an infinite field,³ we say that a property generically holds true in $\mathbb{K}^N$ if there is a nonidentically zero polynomial $P(x_1, \dots, x_N)$ such that the property holds true for all $x \in \mathbb{K}^N$ such that $P(x_1, \dots, x_N) \neq 0$. In particular, if $\mathbb{K}$ is the field of real or complex numbers, the set of points that do not satisfy the property has measure 0.

We also use “generic linear combinations” in the presentation of the decomposition algorithm in order to emphasize the connection with the uniqueness theorem. It is more algorithmically realistic to choose the coefficients $a_1, \dots, a_k$ from a finite set. This yields a randomized algorithm, and its probability of error can be made as small as desired by increasing the size of this set (Remark 4.3).

³If $\mathbb{K}$ is finite, we can replace it by its algebraic closure.

2.1 The Generalized Eigenvalue Problem

Definition 2.1. For two matrices $A, B \in M_{m,n}(\mathbb{K})$, a nonzero vector $v \in \mathbb{K}^n$ is a (generalized) eigenvector of the pair (A, B) if there exists $\lambda \in \mathbb{K}$ such that $Av = \lambda Bv$. Here, λ is the corresponding (generalized) eigenvalue.

The set of matrices of the form $A - \lambda B$ where λ ranges over $\mathbb{K}$ is traditionally called a (linear) *matrix pencil*. We use the same notation (A, B) to denote a pair of matrices or the corresponding matrix pencil.

The (generalized) eigenspace associated to the (generalized) eigenvalue λ is the set of vectors $v \in \mathbb{K}^n$ such that $Av = \lambda Bv$. This is indeed a linear subspace, like a traditional eigenspace. Despite this similarity, generalized eigenvectors and eigenvalues behave somewhat differently than the traditional ones. For instance, a nonzero vector in $\ker A \cap \ker B$ can be viewed as an eigenvector associated to the eigenvalue λ for *any* $\lambda \in \mathbb{K}$. These vectors will therefore belong to the eigenspace associated to λ for all $\lambda \in \mathbb{K}$. This motivates the following definition.

Definition 2.2. We say that an eigenvalue λ of the pair (A, B) is nontrivial if there exists an eigenvector v associated to λ such that $Bv \neq 0$.

In particular, if $\ker B = \{0\}$ all eigenvalues are nontrivial. If $\ker B \neq \{0\}$, the pencil (A, B) is said to have *eigenvalues at infinity*. If $\ker A \cap \ker B \neq \{0\}$ the pencil is said to be degenerate. The pencil (A, B) can also be viewed as a tensor of format $m \times n \times 2$. As pointed out in [BCS97, Exercise 19.9], a pencil is nondegenerate if and only if the 2-slices of the corresponding tensor are linearly independent (the corresponding tensor is said to be *2-concise* in this case). We will not use this fact in the sequel.

When $m = n$ and B is invertible, the eigenvalues/eigenvectors of the pair (A, B) are those of the matrix $B^{-1}A$. Applying this observation requires the computation of a matrix inverse. An inverse-free algorithm for matrix pencil diagonalization can be found in [DDS23]. In full generality, the generalized eigenvalue problem can be solved by computing the Kronecker normal form of the matrix pencil (see e.g. [VD79]). In our case, however, there is a much simpler solution due to the special structure of the matrices involved (see Section C and especially Proposition 4.1 in Section 4.1).

2.2 Facts about matrix-vector decompositions

We now establish basic facts about matrix-vector decompositions (recall that this notion was introduced in Definition 1.2).

Proposition 2.3. The smallest rank of a matrix-vector decomposition is equal to the tensor rank of T .

Proof. Suppose that T has a matrix-vector decomposition of rank r . Writing down each matrix M_l in this decomposition as a sum of $\text{rank}(M_l)$ matrices

of rank one shows that $\text{rank}(T) \leq r$. Assume conversely that T admits a decomposition as a sum of r rank-1 tensors as in (1). If there are no colinear vectors in the list $w_1, \dots, w_\ell$, (1) is already a matrix-vector decomposition with $M_i = u_i \otimes v_i = u_i v_i^T$. If some of these vectors are colinear, we can factorize using the rule:

$$\sum_{\ell} u_{\ell} \otimes v_{\ell} \otimes (\lambda_{\ell} w) = \left(\sum_{\ell} \lambda_{\ell} u_{\ell} \otimes v_{\ell} \right) \otimes w.$$

This results in a matrix-vector decomposition of rank at most r . $\square$

Let $r = \text{rank}(T)$. We say that T has a unique matrix-vector decomposition of rank r if up to permutation, the terms $M_{\ell} \otimes w_{\ell}$ are the same in all matrix-vector decompositions of T of rank r . We will show in Proposition 2.5 that uniqueness of decomposition as a sum of rank-1 tensors implies uniqueness of matrix-vector decompositions. The proof relies on the following well-known lemma (see e.g. [LRA93]).

Proposition 2.4. *The decomposition of a matrix of rank 2 as the sum of two matrices of rank 1 is never unique.*

For completeness, we give a proof of this lemma below (note that it applies to an arbitrary field).

Proof. Let $M = u_1 v_1^T + u_2 v_2^T$ be a matrix of rank 2. Since $M = u_1(v_1 + v_2)^T + (u_2 - u_1)v_2^T$, it suffices to show that

$$\{u_1 v_1^T, u_2 v_2^T\} \neq \{u_1(v_1 + v_2)^T, (u_2 - u_1)v_2^T\}.$$

Suppose that $u_1 v_1^T = u_1(v_1 + v_2)^T$. This implies $u_1 v_2^T = 0$, i.e., $u_1 = 0$ or $v_2 = 0$. This is impossible since M would then be of rank less than 2. The equality $u_1 v_1^T = (u_2 - u_1)v_2^T$ is also impossible. Indeed, the columns of $u_1 v_1^T$ are colinear to u_1 and those of $(u_2 - u_1)v_2^T$ are colinear to $u_2 - u_1$. Then u_1 would be colinear to u_2 , and M would again be of rank less than 2. $\square$

Proposition 2.5. *Suppose that a tensor of rank r admits a unique decomposition*

$$T = \sum_{\ell=1}^r (u_{\ell} \otimes v_{\ell}) \otimes w_{\ell} \tag{4}$$

as a sum of r tensors of rank one. This is also the unique matrix-vector decomposition of T of rank r . In particular, there are no colinear vectors in the list $w_1, \dots, w_{\ell}$.

Proof. We first show that there are no colinear vectors in this list. Assume for instance that w_1 and w_2 are colinear. We can assume that $w_1 = w_2$ by scaling v_1 and w_1 (or v_2 and w_2) if necessary. The sum of the first two

terms in (4) is equal to $M \otimes w_1$ where $M = u_1 \otimes v_1 + u_2 \otimes v_2$. This matrix must be of rank 2 by the minimality of (4). By Lemma 2.4, there is another decomposition $M = u'_1 \otimes v'_1 + u'_2 \otimes v'_2$. Replacing the first two terms of (4) by $u'_1 \otimes v'_1 \otimes w_1$ and $u'_2 \otimes v'_2 \otimes w_1$ yields a rank r decomposition of T which is not equivalent to (4).

We have thus shown that there are no colinear vectors among $w_1, \dots, w_\ell$; hence (4) is a bona fide matrix-vector decomposition. It remains to show that any other matrix-vector decomposition of rank r , say,

$$T = \sum_{\ell} M_{\ell} \otimes w'_{\ell} \quad (5)$$

must be equivalent to (4). We can expand each M_{ℓ} as a sum of $\text{rank}(M_{\ell})$ matrices of rank 1. From these expansions and (5) we obtain a decomposition of T as a sum of r tensors of rank 1. This new decomposition must be equivalent to our first decomposition (4), which is assumed to be unique. But:

- (i) there are no colinear w_{ℓ} in (4);
- (ii) whereas the expansion of M_{ℓ} yields $\text{rank}(M_{\ell})$ rank 1 tensors with the same third-mode vector w'_{ℓ} in our new decomposition of T as a sum of rank 1 tensors.

Hence $\text{rank}(M_{\ell}) = 1$ for all ℓ . It follows that (4) and (5) are equivalent as decompositions as sums of rank 1 tensors, and also as matrix-vector decompositions. $\square$

3 Uniqueness Theorems

After some preliminaries, we establish the equivalence between Theorems 1.3 and 1.4 in Section 3.1. We then prove Theorem 1.4 in Section 3.2. The proofs are based on the notion of (generalized) eigenvalues and eigenvectors for a matrix pencil, which can be found in Section 2.1.

3.1 Equivalence of the uniqueness theorems

Recall that a tensor $T \in \mathbb{K}^{m \times n \times p}$ can be cut into p "slices" $Z_1, \dots, Z_p$ where each slice is a $m \times n$ matrix. These are the 3-slices of T . One can also cut T in the two other directions into its 1-slices and 2-slices. In this paper we will only work with the 3-slices, and henceforth the term "slices" will refer to the 3-slices. It follows immediately from Definition 1.2 that the slices are linear combinations of the matrices M_{ℓ} occurring in a decomposition of T , namely,

$$Z_k = \sum_{\ell=1}^q w_{\ell k} M_{\ell}. \quad (6)$$

Moreover, it is well known that each slice can be expressed as a product of three matrices defined from the vectors occurring in a decomposition of T in its traditional form (1). Namely, if $\text{rank } T = r$, we have

$$Z_k = U D_k V^T, \quad (7)$$

where U is the $m \times r$ matrix having the u_i as column vectors, V is the $r \times n$ matrix having the v_i as column vectors, and $D_k = \text{diag}(w_{1k}, \dots, w_{rk})$. This notation is consistent with e.g. [Moi18]; note however that the transposed notation $Z_k = U^T D_k V$ was used in [Koi25] (i.e., U and V were defined as the matrices having the u_i and v_i as *row vectors*). We record the following simple consequence of (6) and (7):

Lemma 3.1. *Any matrix in the span of the slices of T is of rank at most $\text{rank}(T)$.*

Proof. Let $r = \text{rank}(T)$. By (7), any matrix Z in the span of the Z_k is of the form $Z = U D V^T$, where D is a linear combination of the D_k . Since $D \in M_r(\mathbb{K})$, we have the upper bound $\text{rank}(Z) \leq r$.

This can also be seen from a matrix-vector decomposition of T . Indeed, T has a matrix-vector decomposition of rank r by Proposition 2.3. By (6), Z is a linear combination of the M_ℓ hence $\text{rank}(Z) \leq \sum_\ell \text{rank}(M_\ell) = r$. $\square$

The following simple lemma, follows directly from (6)

Lemma 3.2. *Let $T \in \mathbb{K}^{m \times n \times p}$ be a tensor with slices $T_1, \dots, T_p$ and a decomposition $T = \sum_{\ell=1}^q M_\ell \otimes w_\ell$. For any $A \in M_m(\mathbb{K})$, the tensor T' with slices $AT_1, \dots, AT_p$ admits the decomposition $T' = \sum_{\ell=1}^q (AM_\ell) \otimes w_\ell$.*

Lemma 3.3. *Let $\mathbb{K}$ be an arbitrary field. Let $(u_1, \dots, u_r)$ and $(v_1, \dots, v_r)$ be two families of vectors of $\mathbb{K}^n$, respectively of rank r_u and r_v . For the matrix $M = \sum_{i=1}^r u_i v_i^T$ we have $\text{rank } M \leq \min(r_u, r_v)$. Moreover, if $r_u = r_v = r$ then $\text{rank } M = r$ as well.*

Proof. In the expression for M we can rewrite each u_i as a linear combination of the elements of a basis $e_1, \dots, e_{r_u}$. This yields an expression for M as a sum of r_u matrices of rank at most 1, namely,

$$M = \sum_{i=1}^{r_u} e_i w_i^T$$

where the w_i are linear combinations of the v_i . Hence $\text{rank } M \leq r_u$, and $\text{rank } M \leq r_v$ by a similar argument.

Assume now that $r_u = r_v = r$. It remains to show that $\text{rank } M = r$. This is equivalent to $\dim \ker M = n - r$. A vector $x \in \mathbb{K}^n$ is in the kernel if and only if $\sum_{i=1}^r (v_i^T x) u_i = 0$. Since the u_i are linearly independent, this is equivalent to $v_i^T x = 0$ for all i . Using now the linear independence of the v_i , it follows that the solution space is of dimension $n - r$ as needed. $\square$

We are now ready to show the equivalence of the two uniqueness theorems.

Proposition 3.4. *Let $T = \sum_{i=1}^r u_i \otimes v_i \otimes w_i$ be a tensor of format $m \times n \times p$ such that:*

- (i) *The vectors u_i are linearly independent.*
- (ii) *The vectors v_i are linearly independent.*
- (iii) *Every vector w_i is nonzero.*

Then T has a matrix-vector decomposition $T = \sum_{\ell=1}^q M_\ell \otimes w'_\ell$ where the linear spaces $\text{Im}(M_\ell)$ are in direct sum, the linear spaces $\text{Im}(M_\ell^T)$ are in direct sum, $r = \sum_{\ell=1}^q \text{rank}(M_\ell)$ and $\{w'_1, \dots, w'_\ell\} \subseteq \{w_1, \dots, w_r\}$.

Proof. Let us group together the w_ℓ that are colinear, like in the proof of Proposition 2.3. This yields a matrix-vector decomposition $T = \sum_{\ell=1}^q M_\ell \otimes w'_\ell$ where $\{w'_1, \dots, w'_\ell\} \subseteq \{w_1, \dots, w_r\}$. By Lemma 3.3, the matrix $M = \sum_{\ell=1}^q M_\ell$ is of rank r . This implies that $r = \sum_{\ell=1}^q \text{rank}(M_\ell)$ and that the linear spaces $\text{Im}(M_\ell)$ are in direct sum. A similar reasoning for $M^T = \sum_{\ell=1}^q M_\ell^T$ shows that the linear spaces $\text{Im}(M_\ell^T)$ are in direct sum. $\square$

Here is a converse to this proposition.

Proposition 3.5. *Suppose that a tensor T of format $m \times n \times p$ has a matrix-vector decomposition of the form:*

$$T = \sum_{\ell=1}^q M_\ell \otimes w'_\ell$$

where the linear spaces $\text{Im}(M_\ell)$ are in direct sum, and the linear spaces $\text{Im}(M_\ell^T)$ are also in direct sum. Then T has a decomposition of the form $T = \sum_{i=1}^r u_i \otimes v_i \otimes w_i$ where $r = \sum_{\ell=1}^q \text{rank}(M_\ell)$ and:

- (i) *The vectors u_i are linearly independent.*
- (ii) *The vectors v_i are linearly independent.*
- (iii) *Every vector w_i is nonzero.*

Moreover, we can assume that $w_i \in \{w'_1, \dots, w'_\ell\}$ for all i .

Proof. We apply the transformation of Proposition 3.4 in reverse, i.e., we write each M_ℓ as a sum of rank 1 matrices:

$$M_\ell = \sum_{k=1}^{\text{rank}(M_\ell)} u_{k\ell} \otimes v_{k\ell}.$$

Then we obtain the desired decomposition of T by expanding each product $M_\ell \otimes w'_\ell$. In the resulting expansion, the u_i are linearly independent since they form a basis of $\bigoplus_\ell \text{Im}(M_\ell)$. A similar reasoning for the M_ℓ^T shows that the v_i are linearly independent. The w_i are nonzero since they are the same vectors as the w'_i (each w'_i is repeated $\text{rank}(M_i)$ times in the list of the w_i). $\square$

3.2 Proof of the uniqueness theorems

In this section, we now prove [Theorem 1.4](#). *Throughout this section, we assume that T is a tensor satisfying the (equivalent) hypotheses of [Theorems 1.3](#) and [1.4](#).* We also assume that our field $\mathbb{K}$ is infinite. This is without loss of generality since uniqueness of decomposition for a field implies uniqueness for all its subfields.

As a main step toward the uniqueness theorems, we show in [Theorem 3.11](#) that the linear spaces $\text{Im } M_\ell$ are the same in all matrix-vector decompositions $T = \sum_{\ell=1}^q M_\ell \otimes \gamma_\ell$ of minimal rank. The uniqueness of the matrix-vector decomposition then follows from a simple direct sum argument.

The next claim follows from a simple argument which appears in [\[Koi25\]](#) before [Proposition 11](#). We recall the proof here for the sake of completeness.

Proposition 3.6. *The span of the slices of T contains a matrix of rank r .*

Proof. By [\(7\)](#), the matrices in $\langle Z_1, \dots, Z_p \rangle$ are exactly the matrices of the form:

$$Z = UDV^T, \quad D = \sum_{k=1}^p c_k D_k \quad (8)$$

where $c_1, \dots, c_p \in \mathbb{K}$. Each entry of D is a linear form in the c_k . Since $w_i \neq 0$, for all i there exists k such that $w_{ik} \neq 0$. As a result, these linear forms are all nonzero. Indeed,

$$D = \text{diag}(\langle c, w_1 \rangle, \dots, \langle c, w_r \rangle) \quad (9)$$

and the k -th coefficient of the i -th linear form is $(D_k)_{ii} = w_{ik} \neq 0$. Since $\mathbb{K}$ is infinite, there exists $c_1, \dots, c_p \in \mathbb{K}$ such that the matrix D in [\(8\)](#) is invertible, and the corresponding $Z = UDV^T$ is of rank r since U, D and V^T are all of full rank r . $\square$

It follows from [Proposition 3.6](#) and [Lemma 3.1](#) that $\text{rank}(T) = r$ as claimed in [Theorem 1.3](#).

Corollary 3.7. *In any matrix-vector decomposition of T of rank r , the linear spaces $\text{Im}(M_\ell)$ are in direct sum.*

Proof. Let $T = \sum_{\ell=1}^q M_\ell \otimes w_\ell$ be a matrix-vector decomposition of T of rank r . By [Proposition 3.6](#) there is a matrix Z of rank r in the span of the slices of T , and by [Eq. \(6\)](#) Z is a linear combination of the M_ℓ . Therefore,

$$r = \dim \operatorname{Im}(Z) \leq \dim\left(\sum_{\ell} \operatorname{Im}(M_\ell)\right) \leq \sum_{\ell} \dim(\operatorname{Im} M_\ell) = r.$$

Here the last equality follows from the definition of the rank of a matrix-vector decomposition. We conclude that

$$\dim\left(\sum_{\ell} \operatorname{Im} M_\ell\right) = \sum_{\ell} \dim(\operatorname{Im} M_\ell),$$

and the spaces $\operatorname{Im}(M_\ell)$ are indeed in direct sum. $\square$

Corollary 3.8. *In any matrix-vector decomposition of T of rank r , the linear spaces $\operatorname{Im}(M_\ell^T)$ are in direct sum.*

Proof. Let $T = \sum_{\ell=1}^q M_\ell \otimes \gamma_\ell$ be a matrix-vector decomposition of T of rank r . Consider the tensor $T' = \sum_{i=1}^r v_i \otimes u_i \otimes w_i$, obtained from T by exchanging u_i and v_i in the decomposition $T = \sum_{i=1}^r u_i \otimes v_i \otimes w_i$. For T' we have the matrix-vector decomposition of rank r :

$$T' = \sum_{\ell=1}^q M_\ell^T \otimes \gamma_\ell.$$

The result therefore follows from [Corollary 3.7](#) applied to T' . $\square$

It remains to show that T has a unique matrix-vector decomposition of rank r . First, we'll show (in [Theorem 3.11](#)) that the linear spaces $\operatorname{Im}(M_\ell)$ are the same in all matrix-vector decompositions of T .

Henceforth, for any $c \in \mathbb{K}^p$ we denote by T_c the linear combination $\sum_{k=1}^p c_k Z_k$, and likewise $D_c := \sum_{k=1}^p c_k D_k$ in [Eq. \(9\)](#).

Corollary 3.9. *For a generic choice of a, b in $\mathbb{K}^p$, T_a and T_b are of rank r ; the kernel of T_a (and T_b) is of dimension $n - r$, and is equal to $\bigcap_{\ell \in [q]} \ker M_\ell$.*

Proof. The first property follows from the proof of [Proposition 3.6](#): T_a and T_b are of rank $\leq r$ for any choice of a and b , and the ranks will be equal to r if a and b avoid a union of r hyperplanes. The kernels are therefore of dimension $n - r$.

By [Eq. \(6\)](#), $T_a = \sum_{\ell} \langle a, \gamma_\ell \rangle M_\ell$. If $M_\ell x = 0$ for all ℓ , this immediately shows that $T_a x = 0$. The converse follows from the fact that $\operatorname{Im}(M_\ell)$ are in direct sum, and for a generic choice of a we have $\langle a, \gamma_\ell \rangle \neq 0$ for all ℓ . $\square$

Lemma 3.10. *Let $T = \sum_{\ell=1}^q M_\ell \otimes \gamma_\ell$ be a matrix-vector decomposition of T of rank r . The following properties hold for a generic choice of $a, b \in \mathbb{K}^p$:*

- (i) *If x is an eigenvector of the pair (T_a, T_b) then there is $\ell \in [q]$ such that $T_b x \in \text{Im } M_\ell$.*
- (ii) *There are exactly q nontrivial eigenvalues for (T_a, T_b) , and if λ is a nontrivial eigenvalue then $\langle a, \gamma_\ell \rangle = \lambda \langle b, \gamma_\ell \rangle$ for some $\ell \in [q]$. The dimension of the corresponding eigenspace V_λ is equal to $n - r + \text{rank}(M_\ell)$, and the image of V_λ by T_b is equal to $\text{Im } M_\ell$.*
- (iii) *If x is an eigenvector associated to the nontrivial eigenvalue $\lambda = \langle a, \gamma_\ell \rangle / \langle b, \gamma_\ell \rangle$, we have $T_a x = \langle a, \gamma_\ell \rangle M_\ell x$ and $T_b x = \langle b, \gamma_\ell \rangle M_\ell x$.*

Proof. Let x be an eigenvector of (T_a, T_b) associated to the eigenvalue λ . If $T_b x = 0$ then certainly $T_b x \in \text{Im } M_\ell$ for any ℓ . For the remainder of the proof of (i) we assume that $T_b x \neq 0$. By Eq. (6), $T_a = \sum_\ell \langle a, \gamma_\ell \rangle M_\ell$ and $T_b = \sum_\ell \langle b, \gamma_\ell \rangle M_\ell$. This leads to

$$\sum_\ell \langle a, \gamma_\ell \rangle M_\ell x = \sum_\ell \lambda \langle b, \gamma_\ell \rangle M_\ell x. \quad (10)$$

Remember from Corollary 3.7 that the spaces $\text{Im } M_\ell$ are in direct sum. As a result, we must have $\langle a, \gamma_\ell \rangle = \lambda \langle b, \gamma_\ell \rangle$ for all the ℓ such that $M_\ell x \neq 0$ in Eq. (10). We claim that there is in fact exactly one ℓ such that $M_\ell x \neq 0$. Note first that one cannot have $M_\ell x = 0$ for all ℓ since $T_b x \neq 0$. To continue the proof of the claim, let us assume that

$$\langle a, \gamma_\ell \rangle = \lambda \langle b, \gamma_\ell \rangle \text{ and } \langle a, \gamma_{\ell'} \rangle = \lambda \langle b, \gamma_{\ell'} \rangle$$

for two distinct indices $\ell \neq \ell'$. This implies that

$$\langle a, \gamma_\ell \rangle \langle b, \gamma_{\ell'} \rangle - \langle a, \gamma_{\ell'} \rangle \langle b, \gamma_\ell \rangle = 0. \quad (11)$$

It remains to show that this bilinear form in a and b is not identically 0: this will show that for a generic choice of a and b , (11) cannot hold. The matrix of this bilinear form is equal to $\gamma_\ell \gamma_{\ell'}^T - \gamma_{\ell'} \gamma_\ell^T$. It is indeed nonzero since γ_ℓ and $\gamma_{\ell'}$ are not colinear by definition of a matrix-vector decomposition. This completes the proof of the claim. We have therefore shown that $T_a x = \langle a, \gamma_\ell \rangle M_\ell x$ and $T_b x = \langle b, \gamma_\ell \rangle M_\ell x$, i.e., $T_a x$ and $T_b x$ both belong to $\text{Im } M_\ell$. This completes the proof of (i) and (iii); we have also shown that if λ is a nontrivial eigenvalue then $\lambda = \langle a, \gamma_\ell \rangle / \langle b, \gamma_\ell \rangle$ for some $\ell \in [q]$ (note that $\langle b, \gamma_\ell \rangle$ is generically nonzero for all ℓ). Next, we compute the dimension of the corresponding eigenspace V_λ .

Suppose for instance that $\lambda = \langle a, \gamma_q \rangle / \langle b, \gamma_q \rangle$. The proof of (i) shows that $x \in V_\lambda$ if and only if $M_\ell x = 0$ for all $\ell \neq q$. Let M be the matrix of this

linear system. We have the block decomposition $M^T = (M_1^T \dots M_{q-1}^T)$. By [Corollary 3.8](#) this implies that

$$\text{rank}(M) = \text{rank}(M^T) = \sum_{\ell=1}^{q-1} \text{rank}(M_\ell^T) = r - \text{rank}(M_q),$$

and this shows that $\dim V_\lambda = n - r + \text{rank}(M_q)$ as claimed.

Finally, we build on this computation of $\dim V_\lambda$ to show that $T_b(V_\lambda) = \text{Im } M_\ell$. From (iii) we already know that $T_b(V_\lambda) \subseteq \text{Im } M_\ell$. In order to show that these two subspaces are equal we will show that they have the same dimension. For this, we'll compute the dimension of the kernel K_λ of T_b viewed as a linear operator on V_λ . Let us continue to assume for instance that $\lambda = \langle a, \gamma_q \rangle / \langle b, \gamma_q \rangle$. Recall that $x \in V_\lambda$ if and only if $M_\ell x = 0$ for all $\ell \neq q$. Therefore, $x \in K_\lambda$ if and only if $M_\ell x = 0$ for all $\ell \in [q]$. This shows that K_λ is actually independent of λ , and is equal (by [Corollary 3.9](#)) to $\ker T_b$. Hence $\dim K_\lambda = n - r$, and $\dim T_b(V_\lambda) = \dim V_\lambda - \dim K_\lambda = \text{rank}(M_q)$. $\square$

Theorem 3.11. *Let $T = \sum_{\ell=1}^q M_\ell \otimes \gamma_\ell$ be a matrix-vector decomposition of rank r . If $T = \sum_{\ell=1}^{q'} M'_\ell \otimes \gamma'_\ell$ is another matrix-vector decomposition of rank r , then $q = q'$ and there is a permutation π such that $\text{Im } M'_\ell = \text{Im } M_{\pi(\ell)}$.*

Proof. Since the intersection of two Zariski open sets is nonempty (and even Zariski open), there exist $a, b \in \mathbb{K}^p$ such that the 3 properties of [Lemma 3.10](#) apply to our two decompositions of T . In particular, by [Lemma 3.10](#), q and q' are both equal to the number of nontrivial eigenvalues of the pair (T_a, T_b) . Moreover, the images of the corresponding eigenspaces are the spaces $\text{Im } M_\ell$, and (applying the lemma to the second decomposition of T) these images are *also* the spaces $\text{Im } M'_\ell$. This shows that the second family of spaces is obtained from the first by a permutation of indices. $\square$

With the above theorem at hand, the proof of [Theorem 1.4](#) then follows from two simple lemmas, which we now state and prove.

Lemma 3.12. *Suppose that $E_1, \dots, E_q \subseteq \mathbb{K}^m$ are in direct sum. The linear spaces of matrices:*

$$\mathcal{E}_1 = \{M \in M_{m,n}(\mathbb{K}); \text{Im}(M) \subseteq E_1\}, \dots, \mathcal{E}_q = \{M \in M_{m,n}(\mathbb{K}); \text{Im}(M) \subseteq E_q\}$$

are also in direct sum.

Proof. Suppose that $M_1 + \dots + M_q = 0$ with $M_i \in \mathcal{E}_i$ for $i = 1, \dots, q$. For any $x \in \mathbb{K}^n$, $M_1 x + \dots + M_q x = 0$ and $M_i x \in E_i$ for all i . This implies $M_i x = 0$ for all i since the E_i are in direct sum. Since this is true for any x , we conclude that $M_i = 0$ for all i . $\square$

Lemma 3.13. *If the vectors $u, u' \in \mathbb{K}^m$ and $v, v' \in \mathbb{K}^n$ satisfy $v_j u = v'_j u'$ for all $j = 1, \dots, n$ we have $u \otimes v = u' \otimes v'$.*

Proof. As matrices, $u \otimes v$ and $u' \otimes v'$ are represented respectively by uv^T and $u'v'^T$. These two matrices are equal since their respective columns are equal. $\square$

We can now complete the proof of the uniqueness theorems. For convenience, we restate [Theorem 1.4](#) here.

Theorem 1.4 (Uniqueness theorem, equivalent formulation). *Suppose that a tensor $T \in \mathbb{K}^{m \times n \times p}$ has a matrix-vector decomposition of the form:*

$$T = \sum_{\ell=1}^q M_{\ell} \otimes w_{\ell} \quad (3)$$

where the linear spaces $\text{Im}(M_{\ell})$ are in direct sum and where the linear spaces $\text{Im}(M_{\ell}^T)$ are also in direct sum. Then $\text{rank}(T) = \sum_{\ell=1}^q \text{rank}(M_{\ell})$, and (3) is the unique matrix-vector decomposition of T of minimum rank.

Proof. Consider two matrix-vector decompositions of rank r :

$$T = \sum_{\ell=1}^q M_{\ell} \otimes \gamma_{\ell} = \sum_{\ell'=1}^{q'} M_{\ell'} \otimes \gamma_{\ell'}.$$

By [Theorem 3.11](#), $q = q'$ and we can renumber the terms in this decomposition so that $\text{Im } M'_{\ell} = \text{Im } M_{\ell}$ for $\ell = 1, \dots, q$. Moreover, by [Eq. \(6\)](#) we have two expressions for each of the slices of T :

$$Z_k = \sum_{\ell=1}^q \gamma_{\ell k} M_{\ell} = \sum_{\ell=1}^q \gamma'_{\ell k} M'_{\ell}.$$

Since $\text{Im } M'_{\ell} = \text{Im } M_{\ell}$ and these spaces are in direct sum, [Lemma 3.12](#) shows that $\gamma_{\ell k} M_{\ell} = \gamma'_{\ell k} M'_{\ell}$ for all ℓ . Since this applies to all k , we have $M_{\ell} \otimes \gamma_{\ell} = M'_{\ell} \otimes \gamma'_{\ell}$ by [Lemma 3.13](#) and we have shown that T has a unique matrix-vector decomposition of minimal rank. $\square$

Remark 3.14. *There is an alternative way of proving the above uniqueness theorem via the uniqueness theorem for indecomposable vector space decompositions (the latter proved in [[GKS20](#), Corollary B.2]), which in turn uses the Krull-Schmidt theorem for modules. In short, in the alternative approach, we can show that the direct sum conditions of the given matrix-vector decomposition imply that the adjoint algebra (see [[GKS20](#), Appendix A]) of the set of 3-slices of our tensor can be block-diagonalizable in a unique way. With the above at hand, we can now prove an analogous statement to [Lemma 3.10](#) and [Theorem 3.11](#) which show that any matrix-vector decomposition of minimum rank should have the same adjoint algebra as the 3-slices. The above two facts imply uniqueness of matrix-vector decompositions.*

It is important to notice that the uniqueness of vector space decompositions alone is not enough to prove our result, but that the above is an alternative way to interpret (and prove) our uniqueness theorem.

4 Decomposition algorithm

In this section we propose and analyze an algorithm which, given an input tensor T , computes the unique matrix-vector decomposition guaranteed by [Theorems 1.3](#) and [1.4](#). We will take the point of view of [Theorem 1.4](#): assuming that there is a matrix-vector decomposition

$$T = \sum_{\ell=1}^q M_{\ell} \otimes w_{\ell} \quad (12)$$

where the linear spaces $\text{Im}(M_{\ell})$ are in direct sum and where the linear spaces $\text{Im}(M_{\ell}^T)$ are also in direct sum, we want to compute that (unique) decomposition. In [Section 4.1](#), we first show how to compute the spaces $\text{Im} M_{\ell}$. In [Section 4.2](#) we describe an algorithm for the case where the subspaces $\text{Im}(M_{\ell})$ have a very simple form, i.e., when they are coordinate subspaces. Finally, in [Section 4.3](#) we combine the algorithms from [Sections 4.1](#) and [4.2](#) to give our main algorithm.

4.1 Computing $\text{Im}(M_{\ell})$ in an arbitrary field

In this section $\mathbb{K}$ can be an arbitrary (infinite) field.⁴ As explained in [Section 1.2](#), we only need to assume that we have access to an algorithm for the computation of roots of polynomials with coefficients in $\mathbb{K}$. For the field of real and complex numbers, we provide in [Section C](#) an alternative algorithm based on the Moore-Penrose inverse, in the same style as the classical simultaneous diagonalization algorithm (see [Section A](#)).

We will compute the linear spaces $\text{Im}(M_1), \dots, \text{Im}(M_q)$ by solving a generalized eigenvalue problem⁵ using the following result.

Proposition 4.1. *Let $A = UA'V^T, B = UB'V^T$ where $U \in M_{m,r}(\mathbb{K})$, $V \in M_{n,r}(\mathbb{K})$ and $A', B' \in M_{r,r}(\mathbb{K})$. Assume that $r \leq \min(m, n)$. Then:*

- (i) *All the minors of size r of $A - \lambda B$ are scalar multiples of $\det(A' - \lambda B')$.*
- (ii) *Assume moreover that U, V, A', B' are all of rank r . Let $M_a \in M_{r,r}(\mathbb{K})$ be any submatrix of A of rank r , and let M_b be the matching submatrix of B (i.e., we select the same rows and columns as for M_a). The nontrivial eigenvalues of the pair (A, B) are exactly the roots of the polynomial $P(\lambda) = \det(M_a - \lambda M_b)$.*

Proof. A submatrix of size r of $A - \lambda B$ is of the form $M_a - \lambda M_b$ where $M_a \in M_{r,r}(\mathbb{K})$ is a submatrix of A and M_b is the matching submatrix of B . Note that

$$M_a = U' A' V', \quad M_b = U' B' V' \quad (13)$$

⁴In [Remark 4.3](#) we explain how the algorithm can be adapted to finite fields.

⁵Recall that this was the point of view of [Sections 2.1](#) and [3.2](#).

where U', V' are $r \times r$ submatrices of U and V^T . This implies:

$$P(\lambda) = \det(M_a - \lambda M_b) = (\det U')(\det V') \det(A' - \lambda B'), \quad (14)$$

and we have proved the first part of the proposition.

For the proof of (ii), we first observe that $\text{rank } A = r$ since U, A', V are all of rank r (and likewise, $\text{rank } B = r$). Hence there exists a submatrix M_a of rank r . Also, we observe that $\ker A = \ker B = \ker V^T$ and this kernel is of dimension $n - r$ since U, V, A', B' are all of full rank r . The polynomial $P(\lambda) = \det(M_a - \lambda M_b)$ is not identically 0 since $P(0) \neq 0$.

Let λ be a nontrivial eigenvalue of (A, B) and V_λ be the corresponding eigenspace. By definition, there is $v \notin \ker B$ such that $(A - \lambda B)v = 0$. Note that $\ker B \subseteq V_\lambda$ since $\ker A = \ker B$. Therefore, V_λ must be of dimension at least $n - r + 1$, i.e., $A - \lambda B$ must be of rank at most $r - 1$. Its submatrix $M_a - \lambda M_b$ must therefore be of rank at most $r - 1$, hence $P(\lambda) = 0$.

Assume conversely that $P(\lambda) = 0$. By Eq. (13), U' and V' are of rank r since M_a is of rank r . Hence $\det(A' - \lambda B') = 0$ by Eq. (14). By part (i) of the proposition, this implies $\text{rank}(A - \lambda B) \leq r - 1$. This in turn implies that λ is a nontrivial eigenvalue of (A, B) by the converse of the argument in the preceding paragraph. Namely, since $\ker(A - \lambda B)$ is of dimension at least $n - r + 1$ and $\ker B$ of dimension $n - r$, there must exist v such that $(A - \lambda B)v = 0$ and $Bv \neq 0$. $\square$

We are now ready to state the algorithm to compute the spaces $\text{Im}(M_\ell)$.

Algorithm 1: Computing the linear spaces $\text{Im}(M_1), \dots, \text{Im}(M_q)$.

Input: a tensor $T \in \mathbb{K}^{m \times n \times p}$ with an unknown matrix-vector decomposition $T = \sum_{\ell=1}^q M_\ell \otimes w_\ell$.

Output: The linear spaces $\text{Im}(M_1), \dots, \text{Im}(M_q)$.

- 1 Compute two generic linear combinations $T_a, T_b \in M_{m,n}(\mathbb{K})$ of the 3-slices of T .
 - 2 Compute $r = \text{rank}(T_a)$ and find a submatrix M_a of T_a of size r and rank r .
 - 3 Let M_b be the matching submatrix of T_b , and
 $P(\lambda) = \det(M_a - \lambda M_b)$.
 - 4 Compute the roots $\lambda_1, \dots, \lambda_q$ of P .
 - 5 Compute the corresponding eigenspaces
 $V_1 = \ker(T_a - \lambda_1 T_b), \dots, V_q = \ker(T_a - \lambda_q T_b)$.
 - 6 Output $T_b(V_1), \dots, T_b(V_q)$.
-

The next theorem proves the correctness of the above algorithm.

Theorem 4.2. Suppose that a tensor $T \in \mathbb{K}^{m \times n \times p}$ has a matrix-vector decomposition of the form:

$$T = \sum_{\ell=1}^q M_{\ell} \otimes w_{\ell}$$

where the linear spaces $\text{Im}(M_{\ell})$ are in direct sum (and where the linear spaces $\text{Im}(M_{\ell}^T)$ are also in direct sum.)

Then Algorithm 1 on input T and with vectors $a, b \in \mathbb{K}^p$ of the linear combinations computed at step 1 being generically chosen, Algorithm 1 correctly outputs the linear spaces $\text{Im}(M_1), \dots, \text{Im}(M_q)$.

Proof. By Proposition 4.1, at step 4 we compute the nontrivial eigenvalues of (T_a, T_b) . In this application of Proposition 4.1, the matrices A', B' are diagonal and of full rank since a and b are generically chosen. There are exactly q nontrivial eigenvalues by Lemma 3.10.(ii). We compute the corresponding eigenspaces at step 5, and another application of Lemma 3.10.(ii) shows that the output of step 6 is correct. $\square$

Remark 4.3. In Algorithm 1 the coefficients of the linear combination can be drawn uniformly at random from a finite set S . The proof of Theorem 4.2 reveals that these coefficients should avoid the zero sets of polynomially many polynomials of polynomially bounded degree. By the Schwartz-Zippel Lemma [Sch80, Zip79], we can make the probability of error smaller than, say, $1/3$ (or any other constant) by taking S of polynomial size. This remark also applies to finite fields: we can take $S \subseteq \mathbb{K}$ if $\mathbb{K}$ is large enough. If not, we can take the elements of S from a field extension.

A similar remark applies to the computation of $\text{Im}(M_1), \dots, \text{Im}(M_q)$ using the Moore-Penrose inverse (Proposition C.3).

4.2 Disjoint rows

The assumption that the spaces $\text{Im}(M_{\ell})$ are in direct sum clearly holds in the special case where the only nonzero rows of M_1 are its first $\text{rank}(M_1)$ rows, the only nonzero rows of M_2 are the next $\text{rank}(M_2)$ rows (namely, rows $1 + \text{rank } M_1$ to $\text{rank } M_2 + \text{rank } M_1$), and so on. In this case we say that T has the *disjoint rows property*. From Eq. (6), the slices $Z_1, \dots, Z_p$ must then have the structure:

$$Z_k = \begin{pmatrix} Z_{k1} \\ \vdots \\ Z_{kq} \\ 0 \end{pmatrix} = \begin{pmatrix} w_{1k} M'_1 \\ \vdots \\ w_{qk} M'_q \\ 0 \end{pmatrix} \quad (15)$$

where M'_{ℓ} is the $\text{rank}(M_{\ell}) \times n$ block where all the nonzero entries of M_{ℓ} are located. The block of zeros at the bottom of Z_k is present only when

$m > \text{rank}(T) = \sum_{\ell=1}^q \text{rank}(M_\ell)$; $Z_{k\ell}$ is the block of $\text{rank}(M_\ell)$ rows of Z_k which matches the corresponding block on the right-hand side of Eq. (15). Note that $Z_{k\ell} = w_{\ell k} M'_\ell$, i.e., the blocks of $Z_1, \dots, Z_p$ that are in same position in Eq. (15) are all proportional, and the $w_{\ell k}$ are the coefficients of proportionality. This leads to a very simple decomposition algorithm:

Algorithm 2: Decomposition of a tensor with disjoint rows.

Input: a tensor $T \in \mathbb{K}^{m \times n \times p}$ with the disjoint rows property, and the values $\text{rank}(M_1), \dots, \text{rank}(M_q)$.

Output: The matrix-vector decomposition $T = \sum_{\ell=1}^q M_\ell \otimes w_\ell$.

- 1 For $\ell = 1$ to q :
 - 2 Find a slice Z_k of T such that $Z_{k\ell} \neq 0$.
 - 3 Set $w_{\ell k} = 1$, $M'_\ell = Z_{k\ell}$.
 - 4 For all $j \neq k$ set $w_{\ell j}$ so that $Z_{j\ell} = w_{\ell j} M'_\ell$.
 - 5 Construct M_ℓ by putting the appropriate number of null rows above and below M'_ℓ .
 - 6 Output the decomposition $T = \sum_{\ell=1}^q M_\ell \otimes w_\ell$.
-

Proposition 4.4. *If T has the disjoint rows property, Algorithm 2 produces a correct matrix-vector decomposition $T = \sum_{\ell=1}^q M_\ell \otimes w_\ell$.*

Proof. At line 2 of the algorithm, we look for a slice Z_k where $Z_{k\ell} \neq 0$. There must be such a slice since it is assumed that T has a matrix decomposition $T = \sum_{\ell=1}^q M_\ell \otimes w_\ell$ with the disjoint rows property. Indeed, $Z_{k\ell} = 0$ for all k implies $w_\ell = 0$ or $M_\ell = 0$. This is not allowed by definition of a matrix-vector decomposition. At line 3 we set $w_{\ell k} = 1$. This is legitimate since w_ℓ and M_ℓ are only unique up to scaling. The correctness of the algorithm then follows from Eq. (15). $\square$

Remark 4.5. *In Algorithm 2, we assumed that $\text{rank}(M_1), \dots, \text{rank}(M_q)$ are given as input to the algorithm. We can do this, since we have computed these values (and more) in Section 4.1.*

4.3 Main algorithm

Our main algorithm will reduce the general case to the case of disjoint rows treated in Section 4.2. This is easy once we have determined $\text{Im}(M_1), \dots, \text{Im}(M_q)$. Indeed, we can apply a linear map which sends each $\text{Im}(M_\ell)$ to E_ℓ where E_1 is the space spanned by the first $\text{rank}(M_1)$ vectors of the canonical basis of $\mathbb{K}^m$, E_2 is spanned by the next $\text{rank}(M_2)$ vectors of this basis, etc. This is justified by Lemma 3.2. We are now in position to describe our main algorithm and prove our main algorithmic result, and we restate the latter here for convenience.

Algorithm 3: Matrix-vector decomposition algorithm.

Input: a tensor $T \in \mathbb{K}^{m \times n \times p}$ with an unknown matrix-vector decomposition $T = \sum_{\ell=1}^q M_\ell \otimes w_\ell$.

Output: The above decomposition.

- 1 Determine the linear spaces $\text{Im}(M_1), \dots, \text{Im}(M_q)$ using Algorithm 1 (or via Proposition C.3).
 - 2 Find $A \in GL_m(\mathbb{K})$ such that A maps $\text{Im}(M_1), \dots, \text{Im}(M_q)$ to the linear spaces $E_1, \dots, E_q$ defined at the beginning of Section 4.3.
 - 3 Let T' be the tensor with slices $AT_1, \dots, AT_p$, where $T_1, \dots, T_p$ are the slices of T .
 - 4 Compute a matrix-vector decomposition $T' = \sum_{\ell=1}^q N_\ell \otimes w_\ell$ with Algorithm 2.
 - 5 Output the decomposition $T = \sum_{\ell=1}^q (A^{-1}N_\ell) \otimes w_\ell$.
-

Theorem 1.5 (Matrix-vector decomposition algorithm). *Suppose that a tensor $T \in \mathbb{K}^{m \times n \times p}$ has a matrix-vector decomposition of the form:*

$$T = \sum_{\ell=1}^q M_\ell \otimes w_\ell$$

where the linear spaces $\text{Im}(M_\ell)$ are in direct sum and where the linear spaces $\text{Im}(M_\ell^T)$ are also in direct sum.

There is a randomized, polynomial-time algorithm (Algorithm 3) such that, on input T as above, it outputs the above matrix-vector decomposition (as usual, the M_ℓ and w_ℓ are determined only up to scaling and permutation).

Proof. At Step 2 we find the required matrix A since the spaces $\text{Im}(M_\ell)$ are in direct sum. The tensor T' defined at step 3 has the disjoint rows property by Lemma 3.2. Thereby, we can decompose it with Algorithm 2 at Step 4. Finally, at step 5 we undo the effect of the multiplications by A . $\square$

Remark 4.6. In Algorithm 3 we could even reduce to a block diagonal structure (i.e., disjoint rows and distinct columns) by multiplying the slices of T from the left and from the right. But as demonstrated in Section 4.2 the disjoint rows property alone is sufficient, so we only multiply from the left.

Remark 4.7. This algorithm departs from the standard version of Jennrich's algorithm (as presented in [Moi18, Section 3.3] and Section A) even in the case where the M_ℓ have rank 1 and when we use the Moore-Penrose inverse (Proposition 4.1) to determine the spaces $\text{Im}(M_\ell)$ at step 1. Indeed, we only compute one pseudo-inverse whereas the standard algorithm computes two pseudo-inverses. Moreover, we do not need to solve an overdetermined system of linear equations to find the vectors w_i (as in step 5 of Algorithm 5). Instead, their components are directly read off as coefficients of proportionality in Algorithm 2.

Remark 4.8. Suppose that $\mathbb{K} = \mathbb{Q}$ and that the coefficients of the linear combinations T_a, T_b in Algorithm 1 are chosen at random from a polynomial size set, as suggested in Remark 4.3. Then Algorithm 3 can be implemented efficiently in the Turing machine model of computation: we obtain a randomized algorithm that runs in time polynomial in the bit size of the input tensor. Indeed, at step 4 of Algorithm 1 we compute the nontrivial generalized eigenvalues of the pair (T_a, T_b) as the roots of $P(\lambda)$. By Lemma 3.10.(ii), these eigenvalues are of the form $\lambda = \langle a, \gamma_q \rangle / \langle b, \gamma_q \rangle$, i.e., they are rational numbers. But it is well known that rational roots of polynomials with rational coefficients can be computed in polynomial time, so we can compute the eigenvalues in polynomial time. The other steps of Algorithms 2, 1 and 3 are standard linear algebraic computations that run in polynomial time.

Remark 4.9. Similarly to Remark 3.14, we note here that one can also compute the minimum matrix-vector decomposition by using the indecomposable vector space decomposition algorithm as a subroutine. By block-diagonalizing the adjoint algebra and computing some generalized eigenvalues, we obtain a decomposition for each of the matrices M_ℓ , and we can proceed as we did above to compute the w_ℓ vectors.

5 The minimum rank problem

In this section we provide an application of our uniqueness theorem and algorithmic result: finding the matrices of minimum rank in certain generic vector spaces of matrices. We begin with the following proposition on the minimum ranks of certain vector spaces.

Proposition 5.1. Let V be a subspace of $M_{m,n}(\mathbb{K})$, and let $M_1, \dots, M_p$ be a basis of V . We assume that the matrices in this basis are ordered by non-decreasing rank ($\text{rank } M_i \leq \text{rank } M_{i+1}$). If the subspaces $\text{Im } M_1, \dots, \text{Im } M_p$ are in direct sum, then

$$\min\{\text{rank } M; M \in V \setminus \text{Span}(M_1, \dots, M_{i-1})\} = \text{rank } M_i, \quad \text{for all } i \geq 1.$$

Proof. Let $r_i := \min\{\text{rank } M; M \in V \setminus \text{Span}(M_1, \dots, M_{i-1})\}$. Since $M_i \in V \setminus \text{Span}(M_1, \dots, M_{i-1})$, we have $r_i \leq \text{rank } M_i$. In order to prove the converse inequality, pick any matrix $M \in V \setminus \text{Span}(M_1, \dots, M_{i-1})$. We can write $M = \sum_{j \in [i-1]} \beta_j M_j + \sum_{j \in I} \alpha_j M_j$ where $I \subseteq \{i, \dots, p\}$ is nonempty, and $\alpha_j \neq 0$ for all $j \in I$. If $x \in \ker M$ then $\sum_{j \in [i-1]} \beta_j M_j x + \sum_{j \in I} \alpha_j M_j x = 0$; this implies $M_j x = 0$ for all $j \in I$ by the direct sum assumption. Thus, $\ker M \subseteq \bigcap_{j \in I} \ker M_j$. In particular, we have $\ker M \subseteq \ker M_j$ for any $j \in I$, so that $\text{rank } M \geq \text{rank } M_j \geq \min_{1 \leq i \leq p} \text{rank } M_i$ by the rank-nullity theorem.

This completes the proof since $\text{rank } M_j \geq \text{rank } M_i$ for $j \geq i$. $\square$

5.1 Uniqueness from rank arguments

With an additional assumption, we can strengthen the conclusion of [Proposition 5.1](#). Namely, we can conclude that the only matrices of minimum rank in $V \setminus \text{Span}(M_1, \dots, M_{i-1})$ are elements of the basis $M_1, \dots, M_p$.

Proposition 5.2. *Let V be a subspace of $M_{m,n}(\mathbb{K})$, and let $M_1, \dots, M_p$ be a basis of V . We assume that the matrices in this basis are ordered by nondecreasing rank ($\text{rank } M_i \leq \text{rank } M_{i+1}$). As in [Proposition 5.1](#) we assume that the subspaces $\text{Im } M_1, \dots, \text{Im } M_p$ are in direct sum; in addition we assume that $\text{rank } N^{i,j} > \text{rank } M_i$ for all $1 \leq i < j \leq p$ where*

$$N^{i,j} = \begin{pmatrix} M_i \\ M_j \end{pmatrix} \quad (16)$$

is a $2m \times n$ matrix. Then for all $i \geq 1$,

$$\min\{\text{rank } M; M \in V \setminus \text{Span}(M_1, \dots, M_{i-1})\} = \text{rank } M_i, \quad (17)$$

and the only matrices M that achieve the minimum are scalar multiples of $M_i, M_{i+1}, \dots, M_p$.

Proof. We already know from [Proposition 5.1](#) that [Eq. \(17\)](#) holds true, so we only have to prove the second assertion. Consider therefore a matrix $M \in V \setminus \text{Span}(M_1, \dots, M_{i-1})$ of same rank as M_i , and let us write $M = \sum_j \alpha_j M_j$. There is at least one nonzero coefficient α_{i_1} in this expression with $i_1 \geq i$. Suppose that there is at least one other nonzero coefficient α_{i_2} (which may be smaller or bigger than i). As we have seen in the proof of [Proposition 2.3](#), this implies $\ker M \subseteq \ker M_{i_1} \cap \ker M_{i_2}$. Note that this intersection of kernels is defined by the linear system $N^{i_1, i_2} x = 0$, and $\text{rank } N^{i_1, i_2} > \text{rank } M_{i_1}$ by assumption. This shows that $\text{rank } M > \text{rank } M_{i_1} \geq \text{rank } M_i$, and M would therefore not be of same rank as M_i . $\square$

In particular, taking $i = 1$ in [Eq. \(17\)](#), we see that the nonzero matrices of smallest rank in V are up to scaling the matrices in the basis $M_1, \dots, M_p$ which have same rank as M_1 . We will use this fact in [Section 5.2](#) to analyze a minrank algorithm, and we will use it again in the appendix to given an alternative proof of the Jennrich uniqueness theorem. The general version of [Proposition 5.2](#) ($i \geq 1$ in [Eq. \(17\)](#)) is only provided for the sake of completeness, and we do not use it in the remainder of the paper.

Remark 5.3. *The extra assumption in [Proposition 5.2](#) is necessary, as this new hypothesis is not always satisfied when the images are in direct sum. Take for instance $M_1 = u_1 v^T$, $M_2 = u_1 v^T$. If u_1, u_2 are linearly independent and $v \neq 0$, the images are in direct sum. But $N^{1,2}$ is of the form uv^T , where u is obtained by stacking v_1 on top of v_2 . So this matrix is not of higher rank than M_1 or M_2 . It is worth noting that the extra assumption will be satisfied when $\text{Im}(M_i)$ are in direct sum and when $\text{Im}(M_i^T)$ are also in direct sum.*

5.2 The minrank algorithm

In this section we analyze the following algorithm for the computation of all matrices of minimum rank in a given subspace of matrices (up to scalar multiples). The algorithm simply computes a matrix-vector decomposition of the tensor formed by making the input basis its 3-slices.

Algorithm 4: Minrank algorithm

Input: a subspace of $V \subseteq M_{m,n}(\mathbb{K})$, given by a basis $Z_1, \dots, Z_p$.

Output: Set of matrices $A_1, \dots, A_s \in V$ of minimum rank

- 1 Construct the tensor $T \in \mathbb{K}^{m \times n \times p}$ with slices $Z_1, \dots, Z_p$.
- 2 Apply Algorithm 3 to compute a matrix-vector decomposition $T = \sum_{\ell=1}^q A_\ell \otimes w_\ell$ of minimum rank, where $A_1, \dots, A_q$ are sorted by nondecreasing rank.
- 3 Let $A_1, \dots, A_s$ be the matrices of minimum rank among $A_1, \dots, A_q$, and let $\rho = \text{rank } A_1 = \dots = \text{rank } A_s$. Declare that

$$\rho = \min_{M \in V, M \neq 0} \text{rank } M,$$

and output matrices $A_1, \dots, A_s$.

Theorem 1.6 (Minrank algorithm). *Suppose that $V \subset M_{m,n}(\mathbb{K})$ is a subspace spanned by a basis $M_1, \dots, M_p$ (the “hidden basis”) where the linear spaces $\text{Im}(M_1), \dots, \text{Im}(M_p)$ are in direct sum, and where the linear spaces $\text{Im}(M_1^T), \dots, \text{Im}(M_p^T)$ are also in direct sum.*

There is a randomized, polynomial-time algorithm (Algorithm 4) such that, when given as input any basis $Z_1, \dots, Z_p$ of V , it correctly finds the hidden basis and it outputs $\rho = \min_{M \in V, M \neq 0} \text{rank } M$, as well as matrices $A_1, \dots, A_s$ which are, up to scalar multiplication, the only matrices of rank ρ in V . Moreover, the matrices A_i are a subset of the hidden basis.

Proof. By Proposition 5.1, $\min_{M \in V, M \neq 0} \text{rank } M = \min_{1 \leq i \leq p} \text{rank } M_i$. Moreover, by the case $i = 1$ of Proposition 5.2, the only nonzero matrices in V of minimum rank are up to scalar multiplication the matrices of minimum rank in the list $M_1, \dots, M_p$. Note that the hypothesis $\text{rank } N^{i,j} > \text{rank } M_i$ in Proposition 5.2 is indeed satisfied due to the assumption that the spaces $\text{Im } M_\ell^T$ are in direct sum: we have

$$\text{rank } N^{i,j} = \text{rank}(N^{i,j})^T = \text{rank}(M_i^T) + \text{rank}(M_j^T) > \text{rank } M_i.$$

The correctness of the algorithm therefore follows from two claims which we establish in the remainder of the proof:

- (i) In the matrix-vector decomposition computed at step 2 of the algorithm, the number of terms q is equal to p .

- (ii) The matrices $A_1, \dots, A_q$ are equal up to permutation and scaling to the matrices $M_1, \dots, M_p$ of the hidden basis.

Since $Z_1, \dots, Z_p$ and $M_1, \dots, M_p$ are two bases of V , there is a change of basis matrix $W \in GL_p(\mathbb{K})$ such that $Z_k = \sum_{\ell=1}^p w_{\ell k} M_\ell$. Hence by Eq. (6), we have the decomposition

$$T = \sum_{\ell} M_\ell \otimes w_\ell, \quad (18)$$

where $w_\ell = (w_{\ell 1}, \dots, w_{\ell p})$. This is a matrix-vector decomposition since the w_ℓ are pairwise linearly independent (in fact, they are linearly independent since W is invertible). By Theorem 1.4, this is the unique matrix-vector decomposition of T of minimal rank. The decomposition produced at step 2 of the algorithm must therefore be the same as Eq. (18) up to permutation and scaling. In particular, the two decompositions have the same number of terms (thereby proving claim (i)) and $A_1, \dots, A_q$ are equal up to permutation and scaling to $M_1, \dots, M_p$ (thereby proving claim (ii)). $\square$

Remark 5.4. The decomposition $T = \sum_{\ell=1}^q A_\ell \otimes w_\ell$ computed at step 2 of the algorithm provides a certificate of correctness of its output. Indeed, from this decomposition we can easily check that $p = q$, that the linear spaces $\text{Im}(A_\ell)$ are in direct sum, and that the linear spaces $\text{Im}(A_\ell^T)$ are also in direct sum. These 3 conditions will be satisfied generically by Theorem 1.6; and whenever they are satisfied, Theorem 1.6 guarantees that the algorithm's output is correct (note in particular that the condition $p = q$ implies that $A_1, \dots, A_q$ is a basis of V).

References

- [BCS97] P. Bürgisser, M. Clausen, and M. A. Shokrollahi. *Algebraic Complexity Theory*. Springer, 1997. 9
- [DDS23] James Demmel, Ioana Dumitriu, and Ryan Schneider. Generalized pseudospectral shattering and inverse-free matrix pencil diagonalization. *arXiv preprint arXiv:2306.03700*, 2023. 9
- [DWW25] Jeshu Dastidar, Tait Weicht, and Alexander Wein. Improving the threshold for finding rank-1 matrices in a subspace. *arXiv preprint arXiv:2504.17947*, 2025. 4
- [GKS20] Ankit Garg, Neeraj Kayal, and Chandan Saha. Learning sums of powers of low-degree polynomials in the non-degenerate case. In *Proc. 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 889–899, 2020. 18
- [Har70] Richard Harshman. Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multimodal factor analysis. *UCLA working papers in phonetics*, 1970. 2, 30

- [Har72] Richard A Harshman. Determination and proof of minimum uniqueness conditions for PARAFAC1. *UCLA working papers in phonetics*, 22(111-117):3, 1972. 2, 30
- [Hås90] Johan Håstad. Tensor rank is NP-complete. *Journal of Algorithms*, 11(4):644–654, 1990. 2
- [HSS19] Samuel Hopkins, Tselil Schramm, and Jonathan Shi. A robust spectral algorithm for overcomplete tensor decomposition. In *Proc. 32nd Annual Conference on Learning Theory (COLT)*, 2019. 2
- [JLV23] Nathaniel Johnston, Benjamin Lovitz, and Aravindan Vijayaraghavan. Computing linear sections of varieties: quantum entanglement, tensor decompositions and beyond. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1316–1336. IEEE, 2023. 2, 4, 7, 8
- [KMW25] Pravesh Kothari, Ankur Moitra, and Alexander Wein. Overcomplete tensor decomposition via Koszul-Young flattenings. In *Symposium on Foundations of Computer Science (FOCS)*, 2025. <https://arxiv.org/abs/2411.14344>. 2, 7
- [Koi25] Pascal Koiran. An efficient uniqueness theorem for overcomplete tensor decomposition. *Proc. ACM-SIAM Annual Symposium on Discrete Algorithms (SODA)*, 2025. <https://arxiv.org/abs/2404.07801>. 2, 5, 7, 12, 14
- [Kru77] Joseph Kruskal. Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics. *Linear Algebra and its Applications*, 18(2):95–138, 1977. 2, 30
- [KS23] Pascal Koiran and Subhayan Saha. Complete decomposition of symmetric tensors in linear time and polylogarithmic precision. In *CIAC (International Conference on Algorithms and Complexity)*, 2023. <https://arxiv.org/abs/2211.07407>. Journal version in TCS special issue on CIAC 2023 (published online in May 2025). 29
- [KS24] Pascal Koiran and Subhayan Saha. Undercomplete decomposition of symmetric tensors in linear time, and smoothed analysis of the condition number. <https://arxiv.org/abs/2403.00643>, 2024. 29, 33
- [LRA93] Sue Leurgans, Robert Ross, and Rebecca Abel. A decomposition for three-way arrays. *SIAM Journal on Matrix Analysis and Applications*, 14(4):1064–1083, 1993. 2, 10, 32
- [Moi18] Ankur Moitra. *Algorithmic aspects of machine learning*. Cambridge University Press, 2018. 6, 12, 23, 29, 30, 32
- [MSS16] Tengyu Ma, Jonathan Shi, and David Steurer. Polynomial-time tensor decompositions with sum-of-squares. In *Proc. 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 438–446, 2016. 2
- [Rho10] John Rhodes. A concise proof of Kruskal’s theorem on tensor decomposition. *Linear Algebra and its Applications*, 432(7):1818–1824, 2010. 30

- [Sch80] J. T. Schwartz. Fast probabilistic algorithms for verification of polynomials identities. *Journal of the ACM*, 27:701–717, 1980. 21
- [VD79] Paul Van Dooren. The computation of Kronecker’s canonical form of a singular pencil. *Linear Algebra and its Applications*, 27:103–140, 1979. 9
- [Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *Symbolic and algebraic computation*, pages 216–226. Springer, 1979. 21

A The simultaneous diagonalization algorithm

In this section we recall the simultaneous diagonalization / Jennrich algorithm following [Moi18, Section 3.3]. It provides an efficient decomposition algorithm for generic tensors of rank $r \leq \min(m, n)$. We assume that K is the field of real or complex numbers since this version of the algorithm uses the Moore-Penrose inverse.

Algorithm 5: decomposition by simultaneous diagonalization (sometimes called "Jennrich’s algorithm").

Input: a tensor $T \in K^{m \times n \times p}$ satisfying the conditions of Theorem 1.1.

Output: the (unique) decomposition $T = \sum_{i=1}^r u_i \otimes v_i \otimes w_i$.

- 1 Compute two generic linear combinations

$$T_a = \sum_{k=1}^p a_k T_k, \quad T_b = \sum_{k=1}^p b_k T_k$$

of the slices of T .

- 2 Compute the nonzero eigenvalues $\lambda_1, \dots, \lambda_r$ and the corresponding eigenvectors $u_1, \dots, u_r$ of $T_a T_b^\dagger$.
 - 3 Compute the nonzero eigenvalues $\mu_1, \dots, \mu_r$ and the corresponding eigenvectors $v_1, \dots, v_r$ of $(T_a^\dagger T_b)^T$.
 - 4 Reorder these eigenvectors and their eigenvalues to make sure that the corresponding eigenvalues are reciprocal (i.e., $\lambda_i \mu_i = 1$).
 - 5 Solve for w_i in the linear system $T = \sum_{i=1}^r u_i \otimes v_i \otimes w_i$, output this decomposition.
-

An analysis of this algorithm can be found in [Moi18, Section 3.3], where it is called “Jennrich’s algorithm.” In particular, it can be shown that with high probability over the choice of the coefficients a_k, b_k , each of the matrices at steps 2 and 3 have exactly r distinct nonzero eigenvalues. Moreover, these eigenvalues are reciprocal (refer to step 4). An optimized version of the simultaneous diagonalization algorithm (and a detailed complexity analysis) for the special case of symmetric tensors can be found in [KS23, KS24].

B Jennrich's uniqueness theorem from rank arguments

Nowadays, the best known proof of the Jennrich uniqueness is probably the spectral one. It has the advantage of yielding an efficient tensor decomposition algorithm based on simultaneous diagonalization (see [Moi18, Chapter 3], which emphasizes the algorithmic point of view). This uniqueness theorem can be traced back (in a slightly less general form) to Harshman [Har70], where it is attributed to Jennrich. Another version of the uniqueness theorem appears in a second paper by Harshman [Har72], and the proof seems closer to a simultaneous diagonalization argument. Jennrich's uniqueness theorem also follows from the (more involved) Kruskal uniqueness theorem [Kru77, Rho10]. In this appendix we give a proof which builds on the rank arguments from Section 5.

Lemma B.1. *If the vectors $w_1, \dots, w_q$ in (2) are linearly independent, the span of the matrices $Z_1, \dots, Z_p$ is equal to the span of $M_1, \dots, M_q$.*

Proof. It follows immediately from (6) that the span of the Z_k is included in the span of the M_ℓ , and for this no hypothesis on the w_ℓ is needed.

For the converse first observe that the entries of T are given by the formula:

$$T_{ijk} = \sum_{\ell=1}^q (M_\ell)_{ij} w_{\ell k}. \quad (19)$$

Let us denote by t_{ij} and m_{ij} the column vectors of size p and q with respective entries $(T_{ijk})_{1 \leq k \leq p}$ and $((M_\ell)_{ij})_{1 \leq \ell \leq q}$. We can rewrite (19) as the matrix-vector product $t_{ij} = W m_{ij}$ where W has $w_1, \dots, w_q$ as column vectors. Since these vectors are linearly independent, there is a $q \times p$ matrix W' such that $W'W = I_q$. From this we obtain $m_{ij} = W' t_{ij}$ and the identity

$$M_\ell = \sum_{k=1}^p w'_{k\ell} Z_k,$$

which is converse to (6). □

Lemma B.2. *Let $M \in M_{2m,n}(K)$ be a matrix of the form*

$$M = \begin{pmatrix} uv^T \\ u'v'^T \end{pmatrix}$$

where $u, u' \in K^m$ and $v, v' \in K^n$. If the two vectors u, u' are linearly independent, and the two vectors v, v' are also linearly independent, then $\text{rank } M = 2$.

Proof. In order to show that $\text{rank } M = 2$, we'll show that $\dim \ker M = n - 2$. A vector $x \in K^n$ belongs to $\ker M$ if and only if $(v^T x)u + (v'^T x)u' = 0$. Since u, u' are linearly independent, this is equivalent to $v^T x = v'^T x = 0$. Since v, v' are linearly independent, the space of solutions of this linear system has dimension $n - 2$. $\square$

Theorem B.3 (Jennrich's uniqueness theorem). *Let $T = \sum_{i=1}^r u_i \otimes v_i \otimes w_i$ be a tensor of format $m \times n \times p$ such that:*

- (i) *The vectors u_i are linearly independent.*
- (ii) *Every pair of vectors in the set $\{v_i; 1 \leq i \leq r\}$ is linearly independent.*
- (iii) *The vectors w_i are linearly independent.*

Then $\text{rank}(T) = r$, and the decomposition of T as a sum of r rank one tensors is essentially unique.

Compared to the usual statement of this uniqueness theorem, we have switched the roles of the v_i and w_i : it is usually assumed that the w_i (instead of the v_i) are pairwise linearly independent (see Theorem 1.1). The proof below hinges on the fact that $u_1 \otimes v_1, \dots, u_r \otimes v_r$ are up to scaling the only matrices of rank 1 in the span of the 3-slices of T . With the usual statement of the uniqueness theorem, one would have to work with the 1-slices or the 2-slices instead of the 3-slices.

Proof. Let V be the span of the 3-slices of T . By Lemma B.1, $V = \text{Span}(u_1 v_1^T, \dots, u_r v_r^T)$. Since the u_i are linearly independent and the v_i nonzero, these r matrices are linearly independent. In particular, $\dim V = r$ and V contains r matrices of rank 1. Furthermore, by Lemma B.2 we can apply Theorem 5.2 to V and to its basis $M_1 = u_1 v_1^T, \dots, M_r = u_r v_r^T$. As a result, taking $i = 1$ in (17), we conclude that $M_1, \dots, M_r$ are up to scaling the only rank-1 matrices in V .

Consider now any other decomposition $T = \sum_{i=1}^{r'} u'_i \otimes v'_i \otimes w'_i$. Note that $V \subseteq \text{Span}(u'_1 v'^T_1, \dots, u'_{r'} v'^T_{r'})$, but we have seen that $\dim V = r$. Hence $r' \geq r$, and we have shown that $\text{rank}(T) = r$.

For the remainder of the proof we will assume that $r' = r$; it remains to show that the rank-1 tensors $u_i \otimes v_i \otimes w_i$ are up to permutation the same as $u'_i \otimes v'_i \otimes w'_i$. We have just seen that $V \subseteq \text{Span}(u'_1 v'^T_1, \dots, u'_r v'^T_r)$. Since $\dim V = r$, this must be an equality: $V = \text{Span}(u'_1 v'^T_1, \dots, u'_r v'^T_r)$, i.e., the matrices $u'_1 v'^T_1, \dots, u'_r v'^T_r$ form a basis of V . But we have seen that $u_1 v_1^T, \dots, u_r v_r^T$ are up to scaling the only rank-1 matrices in V . We conclude that $u_i v_i^T = u'_i v'^T_i$, up to scaling and permutation. Finally, we observe that there is a unique way of writing each 3-slice of T as a linear combination of the $u_i v_i^T$ since these matrices form a basis of V . This establishes the uniqueness of the vectors $w_1, \dots, w_r$, and completes the proof. $\square$

C Computing $\text{Im}(M_\ell)$ with the Moore-Penrose inverse

In this section we assume that K is the field of real or complex numbers. We recall the following properties of the Moore-Penrose inverse.

Proposition C.1. *Consider two matrices $A \in M_{m,n}(K)$, $B \in M_{n,p}(K)$. Their Moore-Penrose inverses $A^\dagger \in M_{n,m}(K)$, $B^\dagger \in M_{p,n}(K)$ satisfy the following properties:*

- (i) *If A has linearly independent columns, $A^\dagger A = I_n$.*
- (ii) *If B has linearly independent rows, $BB^\dagger = I_m$.*
- (iii) *If A has linearly independent columns or B has linearly independent rows, then $(AB)^\dagger = B^\dagger A^\dagger$.*

We will switch back and forth between the point of views of Theorems 1.3 and 1.4. Recall that to go from the latter (the matrix-vector point of view) to the former (the viewpoint of “ordinary rank decompositions”), we just need to write each matrix M_ℓ as the following sum of $\text{rank}(M_\ell)$ matrices of rank 1:

$$M_\ell = \sum_i u_i \otimes v_i. \quad (20)$$

Note that $\text{Im}(M_\ell)$ is the span of the vectors u_i occurring in (20). As in Section 3.2, for any $c \in K^p$ we denote by T_c the linear combination of slices $\sum_{k=1}^p c_k Z_k$, and D_c denotes the linear combination $\sum_{k=1}^p c_k D_k$ in (9). By (7), $T_c = U D_c V^T$. The computations that follow are reminiscent of the classical treatment of Jennrich’s algorithm as in, e.g., [Moi18] (see also [LRA93]).

Lemma C.2. *For any $a \in K^p$ and a generically chosen $b \in K^p$, $T_a T_b^\dagger = U D_a D_b^{-1} U^\dagger$.*

Proof. For a generically chosen $b \in K^p$, $\text{rank}(D_b) = r$ as shown in the proof of Proposition 3.6. This implies that $D_b V^T$ has linearly independent rows. Since U has linearly independent columns, Proposition C.1.(iii) implies $T_b^\dagger = (D_b V^T)^\dagger U^\dagger$. By the same token, since D_b has linearly independent columns and V^T linearly independent rows, $(D_b V^T)^\dagger = (V^T)^\dagger D_b^\dagger = (V^T)^\dagger D_b^{-1}$.

To conclude, we multiply by $T_a = U D_a V^T$ and use Proposition C.1.(ii): $V^T (V^T)^\dagger = I_r$ since the r rows of V are linearly independent. $\square$

Proposition C.3. *For generically chosen $a, b \in K^p$, $T_a T_b^\dagger$ has exactly q distinct nonzero eigenvalues and the corresponding eigenspaces are $\text{Im}(M_1), \dots, \text{Im}(M_q)$.*

Proof. We first check that the columns of U (i.e., the u_i) are eigenvectors of $T_a T_b^\dagger$: by the previous lemma, $T_a T_b^\dagger U = (U D_a D_b^{-1} U^\dagger) U = U D_a D_b^{-1}$. Here we use the fact that $U^\dagger U = I_r$ since the r columns of U are linearly independent (Proposition C.1.(i)). Note that u_i is associated to some eigenvalue $\lambda = \langle w_\ell, a \rangle / \langle w_\ell, b \rangle$ where $w_1, \dots, w_q$ are the q distinct “third mode vectors” occurring in a decomposition of T . This eigenvalue is nonzero for a generic $a \in K^p$. Moreover, by (20), $\text{Im } M_\ell$ is included in the corresponding eigenspace V_λ . For generically chosen $a, b \in K^p$, these q eigenvalues are distinct. Like in the analysis of the standard version of Jennrich’s algorithm, this follows from the fact that the w_ℓ are pairwise linearly independent (for a detailed argument, see the proof of Lemma 3.10.(ii)).

In order to complete the proof of the proposition, we still need to derive the converse inclusion ($V_\lambda \subseteq \text{Im } M_\ell$) and we need to show that 0 is the only possible other eigenvalue. We will in fact show that 0 is an eigenvalue of multiplicity $m - r$, which achieves these two goals at once. For this, recall that $\text{rank } U^\dagger = \text{rank } U = r$, hence $\dim \ker U^\dagger = m - r$. Moreover, $\ker U^\dagger \subseteq \ker T_a T_b^\dagger$ by Lemma C.2. This shows that 0 has multiplicity at least $m - r$ as an eigenvalue. This is in fact the exact value of the multiplicity since we have already found other eigenvalues (the $\langle w_\ell, a \rangle / \langle w_\ell, b \rangle$) whose multiplicities sum at least to r . $\square$

It is possible to eliminate the Moore-Penrose inverse from the above computations (and also from the ordinary Jennrich algorithm), and to replace it by the ordinary matrix inverse. We briefly sketch how to do so. First we can obtain the span of the u_i as $\text{Im}(T_a)$ for a generic $a \in K^p$. We can likewise obtain the span of the v_i as $\text{Im}(T_a^T)$. Then we can perform a change of basis to reduce to the case of a tensor T' of format $r \times r \times p$ and rank r . For such a tensor, the Moore-Penrose inverse can be replaced by the ordinary inverse. For the symmetric version of Jennrich’s algorithm, this approach is worked out in detail (with numerical error bounds) in [KS24], see in particular Section 1.3.1 of that paper.

The modifications above result in an algorithm that is applicable to an arbitrary field (not just to the real and complex numbers). We have already presented such an algorithm in Section 4.1. One difference is that the algorithm of Section 4.1 computes the spaces $\text{Im}(M_\ell)$ “in one go” whereas the algorithm that we have just sketched first computes the span of the u_i (i.e., the direct sum of the $\text{Im}(M_\ell)$) before computing the spaces $\text{Im}(M_\ell)$ themselves in a second stage.