A nearly optimal algorithm to decompose binary forms

Matías R Bender, Jean-Charles Faugère, Ludovic Perret, Elias Tsigaridas

A nearly optimal algorithm to decompose binary forms

Matías R. Bender

Sorbonne Université, CNRS, INRIA, *Laboratoire d'Informatique de Paris 6,* LIP6, *Équipe* POLSYS, *4 place Jussieu, F-75005, Paris, France*

Jean-Charles Faugère

Sorbonne Université, CNRS, INRIA, *Laboratoire d'Informatique de Paris 6,* LIP6, *Équipe* POLSYS, *4 place Jussieu, F-75005, Paris, France*

Ludovic Perret

Sorbonne Université, CNRS, INRIA, *Laboratoire d'Informatique de Paris 6,* LIP6, *Équipe* POLSYS, *4 place Jussieu, F-75005, Paris, France*

Elias Tsigaridas

Sorbonne Université, CNRS, INRIA, *Laboratoire d'Informatique de Paris 6,* LIP6, *Équipe* POLSYS, *4 place Jussieu, F-75005, Paris, France*

Abstract

Symmetric tensor decomposition is an important problem with applications in several areas, for example signal processing, statistics, data analysis and computational neuroscience. It is equivalent to Waring's problem for homogeneous polynomials, that is to write a homogeneous polynomial in n variables of degree D as a sum of D-th powers of linear forms, using the minimal number of summands. This minimal number is called the *rank* of the polynomial/tensor. We focus on decomposing binary forms, a problem that corresponds to the decomposition of symmetric tensors of dimension 2 and order D, that is, symmetric tensors of order D over the vector space $\mathbb{K}^2$. Under this formulation, the problem finds its roots in invariant theory where the decompositions are related to canonical forms.

We introduce a *superfast* algorithm that exploits results from *structured linear algebra*. It achieves a *softly linear* arithmetic complexity bound. To the best of our knowledge, the previously known algorithms have at least quadratic complexity bounds. Our algorithm computes a symbolic decomposition in $O(\mathtt{M}(D)\log(D))$ arithmetic operations, where $\mathtt{M}(D)$ is the complexity of multiplying two polynomials of degree D. It is deterministic when the decomposition is unique. When the decomposition is not unique, it is randomized. We also present a Monte Carlo

variant as well as a modification to make it a Las Vegas one.

From the symbolic decomposition, we approximate the terms of the decomposition with an error of $2^{-\varepsilon}$, in $O\big(D\log^2(D)\big(\log^2(D)+\log(\varepsilon)\big)\big)$ arithmetic operations. We use results from Kaltofen and Yagati (1989) to bound the size of the representation of the coefficients involved in the decomposition and we bound the algebraic degree of the problem by $\min(\mathit{rank}, D-\mathit{rank}+1)$. We show that this bound can be tight. When the input polynomial has integer coefficients, our algorithm performs, up to poly-logarithmic factors, $\widetilde{O}_B(D\ell+D^4+D^3\tau)$ bit operations, where τ is the maximum bitsize of the coefficients and $2^{-\ell}$ is the relative error of the terms in the decomposition.

Keywords: Decomposition of binary forms; Tensor decomposition; Symmetric tensor; Symmetric tensor rank; Polynomial Waring's problem; Waring rank; Hankel matrix; Algebraic degree; Canonical form;

1. Introduction

The problem of decomposing a symmetric tensor consists in writing it as the sum of rank-1 symmetric tensors, using the minimal number of summands. This minimal number is known as the rank of the symmetric tensor[1]. The symmetric tensors of rank-1 correspond to, roughly speaking, the D-th outer-product of a vector. The decomposition of symmetric tensor is a common problem which appears in divers areas such as signal processing, statistics, data mining, computational neuroscience, computer vision, psychometrics, chemometrics, among others. For a modern introduction to the theory of tensor, their decompositions and applications we refer to e.g., Comon (2014); Landsberg (2012).

There is an equivalence between decomposing symmetric tensors and solving Waring's problem for homogeneous polynomials, e.g., Comon et al. (2008); Helmke (1992). Given a symmetric tensor of dimension n and order D, that is a symmetric tensor of order D over the vector space

Email addresses: matias.bender@inria.fr (Matías R. Bender), jean-charles.faugere@inria.fr (Jean-Charles Faugère), ludovic.perret@lip6.fr (Ludovic Perret), elias.tsigaridas@inria.fr (Elias Tsigaridas)

URL: http://www-polsys.lip6.fr/~bender/ (Matías R. Bender), http://www-polsys.lip6.fr/~jcf/ (Jean-Charles Faugère), http://www-polsys.lip6.fr/~perret/ (Ludovic Perret), http://www-polsys.lip6.fr/~elias/ (Elias Tsigaridas)

[1] Some authors, e.g., Comon et al. (2008), refer to this number as the symmetric rank of the tensor.

$\mathbb{K}^n$, we can construct a homogeneous polynomial in n variables of degree D. We can identify the symmetric tensors of rank-1 with the D-th power of linear forms. Hence, to decompose a symmetric tensor of order D is equivalent to write the corresponding polynomial as a sum of D-th powers of linear forms using the minimal numbers of summands. This minimal number is the rank of the polynomial/tensor.

Under this formulation, symmetric tensor decomposition dates back to the origin of modern (linear) algebra as a part of Invariant Theory. In this setting, the decomposition of generic symmetric tensors corresponds to canonical forms (Sylvester, 1904, 1851; Gundelfinger, 1887). Together with the theory of apolarity, this problem was of great importance because the decompositions provide information about the behavior of the polynomials under linear change of variables (Kung and Rota, 1984).

Binary Form Decomposition. We study the decomposition of symmetric tensors of order D and dimension 2. In terms of homogeneous polynomials, we consider a binary form

$$f(x,y) := \sum_{i=0}^{D} \binom{D}{i} a_i x^i y^{D-i}, \tag{1}$$

where $a_i \in \mathbb{K} \subset \mathbb{C}$ and $\mathbb{K}$ is some field of characteristic zero. We want to compute a decomposition

$$f(x,y) = \sum_{j=1}^{r} (\alpha_j x + \beta_j y)^D, \tag{2}$$

where $\alpha_1, \ldots, \alpha_r, \beta_1, \ldots, \beta_r \in \overline{\mathbb{K}}$, with $\overline{\mathbb{K}}$ being the algebraic closure of $\mathbb{K}$, and r is minimal. We say that a decomposition *unique* if, for all the decompositions, the set of points $\{(\alpha_j, \beta_j) : 1 \leq j \leq r\} \subset \mathbb{P}^1(\overline{\mathbb{K}})$ is unique, where $\mathbb{P}^1(\overline{\mathbb{K}})$ is the projective space of $\overline{\mathbb{K}}$ (Reznick, 2013a).

Previous work. The decomposition of binary forms, Equation (2), has been studied extensively for $\mathbb{K} = \mathbb{C}$. More than one century ago Sylvester (1851, 1904) described the necessary and sufficient conditions for a decomposition to exist, see Section 2.1. He related the decompositions to the kernel of Hankel matrices. For a modern approach of this topic, we refer to Kung and Rota (1984); Kung (1990); Reznick (2013a); Iarrobino and Kanev (1999). Sylvester's work was extended to a more general kind of polynomial decompositions that we do not consider in this work, e.g., Gundelfinger (1887); Reznick (1996); Iarrobino and Kanev (1999).

Sylvester's results lead to an algorithm (Algorithm 1) to decompose binary forms (see Comon and Mourrain, 1996, Sec. 3.4.3). In the case where the binary form is of odd degree, then we can compute the

decompositions using Berlekamp-Massey algorithm (see Dür, 1989). When the decomposition is unique, the Catalecticant algorithm, which also works for symmetric tensors of bigger dimension (Iarrobino and Kanev, 1999; Oeding and Ottaviani, 2013), improves Sylvester's work. For an arbitrary binary form, Helmke (1992) presented a randomized algorithm based on Padé approximants and continued fractions, in which he also characterized the different possible decompositions. Unfortunately, all these algorithms have complexity at least quadratic in the degree of the binary form.

Besides the problem of computing the decomposition(s) many authors considered the subproblems of computing the rank and deciding whether there exists a unique decomposition, e.g., Sylvester (1851, 1904); Helmke (1992); Comas and Seiguer (2011); Bernardi et al. (2011). For example, Sylvester (1851, 1904) considered generic binary forms, that is binary forms with coefficients belonging to a dense algebraic open subset of $\overline{\mathbb{K}}^{D+1}$ (Comon and Mourrain, 1996, Section 3), and proved that when the degree is $2k$ or $2k+1$, for $k \in \mathbb{N}$, the rank is $k+1$ and that the minimal decomposition is unique only when the degree is odd. In the non-generic case, Helmke (1992); Comas and Seiguer (2011); Iarrobino and Kanev (1999), among others, proved that the rank is related to the kernel of a Hankel matrix and that the decomposition of a binary form of degree $2k$ or $2k-1$ and rank r, is unique if and only if $r \leq k$. With respect to the problem of computing the rank there are different variants of algorithms, e.g., Comas and Seiguer (2011); Comon et al. (2008); Bernardi et al. (2011). Even though there are not explicit complexity estimates, by exploiting recent superfast algorithms for Hankel matrices (Pan, 2001), we can deduce a nearly-optimal arithmetic complexity bound for computing the rank using the approach of Comas and Seiguer (2011).

For the general problem of symmetric tensor decomposition, Sylvester's work was successfully extended to cases in which the decomposition is unique, e.g., Brachat et al. (2010); Oeding and Ottaviani (2013). There are also homotopy techniques to solve the general problem, e.g., to decompose generic symmetric tensors (Bernardi et al., 2017) or, when there is a finite number of possible decompositions and we know at least one of them, to compute all the other decompositions (Hauenstein et al., 2016). There are no complexity estimations for these methods. Besides tensor decomposition, there are other related decompositions for binary forms and univariate polynomials that we do not consider in this work, e.g., Reznick (1996, 2013b); Giesbrecht et al. (2003); Giesbrecht and Roche (2010); García-Marco et al. (2017).

Formulation of the problem. Instead of decomposing the binary form as in Equation (2), we compute $\lambda_1 \dots \lambda_r$, $\alpha_1 \dots \alpha_r$, $\beta_1 \dots \beta_r \in \overline{\mathbb{K}}$, where r is minimal, such that,

$$f(x,y) = \sum_{j=1}^{r} \lambda_j (\alpha_j x + \beta_j y)^D. \tag{3}$$

Since every λ_j belongs to the algebraic closure of the field $\mathbb{K}$, the problems are equivalent. This approach allows us to control the algebraic degree (Bajaj, 1988; Nie et al., 2010) of the parameters λ_j, α_j, and β_j in the decomposition (Section 4.1).

Note that if the field is not algebraically closed and we force the parameters to belong to the base field, that is $\lambda_j, \alpha_j, \beta_j \in \mathbb{K}$, the decompositions induced by Equation (2) and Equation (3) are not equivalent. We do not consider the latter case and we refer to Helmke (1992); Reznick (1992); Comon et al. (2008); Boij et al. (2011); Blekherman (2015) for $\mathbb{K} = \mathbb{R}$, and to Reznick (1996, 2013a); Reznick and Tokcan (2017) for $\mathbb{K} \subset \mathbb{C}$.

Main results. We extend Sylvester's algorithm to achieve a nearly-optimal complexity bound in the degree of the binary form. By considering structural properties of the Hankel matrices, we restrict the possible values for the rank of the decompositions and we identify when the decomposition is unique. We build upon Helmke (1992) and we use the Extended Euclidean Algorithm to deduce a better complexity estimate than what was previously known. Similarly to Sylvester's algorithm, our algorithm decomposes successfully any binary form, without making any assumptions on the input.

First, we focus on *symbolic decompositions*, that is a representation of the decomposition as a sum of a rational function evaluated at the roots of a univariate polynomial (Definition 36). We introduce an algorithm to compute a symbolic decomposition of a binary form of degree D in $O(\mathtt{M}(D)\log(D))$, where $\mathtt{M}(D)$ is the arithmetic complexity of polynomial multiplication (Theorem 43). When the decomposition is unique, the algorithm is deterministic and this is a worst case bound. When the decomposition is not unique, our algorithm makes some random choices to fulfill certain genericity assumptions; thus the algorithm is a Monte Carlo one. However, we can verify if the genericity assumptions hold within the same complexity bound, that is $O(\mathtt{M}(D)\log(D))$, and hence we can also deduce a Las Vegas variant of the algorithm.

Following the standard terminology used in structured matrices (Pan, 2001), our algorithm is *superfast* as its arithmetic complexity matches the size of the input up to poly-logarithmic factors. The symbolic decomposition allow us to approximate the terms in a decomposition,

with a relative error of $2^{-\varepsilon}$, in $O\big(D\log^2(D)\big(\log^2(D)+\log(\varepsilon)\big)\big)$ arithmetic operations (Pan, 2002; McNamee and Pan, 2013). Moreover, we can deduce for free the rank and the border rank of the tensor, see (Comas and Seiguer, 2011, Section 1).

Using results from Kaltofen and Yagati (1989), we bound the algebraic degree of the decompositions by $\min(\text{rank}, D-\text{rank}+1)$ (Theorem 28). Moreover, we prove lower bounds for the algebraic degree of the decomposition and we show that in certain cases the bound is tight (Section 4.1.3). For polynomials with integer coefficients, we bound the bit complexity, up to polylogarithmic factors, by $\widetilde{O}_B(D\ell+D^4+D^3\tau)$, where τ is the maximum bitsize of the coefficients of the input binary form and $2^{-\ell}$ is the error of the terms in the decomposition (Theorem 45). This Boolean worst case bound holds independently of whether the decomposition is unique or not.

This work is an extension of the conference paper (Bender et al., 2016). With respect to the conference version, our main algorithm (Algorithm 3) omits an initial linear change of coordinates as we now rely on fewer genericity assumptions. In contrast with our previous algorithm, we present an algorithm which is deterministic when the decomposition is unique (Theorem 43). When the decomposition is not unique, our algorithm is still randomized but we present bounds for the number of bad choices that it could make (Proposition 29). With respect to the algebraic degree of the problem, we study the tightness of the bounds that we proposed in the conference paper (Theorem 27). We introduce explicit lower bounds showing that our bounds can be tight (Section 4.1.3).

Organization of the paper. First, we introduce the notation. In Section 2, we present the preliminaries that we need for introducing our algorithm. We present Sylvester's algorithm (Section 2.1), we recall some properties of the structure of the kernel of the Hankel matrices (Section 2.2), we analyze its relation to rational reconstructions of series/polynomials (Section 2.3), and we present the Extended Euclidean Algorithm (Section 2.4). Later, in Section 3, we present our main algorithm to decompose binary forms (Algorithm 3) and its proof of correctness (Section 3.3). This algorithm uses Algorithm 4 to compute the kernel of a family of Hankel matrices, which we consider in Section 3.1. Finally, in Section 4, we study the algebraic degree of the problem (Section 4.1), we present tight bounds for it (Section 4.1.3), and we analyze the arithmetic (Section 4.2) and bit complexity of Algorithm 3 (Section 4.3).

Notation. We denote by O, respectively O_B, the arithmetic, respectively bit, complexity and we use $\widetilde{O}$, respectively $\widetilde{O}_B$, to ignore (poly-)logarithmic factors. $\mathtt{M}(n)$ is the arithmetic complexity of multiplying two polynomial of degree n. Let $\mathbb{K}$ be a zero characteristic subfield of $\mathbb{C}$, and $\overline{\mathbb{K}}$ its algebraic closure. If $v = (v_0, \ldots, v_n)^\mathsf{T}$ then $P_v = P_{(v_0,\ldots,v_n)} := \sum_{i=0}^{n} v_i x^i y^{n-i}$. Given a binary form $f(x,y)$, we denote by $f(x)$ the univariate polynomial $f(x) := f(x,1)$. By $f'(x)$ we denote the derivative of $f(x)$ with respect to x. For a matrix M, $\mathtt{rk}(M)$ is its rank and $\mathtt{Ker}(M)$ its kernel.

2. Preliminaries

2.1. An algorithm based on Sylvester's theorem

Sylvester's theorem (Theorem 2) relates the minimal decomposition of a binary form to the kernel of a Hankel matrix. Moreover, it implies an (incremental) algorithm for computing the minimal decomposition. The version that we present in Algorithm 1 comes from Comon and Mourrain (1996, Section 3.2).

Definition 1. *Given a vector $a = (a_0, \ldots, a_D)^\mathsf{T}$, we denote by $\{H_a^k\}_{1\le k\le D}$ the family of Hankel matrices indexed by k, where $H_a^k \in \mathbb{K}^{(D-k+1)\times(k+1)}$ and*

$$H_a^k := \begin{pmatrix} a_0 & a_1 & \cdots & a_{k-1} & a_k \\ a_1 & a_2 & \cdots & a_k & a_{k+1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ a_{D-k-1} & a_{D-k} & \cdots & a_{D-2} & a_{D-1} \\ a_{D-k} & a_{D-k+1} & \cdots & a_{D-1} & a_D \end{pmatrix}. \tag{4}$$

We may omit the index a in H_a^k when it is clear from the context.

Theorem 2 (Sylvester, 1851)**.** *Let $f(x,y) = \sum_{i=0}^{D} \binom{D}{i} a_i x^i y^{D-i}$ with $a_i \in \mathbb{K} \subseteq \mathbb{C}$. Also, consider a non-zero vector $c = (c_0, \ldots, c_r)^\mathsf{T} \in \mathbb{K}^{r+1}$, such that the polynomial*

$$P_c = \sum_{i=0}^{r} c_i x^i y^{r-i} = \prod_{j=1}^{r} (\beta_j x - \alpha_j y)$$

is square-free and $\alpha_j, \beta_j \in \overline{\mathbb{K}}$, for all $1 \le j \le r$. Then, there are $\lambda_1, \ldots \lambda_r \in \overline{\mathbb{K}}$ such that we can write $f(x,y)$ as

$$f(x,y) = \sum_{j=1}^{r} \lambda_j (\alpha_j x + \beta_j y)^D,$$

if and only if $(c_0, \ldots, c_r)^\mathsf{T} \in \mathtt{Ker}(H_a^r)$.

Algorithm 1 INCRDECOMP (Comon and Mourrain, 1996, Figure 1)

1. $r := 1$
2. Get a random $c \in \mathtt{Ker}(H^r)$
3. If P_c is not square-free, $r := r+1$ and GO TO 2
4. Write P_c as $\prod_{j=1}^{r}(\beta_j x - \alpha_j y)$
5. Solve the transposed Vandermonde system:

$$\begin{pmatrix} \beta_1^D & \cdots & \beta_r^D \\ \beta_1^{D-1}\alpha_1 & \cdots & \beta_r^{D-1}\alpha_r \\ \vdots & \ddots & \vdots \\ \alpha_1^D & \cdots & \alpha_r^D \end{pmatrix} \begin{pmatrix} \lambda_1 \\ \vdots \\ \lambda_r \end{pmatrix} = \begin{pmatrix} a_0 \\ \vdots \\ a_D \end{pmatrix}. \tag{5}$$

6. Return $\sum_{j=1}^{r} \lambda_j (\alpha_j x + \beta_j y)^D$

For a proof of Theorem 2 we refer to Reznick (2013a, Theorem 2.1 & Corollary 2.2). Theorem 2 implies Algorithm 1. This algorithm will execute steps 2 and 3 as many times as the rank. At the i-th iteration it computes the kernel of H^i. The dimension of this kernel is $\leq i$ and each vector in the kernel has $i+1$ coordinates. As the rank of the binary form can be as big as the degree of the binary form, a straightforward bound for the arithmetic complexity of Algorithm 1 is at least cubic in the degree.

We can improve the complexity of Algorithm 1 by a factor of D by noticing that the rank of the binary form is either $\mathtt{rk}(H^{\lceil \frac{D}{2} \rceil})$ or $D - \mathtt{rk}(H^{\lceil \frac{D}{2} \rceil}) + 2$ (Comas and Seiguer, 2011, Section 3) (Helmke, 1992, Theorem B). Another way to compute the rank is by using minors (Bernardi et al., 2011, Algorithm 2).

The bottleneck of the previous approaches is that they have to compute the kernel of a Hankel matrix. However, even if we know that the rank of the binary form is r, the dimension of the kernel of H^r can still be as big as $O(D)$; the same bound holds for the length of the vectors in the kernel. Hence, the complexity is lower bounded by $O(D^2)$.

Our approach avoids the incremental construction. We exploit the structure of the kernel

of the Hankel matrices and we prove that the rank has only two possible values (Lemma 17), see also (Comas and Seiguer, 2011, Section 3), (Helmke, 1992, Theorem B), or (Bernardi et al., 2011). Moreover, we use a compact representation of the vectors in the kernel. We describe them as a combination of two polynomials of degree $O(D)$.

2.2. Kernel of the Hankel matrices

The Hankel matrices are among the most studied structured matrices (Pan, 2001). They are related to polynomial multiplication. We present results about the structure of their kernel. For details, we refer to Heinig and Rost (1984, Chapter 5).

Proposition 3. *Matrix-vector multiplication of Hankel matrices is equivalent to polynomial multiplication. Given two binary forms* $A := \sum_{i=0}^{D} a_i x^i y^{D-i}$ *and* $U := \sum_{i=0}^{k} u_i x^i y^{k-i}$*, consider* $R := \sum_{i=0}^{D+k} r_i x^i y^{D+k-i} = A \cdot U$*. If we choose the monomial basis* $\{y^{D+k}, \ldots, x^{D+k}\}$*, then the equality* $A \cdot U = R$ *is equivalent to Equation* (6)*, where the central submatrix of the left matrix is* $H^k_{(a_0,\ldots,a_D)}$ *(Definition 1).*

$$\left(\begin{array}{ccccc}
 & & & & a_0 \\
 & & & a_0 & a_1 \\
 & & \iddots & \iddots & \vdots \\
 & a_0 & \cdots & a_{k-2} & a_{k-1} \\
\hline
a_0 & a_1 & \cdots & a_{k-1} & a_k \\
a_1 & a_2 & \cdots & a_k & a_{k+1} \\
\vdots & \vdots & \iddots & \vdots & \vdots \\
a_{D-k} & a_{D-k+1} & \cdots & a_{D-1} & a_D \\
\hline
a_{D-k+1} & a_{D-k+2} & \cdots & a_D & \\
\vdots & \iddots & \iddots & & \\
a_{D-1} & a_D & & & \\
a_D & & & &
\end{array}\right)
\left(\begin{array}{c} u_k \\ \vdots \\ u_1 \\ u_0 \end{array}\right)
=
\left(\begin{array}{c} r_0 \\ r_1 \\ \vdots \\ r_{k-1} \\ \hline r_k \\ r_{k+1} \\ \vdots \\ r_D \\ \hline r_{D+1} \\ \vdots \\ r_{D+k-1} \\ r_{D+k} \end{array}\right). \tag{6}$$

Consider a family of Hankel matrices $\{H^k_a\}_{1\le k\le D}$ as in Definition 1. There is a formula for the dimension of the kernel of each matrix in the family $\{H^k_a\}_{1\le k\le D}$ that involves two numbers, N^a_1 and N^a_2. To be more specific, the following holds: