Towards Efficient Algorithms for Approximately Solving (Overdetermined) Systems of Polynomial Equations

N. Govindarajan, R. Widdershoven, L. De Lathauwer

Abstract

We revisit the age-old problem of solving a system of multivariate polynomial equations. This problem can be viewed as a simultaneous generalization of solving linear systems and finding roots of univariate polynomials. It is well-known that the aforementioned special cases have widespread applications in science and engineering. It should come as no surprise that the same is true about the more general version of this problem. This is particularly the case in the noisy overdetermined setting, which extends the familiar engineering notion of solving a system of linear equations in a "least-squares" fashion to the polynomial case.

Classically, solving a system of polynomial equations belongs to the field of computational algebraic geometry. The literature advocates two major approaches to the problem. The first approach involves homotopy continuation, where the roots of the desired system are found by continuous deformation of a starting system for which the roots are already known. The second approach, which is more in line with our work, reduces the problem to an eigenvalue problem. The algebraic approach has its origins in resultant theory, tracing back to original contributions by B'ezout, Sylvester, Cayley, and Macaulay. The main idea behind this line of attack is to unveil the structure of the quotient algebra of the ideal generated by the polynomial system. Solutions of the system are subsequently extracted from the eigen-structure of the generated multiplication tables [1].

Up to this point, the literature has primarily focused on the noiseless square case. Herein, it is assumed that the coefficients of the polynomials are known with full precision and the number of equations equals the number of unknowns. On the contrary, a critical component of engineering applications is the estimation of system parameters from an overcomplete set of equations corrupted by noise. In the case of linear systems, numerical linear algebra already provides effective methods to deal with such problems. Analogous methods to treat the more general polynomial case are, however, far fewer, and relatively underdeveloped.

Our work hopes to fill this gap by taking a fresh perspective on the problem. The cornerstone of our proposed framework is the (tensor-based) Macaulay method [4]. This method rests on the idea that a basis for the quotient algebra can be formed by computing a numerical null space of a resultant map. The classical Macaulay matrix, which generalizes the Sylvester resultant matrix of two univariate polynomials to the multivariate case, is a canonical example of such a map. The fact that the null space is obtained through numerical means is an essential ingredient in enabling the solvability of polynomial systems in an approximate sense [5].

This core key feature of the methodology is best explained through an example. Suppose that one wants to solve the overdetermined linear system

$$
\left[\begin{array}{c|cc}
-3 & -1 & -2 \\
-2 & -1 & 1 \\
1 & 7 & 1
\end{array}\right]
\left[\begin{array}{c}
1 \\ \hline x \\ y
\end{array}\right]
=
\left[\begin{array}{c}
0 \\ 0 \\ 0
\end{array}\right]
$$

in a total least-squares sense. As it is well-known, the solution is trivially found by retrieving the smallest right-singular vector of the matrix on the left-hand side of the expression. Since the corresponding singular value is strictly positive, the right-singular vector (after normalization)

will only yield an approximate solution of the system. Interesetingly, a similar strategy may be employed to find approximate solutions for polynomial systems. For example, the overdetermined polynomial system

$$
\begin{array}{c} p_1(x,y) \\ p_2(x,y) \\ p_3(x,y) \end{array}
\left[\begin{array}{cc|cc|ccc} -3 & -1 & -2 & 4 & 6 & 7 \\ -2 & -1 & 1 & 3 & -7 & 5 \\ 1 & 7 & 1 & -8 & 3 & 1 \end{array}\right]
\left[\begin{array}{c} 1 \\ x \\ y \\ x^2 \\ xy \\ y^2 \end{array}\right]
= \left[\begin{array}{c} 0 \\ 0 \\ 0 \end{array}\right]
\tag{1}
$$

has an exact solution at $(x, y) = (1, 0)$, which is also the only solution of the system. This single solution disappears under small perturbations of the coefficients, but one can still view $(x, y) = (1, 0)$ as an approximate solution of the noise-perturbed system. After all, the Vandermonde vector evaluated at this point, i.e., $\begin{bmatrix} 1 & 1 & 0 & 1 & 0 & 0 \end{bmatrix}^\top$, is still an approximate null vector for the noise perturbed matrix. It can therefore be subsequently used to derive approximate solutions.

Unlike the linear system, there are several complications that one has to treat in the polynomial case. Firstly, the existence of other artificial null vectors in (1) make the retrieval of the Vandermonde solution vector impossible. This first complication is resolved by adding additional equations to the system:

$$
\begin{array}{c} p_1(x,y) \\ p_2(x,y) \\ p_3(x,y) \\ xp_1(x,y) \\ xp_2(x,y) \\ xp_3(x,y) \\ yp_1(x,y) \\ yp_2(x,y) \\ yp_3(x,y) \end{array}
\left[\begin{array}{cc|ccc|cccc} -3 & -1 & -2 & 4 & 6 & 7 & 0 & 0 & 0 & 0 \\ -2 & -1 & 1 & 3 & -7 & 5 & 0 & 0 & 0 & 0 \\ 1 & 7 & 1 & -8 & 3 & 1 & 0 & 0 & 0 & 0 \\ 0 & -3 & 0 & -1 & -2 & 0 & 4 & 6 & 7 & 0 \\ 0 & -2 & 0 & -1 & 1 & 0 & 3 & -7 & 5 & 0 \\ 0 & 1 & 0 & 7 & 1 & 0 & -8 & 3 & 1 & 0 \\ 0 & 0 & -3 & 0 & -1 & -2 & 0 & 4 & 6 & 7 \\ 0 & 0 & -2 & 0 & -1 & 1 & 0 & 3 & -7 & 5 \\ 0 & 0 & 1 & 0 & 7 & 1 & 0 & -8 & 3 & 1 \end{array}\right]
\left[\begin{array}{c} 1 \\ x \\ y \\ x^2 \\ xy \\ y^2 \\ x^3 \\ x^2y \\ xy^2 \\ y^3 \end{array}\right]
= \left[\begin{array}{c} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{array}\right].
$$

The above matrix presents an example of a Macaulay matrix at a certain degree. This Macaulay matrix contains the Vandermonde Vector

$$
\begin{bmatrix} 1 & | & 1 & 0 & | & 1 & 0 & 0 & | & 1 & 0 & 0 & 0 \end{bmatrix}^\top
$$

as the only null vector up to scaling ambiguity. Secondly, a polynomial system can, in general, admit multiple solutions. Consequently, the numerically obtained (approximate) null basis will not be immediately in Vandermonde form. This second complication is resolved by performing an additional unmixing operation to retrieve the Vandermonde basis. This Vandermonde basis reconstruction can be viewed as a multi-dimensional harmonic retrieval problem [3], and can be effectively solved by computing a canonical polyadic decomposition of a tensor formed from the obtained null space basis [4].

In general, solutions for a polynomial system can be determined in two computational steps: (i) compute the null space of the Macaulay matrix, and (ii) compute a polyadic decomposition of a tensor to retrieve the solutions. In the case of overdetermined systems that only possess a handful of solutions, the first step is a major computational bottleneck. Macaulay matrices grow very rapidly in size for even moderately-sized polynomial systems. This makes the null space computation

prohibitively expensive. To get an impression of the complexity, a polynomial in N variables of degree D contains $\binom{N+D}{N}$ monomial terms. The Macaulay matrix at degree $D_{\text{mac}} \geq D$ of S such polynomials is the matrix that is constructed from the polynomial coefficients of the system $\{p_s\}_{s=1}^{S}$ such that its rows span the set of polynomials

$$\left\{ p := \sum_{s=1}^{S} h_s \cdot p_s : \quad \deg(p) \leq D_{\text{mac}} \right\}. \tag{2}$$

Both the numbers of rows and columns of this matrix grow combinatorially in size with respect to D_{mac}, D, and N.

Fortunately, the Macaulay matrix is highly structured. In the monomial basis, the Macaulay matrix possesses multilevel Toeplitz structures. It also has recursive properties, and in certain cases, the matrix can be highly sparse. Furthermore, the Chebyshev variant of the Macaulay matrix is multilevel Toeplitz-plus-Hankel. Subsequently, much of our research has been dedicated to developing efficient null-space computation algorithms that exploit the structures in the matrix. Our efforts in this area have led to some interesting work; some of which is still ongoing research:

1. In [2] we considered exploiting the low-displacement rank structure of the Macaulay matrix to compute the null space through a rank-revealing LU factorization using the Gohberg-Kailath-Olshevsky (GKO) algorithm. Although this approach reduced the memory and time complexity of the null-space computation significantly, the savings do not scale well for polynomial systems in many variables.

2. In recent (almost completed) work [6], we introduced an alternative approach that exploits the shift structure of the Macaulay matrix in a different way. This method scales more gracefully with the number of variables. The method relies on computing the null spaces of nested subblocks of increasing size by using the fact that the same subblocks are repeated throughout the Macaulay matrix and that the null space of two stacked (block-)rows equals the intersection of their individual null spaces.

3. Currently, we are also investigating an approach to compute the null space with a Krylov-based iterative technique. This option seems to be particularly attractive for overdetermined systems that only have a few (approximate) solutions. This is because, in this case, the Macaulay matrix will have a relatively small (numerical) null space. In the monomial basis, the Macaulay matrix has fast matrix-vector product with the help of fast Fourier transforms. Interestingly, one can also obtain a fast matrix-vector product for the Macaulay matrix in the Chebyshev basis using fast cosine transforms. A big challenge is to make the Krylov-based method converge in a reasonable number of iterations. Finding good preconditioners may be essential.

At the Householder Symposium, our goal is to share this research with our colleagues.

References

[1] David A Cox. Stickelberger and the eigenvalue theorem. In *Commutative Algebra*, pages 283–298. Springer, 2021.

[2] Nithin Govindarajan, Raphaël Widdershoven, Shivkumar Chandrasekaran, and Lieven De Lathauwer. A fast algorithm for computing macaulay null spaces of bivariate polynomial systems. *SIAM J. Matrix Anal. Appl.*, 45(1):368–396, 2024.

[3] M. Sørensen and L. De Lathauwer. Multidimensional harmonic retrieval via coupled canonical polyadic decomposition — Part I: Model and identifiability. *IEEE Transactions on Signal Processing*, 65(2):517–527, 1 2017.

[4] Jeroen Vanderstukken and Lieven De Lathauwer. Systems of polynomial equations, higher-order rensor decompositions and multidimensional harmonic retrieval: A unifying framework. Part I: The canonical polyadic decomposition. *SIAM J. Matrix Anal. Appl.*, 42(2):883–912, 2021.

[5] Raphaël Widdershoven, Nithin Govindarajan, and Lieven De Lathauwer. Overdetermined systems of polynomial equations: tensor-based solution and application. In *2023 31st European Signal Processing Conference (EUSIPCO)*, pages 650–654. IEEE, 2023.

[6] Raphaël Widdershoven, Nithin Govindarajan, and Lieven De Lathauwer. Fast macaulay null space through the intersection of shifted null spaces. *Technical Report 25-56, ESAT-STADIUS, KU Leuven (Leuven, Belgium)*, 2025.