

Efficient Computation of Macaulay Matrix Null Spaces Through Exploiting Shift-Invariant Structures

Nithin Govindarajan

with Raphaël Widdershoven, Shiv Chandrasekaran, and Lieven De Lathauwer

June 15th 2023

Overview

Macaulay matrices and their null spaces: a computational challenge

The low displacement property of Macaulay matrices

The fast algorithm

Analysis of complexity

Numerical experiments

Generalizations

Conclusions

The problem that we wish to solve

Let $S \geq N$, and consider system of the multivariate polynomials

$$\Sigma : \begin{cases} p_1 = p_1(x_1, x_2, \dots, x_N) \\ \vdots \\ p_S = p_S(x_1, x_2, \dots, x_N) \end{cases}, \quad (1)$$

with $\deg(p_s) = d_s$.

Goal:

Find all roots $\left\{ (t^{(r)}, x_1^{(r)}, \dots, x_N^{(r)}) \right\}_{r=1}^R$ of the *homogenized system* Σ_h .

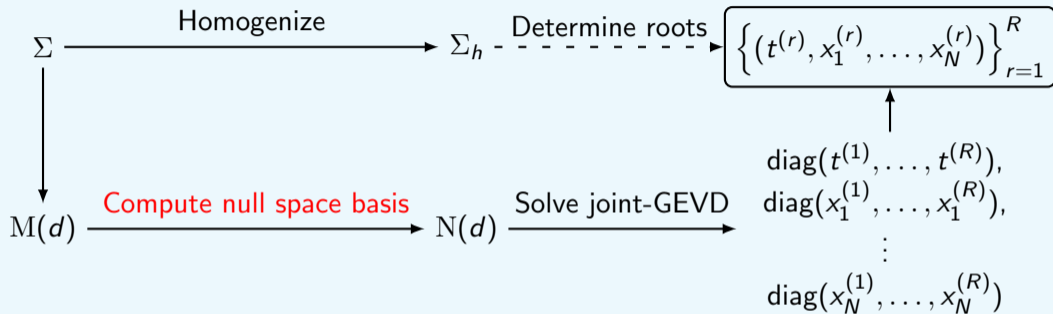
NOTE: In practice, we may be only interested in the affine roots...

The Macaulay matrix $M(d)$ and its right null space

The rows of $M(d)$ span the set of polynomial combinations

$$\left\{ \sum_{s=1}^S h_s \cdot p_s : \deg(h_s) = d - d_s \right\}.$$

Null space computation is a major computational bottleneck in *many* algorithms!



(Vanderstukken and De Lathauwer 2021)

Overview

Macaulay matrices and their null spaces: a computational challenge

The low displacement property of Macaulay matrices

The fast algorithm

Analysis of complexity

Numerical experiments

Generalizations

Conclusions

Focus: (Possibly overdetermined) bivariate polynomial systems

For simplicity of exposition, we assume $\deg(p_s) = d_\Sigma$, i.e.,

$$\Sigma : \left\{ \begin{array}{l} p_1(x, y) = \sum_{i=0}^{d_\Sigma} \sum_{j=0}^{d_\Sigma-i} c_{1ij} x^i y^j = 0 \\ \vdots \\ p_s(x, y) = \sum_{i=0}^{d_\Sigma} \sum_{j=0}^{d_\Sigma-i} c_{sij} x^i y^j = 0 \end{array} \right.$$

In lex ordering Macaulay is almost Toeplitz block-(block-)Toeplitz!

$$M(4) = \begin{array}{c} \begin{array}{l} p_1 \\ p_2 \\ xp_1 \\ xp_2 \\ x^2p_1 \\ x^2p_2 \end{array} \\ \hline \begin{array}{l} yp_1 \\ yp_2 \\ xyp_1 \\ xyp_2 \\ y^2p_1 \\ y^2p_2 \end{array} \end{array} \begin{bmatrix} \begin{array}{ccccc} 1 & x & x^2 & x^3 & x^4 \\ \begin{array}{c} 1 \\ 9 \end{array} & \begin{array}{c} 6 \\ 1 \end{array} & \begin{array}{c} 4 \\ 3 \end{array} & & \\ \begin{array}{c} 1 \\ 9 \end{array} & \begin{array}{c} 6 \\ 1 \end{array} & \begin{array}{c} 4 \\ 3 \end{array} & & \\ & \begin{array}{c} 1 \\ 9 \end{array} & \begin{array}{c} 6 \\ 1 \end{array} & \begin{array}{c} 4 \\ 3 \end{array} & \\ & & \begin{array}{c} 1 \\ 9 \end{array} & \begin{array}{c} 6 \\ 1 \end{array} & \begin{array}{c} 4 \\ 3 \end{array} \\ & & & \begin{array}{c} 1 \\ 9 \end{array} & \begin{array}{c} 6 \\ 1 \end{array} & \begin{array}{c} 4 \\ 3 \end{array} \end{array} & \begin{array}{cccc} y & xy & x^2y & x^3y \\ \begin{array}{c} 2 \\ 8 \end{array} & \begin{array}{c} 5 \\ 7 \end{array} & & \\ \begin{array}{c} 2 \\ 8 \end{array} & \begin{array}{c} 5 \\ 7 \end{array} & & \\ & \begin{array}{c} 2 \\ 8 \end{array} & \begin{array}{c} 5 \\ 7 \end{array} & \\ & & \begin{array}{c} 2 \\ 8 \end{array} & \begin{array}{c} 5 \\ 7 \end{array} \end{array} & \begin{array}{ccc} y^2 & xy^2 & x^2y^2 \\ \begin{array}{c} 3 \\ 2 \end{array} & & \\ & \begin{array}{c} 3 \\ 2 \end{array} & \\ & & \begin{array}{c} 3 \\ 2 \end{array} \end{array} & \begin{array}{cc} y^3 & xy^3 \\ \begin{array}{c} 3 \\ 2 \end{array} & \\ & \begin{array}{c} 3 \\ 2 \end{array} \end{array} & \begin{array}{c} y^4 \\ \begin{array}{c} 3 \\ 2 \end{array} \end{array} \end{bmatrix}$$

The Macaulay matrix for the general bivariate case

Let $\Delta d := d - d_\Sigma$. Then,

$$M(d) := \begin{bmatrix} M_{0,0} & M_{1,0} & \cdots & M_{d_\Sigma,0} & & & \\ & M_{0,1} & M_{1,1} & \cdots & M_{d_\Sigma,1} & & \\ & & \ddots & \ddots & & \ddots & \\ & & & M_{0,\Delta d} & M_{1,\Delta d} & \cdots & M_{d_\Sigma,\Delta d} \end{bmatrix} \in \mathbb{C}^{\frac{S}{2}(\Delta d+1)(\Delta d+2) \times \frac{1}{2}(d+1)(d+2)},$$

with

$$M_{i,j} := \begin{bmatrix} \mathbf{c}_{0i} & \mathbf{c}_{1i} & \cdots & \mathbf{c}_{(d_\Sigma-i)i} & & & \\ & \mathbf{c}_{0i} & \mathbf{c}_{1i} & \cdots & \mathbf{c}_{(d_\Sigma-i)i} & & \\ & & \ddots & \ddots & & \ddots & \\ & & & \mathbf{c}_{0i} & \mathbf{c}_{1i} & \cdots & \mathbf{c}_{(d_\Sigma-i)i} \end{bmatrix} \in \mathbb{C}^{S(\Delta d+1-j) \times (d+1-i-j)}.$$

Intermezzo: The generalized shift matrix

Let

$$Z_{p,\varphi} := \begin{bmatrix} & & & \varphi \\ & & & \\ & & & \\ & & \dots & \\ & & & 1 \\ & & & \end{bmatrix} \in \mathbb{C}^{p \times p}.$$

Eigen-decomposition of $Z_{p,\varphi}$

Denote $\omega_p := \exp(-2\pi i/p)$ and $F_p \in \mathbb{C}^{p \times p}$ the (unitary) DFT matrix. Then

$$Z_{p,\varphi} = (D_{p,\varphi} F_p)(\varphi^{1/p} \Omega_p)(D_{p,\varphi} F_p)^{-1},$$

where $D_{p,\varphi} := \text{diag}(1, \varphi^{-1/p}, \dots, \varphi^{-(p-1)/p})$, $\Omega_p := \text{diag}(1, \bar{\omega}_p, \dots, \bar{\omega}_p^{p-1})$.

Toeplitz matrices and their low “displacement rank” properties

Consider the so-called displacement equation

$$\begin{aligned} Z_{4,1} \begin{bmatrix} t_0 & t_1 & t_2 & t_3 \\ t_{-1} & t_0 & t_1 & t_2 \\ t_{-2} & t_{-1} & t_0 & t_1 \\ t_{-3} & t_{-2} & t_{-1} & t_0 \end{bmatrix} - \begin{bmatrix} t_0 & t_1 & t_2 & t_3 \\ t_{-1} & t_0 & t_1 & t_2 \\ t_{-2} & t_{-1} & t_0 & t_1 \\ t_{-3} & t_{-2} & t_{-1} & t_0 \end{bmatrix} Z_{4,\phi} \\ = \\ \underbrace{\begin{bmatrix} t_{-3} - t_1 & t_{-2} - t_2 & t_{-1} - t_3 & t_0 - \phi t_0 \\ 0 & 0 & 0 & t_3 - \phi t_{-1} \\ 0 & 0 & 0 & t_2 - \phi t_{-2} \\ 0 & 0 & 0 & t_1 - \phi t_{-3} \end{bmatrix}}_{\text{rank is only two!}} \end{aligned}$$

NOTE: $\phi \in \mathbb{C}$ is chosen such that the operator $\mathcal{D} : T \mapsto Z_{4,1}T - TZ_{4,\phi}$ remains invertible!

The key observation that shall allow for a faster algorithm!

Consider the displacement operator

$$\mathcal{D}\{M(d)\} = \begin{bmatrix} Z_{d+1,1} \otimes I_S & & \\ & \ddots & \\ & & Z_{1,1} \otimes I_S \end{bmatrix} M(d) - M(d) \begin{bmatrix} Z_{d+1,\varphi_{d+1}} & & \\ & \ddots & \\ & & Z_{1,\varphi_1} \end{bmatrix}.$$

$M(d)$ has relative “low” displacement rank too!

Dimensions of $M(d) \in \mathbb{C}^{\frac{S}{2}(\Delta d+1)(\Delta d+2) \times \frac{1}{2}(d+1)(d+2)}$ grow *quadratically* w.r.t. d , but

$$\text{rank } \mathcal{D}\{M(d)\} \leq S(\Delta d + 1) = S(d + 1 - d_\Sigma).$$

grows only *linearly* with d .

Overview

Macaulay matrices and their null spaces: a computational challenge

The low displacement property of Macaulay matrices

The fast algorithm

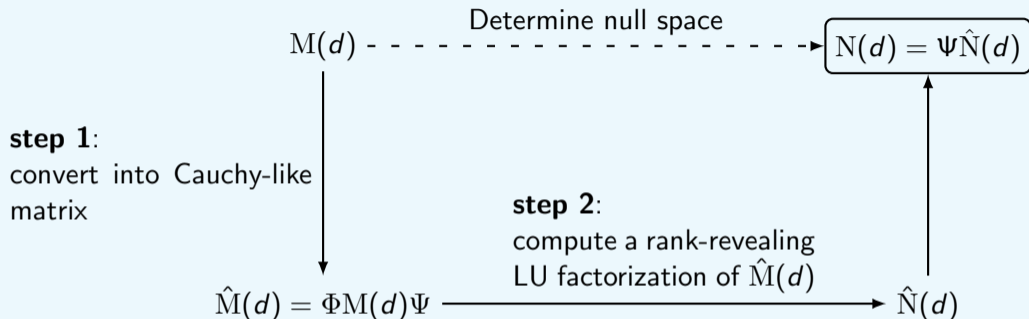
Analysis of complexity

Numerical experiments

Generalizations

Conclusions

An overview of the fast algorithm



Both steps can be done *fast*!

Step 1: convert into Cauchy-like matrix

Apply unitary transformations Φ and Ψ such that $\Phi M(d)\Psi =: \hat{M}(d)$ is Cauchy-like, i.e., its entries are of the form

$$\left[\hat{M}(d)\right]_{ij} := [\Phi M(d)\Psi]_{ij} = \frac{\mathbf{u}_i^* \mathbf{v}_j}{\mu_i - \nu_j}, \quad \mathbf{u}_i, \mathbf{v}_j \in \mathbb{C}^{S(\Delta d+1)}.$$

Step 1 *done fast*: use displacement rank theory!

Best explained through the simpler Toeplitz case...

$$Z_{n,1}T_n - T_n Z_{n,\varphi} = UV^*$$

↓

$$(F_n \Omega_n F_n^*) T_n - T_n \left((D_{n,\varphi} F_n) (\varphi^{1/n} \Omega_n) (D_{n,\varphi} F_n)^{-1} \right) = UV^*$$

↓

$$\Omega_n F_n^* T_n D_{n,\varphi} F_n - F_n^* T_n D_{n,\varphi} F_n \left(\varphi^{1/n} \Omega_n \right) = (F_n^* U) F_n^* D_{n,\varphi}^* V^*$$

↓

$$\text{diag}(\mu)C - C\text{diag}(\nu) = RS^* =: G$$

↓

Displacement equation for *Cauchy-like* matrix!

Step 2: compute a rank-revealing LU factorization of $\hat{M}(d)$

Let $r(d) := \text{rank } M(d)$. Compute a rank-revealing LU (RRLU) factorization

$$\begin{aligned}\Pi_1 \hat{M}(d) \Pi_2 &= \begin{bmatrix} I_{r(d)} & & \\ \hat{M}_{21} \hat{M}_{11}^{-1} & & \\ & & I_{d_{\Sigma}^2} \end{bmatrix} \begin{bmatrix} \hat{M}_{11} & & \\ & \hat{M}_{22} - \hat{M}_{21} \hat{M}_{11}^{-1} \hat{M}_{12} & \\ & & \end{bmatrix} \begin{bmatrix} I_{r(d)} & \hat{M}_{11}^{-1} \hat{M}_{12} \\ & I_{d_{\Sigma}^2} \end{bmatrix} \\ &\approx \begin{bmatrix} \hat{M}_{11} \\ \hat{M}_{21} \end{bmatrix} \begin{bmatrix} I_{r(d)} & \hat{M}_{11}^{-1} \hat{M}_{12} \end{bmatrix}\end{aligned}$$

Expression for the null space $N(d)$

$$N(d) = \Psi \Pi_2 \begin{bmatrix} \tilde{N} \\ I_{d_{\Sigma}^2} \end{bmatrix}, \quad \tilde{N} := -\hat{M}_{11}^{-1} \hat{M}_{12}.$$

NOTE: Gaussian elimination on Macaulay matrix \equiv polynomial reductions for e.g., Grobner Basis reductions (Eder and Faugère 2017)

Step 2 *done fast*: Apply Schur algorithm on Cauchy-like matrix (Heinig 1995)

$$\text{diag}(\boldsymbol{\mu}) \begin{bmatrix} A & G^* \\ F & B \end{bmatrix} - \begin{bmatrix} A & G^* \\ F & B \end{bmatrix} \text{diag}(\boldsymbol{\nu}) = \begin{bmatrix} R_1 \\ R_2 \end{bmatrix} \begin{bmatrix} S_1 \\ S_2 \end{bmatrix}^*$$

↓

$$\text{diag}(\boldsymbol{\mu}) \begin{bmatrix} A & \\ & B - FA^{-1}G^* \end{bmatrix} - \begin{bmatrix} A & \\ & B - FA^{-1}G^* \end{bmatrix} \text{diag}(\boldsymbol{\nu}) = \begin{bmatrix} R_1 \\ \tilde{R}_2 \end{bmatrix} \begin{bmatrix} S_1 \\ \tilde{S}_2 \end{bmatrix}^*,$$

where $\tilde{R}_2 = R_2 - FA^{-1}R_1$ and $\tilde{S}_2 = S_2 - G(A^*)^{-1}S_1$,

Main idea of Schur algorithm

Perform Gauss elimination on the *generators* instead of the dense matrix itself!

Step 2 done fast: approximate total pivoting through QR-decomposition of generators

Recall

$$\text{diag}(\boldsymbol{\mu})C - C\text{diag}(\boldsymbol{\nu}) = RS^* =: G \in \mathbb{C}^{n \times n}.$$

(Gu 1998, Lemma 3.1)

Let j_{\max} denote the column with largest 2-norm in G . Then,

$$\max_{1 \leq i \leq n} |c_{ij_{\max}}| \geq \frac{1}{K\sqrt{n}} \max_{1 \leq i, j \leq n} |c_{ij}|,$$

with $K := \max_{1 \leq i, j, l, j' \leq n} |\mu_i - \nu_j| / |\mu_l - \nu_{j'}|$.

Full QR-decomposition of R expensive! \rightarrow Use *fancy* QR updating techniques!

Overview

Macaulay matrices and their null spaces: a computational challenge

The low displacement property of Macaulay matrices

The fast algorithm

Analysis of complexity

Numerical experiments

Generalizations

Conclusions

We are able to reduce the flop complexity from $\mathcal{O}(d_\Sigma^6)$ to $\mathcal{O}(d_\Sigma^5)$!

Assumptions:

- Σ has a consistent set of equations \rightarrow number of roots = d_Σ^2 ,
- $S \ll d_\Sigma$.

A quick complexity overview for each step

- Step 1: $\mathcal{O}(S \cdot d_\Sigma \cdot \Delta d \cdot d \log d)$
- Step 2: $\mathcal{O}(r(d) \cdot S^2 d^3)$

$d \leq 2d_\Sigma - 2$ to find a null space containing all system roots

$$\downarrow \\ \mathcal{O}(d_\Sigma^5)$$

NOTE: We can overcome the restriction $S \ll d_\Sigma$ by random sampling of the rows!

Overview

Macaulay matrices and their null spaces: a computational challenge

The low displacement property of Macaulay matrices

The fast algorithm

Analysis of complexity

Numerical experiments

Generalizations

Conclusions

The error metric used to assess performance

Let $Q \in \mathbb{C}^{n(d) \times d_{\Sigma}^2}$ be an orthonormal basis for $\text{col } N(d)$, then define

$$\epsilon := \frac{\|M(d)Q\|_2}{\|M(d)\|_2} \geq \frac{\sigma_{r(d)+1}}{\sigma_1} =: \epsilon_{\min},$$

Algorithm stability: error grows linearly with problem size

Median error over 100 runs for *square* systems with different methods and degrees.


	d_{Σ}				
	2	4	8	16	32
SVD on $M(d)$	2.23e-16	3.75e-16	5.70e-16	7.94e-16	9.51e-16
SVD on $\hat{M}(d)$	2.57e-16	4.77e-16	7.54e-16	9.97e-16	1.15e-15
GECP on $M(d)$	1.40e-16	3.11e-16	8.33e-16	1.02e-14	1.40e-13
GECP on $\hat{M}(d)$	2.08e-16	4.65e-16	1.03e-15	9.73e-15	1.21e-13
GECP on \mathcal{C}	4.35e-16	1.51e-15	1.35e-14	1.72e-13	2.81e-12
GEAP on \mathcal{C}	4.21e-16	3.63e-15	3.88e-14	3.19e-13	4.48e-12

Sources of error:

- switching to LU instead of an SVD
- working with the compact Cauchy representation \mathcal{C}
- switching to approximate pivoting ← Surprisingly not so bad!

Our experiments indicate that the flop complexity is indeed $O(d_\Sigma^5)$

The measurements are the median of an adapted number of runs after warmup.



Overview

Macaulay matrices and their null spaces: a computational challenge

The low displacement property of Macaulay matrices

The fast algorithm

Analysis of complexity

Numerical experiments

Generalizations

Conclusions

The algorithm generalizes to Chebyshev systems!

$$\Sigma : \left\{ \begin{array}{l} p_1(x, y) := \sum_{i=0}^{d_\Sigma} \sum_{j=0}^{d_\Sigma-i} b_{1ij} T_i(x) T_j(y) = 0 \\ \vdots \\ p_S(x, y) := \sum_{i=0}^{d_\Sigma} \sum_{j=0}^{d_\Sigma-i} b_{Sij} T_i(x) T_j(y) = 0 \end{array} \right.$$

Key ideas to arrive to an $\mathcal{O}(d_\Sigma^5)$ algorithm:

- *Toeplitz-plus-Hankel* instead of just Toeplitz.
- Apply same techniques but with a *modified* displacement equation.

The algorithm does not nicely extend for general N -dimensional systems

- Displacement rank theory does not generalize nicely to higher dimensions ☹️
- Diminishing returns: $\mathcal{O}(d^{3N})$ to $\mathcal{O}(d^{3N-1})$
- Open problem: how to exploit multi-level Toeplitz structures?

Overview

Macaulay matrices and their null spaces: a computational challenge

The low displacement property of Macaulay matrices

The fast algorithm

Analysis of complexity

Numerical experiments

Generalizations

Conclusions





Conclusions


- An asymptotically *faster* algorithm for Macaulay null space computation.
- Generalizes to Chebyshev systems
- Generic n -dimensional systems still a challenge! (assume additional structure?)

To learn more, see our pre-print:

Govindarajan, Nithin et al. (2023). “A fast algorithm for computing Macaulay nullspaces of bivariate polynomial systems”. In: Technical Report 23-16, ESAT-STADIUS, KU Leuven (Leuven, Belgium).

References I

-  Eder, Christian and Jean-Charles Faugère (2017). “A survey on signature-based algorithms for computing Gröbner bases”. In: *J. Symb. Comput.* 80, pp. 719–784.
-  Gu, Ming (1998). “Stable and Efficient Algorithms for Structured Systems of Linear Equations”. In: *SIAM J. Matrix Anal. Appl.* 19.2, pp. 279–306.
-  Heinig, Georg (1995). “Inversion of Generalized Cauchy Matrices and other Classes of Structured Matrices”. In: *Linear Algebra for Signal Processing*. Ed. by A. Bojanczyk and G. Cybenko. Vol. 69. The IMA Volumes in Mathematics and its Applications. New York, NY, USA: Springer.
-  Vanderstukken, Jeroen and Lieven De Lathauwer (2021). “Systems of Polynomial Equations, Higher-Order Rensor Decompositions and Multidimensional Harmonic Retrieval: A Unifying Framework. Part I: The Canonical Polyadic Decomposition”. eng. In: *SIAM J. Matrix Anal. Appl.* 42.2, pp. 883–912.

-  Vanderstukken, Jeroen et al. (2021). “Systems of Polynomial Equations, Higher-Order Tensor Decompositions, and Multidimensional Harmonic Retrieval: A Unifying Framework. Part II: The Block Term Decomposition”. In: *SIAM J. Matrix Anal. Appl.* 42.2, pp. 913–953.