Non-commutative rank of linear matrices, related structures and applications

Gábor Ivanyos
MTA SZTAKI

SIAM AG 2019, 9-13 July 2019, Bern.

Based on joint works with Youming Qiao and K. V. Subrahmanyam

Commutative and noncommutative rank

- linear matrix: $A(x) = A(x_1, \ldots, x_k) = A_1x_1 + \ldots + A_kx_k$

  $\sim$ matrix space $\mathcal{A} = \langle A_1, \ldots, A_k \rangle$; $\quad A_1, \ldots, A_k \in F^{n\times n}$
- (commutative) rank $\operatorname{rk} A(x)$: as a matrix over $F(x_1, \ldots, x_n)$

  max rank from $\mathcal{A}$ (if F "large enough")
- Task: compute $\operatorname{rk} A(x)$ (attributed to Edmonds 1967)

  an instance of PIT, $\in RP$, not known to be in P

  "derandomization" would have remarkable consequences in complexity theory (Kabanets, Impagliazzo 2003)
- noncommutative rank $\operatorname{ncrk} A(x)$: as a matrix over the free skewfield

  max rank from $\mathcal{A} \otimes_F D$; ("D-span" of A_js; D: *some* skewfield)

  (Gaussian elim. and consequences to rank remain valid for skewfields)

Commutative vs. noncommutative rank

- $\operatorname{rk} A(x) \leq \operatorname{ncrk} A(x)$
- Example for $<$: A_1, A_2, A_3 a basis for the skew-symmetric 3 by 3 real matrices

  $\operatorname{rk} A(x) = 2$; $\operatorname{ncrk} A(x) = 3$ (over the quaternions)
- which one is easier to compute?
  - ncrk is a proper relaxation of rk
  - however its definition is more complicated

    uses a difficult object or a (possibly) infinite family of skewfields

    (can be pulled down to exp size)

    $\Rightarrow$???? randomized poly alg?????

Commutative vs. noncommutative rank

- $\operatorname{rk} A(x) \leq \operatorname{ncrk} A(x)$
- Example for $<$: A_1, A_2, A_3 a basis for the skew-symmetric 3 by 3 real matrices

  $\operatorname{rk} A(x) = 2$; $\operatorname{ncrk} A(x) = 3$ (over the quaternions)
- which one is easier to compute?
  - ncrk is a proper relaxation of rk
  - however its definition is more complicated

    uses a difficult object or a (possibly) infinite family of skewfields

    (can be pulled down to exp size)

  $\Rightarrow$???? randomized poly alg?????
- ncrk is "easier":

  computable even in **deterministic polynomial** time!

  (Garg, Gurvits, Oliveira, Wigderson 2015-2016; IQS 2015-2018)

The nc rank as a rank of a large matrix

- Can assume D: central of dimension d^2 over F
  - $D \otimes L \cong L^{d \times d} (= M_d(L))$ for some L
  - both D and $F^{d \times d}$ embedded in $L^{d \times d}$
- gives switching procedures

$$\mathcal{A} \otimes D \longleftrightarrow \mathcal{A} \otimes F^{d \times d} \subseteq F^{nd \times nd}$$

$$\text{rank } r \text{ over } D \longrightarrow \text{rank} \geq r \cdot d \text{ over } F$$

$$\text{rank} \geq \lceil R/d \rceil \text{ over } D \longleftarrow \text{rank } R \text{ over } F$$

- composition ($\leftarrow$, then $\rightarrow$): "rounding up" the rank of a matrix in $\mathcal{A} \otimes F^{d \times d}$ to a multiple of d

  IQS 2015: can be done in deterministic poly time (for suitable D)

Remark: determinants of matrices in $\mathcal{A} \otimes F^{d \times d}$ $\sim$ invariants of $SL_n \times SL_n$

Inflated matrix spaces

- $\mathcal{A} \otimes F^{d\times d}$: *inflated* matrix space (d: infl. factor)
  n by n matrices with entries from $F^{d\times d}$
- based on the rounding, Derksen-Makam 2015–2017, a reduction tool to show

  $$\text{ncrk}\, A(x) = \frac{1}{d} \text{max rank in } \mathcal{A} \otimes F^{d\times d}$$

  for some $d \leq n-1$.
- $\Rightarrow \exists$ randomized poly time alg for ncrk

Constructive deterministic results

- IQS 2015-2018: a deterministic poly time algorithm
  - computes a matrix of rank $d \cdot \operatorname{ncrk} A(x)$ in $\mathcal{A} \otimes F^{d \times d}$
    $d \leq n-1$ (or $d \leq n \log n$ if F is too small)
  - computes a witness for that ncrk cannot be larger
  - uses analogues of the alternating paths for matchings if graphs
    + an efficient implementation of the DM reduction tool
- Garg, Gurvits, Oliveira, Wigderson 2015-2016:
  - different approach for $\operatorname{char} F = 0$
    (not through such witnesses)

The witnesses: shrunk subspaces (a Hall-like obstacles)

- ℓ-shrunk subspace: $U \leq F^n$ mapped to a subspace of dimension $\dim U - \ell$ by $\mathcal{A}$

$$\mathcal{A} \leq \begin{pmatrix} * & * & 0 & 0 & 0 \\ * & * & 0 & 0 & 0 \\ * & * & 0 & 0 & 0 \\ * & * & * & * & * \end{pmatrix} \text{ alias } \begin{pmatrix} * & * & & & \\ * & * & & & \\ * & * & & & \\ * & * & * & * & * \end{pmatrix}$$

$\exists$ ℓ-shrunk subsp. $\Rightarrow$ the max rank in $\mathcal{A}$ is at most $n - \ell$

- Inheritance: $U \otimes F^{d \times d}$ mapped to a subspace of dim less by $\ell \cdot d \Rightarrow$ max rank in $\mathcal{A} \otimes F^{d \times d}$ is at most $nd - \ell d$.
- $\Rightarrow$ ncrk $\leq n - \ell$
- $\sim$ a characterization of the nullcone of invariants $SL_n \times SL_n$ (by Hilbert-Mumford)

Wong sequence

- attempt to find a shrunk subspace (from Fortin, Reutenauer 2004, also I, Karpinski, Qiao, Santha 2013-2015)
- Assume we have $B \in \mathcal{A}$ with $\mathrm{rk}\, B = \mathrm{ncrk}$, $\ell = n - \mathrm{ncrk}$, U ℓ-shrunk. Then

$$U \geq \ker B \text{ and } \mathcal{A}U = \mathrm{Im}\, B.$$

- Wong sequence ($\sim$ alternating forest in bipartite graph matching):

$$U_1 = \ker B;\ U_{i+j} = B^{-1}(\mathcal{A}U_j) \qquad \text{(inverse image for } B\text{)}$$

  - Either stabilizes in $\mathrm{Im}\, B$: gives an ℓ-shrunk subspace
  - or "escapes" : $\mathcal{A}U_j \not\subseteq \mathrm{Im}\, B$: ($\sim$ $\exists$ augmenting path)

Escaping Wong sequence $\sim$ augmenting path

- sequence $i_1, \ldots, i_s$ – with s smallest – s.t.

$$A_{i_s} B^{-1}(A_{i_{s-1}} B^{-1}(\ldots B^{-1}(A_{i_1} \ker B))) \not\subseteq \operatorname{Im} B$$

- Put $A_1' = B' = B \otimes I_d$, $A_2' = \sum A_{i_j} \otimes E_{j,j+1} \in \mathcal{A} \otimes F^{d \times d}$; $\mathcal{A}' = \langle A_1', A_2' \rangle$ (d large enough)
- Then the Wong seq. escapes $\operatorname{Im} B'$ and $C' = B' + \lambda A_2'$ has rank $> d \cdot \operatorname{rk} B$ for some λ
- Round up the rank of C' in $\mathcal{A} \otimes F^{d \times d}$ to a multiple of d

The iterative algorithm

- iterate the above "scaled" rank incrementation procedure (with iteratively "inflating" $\mathcal{A}$)
- combine with the reduction tool to keep final "inflation" factor small.
- Result: $A \in \mathcal{A} \otimes F^{d \times d}$ of rank $d \cdot \text{ncrk}$; and a maximally (by $(n - d \cdot \text{ncrk}\,)))$ shrunk subspace (of F^{nd}) for $\mathcal{A} \otimes F^{d \times d}$. $(d \leq n - 1.)$
- Use converse of inheritance to obtain a maximally (by $n - \text{ncrk}\,$) shrunk subspace of F^n for $\mathcal{A}$.
- Remarks:

  (1) Actually, *the smallest* maximally shrunk subspace found. ((0) if ncrk $= n$.)

  (2) The largest one can also be found (duality)

The echelon structure

- In bases resp. smallest and largest maximally shrunk subspaces:

$$\mathcal{A} \subseteq \begin{pmatrix} * & * & * & * & * & * & *] \\ & & \bullet & \bullet & \bullet & * & * \\ & & \bullet & \bullet & \bullet & * & * \\ & & \bullet & \bullet & \bullet & * & * \\ & & & & & * & * \\ & & & & & * & * \\ & & & & & * & * \end{pmatrix}$$

- The "middle diagonal block" of $\mathcal{A}$ (filled with $\bullet$) is of full ncrk . Can be:
  - $n \times n$ (if $\operatorname{ncrk} \mathcal{A} = n$)
  - 0×0 (unique maximally shrunk subspace)
  - Further maximally shrunk subspaces can be found by block triangularizing the $\bullet$-block.

Block triangularization in the full ncrk case

- $\sim$ finding flag of 0-shrunk subspaces U ($\dim \mathcal{A}U = \dim U$)
- If $I \in \mathcal{A}$ then (as $\mathcal{A}W \geq W$) equivalent to $\mathcal{A}U = U$.
  - U: a submodule for $\mathcal{A}$,
  - for many F, $\exists$ good algorithms
- If $A \in \mathcal{A}$ of full rank found, $I \in A^{-1}\mathcal{A}$.
- In the general case,
  - Find $A \in \mathcal{A} \otimes F^{d \times d}$ of full rank,
  - Block triangularize $\mathcal{A} \otimes F^{d \times d}$ as above
  - Pull back by "reverse inheritance"
- Applicable in multivariate cryptography e.g, for breaking Patarin's balanced Oil and Vinegar scheme.

On Wong sequences and the commutative rank

Wong sequence: $U_1 = \ker B$; $U_{i+j} = B^{-1}(\mathcal{A}U_j)$.

- Bläser, Jindal & Pandey (2017): deterministic rank approximation scheme based on the speed/length

In extreme cases, ncrk $=$ rk

- Immediately escaping case: length 1
  - $\operatorname{rk}(B + \lambda A_i) > \operatorname{rk} B$ for some i and λ: $\longrightarrow$ "blind" rank incrementing algorithm
  - holds for $\mathcal{A} = Hom(V_1, V_2)$ where V_1, V_2 semisimple modules
  - holds when $\mathcal{A}$ simultaneously diagonalizable
- Slim Wong sequence $\dim U_{j+1} = \dim U_j + 1$
  - $\operatorname{rk}(B + \lambda \sum_{j=1}^{k} A_j) > \operatorname{rk} B$ for some λ
  - holds for $k = 2$
  - can be enforced if $\mathcal{A}$ spanned by rank 1 matrices (even if they are not given explicitly)