Limitations of algebraic lower bound proofs

Markus Bläser

Saarland University

with Christian Ikenmeyer, Gorav Jindal, and Vladimir Lysikov

How to prove a lower bound?

- presumably hard function f
- class of easy functions $\mathcal{E}$
- lower bound: a proof for $f \notin \mathcal{E}$.

Approach

Exhibit a property $\mathcal{P}$ such that

- f does not have property $\mathcal{P}$ and
- all $g \in \mathcal{E}$ have property $\mathcal{P}$.

- $\mathcal{P}$ should be easier that "being in $\mathcal{E}$".

Natural proofs

Definition (Razborov & Rudich)

A property $\mathcal{P}$ of Boolean functions is *natural* if it has the following properties:

Usefulness: If $f : \{0,1\}^n \to \{0,1\}$ has $\mathrm{poly}(n)$-sized circuits, then $f \in \mathcal{P}$.

Constructivity: Given f by a truthtable of size $N = 2^n$, we can decide $f \in \mathcal{P}$ in time $\mathrm{poly}(N)$.

Largeness: A random function is not in $\mathcal{P}$ with probability at least $1/\mathrm{poly}(N) = 2^{-O(n)}$.

The Razborov–Rudich barrier

- A function $f : \{0,1\}^n \times \{0,1\}^\ell \to \{0,1\}$ is *pseudorandom* if when sampling the key $k \in \{0,1\}^\ell$ uniformly at random, the resulting distribution $f(\,.\,,k)$ is computationally indistinguishable from a truly random function.
- If oneway functions exists, so do pseudorandom functions.

Theorem (Razborov & Rudich)

A natural property $\mathcal{P}$ distinguishes a pseudorandom function having $\mathrm{poly}(n)$-size circuits from a truly random function in time $2^{O(n)}$.

Conclusion

If you believe in private key cryptography, then no natural proof will show superpolynomial circuit lower bounds.

Valiant's classes

Definition

A sequence of polynomials $(f_n) \in K[X]$ is a *p-family* if for all n,

1. $f_n \in K[X_1, \dots, X_{p(n)}]$ for some polynomially bounded function p and
2. $\deg f_n \leq q(n)$ for some polynomially bounded function q.

Definition

The class VP consists of all p-families (f_n) such that each f_n is computed by an algebraic circuit of polynomial size.

Example:

$$\mathrm{det}_n X = \sum_{\pi \in S_n} \mathrm{sgn}(\pi) X_{1,\pi(1)} \dots X_{n,\pi(n)}$$

Rank of bilinear maps

- bilinear forms $b_1(X, Y), \ldots, b_n(X, Y)$
- in variables $X = \{x_1, \ldots, x_k\}$ and $Y = \{y_1, \ldots, y_m\}$.

Definition

The *rank* of $b_1, \ldots, b_n$ is the smallest number r of bilinear products $P_1, \ldots, P_r$, that is, $P_i = L_i(X) \cdot Q_i(Y)$ with linear forms L_i and Q_i such that

$$b_1, \ldots, b_n \in \langle P_1, \ldots, P_r \rangle.$$

Example:

- Matrix multiplication: $b_{i,j} = \sum_{h=1}^{n} X_{i,h} Y_{h,j}$
- (Univariate) polynomial multiplication: $b_i = \sum_{j=0}^{i} X_j Y_{i-j}$

Tensor rank

Write

$$\sum_{\nu=1}^{n} b_\nu z_\nu = \sum_{h=1}^{k} \sum_{i=1}^{m} \sum_{j=1}^{n} t_{h,i,j} x_h y_i z_j.$$

$$t = (t_{h,i,j}) \in K^k \otimes K^m \otimes K^n$$

is the tensor corresponding to $b_1, \ldots, b_n$.

Definition

$u \otimes v \otimes w = (u_h v_i w_j) \in K^k \otimes K^m \otimes K^n$ is called a *rank-one tensor*.

Definition (Rank)

$R(t)$ is the smallest r such that there are rank-one tensors $t_1, \ldots, t_r$ with $t = t_1 + \cdots + t_r$.

Algebraic proofs

Boolean world

Objects: Boolean functions

Representations: Table of values

Lower bound proofs: Boolean property

Algebraic world

Objects: Polynomials

Representations: List of coefficients

Lower bound proofs: Polynomials (in the lists of coefficients)

Lower bound proofs can only prove that a list of coefficients is not contained in an algebraic variety!

Border complexity

Polynomial multiplication mod X^2:

$$(a_0 + a_1X)(b_0 + b_1X) = \underbrace{a_0b_0}_{f_0} + (\underbrace{a_1b_0 + a_0b_1}_{f_1})X + a_1b_1X^2$$

1	0
0	0

0	1
1	0

Observation

$R(t) = 3$

However, t can be approximated by tensors of rank 2.

$$t(\epsilon) = (1, \epsilon) \otimes (1, \epsilon) \otimes (0, \tfrac{1}{\epsilon}) + (1, 0) \otimes (1, 0) \otimes (1, -\tfrac{1}{\epsilon})$$

1	0
0	0

0	1
1	ϵ

Waring rank

$K[X_1, \ldots, X_n]_d$ denotes all homogeneous degree d polynomials

Definition (Waring rank)

The *Waring rank* of a polynomial $p \in K[X_1, \ldots, X_m]_d$ is the smallest number of homogeneous linear forms $\ell_1, \ldots, \ell_r$ such that

$$p = \ell_1^d + \cdots + \ell_r^d.$$

- finite quantity
- also called symmetric rank

Example:

$$4XY = (X+Y)^2 - (X-Y)^2$$

The discriminant

Simplest case: $\mathbb{C}[X, Y]_2$

- Elements are of the form $aX^2 + bXY + cY^2$
- Algebraic proofs are elements in $\mathbb{C}[a, b, c]$

Fact

There are $\alpha, \beta \in \mathbb{C}$ *such that* $aX^2 + bXY + cY^2 = (\alpha X + \beta Y)^2$ *iff* $b^2 - 4ac = 0$.

- The Waring rank of $4XY$ is 2.
- The set of polynomials of Waring rank 1 equals the set of polynomials of border Waring rank 1.

Natural proofs, rephrased

Boolean:

If private key cryptography works, then

- proofs are hard to compute or
- work only for very specific functions

Algebraic:

If ????, then

- proofs have high circuit complexity

Algebraic natural proofs

Definition (Forbes, Shpilka & Volk)

Let $M \subseteq K[X]$ be a set of monomials.
Let $\mathcal{C} \subseteq \langle M \rangle$ and let $\mathcal{D} \subseteq K[T_m : m \in M]$.

A polynomial $D \in \mathcal{D}$ is an *algebraic $\mathcal{D}$-natural proof against $\mathcal{C}$*, if

1. D is a nonzero polynomial and
2. for all $f \in \mathcal{C}$, $D(f) = 0$, that is, D vanishes on the coefficient vectors of all polynomials in $\mathcal{C}$.

A similar concept was defined independently by Grochow, Kumar, Saks & Saraf.

Succinct hitting sets

Definition

A *hitting set* for $\mathcal{P} \subseteq K[X_1, \ldots, X_\mu]$ is a set $\mathcal{H} \subseteq K^\mu$ such that for all $p \in \mathcal{P}$, there is an $h \in \mathcal{H}$ such that $p(h) \neq 0$.

Definition (Succinct hitting sets)

Let $M \subseteq K[X]$ be a set of monomials.
Let $\mathcal{C} \subseteq \langle M \rangle$ and let $\mathcal{D} \subseteq K[T_m : m \in M]$.

$\mathcal{C}$ is a *$\mathcal{C}$-succinct hitting set for $\mathcal{D}$* if $\mathcal{C}$ viewed as a set of vectors of coefficients of length $|M|$ is a hitting set for $\mathcal{D}$.

The succinct hitting set barrier

Theorem

Let $M \subseteq K[X]$ *be a set of monomials.*
Let $\mathcal{C} \subseteq \langle M \rangle$ *and let* $\mathcal{D} \subseteq K[T_m : m \in M]$.

There are algebraic $\mathcal{D}$*-natural proofs against* $\mathcal{C}$ *iff* $\mathcal{C}$ *is not a* $\mathcal{C}$*-succinct hitting set for* $\mathcal{D}$.

Corollary

Let $\mathcal{C} \subseteq K[X_1, \ldots, X_n]$ *with degree* $\leq d$ *and computable by* $\mathrm{poly}(n, d)$*-size circuits.*
Then there is an algebraic $\mathrm{poly}(N_{n,d})$*-natural proof against* $\mathcal{C}$ *iff there is no* $\mathrm{poly}(n, d)$*-succinct hitting set for* $\mathrm{poly}(N_{n,d})$*-size circuits in* $N_{n,d}$ *variables.*

The succinct hitting set barrier (2)

Typical regime:

- $N_{n,d} = \binom{n+d}{d}$
- $d = \mathrm{poly}(n) \longrightarrow \mathrm{poly}(n) = \mathrm{poly}\log(N_{n,d})$

Conjecture/Wish/Fear

There are $\mathrm{poly}\log(N)$-succinct hitting sets for $\mathrm{poly}(N)$-size circuits.

Tensor rank

Definition

1. A tensor $t \in K^{k \times m \times n}$ has rank-one if $t = u \otimes v \otimes w := (u_h v_i w_j)$ for $u \in K^k$, $v \in K^m$, and $w \in K^n$.
2. The rank $R(t)$ of a tensor $t \in K^{k \times m \times n}$ is the smallest number r of rank-one tensors $s_1, \ldots, s_r$ such that $t = s_1 + \cdots + s_r$.
3. S_r denotes the set of all tensors of rank r.

Definition

$D \in K[X_1, \ldots, X_{kmn}]$ is a $\mathrm{poly}(k, m, n)$-natural proof against S_r if

- D is nonzero,
- D vanishes on S_r, and
- D is computed by circuits of size $\mathrm{poly}(k, m, n)$.

Tensor rank (2)

Good news:

Theorem (Håstad)

Tensor rank is NP*-hard.*

Theorem (Shitov; Schaefer & Stefankovic)

Tensor rank is as hard as the existential theory over $\mathbb{K}$*.*

Bad news:

- S_r is not the zero set of a set of polynomials.
- When D vanishes on S_r, it also vanishes on its closure $\overline{S_r}$.
- $X_r := \overline{S_r}$ is the set of tensors of *border rank* $\leq r$.
- X_r contains tensors of rank $> r$.

(Generalized) matrix completion

Definition

Let $A_0, A_1, \dots, A_m \in K^{n \times n}$. The *completion rank* of $A_0, A_1, \dots, A_m$ is the minimum number r such that there are scalars $\lambda_1, \dots, \lambda_m$ with

$$\mathrm{rk}(A_0 + \lambda_1 A_1 + \cdots + \lambda_m A_m) \leq r.$$

We denote the completion rank by $\mathrm{CR}(A_0, A_1, \dots A_m)$.

- Can also be phrased in terms of an affine linear matrix $A_0 + X_1 A_1 + \cdots + X_m A_m$.

(Generalized) matrix completion (2)

- The set of all $(m+1)$-tuples of $n \times n$-matrices together with m scalars $\lambda_1, \dots, \lambda_m$

$$(A_0, A_1, \dots, A_m, \lambda_1, \dots, \lambda_m) \in K^{(m+1)n^2+m}$$

  such that

$$\mathrm{rk}(A_0 + \lambda_1 A_1 + \cdots + \lambda_m A_m) \leq r$$

  is a closed set, since it is defined by vanishing of all $(r+1) \times (r+1)$-minors.
- Denote this set by $P_r^{m,n}$.
- Let $C_r^{m,n}$ be the projection of $P_r^{m,n}$ onto the first $(m+1)n^2$ components, that is, $C_r^{m,n}$ is the set of all $(A_0, A_1, \dots, A_m)$ with $\mathrm{CR}(A_0, A_1, \dots, A_m) \leq r$.
- $C_r^{m,n}$ is not closed.

Example

- Let

$$A_0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \qquad \text{and} \qquad A_1 = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

  $\mathrm{CR}(A_0, A_1) = 2.$

- Let

$$\underbrace{\begin{pmatrix} 1 & 0 \\ \epsilon & 1 \end{pmatrix}}_{=:A_{0,\epsilon}} + \frac{1}{\epsilon} \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 1/\epsilon \\ \epsilon & 1 \end{pmatrix}.$$

  $\mathrm{CR}(A_{0,\epsilon}, A_1) = 1$ for every $\epsilon \neq 0$.

- $(A_{0,\epsilon}, A_1)$ converges to (A_0, A_1) in the Euclidean topology.
- (A_0, A_1) is contained in the Euclidean closure of C_1.

Closure

Example:

- Let B be any rank-one matrix.
- The completion rank of (I, B) is at least $n - 1$.
- We can approximate B by $B + \epsilon I$.
- But $I - \frac{1}{\epsilon}(B + \epsilon I)$ has rank 1.

Conclusion:

- The rank of the approximating matrices should not be larger than the rank of the matrix itself.
- We take the closure in $K^{n \times n} \times K^{n \times n}_{r_1} \times \cdots \times K^{n \times n}_{r_m}$, where $K^{n \times n}_{\rho}$ denotes the closed set of matrices of rank at most ρ and $r_i = \mathrm{rk}(A_i)$.

Border completion rank

Definition

Let $A_0, A_1, \ldots, A_m \in K^{n\times n}$. The *border completion rank* of $A_0, A_1, \ldots, A_m$ is the minimum number r such that there are approximations $\tilde{A}_i \in K(\epsilon)^{n\times n}_{\mathrm{rk}(A_i)}$ with $\tilde{A}_i = A_i + O(\epsilon)$, $0 \le i \le m$, and rational functions $\lambda_1, \ldots, \lambda_m \in K(\epsilon)$ with

$$\mathrm{rk}(\tilde{A}_0 + \lambda_1\tilde{A}_1 + \cdots + \lambda_m\tilde{A}_m) \le r.$$

We denote the border completion rank by $\underline{\mathrm{CR}}(A_0, A_1, \ldots A_m)$.

Hardness of completion rank

- ϕ formula in 2-CNF over the variables $x_1, \ldots, x_t$ with clauses $c_1, \ldots, c_s$.
- Given b, it is NP-hard to decide whether there is an assignment satisfying at least b clauses.

Clause gadget: $c_i = L_1 \vee L_2$

$$\begin{pmatrix} 1-\ell_1 & 1 \\ 0 & 1-\ell_2 \end{pmatrix}$$

- ℓ_j in the matrix is x_k if the literal $L_j = x_k$ and it is $1 - x_k$ if $L_j = \neg x_k$, $j = 1, 2$.

Observation

The clause gadget has rank 1 *iff at least one of the literals* ℓ_1, ℓ_2 *is set to be* 1. *Otherwise, it has rank* 2.

Hardness of completion rank (2)

- All clause gadgets are blocks of our desired block diagonal matrix.
- We get a matrix $A_0 + x_1 A_1 + \cdots + x_t A_t$ with affine linear forms as entries

Proposition

$\mathrm{CR}(A_0, A_1, \ldots, A_t) \leq 2s - b$ *iff* b *clauses of* ϕ *can be satisfied.*

Thus the problem $\mathrm{CR}(A_0, A_1, \ldots, A_t) \overset{?}{\leq} k$ *is* NP*-hard.*

Hardness of border completion rank

Observation

Each A_i, $i \geq 1$, is a diagonal matrix with diagonal entries ± 1. If the j^{th} diagonal entry of A_i is nonzero, then the j^{th} diagonal entry of any other A_k is zero, $i, k \geq 1$.

Let $\tilde{A}_0, \tilde{A}_1, \ldots, \tilde{A}_t$ be approximations to $A_0, A_1, \ldots, A_t$, that is, $\tilde{A}_i = A_i + O(\epsilon)$.

Lemma

There are (invertible) matrices $S = I_n + O(\epsilon)$ and $T = I_n + O(\epsilon)$ such that $S \cdot (\tilde{A}_0 + \lambda_1 \tilde{A}_1 + \cdots + \lambda_t \tilde{A}_t) \cdot T = \widehat{A_0} + \lambda_1 A_1 + \cdots + \lambda_t A_t$ for some $\widehat{A_0} = A_0 + O(\epsilon)$.

Hardness of border completion rank (2)

Lemma

$\underline{\mathrm{CR}}(A_0, A_1, \ldots, A_t) \leq 2s - b$ *iff* b *clauses of* ϕ *can be satisfied.*

- $\Leftarrow$ follows from hardness proof for CR.
- Assume there are $\lambda_i = a_{i,0}\epsilon^{d_i} + a_{i,1}\epsilon^{d_i+1} + \ldots$ with $a_{i,0} \neq 0$ such that $\mathrm{rk}(\tilde{A}_0 + \lambda_1 A_1 + \cdots + \lambda_t A_t) \leq 2s - b$.
- λ_i induce an assignment to the x_i and thus to literals ℓ_j.
- A clause gadget looks like

  $$\begin{pmatrix} 1 + O(\epsilon) - \ell_1 & 1 + O(\epsilon) \\ O(\epsilon) & 1 + O(\epsilon) - \ell_2 \end{pmatrix}$$

  To have rank 1, $\ell_1 = 1 + O(\epsilon)$ or $\ell_2 = 1 + O(\epsilon)$. We call such clauses "ϵ-satisfied".
- If we have at least b "ϵ-satisfied" clauses, then we substitute $\epsilon = 0$ in corresponding λ_i and get an exact assignment.
- If there are $< b$ ϵ-satisfied clauses, then $\underline{\mathrm{CR}}(A_0, A_1, \ldots, A_t) > 2s - b$.

Algebraic natural proofs for border completion rank

Let $t \in K^{n \times n \times (m+1)}$. An *algebraic* $\mathrm{poly}(n)$-*natural proof* for the border completion rank of t being $> r$ is a polynomial $P \in K[X_{h,i,j} | 1 \leq h, i \leq n,\ \ 1 \leq j \leq m]$ such that

1. $P(t) \neq 0$,
2. $P(s) = 0$ for every $s \in K^{n \times n \times (m+1)}$ with $\underline{\mathrm{CR}}(s) \leq r$.
3. P is computed by a constant-free algebraic circuit of size $\mathrm{poly}(n)$.

Universal tensors

Observation

Let $U_{i,j}, V_{i,j}$, $1 \leq i \leq \rho$, $1 \leq j \leq n$ *be indeterminates. If we substititute arbitrary constants for the indeterminates in* $\sum_{i=1}^{\rho}(U_{i,1}, \dots, U_{i,n})^T(V_{i,1}, \dots, V_{i,n})$, *then we get all matrices in* $K_{\rho}^{n\times n}$

Lemma

Let $Q_0, Q_1, \dots, Q_t$ *be polynomial matrices as in the observation above having ranks* $r_0, \dots, r_t$, *respectively. We use fresh variables for each* Q_i.
Let $g := (Q_0 - Z_0 Q_1 - \cdots - Z_t Q_t, Q_1, \dots, Q_t)$, *where* $Z_1, \dots, Z_t$ *are new variables. If we substitute arbitrary constants for the indeterminates, then we get all tensors of completion rank* $\leq r_0$ *with the* i^{th} *slice having rank* $\leq r_i$, $1 \leq i \leq t$.

Main result

Theorem

For infinitely many n, there is an m, a tensor $t \in K^{n \times n \times m}$ and a value r such that there is no algebraic $\mathrm{poly}(n)$-natural proof for the fact that $\underline{\mathrm{CR}}(t) > r$ unless coNP $\subseteq \exists$BPP.

- Let ϕ be a formula in 2-CNF and let $b \in \mathbb{N}$. We want to check whether every assignment satisfies $< b$ clauses of ϕ. This problem is coNP-hard.
- Let $T_\phi = (A_0, \dots, A_t)$ be the tensor constructed above.
- Guess a circuit C of polynomial size computing some P.
- Decide whether $P(g) = 0$ using polynomial identity testing.
- Check whether $P(T_\phi) \neq 0$. If yes, then accept. Otherwise reject.

Relation to tensor (border) rank

Theorem (Derksen)

If $t = (A_0, A_1, \ldots, A_m)$ *is a concise tensor such that* $\mathrm{rk}(A_1) = \cdots = \mathrm{rk}(A_m) = 1$*. Then*

$$R(t) = \mathrm{CR}(t) + m.$$

Proposition

If $t = (A_0, A_1, \ldots, A_m)$ *is a tensor such that* $\mathrm{rk}(A_1) = \cdots = \mathrm{rk}(A_m) = 1$*. Then*

$$\underline{R}(t) \leq \underline{\mathrm{CR}}(t) + m.$$

Approximability of tensor rank

With similar methods, we can prove that tensor is hard to approximate.

Fun fact:

- Johan Håstad asked me this question about 15 years ago on the train back from Oberwolfach.
- Independently of our work, hardness of approximation was recently shown by Song, Woodruff and Zhong (under ETH)
- and Joseph Swernofsky

Tensor rank is hard to approximate

- Let ϕ be a formula in 3-CNF with t variables and s clauses such that every variable appears in a constant number c of clauses. Note that $s = O(t)$.
- We construct a matrix completion problem as before.
- We will have variable gadgets and clause gadgets.
- They will appear as blocks on the main diagonal.
- **Problem:** Everything needs to be of rank 1.

Variable gadget

$$\begin{pmatrix} 1 & x & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & u & 0 & u-u_1 & 0 & u-u_2 & 0 & 0 \\ 0 & u-u_3 & 1 & u & 0 & u-u_4 & 0 & 0 \\ 0 & 0 & 1 & v & 0 & 0 & 0 & 2v-v_1 \\ 0 & u-u_5 & 0 & u-u_6 & 1 & u & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & w & 2w-w_1 & 0 \\ 0 & 0 & 0 & v-v_2 & 0 & 0 & 1 & 2(v-1/2) \\ 0 & 0 & 0 & 0 & 0 & w-w_2 & 2(w-1/2) & 1 \end{pmatrix}$$

Lemma

1. *If x is set to 0 or 1, then the local variables in the variable gadget can be set such that the resulting matrix has rank 4.*
2. *If the variables are set in such a way that the rank of the variable gadget is 4, then x is set to 0 or 1.*

Variable gadget

$$
\begin{pmatrix}
1 & x & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & u & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & u & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & v & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & u & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & w & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 2(v-1/2) \\
0 & 0 & 0 & 0 & 0 & 0 & 2(w-1/2) & 1
\end{pmatrix}
$$

Lemma

1. *If x is set to 0 or 1, then the local variables in the variable gadget can be set such that the resulting matrix has rank 4.*
2. *If the variables are set in such a way that the rank of the variable gadget is 4, then x is set to 0 or 1.*

Clause gadget

$$\begin{pmatrix} 1 & x & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & u & 0 & 0 & 0 & 0 & s(u)-u_1 & 0 & 0 \\ 0 & 0 & 1 & y & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & v & 0 & 0 & 0 & s(v)-v_1 & 0 \\ 0 & 0 & 0 & 0 & 1 & z & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & w & 0 & 0 & s(w)-w_1 \\ 0 & u-u_2 & 0 & 0 & 0 & 0 & 1-\ell(u) & 1 & 0 \\ 0 & 0 & 0 & v-v_2 & 0 & 0 & 0 & 1-\ell(v) & 1 \\ 0 & 0 & 0 & 0 & 0 & w-w_2 & 0 & 0 & 1-\ell(w) \end{pmatrix}$$

- $\ell(u) = u$ if x appears positive in the clause and $\ell(u) = 1 - u$ otherwise.
- $s(u) = -u$ if x appears positive in the clause and $s(u) = u$ otherwise.

Hardness of approximation

Lemma

Assume that ϕ is either satisfiable or any assignment satisfies at most $(1-\epsilon)$ of the clauses for some $\epsilon > 0$.

1. *If ϕ is satisfiable, then the tensor rank of T_ϕ is $4t + 5s$.*
2. *If ϕ is not satisfiable, then the tensor rank of T_ϕ is at least $4t + 5s + \delta t$ for some constant $\delta > 0$.*

Theorem

Tensor rank is NP*-hard to approximate.*

The permanent

Permanent:

$$\operatorname{per}_n X = \sum_{\pi \in S_n} X_{1,\pi(n)} \cdots X_{n,\pi(n)}$$

- complete for VNP
- evaluation at $\{0,1\}$-matrices is #P-hard under Turing reductions.

Consider the hypersurface

$$\mathcal{Z}_n = \{M \in K^{n \times n} \mid \operatorname{per} M = 0\}.$$

How hard is it to prove that some M is not in $\mathcal{Z}_n$?

Matrices with permanent zero

Let X be an $n \times n$ matrix. Construct a matrix Z as follows:

$$\begin{cases} z_{ij} = x_{ij} & \text{for } i \leq n-1, \\ z_{nj} = x_{nj} \operatorname{per} X_{nn} & \text{for } j \leq n-1, \\ z_{nn} = -\sum_{j=1}^{n-1} x_{nj} \operatorname{per} X_{nj}, & \end{cases}$$

where X_{ij} is the matrix obtained from X by removing the i^{th} row and the j^{th} column.

Observation

We have $\operatorname{per} Z = 0$. *Moreover, any matrix with* $\operatorname{per} Z = 0$ *and* $\operatorname{per} Z_{nn} \neq 0$ *can be obtained in this way.*

Natural proofs for matrices with permanent zero

Theorem

Let $\mathcal{Z}_n \subseteq \mathsf{K}^{n\times n}$ be the set of matrices with permanent 0.
If $\mathcal{Z}_n$ has algebraic VP^0-natural proofs, then $\mathsf{P}^{\#\mathsf{P}} \subseteq \exists\mathsf{BPP}$.

- Construct iteratively a polynomial size circuit computing per_k.
- Using the circuit for per_{k-1} compute a small circuit computing Z_k.
- Guess a polynomial size circuit C_k vanishing on $\mathcal{Z}_k$
- Verify this by checking $C_k(Z_k) = 0$.
- By Hilbert's Nullstellensatz, per_k^e divides C_k.
- Compute a small circuit for per_k using Kaltofen's factoring algorithm.

GCT breaks the algebraic natural proofs barrier

- $\mathcal{Z} \subseteq \mathbb{C}^{n \times n}$ all matrices with permanent 0.
- $\mathrm{GL}_n \times \mathrm{GL}_n$ acts on $\mathbb{C}^{n \times n}$ via left-right multiplication:

$$(g_1, g_2) \cdot A := g_1 A (g_2)^T.$$

- Let $Q_n \subseteq \mathrm{GL}_n$ denote the group of monomial matrices, i.e., matrices with nonzero determinant that have a single nonzero entry in each row and column.
- $\mathcal{Z}$ is closed under the action of the group $G := Q_n \times Q_n \subseteq \mathrm{GL}_n \times \mathrm{GL}_n$, which means that if $A \in Z$, then $gA \in Z$ for all $g \in G$.

The GCT framework

- Assume that $A \in \mathcal{Z}$.
- $GA := \{gA \mid g \in G\}$ is contained in $\mathcal{Z}$
- $\overline{GA} \subseteq \mathcal{Z}$ as a subvariety.
- For a Zariski-closed subset $Y \subseteq \mathbb{C}^{n \times n}$ let $I(Y) \subseteq \mathbb{C}[\mathbb{C}^{n \times n}]$ denote the *vanishing ideal* of Y.
- $I(Y)_d$ the homogeneous degree d component of $I(Y)$. (inherits grading)
- *Coordinate ring* $\mathbb{C}[Y]$ of Y is the quotient $\mathbb{C}[Y] := \mathbb{C}[\mathbb{C}^{n \times n}]/I(Y)$, inherits the grading $\mathbb{C}[Y]_d := \mathbb{C}[\mathbb{C}^{n \times n}]_d / I(Y)_d$.
- Since $\overline{GA} \subseteq \mathcal{Z}$, $I(\mathcal{Z})_d \subseteq I(\overline{GA})_d$ for all d.
- Canonical surjection by restriction: $\mathbb{C}[\mathcal{Z}]_d \twoheadrightarrow \mathbb{C}[\overline{GA}]_d$

Representations

Definition

- An H-*representation* is a finite dimensional vector space V with a group homomorphism $\rho : H \to \mathrm{GL}(V)$.
  We write gf for $(\rho(g))(f)$.
- A linear map $\varphi : V_1 \to V_2$ between two H-representations is called *equivariant* if for all $g \in H$ and $f \in V_1$, $\varphi(gf) = g\varphi(f)$.
- A bijective equivariant map is called an H-isomorphism.
- Two H-representations are called *isomorphic* if an H-isomorphism exists from one to the other.
- A linear subspace of an H-representation that is closed under the action of H is called a *subrepresentation*.
- An H-representation whose only subrepresentations are itself and 0 is called *irreducible*.

Representations (2)

- *Canonical pullback*: $(gf)(B) := f(g^T B)$ for $g \in G$, $f \in \mathbb{C}[Y]$, $B \in \mathbb{C}^{n \times n}$.
- Turns $\mathbb{C}[\mathcal{Z}]_d$ and $\mathbb{C}[\overline{GA}]_d$ into G-representations.
- G is *linearly reductive*, which means that every G-representation V decomposes into a direct sum of irreducible representations.
- For each type λ the *multiplicity* $\mathrm{mult}_\lambda(V)$ of λ in V is unique.

Lemma (Schur)

For an equivariant map $\varphi : V \to W$, the image $\varphi(V)$ is a G-representation and $\mathrm{mult}_\lambda(V) \geq \mathrm{mult}_\lambda(\varphi(V))$.

- The map $\mathbb{C}[\mathcal{Z}]_d \twoheadrightarrow \mathbb{C}[\overline{GA}]_d$ is equivariant, thus

$$\mathrm{mult}_\lambda(\mathbb{C}[\mathcal{Z}]_d) \geq \mathrm{mult}_\lambda(\mathbb{C}[\overline{GA}]_d).$$

- A λ that violates this is an *obstruction* and proves "$A \notin \mathcal{Z}$".

Main result

Theorem

Let $G := Q_n \times Q_n$ *and* $\nu := (((1^n),(n)),((1^n),(n)))$*. Then*

- $\mathrm{mult}_\nu(\mathbb{C}[Z]_n) = 0$ *and*
- $\mathrm{mult}_\nu(\mathbb{C}[\overline{GA}]_n) = \begin{cases} 0 & \textit{if } A \in Z \\ 1 & \textit{otherwise} \end{cases}.$

- Subrepresentation is $\langle \mathrm{per} \rangle$ with $\mathrm{mult}_\nu\, \mathbb{C}[\mathbb{C}^{n\times n}]_n = 1$.
- $\mathrm{mult}_\nu(\mathrm{I}(\mathcal{Z})_n) = 1$ and thus $\mathrm{mult}_\nu(\mathbb{C}[\mathcal{Z}]_n) = 0$.
- For $A \in \mathcal{Z}$, $\overline{GA} \subseteq \mathcal{Z}$. Therefore $\mathrm{mult}_\nu(\mathbb{C}[\overline{GA}]_n) = 0$.
- For $A \notin \mathcal{Z}$, $\mathrm{mult}_\nu(\mathrm{I}(\overline{GA}))_n) = 0$ and therefore $\mathrm{mult}_\nu(\mathbb{C}[\overline{GA})]_n = 1)$.

Thank You!