Worst-case to average-case reductions via additive combinatorics

[STOC '22]

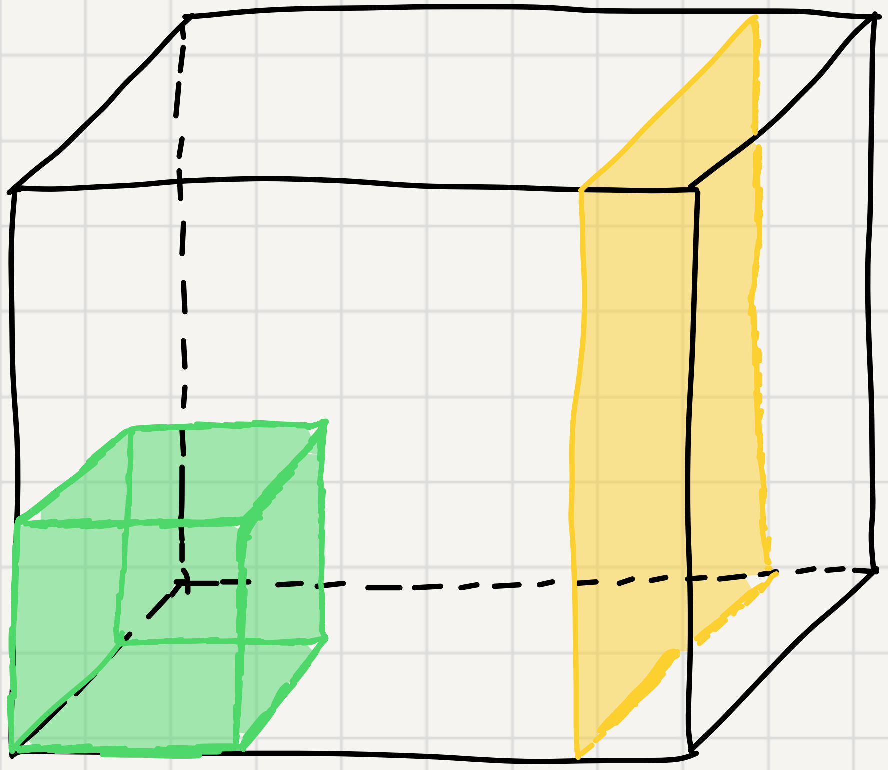

Vahid Asadi, Alexander Golovnev, Tom Gur, Igor Shinkar

Plan

① What we do, and why?

② A naive approach + challenges

③ The key idea: local-correction via additive combinatorics

④ Sketch of proof

① What we do, and why?

"Boosting knowledge" via average-to-worst case reductions

"Boosting knowledge" via average-to-worst case reductions

Suppose we know how to solve a problem on few instances

Can we derive how to solve all of them?

Example– Matrix multiplication

Problem: Given $A, B \in \mathbb{F}^{n\times n}$, compute $A \cdot B$.

Suppose ALG s.t.

$$\Pr_{A,B\in\mathbb{F}^{n\times n}}[ALG(A,B) = A\cdot B] \geq \alpha$$

$\alpha = 2/3$

average-case

$$\boxed{\begin{array}{c}\alpha = 0.01 \\ \alpha = o(1)\end{array}}$$

Can we boost α to 1?

Two perspectives

Worst-case to average-case reductions

OPTIMIST

"A new paradigm for designing algorithms!"

PESSIMIST

"Show lower bounds even for weak average-case complexity"

This work focuses on the 1% regime.
and even $o(1)$!

Our contribution

A new framework for worst-case to average-case reductions via local-correction lemmas based on additive combinatorics

Applications

1) Algorithms for matrix multplication

2) Data structures for all linear problems

3) Online matrix-vector multiplication

4) weak average case: polygnomial evaluation

② naive approach + challenges

via a concrete example

This talk

Illustrate the framework via matrix multiplication

Theorem: If there exists ALG running in time T

s.t. $\Pr_{A,B\in\mathbb{F}^{n\times n}}[ALG(A,B) = A\cdot B] \geq \alpha$,

then there exists ALG' running in time $O(T)$

s.t. for all $A,B\in\mathbb{F}^{n\times n}$, w.p. $1-\delta$

$$ALG(A,B) = A\cdot B$$

Remark: $O(T)$ hides a factor of roughly $1/\delta\cdot\alpha$

A trivial special case: high-agreement regime

Suppose $\Pr_{A,B\in\mathbb{F}^{n\times n}}[\mathrm{ALG}(A,B)=A\cdot B]\geq 0.99$ → Same holds for $\alpha\geq\frac{3}{4}+\varepsilon$

Idea: BLR-type local correction

1) Sample $R,S\sim U(\mathbb{F}^{n\times n})$

2) Write $A=R+(A-R)$, $B=S+(B-S)$ → Each component is uniformly distributed

3) Compute $M=\mathrm{ALG}(R,S)+\mathrm{ALG}(A-R,S)+\mathrm{ALG}(R,B-S)+\mathrm{ALG}(A-R,B-S)$

Note that $\Pr[M=A\cdot B]\geq 1-4\cdot 0.01>9/10$

The challenge: low-agreement regime

In the 1% regime ($\Pr[\mathrm{ALG}(A,B) = A \cdot B] \geq 0.01$)

this approach completely breaks

Even for $\alpha < 0.75$

Example:

Suppose $\mathrm{ALG}(A,B) = A \cdot B \iff A_{11} = 0$ (o/w $\mathrm{ALG}(A,B) = \bar{0}$)

Here $\alpha = 1/2$, yet no such decomposition self-correct $A \cdot B$ if $A_{11} = 1$

$$\begin{bmatrix} 1 & \\ & \end{bmatrix} = \begin{bmatrix} 0 & \\ & \end{bmatrix} + \begin{bmatrix} 1 & \\ & \end{bmatrix}$$

Is all hope lost?

③ The key idea:

local-correction via additive combinatorics

Local correction via additive combinatorics

A-C studies approximate notions of algebraic structures via the perspective of combinatorics, number theory, harmonic analysis.

The sumset of a set X is defined as

$X+X=\{x_1+x_2 : x_1,x_2 \in X\}$. Generally: $t\cdot X=\{\sum_{i=1}^{t} x_i : x_1,\dots,x_t \in X\}$.

These notions quantify a combinatorial analogue of approximate subgroup structure.

Small sumsets imply approximate closure.

Bogolyubov's lemma

Let $X \subseteq \mathbb{F}_2^n$ of density $\frac{|X|}{2^n} \geq \alpha$. Then, there exists a subspace $V \subseteq 4X$ of dimension $\dim(V) \geq n - 1/\alpha^2$.

key idea: Use Bogolyubov's lemma for local correction!

How? Suppose $\Pr_{A,B \in \mathbb{F}^{n \times n}}[ALG(A,B) = A \cdot B] \geq \alpha$

Denote $X = \{(A,B) : ALG(A,B) = A \cdot B\}$. Note $\mu(X) \geq \alpha$

Hence, there exist a large subspace V s.t. $v \in V$ decomposes to $v = x_1 + x_2 + x_3 + x_4$, $x_1, \dots, x_4 \in X$

Problems with Bogolyubov local correction

Denote $X = \{(A,B) : ALG(A,B) = A \cdot B\}$. Note $\mu(X) \geq \alpha$.

Hence, there exist a large subspace V s.t. $v \in V$ decomposes to $V = x_1 + x_2 + x_3 + x_4$, $x_1, \dots, x_4 \in X$

☐ How do we obtain the decomposition?

☐ How to deal with inputs outside of V?

Obtaining the decomposition

All we shall need is a probabilistic Bogolyubov lemma.

Lemma: Let $X \subseteq \mathbb{F}_2^n$ s.t. $|X|/2^n \geq \alpha$.

There exist a subspace V of dim $n - \frac{1}{\alpha^2}$

s.t. $\forall v \in V \quad \Pr_{x_1,x_2,x_3}[x_1, x_2, x_3, v - x_1 - x_2 - x_3 \in X] \geq \alpha^5$.

This will suffice for matrix multiplication over $\mathbb{F}_2^n$, and constant $\alpha > 0$.

④ Sketch of proof

Matrix multiplication – Sketch of proof

Problem: Given $A, B \in \mathbb{F}^{n \times n}$, compute $A \cdot B$.

Suppose ALG s.t.

$$\Pr_{A,B \in \mathbb{F}^{n \times n}}[\mathrm{ALG}(A,B) = A \cdot B] \geq \alpha$$

Goal: boost α to 1.

For simplicity, assume:

1) ALG is deterministic.
2) $\mathbb{F} = \mathbb{F}_2$.
3) α is a constant.
4) The input (A,B) satisfies $\Pr_{B'}[\mathrm{ALG}(A,B') = A \cdot B'] \geq \alpha$.

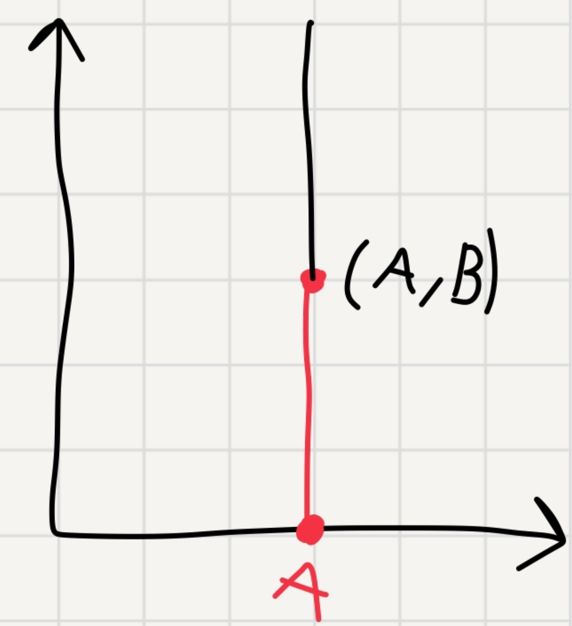

Two simple facts

1) Given a potentially wrong output $A \cdot B$, we can efficiently check the solution via Freivald's algorithm.

Lemma: Given $A, B, C \in \mathbb{F}^{n \times n}$, there exists a prob. alg. verifing $A \cdot B = C$ with high probability, in time $O(n^2)$

2) Denoting by $X = \{B' \in \mathbb{F}^{n \times n} : ALG(A, B') = A \cdot B'\}$ the "good" B's, if $B \in X$, then ALG is successful.

Hence, the goal is to locally correct $B \notin X$.

Local correction via Bogolybov's lemma

First idea: use Bogolyubov's lemma to locally-correct inputs that lie in a large subspace.

Lemma: Let $X \subseteq \mathbb{F}_2^n$ s.t. $|X|/2^n \geq \alpha$.

There exist a subspace V of dim $n - \frac{1}{\alpha^2}$

s.t. $\forall v \in V$ $\Pr_{x_1,x_2,x_3}[x_1, x_2, x_3, v - x_1 - x_2 - x_3 \in X] \geq \alpha^5$.

Given $X = \{B' \in \mathbb{F}^{n \times n} : ALG(A, B') = A \cdot B'\}$ we obtain $V \subseteq \mathbb{F}^n$ with $\dim(V) \geq n^2 - 1/\alpha^2$ s.t. $\forall$ $B' \in V$

$\Pr[M_1, M_2, M_3, M_4 \in X] \geq \alpha^5$, where $M_4 = B' - M_1 - M_2 - M_3$

Local correction via Bogolybov's lemma

Given $X = \{B' \in \mathbb{F}^{n\times n} : ALG(A,B') = A\cdot B'\}$ we obtain $V \subseteq \mathbb{F}^n$ with $\dim(V) \geq n^2 - 1/\alpha^2$ s.t. $\forall\ B' \in V$

$$\Pr[M_1, M_2, M_3, M_4 \in X] \geq \alpha^s, \text{ where } M_4 = B' - M_1 - M_2 - M_3$$

If this event occurs, then

$$\sum_{i=1}^{4} ALG(A, M_i) = \sum_{i=1}^{4} A\cdot M_i = A\cdot\left(\sum_{i=1}^{4} M_i\right) = A\cdot B'$$

as required.

But success probability α^s is far smaller than desired...

We amplify $O(1/\alpha^s)$ times via Freivald's algorithm!

And what about $B' \notin V$?

Local correction via Bogolybov's lemma

Recap: good inputs $X = \{B' \in \mathbb{F}^{n \times n} : ALG(A, B') = A \cdot B'\}$

Bogolybov subspace $V \subseteq 4X$

Case 1: If $B' \in X$, just run $ALG(A, B')$

Case 2: If $B' \in V$, locally correct via Bogolyubov's lemma

Case 3: If $B' \notin V$, um... did we really gain anything?

We started with X of density α

V has smaller density... **but it has structure!**

Low-rank random matrix shifts

<u>Observation:</u> If $A \in \mathbb{F}^{n \times n}$ has rank K, then AB can be computed in time $O(K \cdot n^2)$ given a rank-K decomposition

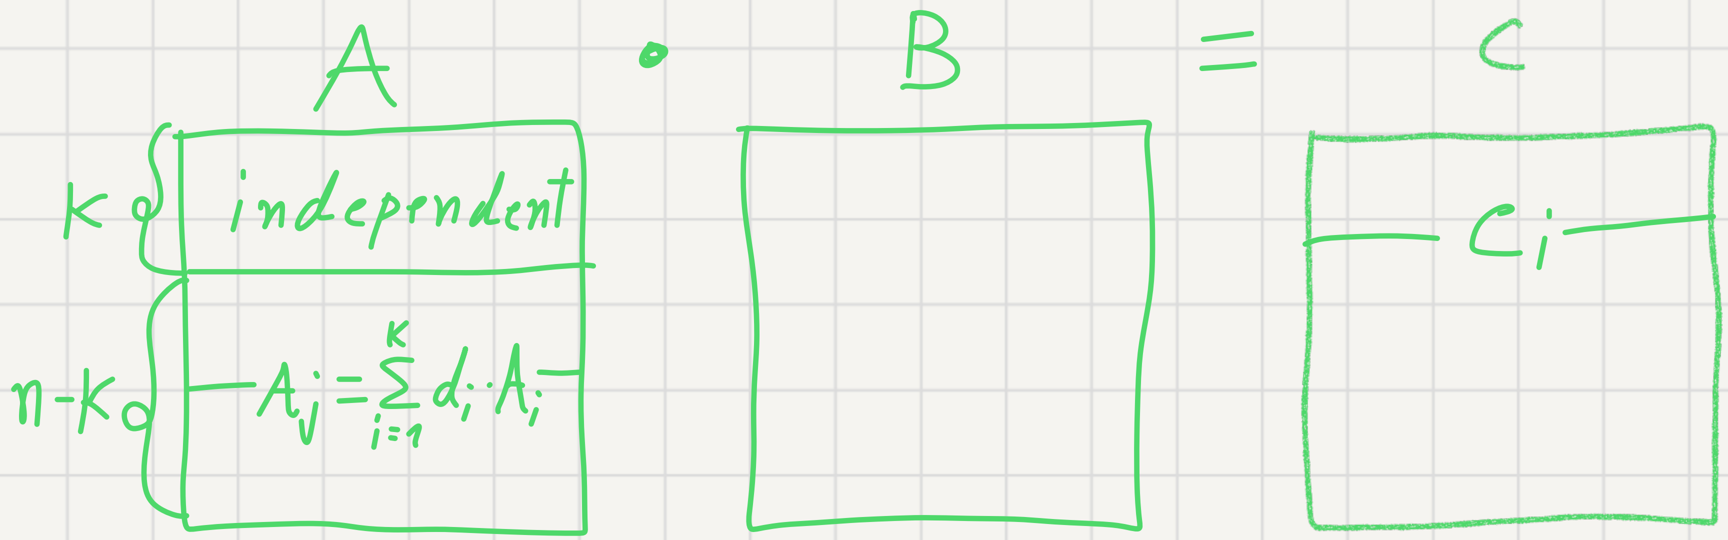

- Multiply the indep. rows in time $O(K \cdot n^2)$
- To compute the rest, let $C_i = A_i \cdot B \quad \forall i \in [K]$, observe $A_j = \sum_{i=1}^{K} d_i \cdot A_i \Rightarrow C_j = \sum_{i=1}^{K} d_i \cdot C_i$, can be computed in time $O(K \cdot n)$ for each $j \in [n]$.

The reduction

$$R_B: \quad \begin{array}{r|c|} \hline k\{ & \text{random independent} \\ \hline n-k\{ & (R_B)_j = \sum_{i=1}^{k} d_i (R_B)_i \\ \hline \end{array}$$

1) Given (A, B), sample $R_B \in \mathbb{F}^{n \times n}$.

2) Observe that if $\dim(V) = n^2 - k$, then $\Pr[B + R_B \in V] \geq \frac{1}{2|\mathbb{F}|^k}$.

3) Compute $ALG_V(A, B + R_B)$.

4) Verify using Freivald's algorithm.

5) Compute $A \cdot R_B$ in time $O(k \cdot n^2)$.

6) Return $ALG_V(A, B + R_B) - A \cdot R_B$.

Summary & open problems

Worst-case to average-case reductions

OPTIMIST

PESSIMIST

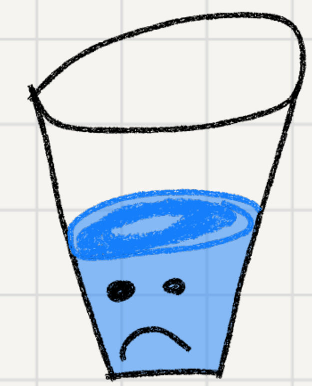

"A new paradigm for designing algorithms!"

"Show lower bounds even for weak average-case complexity"

This work focuses on the 1% regime.
and even $o(1)$

Summary & open problems

A new framework for worst-case to average-case reductions via local-correction lemmas based on additive combinatorics

Applications

1) Algorithms for matrix multplication

2) Data structures for all linear problems

3) Online matrix-vector multiplication

4) weak average case: polynomial evaluation

Summary & open problems

- Can the framework be extended to other models? E.g., communication complexity, property testing, and PAC learning?

- Can we obtain average-to-worst case reductions for all linear problems for both circuits and uniform algorithms?

Our work reduces the above to efficient verification for the problem.

Thank you!