A note on the discrepancy of matrices with bounded row and column sums

Nicholas J. A. Harvey[1]

Department of Computer Science, University of British Columbia, Vancouver, Canada

Abstract

A folklore result uses the Lovász local lemma to analyze the discrepancy of hypergraphs with bounded degree and edge size. We generalize this result to the context of real matrices with bounded row and column sums.

1. Introduction

In combinatorics, discrepancy theory is the study of red-blue colorings of a hypergraph's vertices such that every hyperedge contains a roughly equal number of red and blue vertices. A classic survey on this topic is [3].

Many combinatorial discrepancy results have a more general form as a geometric statement about discrepancy of real vectors [3, §4]. Some examples include the Beck-Fiala theorem [2] and Spencer's "six standard deviations" theorem [9]. One exception is the following folklore result on the discrepancy of hypergraphs of bounded degree and edge size [10, pp. 693] [4, Proposition 12].

Theorem 1. *Let H be a hypergraph of maximum degree Δ and maximum edge size R. Then there is a red-blue coloring of the vertices such that, for every edge e, the numbers of red and blue vertices in e differ by at most $2\sqrt{R\ln(R\Delta)}$.*

The proof is a simple consequence of the Lovász local lemma.

We show that this theorem also has a more general form as a geometric statement about discrepancy of real vectors. Theorem 2 recovers Theorem 1 (up to constants) by letting $V_{i,j} \in \{0,1\}$ indicate whether vertex j is contained in edge i. As usual, let $[m] = \{1,\ldots,m\}$ and let $\|\cdot\|_p$ denote the ℓ_p-norm.

Theorem 2. *Let V be an $m \times n$ real matrix, let v^i denote its i^{th} row and let v_j denote its j^{th} column. Suppose that*

- $\|v^i\|_1 \leq R$,
- $|V_{i,j}| \leq 1$ *for all i, j, and*
- $\|v_j\|_1 \leq \Delta$ *for all j.*

Assume that $R \geq 4$ and $\Delta \geq 2$. There exists $y \in \{-1,+1\}^n$ with $\|Vy\|_\infty \leq O(\sqrt{R\log(R\Delta)})$.

It is worth calling attention to the main principle underlying this theorem. Ordinarily, one thinks of applying the local lemma in scenarios where the dependencies between the events are described by a sparse graph. In our scenario, that would usually be taken to mean that each column has small support. The main principle of our theorem is that, in our scenario, sparsity can be measured by the ℓ_1-norm of the columns rather than by the support size (the so-called ℓ_0-norm).

Email address: nickhar@cs.ubc.ca (Nicholas J. A. Harvey)

[1]Supported by an NSERC Discovery Grant and a Sloan Foundation Fellowship. Phone: (604) 827-5941.

2. The Proof

In order to prove Theorem 2, it is convenient to eliminate negative values and to rescale the entries. Our main technical result is as follows. Let $\lg x$ denote the base-2 logarithm of x.

Theorem 3. *Let A be a non-negative real matrix of size $m \times n$, and let $a_1, \ldots, a_n \in \mathbb{R}^m_{\geq 0}$ denote its columns. Suppose that*

- $\|\sum_j a_j\|_\infty \leq 1$,
- $A_{i,j} \leq \beta$ *for all* i, j, *and*
- $\|a_j\|_1 \leq \delta$ *for all* j.

Assume that $\beta \leq \min\{\delta/2, 1/4\}$. Define $\alpha := \sqrt{\lg(\delta/\beta^2)} > \sqrt{2}$. Then there exists a vector $y \in \{-1,+1\}^n$ such that

$$\|Ay\|_\infty \;\leq\; 16\alpha\sqrt{\beta}.$$

Theorem 2 is derived from Theorem 3 by rescaling the vectors and separately considering the positive and negative coordinates.

Proof of Theorem 2. Define $v_j^+, v_j^- \in \mathbb{R}^n$ by

$$\begin{aligned}(v_j^+)_i &= \max\{(v_j)_i/R, 0\}\\ (v_j^-)_i &= \max\{-(v_j)_i/R, 0\}.\end{aligned}$$

Let $a_j \in \mathbb{R}^{2n}$ be the vector obtained by concatenating v_j^+ and v_j^-, and let A be the non-negative matrix whose j^{th} column is a_j. Then $0 \leq A_{i,j} \leq 1/R$, $\|a_j\|_1 = \|v_j\|_1/R \leq \Delta/R$, and $\|\sum_j a_j\|_\infty \leq 1$. Applying Theorem 3 with $\delta = \Delta/R$ and $\beta = 1/R$, there must exist a vector $y \in \{-1,+1\}^m$ with

$$\|Ay\|_\infty \;\leq\; 16\alpha\sqrt{\beta} \;=\; O(\sqrt{\lg(R\Delta)/R}).$$

Since $\|Vy\|_\infty \leq 2R\|Ay\|_\infty$, we obtain that $\|Vy\|_\infty = O(\sqrt{R\lg(R\Delta)})$. □

We now turn to the proof of Theorem 3. Suppose we choose the vector $y \in \{-1,+1\}^n$ uniformly at random. The discrepancy of row i is the value $|\sum_j A_{i,j}y_j|$. Our goal is to bound $\|Ay\|_\infty = \max_i |\sum_j A_{i,j}y_j|$, which is the maximum discrepancy of any row.

One challenge in analyzing $\|Ay\|_\infty$ is that the entries of A can have various magnitudes, so one row's discrepancy can have various degrees of dependence on all other rows' discrepancy. The Lovász local lemma treats events as either dependent or independent, and cannot easily deal with mild amounts of dependence.

A natural approach to address this issue is to partition each row of A into sets whose entries all have roughly the same magnitude. Define $b := \lfloor -\lg\beta \rfloor \geq 2$, so that every entry of every A is at most 2^{-b}. For $k \geq b$, let

$$S_{i,k} \;=\; \{\, j \,:\, \lfloor -\lg A_{i,j} \rfloor = k \,\}$$

be the locations of the entries in row i that take values in $(2^{-(k+1)}, 2^{-k}]$.

To bound the discrepancy of row i, we will actually bound the discrepancy of each set $S_{i,k}$ (i.e., $|\sum_{j \in S_{i,k}} A_{i,j}y_j|$). The amount of discrepancy of $S_{i,k}$ that we will allow is a quantity T_k that is carefully chosen such that the total discrepancy of row i will be small, but yet the Lovász local lemma can still be applied. By the triangle inequality, the total discrepancy of row i is at most the sum of the discrepancies of each $S_{i,k}$.

Define

$$\epsilon \;:=\; 8\alpha\sqrt{\beta} \;>\; 8\sqrt{\beta}. \tag{1}$$

Let $\mathcal{E}_{i,k}$ be the event that the discrepancy of $S_{i,k}$ exceeds

$$T_k \;:=\; \epsilon \sum_{j\in S_{i,k}} A_{i,j} + \alpha 2^{-k/2}. \tag{2}$$

The definition of T_k is chosen such that $\mathcal{E}_{i,k}$ is unlikely to occur; the first term ensures this in the case that $|S_{i,k}|$ is large, and the second term ensures this in the case that $|S_{i,k}|$ is small. As shown below in (4), each of the two terms contributes $O(\alpha\sqrt{\beta})$ to the overall discrepancy of row i.

The probability of $\mathcal{E}_{i,k}$ can be analyzed by a basic Hoeffding bound: if $\{X_i\}_{i\le\ell}$ are independent random variables, each $X_i \in [-1,+1]$, and $X = X_1 + \cdots + X_\ell$, then $\Pr[\,|X| > a\,] \le 2e^{-a^2/2\ell}$. Applying this bound to the discrepancy of $S_{i,k}$, we get that

$$\begin{aligned}
\Pr[\,\mathcal{E}_{i,k}\,] \;&\le\; 2\exp\big(-(T_k 2^k)^2/2|S_{i,k}|\big) \\
&<\; 2\exp\left(-\frac{\epsilon^2}{2|S_{i,k}|}\Big(2^k{\textstyle\sum_{j\in S_{i,k}}} A_{i,j}\Big)^2 - \frac{2\epsilon}{2|S_{i,k}|}\alpha 2^{k/2}\Big(2^k{\textstyle\sum_{j\in S_{i,k}}} A_{i,j}\Big)\right) \\
&\le\; 2\exp\Big(-\frac{\epsilon^2}{8}|S_{i,k}| - \frac{\epsilon}{2}\alpha 2^{k/2}\Big) \;=:\; p_{i,k},
\end{aligned} \tag{3}$$

where the last inequality uses $\sum_{j\in S_{i,k}} A_{i,j} \ge 2^{-(k+1)}|S_{i,k}|$.

2.1. Discrepancy assuming no events occur

Suppose that none of the events $\mathcal{E}_{i,k}$ happen. Then the total discrepancy of row i is at most

$$\begin{aligned}
\sum_{k\ge b} T_k \;&=\; \epsilon\sum_{k\ge b}\sum_{j\in S_{i,k}} A_{i,j} + \alpha\sum_{k\ge b} 2^{-k/2} \\
&\le\; \epsilon + \alpha\sum_{k\ge b} 2^{-k/2} \qquad \text{(since we assume } {\textstyle\sum_{j=1}^n} A_{i,j} \le 1) \\
&=\; \epsilon + \alpha\frac{2^{-b/2}}{1-2^{-1/2}} \\
&\le\; \epsilon + 4\alpha\sqrt{2\beta} \qquad \text{(since } 2^{-b} \le 2^{-(\lg(1/\beta)-1)} = 2\beta) \\
&\le\; 16\alpha\sqrt{\beta}.
\end{aligned} \tag{4}$$

2.2. Avoiding the events

We will use the local lemma to show that, with positive probability, none of the events $\mathcal{E}_{i,k}$ occur. To do so, we must show that these events have limited dependence. Consider $\mathcal{E}_{i,k}$, which is the event that the elements in row i of value roughly 2^{-k} have large discrepancy. This event depends only on the random values $\{\, y_j \,:\, j \in S_{i,k} \,\}$. We will bound the total failure probability of the events that depend on those random values.

The local lemma can be stated as follows [1, Theorem 5.1.1]:

Theorem 4. *Let $\mathcal{E}_1,\ldots,\mathcal{E}_n$ be events in a probability space. Let $\Gamma(\mathcal{E}_i)$ be a set of events (other than $\mathcal{E}_i$ itself) chosen such that $\mathcal{E}_i$ is mutually independent of all events outside of $\Gamma(\mathcal{E}_i)\cup\{\mathcal{E}_i\}$. Suppose one can associate a value $x(\mathcal{E}_i) \in (0,1)$ with each event $\mathcal{E}_i$ such that*

$$\Pr[\,\mathcal{E}_i\,] \;\le\; x(\mathcal{E}_i)\cdot \prod_{\mathcal{F}\in\Gamma(\mathcal{E}_i)} \big(1 - x(\mathcal{F})\big). \tag{5}$$

Then, with positive probability, no event $\mathcal{E}_i$ occurs.

The value that we assign to $\mathcal{E}_{i,k}$ is

$$x(\mathcal{E}_{i,k}) \;:=\; 2\exp\big(-\epsilon^2|S_{i,k}|/16 - \epsilon\alpha 2^{k/2}/2\big). \tag{6}$$

This choice deserves some explanation. When applying the local lemma, it is common to use either its "symmetric" form [1, Corollary 5.1.2], in which the $x(\cdot)$ values are all equal, or its "asymmetric" form [7, pp. 221], in which $x(\mathcal{E}_i)$ is taken to be $2\Pr[\,\mathcal{E}_i\,]$ and this quantity must be sufficiently small. Instead, we define $x(\mathcal{E}_{i,k})$ to be approximately the square root of $p_{i,k}$, which denotes our upper bound on $\Pr[\,\mathcal{E}_{i,k}\,]$ from (3). This choice turns out to be appropriate because, even though there will be substantial dependence amongst the events, the product $\prod_{\mathcal{F}\in\Gamma(\mathcal{E}_{i,k})}(1-x(\mathcal{F}))$ in (5) that controls the amount of dependence can also be bounded by roughly $\sqrt{p_{i,k}}$, so (5) is satisfied. In contrast, the product is bounded by a constant in both the "symmetric" and "asymmetric" forms of the local lemma.

Claim 5. *$x(\mathcal{E}_{i,k}) < 1/2$ for every $i \in [m]$ and $k \geq b$.*

Proof. By (1) we have $\epsilon > 8\sqrt{\beta}$, so

$$\epsilon 2^{k/2} \;\geq\; \epsilon\sqrt{2^b} \;\geq\; \epsilon\sqrt{2^{\lg(1/\beta)-1}} \;=\; \epsilon\sqrt{1/2\beta} \;>\; 4.$$

It follows that $x(\mathcal{E}_{i,k}) \leq 2\exp(-\epsilon 2^{k/2}/2) < 2\exp(-\sqrt{2}) < 1/2$. □

Our next step is to characterize $\Gamma(\mathcal{E}_{i,k})$, the events that are dependent on $\mathcal{E}_{i,k}$. We let $\mathcal{C}_{j,k}$ be the events corresponding to all entries of value roughly 2^{-k} in the j^{th} column.

$$\mathcal{C}_{j,k} \;:=\; \{\, \mathcal{E}_{i,k} \;:\; \lfloor -\lg A_{i,j} \rfloor = k \,\} \qquad (\text{for } j \in [n],\; k \geq b)$$

Next, $\mathcal{Y}_j$ contains all events corresponding to all entries in the j^{th} column. In other words, $\mathcal{Y}_j$ is the set of all events that depend on the random variable y_j.

$$\mathcal{Y}_j \;:=\; \bigcup_{k\geq b} \mathcal{C}_{j,k} \;=\; \{\, \mathcal{E}_{i,\lfloor -\lg A_{i,j}\rfloor} \;:\; i \in [m] \,\} \qquad (\text{for } j \in [n])$$

Finally, since $\mathcal{E}_{i,k}$ depends only on the random variables $\{\, y_j \;:\; j \in S_{i,k} \,\}$, the set $\Gamma(\mathcal{E}_{i,k})$ consists of all events that depend on any of those labels.

$$\Gamma(\mathcal{E}_{i,k}) \;=\; \bigcup_{j\in S_{i,k}} \mathcal{Y}_j.$$

Claim 6. *For every event $\mathcal{E}_{i,k}$, inequality (5) is satisfied.*

Proof. The main goal of the proof is to give a good lower bound for $\prod_{\mathcal{F}\in\Gamma(\mathcal{E}_{i,k})}(1-x(\mathcal{F}))$. Claim 5 shows that $x(\mathcal{F}) \leq 1/2$, so

$$\prod_{\mathcal{F}\in\Gamma(\mathcal{E}_{i,k})}(1-x(\mathcal{F})) \;\geq\; \prod_{\mathcal{F}\in\Gamma(\mathcal{E}_{i,k})}\exp(-2x(\mathcal{F})) \;=\; \exp\Bigg(-2\sum_{\mathcal{F}\in\Gamma(\mathcal{E}_{i,k})} x(\mathcal{F})\Bigg). \tag{7}$$

So it suffices to give a good upper bound for $\sum_{\mathcal{F}\in\Gamma(\mathcal{E}_{i,k})} x(\mathcal{F})$.

First we derive a technical inequality that is quite loose, but suffices for our proof.

$$\begin{aligned}
\epsilon\cdot\alpha 2^{k/2}/2 \;&=\; 8\alpha\sqrt{\beta}\cdot\alpha 2^{k/2}/2 \qquad (\text{by } (1))\\
&=\; \alpha^2\cdot 2\sqrt{\beta}\cdot 2^{1+b/2+(k-b)/2}\\
&=\; \lg(\delta/\beta^2)\cdot\big(2\sqrt{\beta}2^{b/2}\big)\cdot 2^{1+(k-b)/2}\\
&>\; \big(b+\lg(\delta/\beta)\big)\cdot 2^{1+(k-b)/2} \qquad (\text{since } \lg(1/\beta)\geq b \text{ and } 2^{b/2}\geq\sqrt{1/2\beta})\\
&\geq\; \big(b+\lg(\delta/\beta)\big) + 2^{1+(k-b)/2} \qquad (\text{since } xy \geq x+y \text{ if } x,y\geq 2)\\
&\geq\; \big(b+\lg(\delta/\beta)\big) + (k-b) \qquad (\text{since } 2^{1+i/2}\geq i \text{ for all } i\geq 0)\\
&=\; k+\lg(\delta/\beta)
\end{aligned} \tag{8}$$

Next, consider all the events that depend on y_j. Then

$$\begin{aligned}
\sum_{\mathcal{F}\in\mathcal{Y}_j} x(\mathcal{F}) &= \sum_{k\geq b}\sum_{\mathcal{F}\in\mathcal{C}_{j,k}} x(\mathcal{F}) \\
&\leq \sum_{k\geq b}\sum_{\mathcal{F}\in\mathcal{C}_{j,k}} \exp(-\epsilon\alpha 2^{k/2}/2) \qquad \text{(by (6))} \\
&\leq \sum_{k\geq b} |\mathcal{C}_{j,k}|\cdot e^{-(k+\lg(\delta/\beta))} \qquad \text{(by (8))} \\
&\leq \sum_{k\geq b} \Big|\left\{\, i \,:\, A_{i,j}\in(2^{-k-1},2^{-k}] \,\right\}\Big| \cdot 2^{-(k+\lg(\delta/\beta))} \\
&\leq (2\delta)\cdot(\beta/\delta) \;=\; 2\beta,
\end{aligned}$$

since the j^{th} column of A sums to at most δ. Therefore

$$\sum_{\mathcal{F}\in\Gamma(\mathcal{E}_{i,k})} x(\mathcal{F}) \;=\; \sum_{j\in S_{i,k}}\sum_{\mathcal{F}\in\mathcal{Y}_j} x(\mathcal{F}) \;\leq\; 2|S_{i,k}|\beta.$$

Combining this with (7), we obtain the lower bound

$$\begin{aligned}
x(\mathcal{E}_{i,k})\cdot\prod_{\mathcal{F}\in\Gamma(\mathcal{E}_{i,k})}(1-x(\mathcal{F})) &\geq x(\mathcal{E}_{i,k})\cdot\exp\Bigg(-2\sum_{\mathcal{F}\in\Gamma(\mathcal{E}_{i,k})} x(\mathcal{F})\Bigg) \\
&\geq 2\exp\big(-\epsilon^2|S_{i,k}|/16-\epsilon\alpha 2^{k/2}/2\big)\cdot\exp\big(-4|S_{i,k}|\beta\big) \\
&= 2\exp\Big(-|S_{i,k}|(\epsilon^2/16+4\beta)-\epsilon\alpha 2^{k/2}/2\Big) \\
&\geq 2\exp\Big(-|S_{i,k}|\epsilon^2/8-\epsilon\alpha 2^{k/2}/2\Big) \\
&= p_{i,k} \;\geq\; \Pr[\,\mathcal{E}_{i,k}\,]
\end{aligned}$$

where the penultimate inequality holds because $\epsilon^2/8\geq\epsilon^2/16+4\beta$, which follows because $\epsilon\geq 8\sqrt{\beta}$ (cf. (1)). This proves (5). □

The previous claim shows that the hypotheses of the local lemma are satisfied. So there exists a vector $y\in\{-1,+1\}^n$ such that none of the events $\mathcal{E}_{i,k}$ hold. As in (4), this implies that every row has discrepancy at most $16\alpha\sqrt{\beta}$. In other words, $\|Ay\|_\infty\leq 16\alpha\sqrt{\beta}$. This completes the proof of Theorem 3.

3. Conclusion

Many discrepancy theorems on hypergraphs have a more general statement about the discrepancy of real-valued matrices [3, §4]. We have provided another occurrence of this phenomenon by proving Theorem 2, which generalizes Theorem 1. Our result also yields a randomized algorithm with running time poly(m,n) for finding the desired vector y by directly applying the Moser-Tardos algorithm [8]. Independently, Harris and Srinivasan [5] have proven a more general result which also implies Theorem 2 and also provides an efficient randomized algorithm. To prove their result they derive a generalized form of the local lemma that can accommodate mild dependencies between events.

We are not aware of any result showing that either Theorem 1 or 2 is optimal. It seems conceivable that the logarithmic factor could be improved.

Conjecture 7. *Let V be an $m\times n$ real matrix with $|V_{i,j}|\leq 1$, $\left\|v^i\right\|_1\leq R$, and $\|v_j\|_1\leq\Delta$ for all $i\in[m], j\in[n]$. There exists $y\in\{-1,+1\}^n$ with $\|Vy\|_\infty\leq O(\sqrt{R\log(2+\Delta/R)})$.*

In the special case that $R = n$ and $\Delta = m$, this would recover the classic result of Spencer [9]. The following conjecture also seems plausible. Let $\|\cdot\|$ denote the spectral norm on the set of Hermitian matrices, i.e., the maximum absolute value of all eigenvalues, and let tr denote the trace functional.

Conjecture 8. *Let $A_1, \ldots, A_n$ be Hermitian, positive semi-definite matrices of size $m \times m$ satisfying $\sum_{i=1}^n A_i \preceq I$, $\|A_i\| \leq \beta$ and $\operatorname{tr} A_i \leq \delta$ for all i. Assume $\beta \leq \delta$. Then there exists $y \in \{-1, +1\}^n$ with $\|\sum_{i=1}^n y_i A_i\| \leq O(\sqrt{\delta \log(2 + \delta/\beta)})$.*

The special case in which the matrices are diagonal would imply Theorem 3. In the special case that the matrices have rank one we have $\delta = \beta$, so this would imply the recent result of Marcus et al. [6].

References

[1] N. Alon and J. Spencer. *The Probabilistic Method.* Wiley, 2000.

[2] József Beck and Tibor Fiala. "Integer-making" theorems. *Discrete Applied Mathematics*, 3(1):1–8, 1981.

[3] József Beck and Vera T. Sós. Discrepancy theory. In R. Graham and M. Grötschel and L. Lovász, editor, *Handbook of Combinatorics*, pages 1405–1446. Elsevier Science B.V., 1995.

[4] Béla Bollobás, David Pritchard, Thomas Rothvoss, and Alex Scott. Cover-decomposition and polychromatic numbers. *SIAM Journal on Discrete Math*, 27(1):240–256, 2013.

[5] David G. Harris and Aravind Srinivasan. The Moser-Tardos framework with partial resampling. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science*, pages 469–478, 2013.

[6] Adam Marcus, Daniel A. Spielman, and Nikhil Srivastava. Interlacing Families II: Mixed Characteristic Polynomials and The Kadison-Singer Problem, June 2013. `http://arxiv.org/abs/1306.3969`.

[7] Michael Molloy and Bruce Reed. *Graph Coloring and the Probabilistic Method.* Springer, 2002.

[8] Robin Moser and Gabor Tardos. A constructive proof of the general Lovász Local Lemma. *Journal of the ACM*, 57(2), 2010.

[9] Joel Spencer. Six standard deviations suffice. *Trans. Amer. Math. Soc.*, 289:679–706, 1985.

[10] Aravind Srinivasan. Improving the discrepancy bound for sparse matrices: better approximations for sparse lattice approximation problems. In *Proceedings of the 8th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 692–701, 1997.