Electronic Colloquium on Computational Complexity, Report No. 5 (2022)

Sunflowers: from soil to oil *

Anup Rao
University of Washington
anuprao@cs.washington.edu

Abstract

A *sunflower* is a collection of sets whose pairwise intersections are identical. In this article, we shall go sunflower-picking. We find sunflowers in several seemingly unrelated fields, before turning to discuss recent progress on the famous sunflower conjecture of Erdős and Rado, made by Alweiss, Lovett, Wu and Zhang.

1 Sunflowers

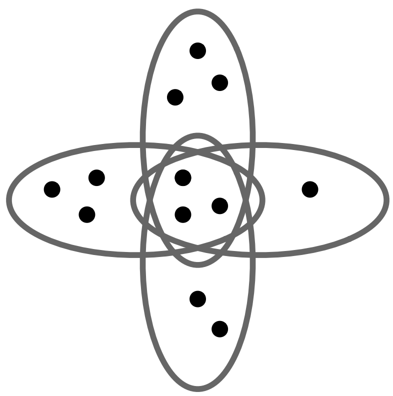

Figure 1: A sunflower with 4 petals.

The moral of Ramsey theory is that large systems can exhibit surprising structure. There are many examples of this kind, starting with the prototypical one: every graph on n vertices either contains a clique[1] on $(1/2) \cdot \log_2 n$ vertices, or an independent set[2] on $(1/2) \cdot \log_2 n$ vertices. Roth's theorem [14] proves that every subset of $\{1, \ldots, n\}$ of density $\Omega(1)$ must contain an arithmetic progression[3]. The Hales-Jewett theorem [9] and Ajtai and Szemeredi's Corner theorem [1] are other examples of this phenomenon.

*This exposition appears as a companion piece to a talk given at the current events bulletin at the Joint Mathmetical Meeting of the AMS in 2022.

[1]Mutually adjacent vertices

[2]Mutually non-adjacent vertices

[3]Three numbers $a, a + d, a + 2d$

ISSN 1433-8092

A *sunflower* with p petals is a collection of p sets whose pairwise intersections are identical. The common intersection is called the *core*. In 1960, Erdős and Rado [4] proved a Ramsey theoretic result concerning sunflowers: every large collection of sets must contain a sunflower. They gave a simple inductive argument showing that every collection of more than $k! \cdot (w-1)^k$ sets of size at most k must contain a sunflower with w petals[4]. There are examples with $\Omega(w)^k$ sets that have no sunflowers, and they conjectured that the correct bound is $O(w)^k$.

The seeds were planted, and the search for sunflowers and sunflower lemmas began in earnest. We begin this article by taking a tour through various fields where sunflowers are essential. We shall see examples relevant to finding arithmetic progressions in sumsets, understanding models of computation such as monotone boolean circuits and data structures, and fundamental questions about the threshold of a monotone function. In each of these arenas, we skip details and zoom in to focus on the role played by sunflowers.

In 2019, Alweiss, Lovett, Wu and Zhang [2] made significant progress towards proving the sunflower conjecture. Subsequent refining by myself [12], Frankston, Kahn, Narayanan and Park [7] and Bell, Chueluecha and Warnke [3] led to the result that every collection of $O(w \log k)^k$ sets of size at most $O(k)$ must contain a sunflower with w petals. We shall return to the ideas that lead to this improved bound near the end of this article.

2 Arithmetic Progressions in Sumsets

In 1992, Erdős and Sárközy [5] used sunflowers to find arithmetic progressions in subset sums. Given a set $T \subseteq \{1, \ldots, n\}$, let $\mathsf{sum}(T)$ denote the quantity $\sum_{x \in T} x$. Given any set $S \subseteq \{1, \ldots, n\}$ of size $|S| \gg \log^2 n$, they proved that there are subsets $T_1, \ldots, T_{w+1} \subseteq S$, with $w \approx |S|/\log^2 n$, such that the sequence

$$\mathsf{sum}(T_1), \mathsf{sum}(T_2), \cdots, \mathsf{sum}(T_{w+1})$$

is an arithmetic progression. Much like the sunflower lemma, this is an example of finding structure in a large system. However, the structure we seek here is an arithmetic progression; what does this have to do with sunflowers? Erdős and Sárközy move between the two structures as follows. First, by counting the number of possible sums that can be obtained by subsets of S, and estimating a binomial coefficient, they show that some $(w \log n)^{\log n}$ subsets of S of size $\log n$ must attain the same sum. By the sunflower lemma, and the choice of parameters, this collection of sets is guaranteed to contain a sunflower. The proof is completed by the following claim, whose proof we leave as an exercise (Figure 2):

[4]Often the sunflower lemma is stated under the assumption that each set is of size *exactly* k rather than at most k. Here we use the more general form because many application rely on this form, and all of the ideas for proving the lemmas carry through.

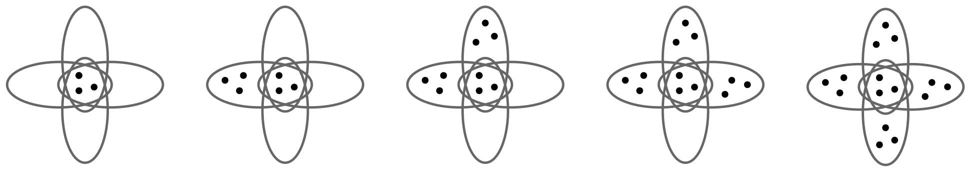

Figure 2: 4 petals induces an arithmetic progression of length 5.

Claim 1. If $S_1, \ldots, S_w \subseteq \{1, \ldots, n\}$ is a sunflower with core C, $|S_1| = \cdots = |S_w|$, and $\mathsf{sum}(S_1) = \mathsf{sum}(S_2) = \cdots = \mathsf{sum}(S_w)$, then

$$\mathsf{sum}(C), \mathsf{sum}(S_1), \mathsf{sum}(S_1 \cup S_2), \ldots, \mathsf{sum}(S_1 \cup \cdots \cup S_w)$$

is an arithmetic progression.

3 Monotone Circuit Lower Bounds

Sunflowers have had a huge impact in theoretical computer science. Perhaps the most well-known example is Razborov's [13] proof from 1985 that there are no small monotone circuits computing the *clique* function. Here, I will give a cartoon description of this clever argument.

A boolean circuit computes with the help of *gates* implementing boolean logic. These logic gates can compute the OR, AND or negation of their inputs. The inputs to the gates are either the outputs of other gates, or input variables. The size of the circuit is the number of wires used, which is the same as the number of connections made between gates. A *monotone* circuit is a boolean circuit that does not have any gates computing negations. The circuit computs a function if there is a gate whose value is equal to the value of the function, for every choice of the input variables.

For a graph G on n vertices, and a set S of vertices, define

$$\mathsf{clique}_S(G) = \begin{cases} 1 & \text{if } G \text{ contains a clique on the vertices of } S, \\ 0 & \text{otherwise.} \end{cases}$$

The function of interest for us is

$$\mathsf{clique}_k(G) = \bigvee_{S \subseteq \{1, \ldots, n\}, |S| = k} \mathsf{clique}_S(G),$$

which computes whether or not the graph contains a clique of size k. Razborov proves that this function requires exponentially large monotone circuits, if $k \approx n^{1/3}$. Razborov's result is one of the few examples where we are able to prove lower bounds on reasonable models of computation: it is a gem of theoretical computer science.

At a high level, sunflowers are used critically to show that any circuit computing clique_k can be used to obtain a smaller circuit with the same ability. Each such step involves a tiny error. We obtain a good approximation to the original circuit that is so simple that we can directly reason that it does not work. This proves that the original circuit does not work either.

Now, let us give a few more details. Let G be a graph on n vertices that contains a uniformly random clique of size k, and no other edges. Let H be a uniformly random $(k-1)$-partite graph. G always contains a clique of size k, while H never contains a clique of size k. A monotone circuit computing the clique function would have to output 1 on G and 0 on H. An input variable to the circuit is the indicator for the presence of an edge, which can be thought of as clique_S for some set S of size 2.

Let us discuss how to approximate the circuit by a simpler circiut. First, we claim that $\mathsf{clique}_S \wedge \mathsf{clique}_T$ can be safely replaced by $\mathsf{clique}_{S\cup T}$. This is because by the choice of G,

$$\mathsf{clique}_S(G) \wedge \mathsf{clique}_T(G) \leq \mathsf{clique}_{S\cup T}(G),$$

and by the choice of H,

$$\mathsf{clique}_S(H) \wedge \mathsf{clique}_T(H) \geq \mathsf{clique}_{S\cup T}(H).$$

Thus, carrying out this approximation preserves the ability of the circuit to distinguish G from H, while reducing the size of the circuit.

Sunflowers play a key role in approximating OR gates, via the following claim:

Claim 2. If $S_1, \ldots, S_w$ form a sunflower with core C, and all sets S_i are of size at most $\sqrt{k}$, then

$$\mathsf{clique}_{S_1}(G) \vee \cdots \vee \mathsf{clique}_{S_w}(G) \leq \mathsf{clique}_C(G),$$

and with high probability over the choice of H,

$$\mathsf{clique}_{S_1}(H) \vee \cdots \vee \mathsf{clique}_{S_w}(H) \geq \mathsf{clique}_C(H).$$

When the input is G, the claim is trivial. When the input is H, the approximation causes a problem if there is a clique on C in H, yet none of the petals constitute a clique. This is extremely unlikely to happen: given that the core is a clique, the events that the petals are also cliques are independent, and the choice of parameters ensures that each occurs with probability $\Omega(1)$. So, one can argue that one of the petals will be a clique with probability $1 - 2^{-\Omega(w)}$.

Thus, if t is large enough, any expression of the type

$$\mathsf{clique}_{S_1} \vee \ldots \vee \mathsf{clique}_{S_t}$$

can be approximated by a smaller expression of the same type—use the sunflower lemma to find a sunflower among the sets and replace it by the core. Repeatedly applying these operations, one can show that any arbitrary small monotone circuit can be approximated by a circuit whose structure is so simple that it is trivial to verify that it cannot distinguish G from H.

4 Lower Bounds for Data Structures

Data structures are a fundamental concept in computer science. They are used to efficiently maintain an object so that the object can be quickly modified and queried. Our next example is a lower bound on the running time of data structures for the problem of maintaining a set and computing its minimum, from my work with Ramamoorthy [11], building on [6, 8].

We showed that any data structure that can maintain a subset of numbers $T \subseteq \{1, \ldots, n\}$ and can quickly and non-adaptively compute the minimum element of T must access $\Omega(\log n / \log \log n)$ locations for one of its operations. Our result is independent of the algorithm used to implement the data structure and the particular encoding of the data (namely T) used, the argument only relies on the sets of locations that the data structure reads and writes to.

A valid data structure for our purposes is one that encodes the set T as a vector $\mathsf{enc}(T) \in \{1, \ldots, n\}^m$. The data structure is associated with a family of subsets of the coordinates of the encoding $S_1, \ldots, S_n \subseteq \{1, \ldots, m\}$ and an algorithm for manipulating $\mathsf{enc}(T)$. For each i, the algorithm is able change $\mathsf{enc}(T)$ to either $\mathsf{enc}(T \cup \{i\})$ or $\mathsf{enc}(T - \{i\})$ and compute the new minimum of the set by reading and writing to the coordinates of $\mathsf{enc}(T)$ given by S_i. Under just these assumptions, we prove that some set S_i must be of size $\Omega(\log n / \log \log n)$.

If all of the sets $S_1, \ldots, S_n$ are of size $\ll \frac{\log n}{\log \log n}$, the choice of parameters implies that there is a sunflower, say $S_1, \ldots, S_w$, with $w \approx (\log n)^{100}$, and core C. Then the key claim is:

Claim 3. If $S_1, \ldots, S_w$ is a sunflower with core C, then every subset of $\{1, \ldots, w\}$ has an encoding as a vector $\mathsf{enc}(W) \in \{1, \ldots, n\}^{|C|}$.

This claim combined with a straightforward counting argument implies that $|C| \geq \Omega(\frac{\log n}{\log \log n})$, proving that one of the sets S_i must be large. To prove the claim, for any set $T \subseteq \{1, \ldots, w\}$, arrive at its encoding by deleting the elements of the set $\{1, \ldots, w\} - T$ from the encoding of $\{1, \ldots, w\}$. The claimed encoding corresponds to the contents of the core at this point, which is a string in $\{1, \ldots, n\}^{|C|}$. T can be recovered from the encoding by computing the minimum of T, then deleting the minimum, then computing the minimum and deleting it, and repeating these operations over and over until the entire set T has been recovered. Because each of these operations only interact on the coordinates of $\mathsf{enc}(T)$ that correspond to the core, the contents of the core are enough to simulate the entire process and determine T.

5 Estimating the Threshold of Monotone Functions

Suppose that we are given a parameter $0 \leq \epsilon \leq 1$ and sample a random graph by including each edge independently with probability ϵ. How can we estimate the probability that the graph contains a perfect matching?

This is a special case of a more general question. Let $X \in \{0, 1\}^n$ be sampled by independently setting each coordinate:

$$X_i = \begin{cases} 1 & \text{with probability } \epsilon, \\ 0 & \text{with probability } 1-\epsilon. \end{cases}$$

Let $f : \{0,1\}^n \to \{0,1\}$ be a *monotone* function, meaning that $x \geq y$ implies that $f(x) \geq f(y)$ ($x \geq y$ if $x_i \geq y_i$ for all i). Because f is monotone, $\mathbb{E}[f(X)]$ is increasing in ϵ. The *threshold* of f is the value of ϵ for which $\mathbb{E}[f(X)] = 1/2$. There are a couple of generic ways to bound $\mathbb{E}[f(X)]$, and these bounds induce other kinds of thresholds that capture something about the structure of f. These ideas were explored extensively by Kahn and Kalai [10], Talagrand [15] and Frankston, Kahn, Narayanan and Park [7].

Every monotone function f admits a minimal collection of binary strings F such that $f(x) = 1$ if and only if there is an element $y \in F$, with $y \leq x$. So, by the union bound:

$$\mathbb{E}[f(X)] \leq \sum_{y \in F} \mathbb{P}[y \leq X], \tag{5.1}$$

The *expectation-threshold* is the value of ϵ for which the right-hand-side of (5.1) is equal to $1/2$. By (5.1), the threshold is always at least as large as the expectation-threshold. When f computes whether or not a graph has a perfect matching, the threshold is $\approx \frac{\log n}{n}$, while the expectation threshold is $\approx \frac{1}{n}$. Kahn and Kalai conjectured that this is in fact the worst possible ratio: the threshold is always at most $O(\log n)$ times larger than the expectation-threshold.

In general, the union bound can be quite far from tight. It is not tight when the events $y \leq X$ have intersections of significant measure. There is a more sophisticated way to get upper bounds on $\mathbb{E}[f(X)]$, as observed by Talagrand [15]—it can be thought of as a fractional variant of the union bound. For $z \in \{0,1\}^n$, let $|z| = \sum_{i=1}^n z_i$. Suppose there is a probability distribution Z on $\{0,1\}^n$ and κ satisfying

$$f(x) \leq \kappa \cdot \mathop{\mathbb{E}}_{Z}[1_{Z \leq x} \cdot \epsilon^{-|Z|}],$$

for all x. Then we obtain the upper bound:

$$\mathop{\mathbb{E}}_{X}[f(X)] \leq \kappa \cdot \mathop{\mathbb{E}}_{X,Z}[1_{Z \leq X} \cdot \epsilon^{-|Z|}] = \kappa, \tag{5.2}$$

since for any fixed z, the probability that $z \leq X$ is exactly $\epsilon^{|z|}$.

The *fractional-expectation-threshold* is the value of ϵ for which there is a distribution Z with $\kappa = 1/2$. The union bound (5.1) can also be proved using (5.2), because if Z is sampled so that

$$\mathbb{P}[Z = z] = \begin{cases} \frac{\mathbb{P}[z \leq X]}{\sum_{w \in F} \mathbb{P}[w \leq X]} & \text{if } z \in F, \\ 0 & \text{otherwise,} \end{cases}$$

then we have

$$f(x) \leq \Big(\sum_{z \in F} \mathbb{P}[z \leq X] \Big) \cdot \mathop{\mathbb{E}}_{Z}[1_{Z \leq x} \cdot \epsilon^{-|Z|}] = \sum_{z \in F} \mathbb{P}[z \leq X],$$

for all x, proving (5.1). So, the bound given by (5.2) is certainly at least as good as the bound given by (5.1). In particular, this implies that the threshold is at least as large as the fractional-expectation-threshold, which in turn is at least the expectation-threshold. But how far apart can these numbers be?

Talagrand conjectured that the fractional-expectation-threshold is within a multiplicative factor of $O(\log n)$ from the threshold, and within an $O(1)$ factor of the expectation-threshold. Frankston, Kahn, Narayanan and Park [7] proved that the fractional-expectation-threshold is within $O(\log n)$ of the threshold, so resolving Talagrand's first conjecture. This allows to compute the threshold for many graph properties, such as perfect matchings, Hamiltonian circuits and bounded degree spanning trees. The ideas used to prove new sunflower lemmas play a key role in their proof, as we shall see below.

Talagrand made an important observation that is ultimately useful to understanding the gap between the threshold and the fractional-expectation-threshold. Suppose that κ is the smallest number for which there is a distribution Z establishing (5.2). Then by von-Neumann's minimax theorem, there is a distribution U on $\{0,1\}^n$ such that for every choice of z,

$$\mathop{\mathbb{E}}_{U}[f(U)] \geq \kappa \cdot \mathop{\mathbb{E}}_{U,z}[1_{z \leq U} \cdot \epsilon^{-|z|}]. \tag{5.3}$$

Without loss of generality, we may assume that U is supported on the min-terms of f, since we can always modify the distribution in this way and preserve the inequality. Moreover, the distribution of U can be thought of as the uniform distribution on a multiset $\mathcal{S}$, by equating binary strings with subsets of $\{1, \dots, n\}$, and finding a close enough rational approximation to the distribution of U. So, after making these changes, we can rewrite (5.3) as:

$$\epsilon^{|z|}/\kappa \geq \mathop{\mathbb{E}}_{U,z}[1_{z \leq U}], \tag{5.4}$$

where here U samples a random element from a multi-set containing the min-terms of f. By (5.4), this multi-set has a very interesting property: It must be *spread*, in the sense that very few of these sets can all contain the same set. The fraction of min-terms containing z is at most $\mathbb{E}_{U,z}[1_{z \leq U}] \leq \epsilon^{|z|}/\kappa \leq r^{-|z|}$, for $r = \kappa/\epsilon$.

So, where are the sunflowers? This time, the sunflowers show up figuratively: our best known method for finding sunflowers involves understanding the probability that a random set contains an element of a spread family. To explain, let us put on hold our study of these thresholds and discuss the ideas needed to prove sunflower lemmas.

6 Finding Sunflowers (in Spread Families)

At last we return to the heart of the matter: how many sets of size k are sufficient to ensure the presence of a sunflower with w petals? Alweiss, Lovett, Wu and Zhang discovered an elementary counting argument that is surprisingly powerful to help answer this question.

Given a collection $\mathcal{S}$ of sets, we shall say that the sets are *r-spread* (for some parameter $r = O(w \log k)$) if for every set T of size at most $k-1$, the fraction of sets in the collection that contain T is at most $r^{-|T|}$. As we saw in the last section, this definition naturally arose in the work of Talagrand, but it is very directly applicable to proving a better sunflower lemma. Suppose $|\mathcal{S}| \geq r^k$. If $\mathcal{S}$ is not r-spread, then there is a set T such that the family $\mathcal{S}' = \{S \in \mathcal{S} : T \subseteq S\}$ has at least $r^{k-|T|}$ sets. In this case, we inductively find a sunflower in the family of sets of size at most $k - |T|$ obtained by deleting T from the sets of $\mathcal{S}'$. Adding T back into this sunflower gives us a sunflower in our original family of sets. So, it only remains to find sunflowers in spread families.

The main technical claim is:

Claim 4. For $r = O((1/\epsilon)\log k)$, if $\mathcal{S}$ is r-spread, and X is a random set sampled by including each element independently with probability ϵ, then X contains a set of $\mathcal{S}$ with probability at least $1/2$.

We note that the Claim holds even if $\mathcal{S}$ is a multi-set, which is useful for the application to understanding the thresholds of monotone functions. Let $X_1, \ldots, X_{2w}$ be a random partition of the universe into $2w$ sets, and set $\epsilon = 1/(2w)$, so $r = O(w \log k)$. Claim 4 implies that w of these sets will contain a set of the family in expectation, and so there must be w mutually disjoint sets: a sunflower with w petals.

To exhibit the key ideas used to prove the claim, let us settle for a weaker goal. Suppose $\mathcal{S} = \{S_1, \ldots, S_{r^k}\}$, and $X \subseteq \{1, \ldots, n\}$ is uniformly random. We will show that there must be a set $S_i \in \mathcal{S}$ such that $|X \cap S_i| \geq 0.99k$ with high probability.

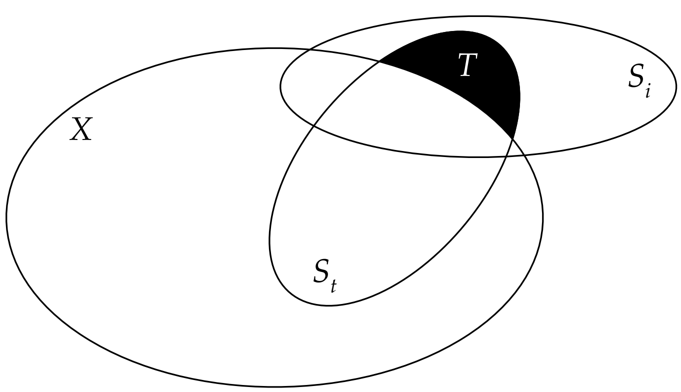

Figure 3: T must be large, and so there are few choices for S_i.

Consider the possible pairs (X, S_i) where X does *not* share $0.99k$ elements with any set of $\mathcal{S}$, and S_i is an arbitrary element of $\mathcal{S}$.

(i). There are at most 2^n choices for $X \cup S_i$.

(ii). Given $X \cup S_i$, let t be the smallest index such that $S_t \subseteq X \cup S_i$. There is exactly one choice for t.

(iii). Let $T = S_t - X$. This is a subset of S_t, so there are at most 2^k choices for T.

(iv). We have now identified a set T that is contained in S_i. Since X does not cover $0.99k$ elements of S_t, $|T| \geq 0.01k$. Because $\mathcal{S}$ is r-spread, there can be at most $|\mathcal{S}| \cdot r^{-|T|}$ choices for S_i.

(v). Finally, there are at most 2^k choices for $X \cap S_i$. Because the sets $X \cup S_i$, S_i and $X \cap S_i$ determine X, all of these choices determine the pair (X, S_i).

In this way, the number of such pairs (X, S_i) is at most $2^n \cdot 2^k \cdot |\mathcal{S}| \cdot r^{-0.01k} \cdot 2^k = |\mathcal{S}| \cdot 2^{n+2k-0.01k \log r}$. Because the number of choices for S_i is $|\mathcal{S}|$, this implies that the probability that a random choice of X fails to intersect a set in $0.99k$ elements is at most $2^{k(2-0.01 \log r)}$, which can be made very small for r a large constant. Methods from information theory along with a few other ideas can be used to prove Claim 4 in its entirety.

7 Estimating the Threshold of Monotone Functions (contd.)

Armed with our ability to reason about random sets containing a set of the spread family (Claim 4), we can now return to finish the story of the gap between the threshold and the fractional-expectation-threshold.

Suppose the fractional-expectation-threshold of a monotone function f is γ. Then recall that we have a collection of min-terms (possibly with repetitions) that correspond to sets that are r-spread, with $r = 1/(2\gamma)$. Let k be the size of the largest min-term, and let C be some large constant. Then if we set $\epsilon = C\gamma \log k$, we get that $r = C(1/(2\epsilon)) \log k$. For some large C, Claim 4 implies that a random set X, where each element is included in X with probability ϵ, will contain one of the min-terms with probability at least $1/2$. If X is viewed as a binary string, this implies that

$$\mathbb{E}[f(X)] \geq 1/2,$$

proving that the threshold of f is at most $O(\gamma \log k)$. Thus, we get that the ratio between the threshold and the fractional-expectation-threshold is at most $O(\log n)$.

8 Conclusion

Sunflowers have had an enormous impact on a surprising number of different fields. They are certain to spring up in new places in the future. The methods of [2] may well lead to even stronger sunflower lemmas, or find applications in places where there are no sunflowers. It is an exciting time to be playing with these concepts!

References

[1] Miklos Ajtai and Endre Szemeredi. *Stud. Sci. Math. Hungar*, (9):9–11, 1974.

[2] Ryan Alweiss, Shachar Lovett, Kewen Wu, and Jiapeng Zhang. Improved bounds for the sunflower lemma. In *Proccedings of the 52nd Annual Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 624–630. ACM, 2020.

[3] Tolson Bell, Suchakree Chueluecha, and Lutz Warnke. Note on sunflowers. *Discret. Math.*, 344(7):112367, 2021.

[4] Paul Erdős and Richard Rado. Intersection theorems for systems of sets. *Journal of London Mathematical Society*, 35:85–90, 1960.

[5] Paul Erdös and András Sárközy. Arithmetic progressions in subset sums. *Discret. Math.*, 102(3):249–264, 1992.

[6] Gudmund Skovbjerg Frandsen, Peter Bro Miltersen, and Sven Skyum. Dynamic word problems. *J. ACM*, 44(2):257–271, 1997.

[7] Keith Frankston, Jeff Kahn, Bhargav Narayanan, and Jinyoung Park. Thresholds versus fractional expectation-thresholds. *CoRR*, abs/1910.13433, 2019.

[8] Anna Gál and Peter Bro Miltersen. The cell probe complexity of succinct data structures. *Theor. Comput. Sci*, 379(3):405–417, 2007.

[9] A. W. Hales and R. I. Jewett. Regularity and positional games. *Transactions of the American Mathematical Society*, 106(2):222–229, 1963.

[10] Jeff Kahn and Gil Kalai. Thresholds and expectation thresholds. *Comb. Probab. Comput.*, 16(3):495–502, 2007.

[11] Sivaramakrishnan Natarajan Ramamoorthy and Anup Rao. Lower bounds on non-adaptive data structures maintaining sets of numbers, from sunflowers. In *33rd Computational Complexity Conference, CCC 2018, June 22-24, 2018, San Diego, CA, USA*, pages 27:1–27:16, 2018.

[12] Anup Rao. Coding for sunflowers. *Discrete Analysis*, 2020.

[13] A. A. Razborov. Some lower bounds for the monotone complexity of some boolean functions. *Soviet Math. Dokl.*, 31(31):354–357, 1985.

[14] K. F. Roth. On Certain Sets of Integers. *Journal of the London Mathematical Society*, s1-28(1):104–109, 01 1953.

[15] Michel Talagrand. Are many small sets explicitly small? In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 13–36. ACM, 2010.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il