BULLETIN OF THE ICA
Volume **105** (2025), 45–53

Algebraic sunflowers

Nathanael Ackerman and Mostafa Mirabi

Abstract. We study sunflowers within the context of finitely generated substructures of ultrahomogeneous structures. In particular, we look at bounds on how large a set system is needed to guarantee the existence of sunflowers of a given size. We show that if we fix the size of the sunflower, the function that takes the size of the substructures in our set system and outputs the size of a set system needed to guarantee a sunflower of the desired size can grow arbitrarily slowly.

1 Introduction

A sunflower, also known as a Δ-system, is a collection of sets such that any two pairs of distinct sets in the collection have a common intersection. When delving into the study of sunflowers, a common focus lies in determining the existence of large sunflowers inside a given collection of sets. The existence of large sunflowers has wide ranging applications in computer science, including in the study of circuit lower bounds, matrix multiplication, pseudo-randomness, and cryptography. For a survey of the connections to computer science see [2].

Additionally, the existence of large sunflowers also has applications in mathematical logic, including in the study of forcing large generic structures. See for example [5], [6], and [1] (where sunflowers are called Δ-systems).

One of the first major results showing the existence of large sunflowers in any collection of finite sets of the same size was the "Sunflower Lemma" of Erdős and Rado.

Lemma 1.1 (Sunflower Lemma [3]). *Let* $\mathrm{SF}\colon \omega \times \omega \to \omega$ *be the function where, for* $n, k \in \omega$, $\mathrm{SF}(n, k)$ *is the minimal value such that any set* $\mathcal{F}$ *of*

Mathematics Subject Classifications: 05D05, 03C13
Key words and phrases: Sunflower Lemma, Sunflower Conjecture, Ultrahomogeneous structures
Corresponding author: Nathanael Ackerman `<nate@aleph0.net>`

Received: 22 October 2023
Accepted: 15 September 2024

sets of size k with $|\mathcal{F}| \geq \mathrm{SF}(n,k)$ has a sunflower of size n. Then for all $n \in \omega$ and $k \geq 3$,

$$\mathrm{SF}(n,k) \leq k!(n-1)^k.$$

The Sunflower Lemma asserts that any sufficiently large collection of sets of a fixed size is guaranteed to contain a large sunflower. Moreover, it provides an upper bound on how large the collection of sets must be.

While the Sunflower Lemma gives an upper bound on the size of the sets needed to have a sunflower, it was clear even when it was proved that it was not optimal. This gave rise to the "Sunflower Conjecture".

Conjecture 1.2 (Sunflower Conjecture [3])**.** For all $n \in \omega$ and $k \geq 3$, there is a constant c_n such that

$$\mathrm{SF}(n,k) \leq c_n^k.$$

Over the years some progress has been made on this conjecture, with an important step forward coming recently in [2]. However, the full conjecture remains open.

Recently a new approach to studying the Sunflower Conjecture emerged in [4] in which the authors studied the Sunflower Conjecture in the context of set systems that satisfy various structural properties. Specifically they considered three structural properties: the case where the set system has finite VC dimension, the case where the set system has finite Littlestone dimension, and the case when the required set system arises naturally from some background structure, e.g., the plane. In [4], the authors were able to provide improved bounds on SF when the set system has finite VC dimension as well as finite Littlestone dimension. However, in the case where the set systems arise from the background structure of the plane, they were able to prove the full Sunflower Conjecture.

In this paper, we aim to study the functions SF under the assumption that the set system comes from some background structure. Specifically, we work in a situation where there is some background ultrahomogeneous algebraic structure $\mathcal{M}$ where every finitely generated substructure of $\mathcal{M}$ is finite. Let $\mathrm{SF}_{\mathcal{M}}$ be the analog of SF restricted to finitely generated substructures of $\mathcal{M}$. Our main result shows that for each $n \in \omega$ and each function $\alpha \colon \omega \to \omega$ there is an ultrahomogeneous structure $\mathcal{M}$ such that, for all $r \geq 3$, $\mathrm{SF}_{\mathcal{M}}(n,r) \leq \alpha(r)$, i.e., for every function, no matter how slowly it grows, we can find an algebraic structure such that the bound on the size of the set needed for an n-size sunflower grows slower than that function. Further, we can choose the structure $\mathcal{M}$ to be totally categorical.

1.1 Notation

Let $n^{(k)}$ be the number of non-repeating sequences of size k on n-elements. Whenever p is a unary function from some set to itself and $k \in \omega$ let p^k be the k-fold composition of p with itself. We also let $\overline{p}(x) = \{p^k(x)\}_{0<k\in\omega}$. If X is a set we let $\mathcal{P}(X)$ denote the collection of subsets of X. If $f\colon X \to Y$ is a function and $X_0 \subseteq X$ we let $f''[X_0] = \{f(x) : x \in X_0\}$, i.e., the image of X_0 under f.

If $\alpha\colon \omega \to \omega$ is a non-decreasing function whose limit is infinity let $\alpha^\circ : \omega \to \omega$ be such that $\alpha^\circ(n) = \min\{k : \alpha(k) \geq n\}$.

Suppose $\mathscr{L}$ is a first-order language and $\mathcal{M}$ is a $\mathscr{L}$-structure. We say $\mathcal{M}$ is an *algebraic structure* when $\mathscr{L}$ has only function symbols. We say $\mathcal{M}$ is *locally finite* if for every finite set A there is a finite substructure of $\mathcal{M}$ containing A. We let $\mathrm{Sub}_k(\mathcal{M})$ be the collection of finitely generated substructures of $\mathcal{M}$ of size at most k. We let $\mathrm{Sub}(\mathcal{M}) = \bigcup_{k\in\omega} \mathrm{Sub}_k(\mathcal{M})$.

We say that an $\mathscr{L}$-structure $\mathcal{M}$ is *ultrahomogeneous* if, whenever A, B are finitely generated substructures of $\mathcal{M}$ and $\sigma\colon A \to B$ is an isomorphism, there is an automorphism of $\mathcal{M}$ extending σ. We say $\mathcal{M}$ has a *strong amalgamation property* or (SAP) if, whenever $i_0\colon A \to B$ and $j_0\colon A \to C$ are embeddings of finitely generated substructures of $\mathcal{M}$, there are embeddings $i_1\colon B \to \mathcal{M}$ and $j_1\colon C \to \mathcal{M}$ such that $i_1 \circ i_0 = j_1 \circ j_0$ and $i_1''[B] \cap j_1''[C] = (i_1 \circ i_0)''[A]$. Note $\mathcal{M}$ has (SAP) precisely when its age does.

An infinite structure is called *totally categorical* whenever any two models of its first-order theory of the same size are isomorphic. Structures that are totally categorical are particularly simple from a model theoretic point of view. For example, note that the theory of sets in the empty language is an example of totally categorical structures.

2 Algebraic sunflowers

We now formalize the functions we want to consider.

Definition 2.1. Suppose $\mathcal{M}$ is a locally finite $\mathscr{L}$-structure and $X \subseteq \mathrm{Sub}(\mathcal{M})$.

- We say X is *uniform* if for all $A, B \in X$ there is an isomorphism between A and B.
- We say X is *strongly uniform* if for all $A, B \in X$ there is an isomorphism from A to B that is the identity on $A \cap B$.
- We say X is a *sunflower* if for all $A_0, A_1, B_0, B_1 \in X$ we have $A_0 \cap A_1 = B_0 \cap B_1$.

Definition 2.2. Suppose $\mathscr{L}$ is a language and $\mathcal{M}$ is a locally finite $\mathscr{L}$-structure. For each $n, k \in \omega$ let $\mathrm{SF}_{\mathcal{M}}(n, k)$ be the least ℓ such that, whenever $X \subseteq \mathrm{Sub}_k(\mathcal{M})$ and $|X| \geq \ell$, there is an $X_0 \subseteq X$ such that $|X_0| = n$ and X_0 is a sunflower.

We omit mention of $\mathcal{M}$ when $\mathscr{L}$ is the empty language and $\mathcal{M}$ is infinite. In this case $\mathrm{SF}(n, k)$ is the size of a set, all of whose elements have size at most k, that guarantees a sunflower of size n. Note that if X is a finite set of elements, all of which have size at most k, there is a set X^* and a bijection $i\colon X \to X^*$ such that

- for all $x \in X^*$, $|x| = k$,
- for all $x \in X$, $x \subseteq i(x)$,
- for all $x \in X$, $\big(i(x) \setminus x\big) \cap \big(\bigcup X\big) = \emptyset$,
- for any $X_0 \subseteq X$, X_0 is a sunflower if and only if $i''[X_0]$ is a sunflower.

In particular, when dealing with sunflowers for sets without structure, there is no difference with considering sets all of the same size or sets of bounded size. This is equivalent to saying that we lose no generality in only considering strongly uniform sunflowers.

Note if there are no sunflowers of size n consisting of elements of size at most k then $\mathrm{SF}_{\mathcal{M}}(n, k) = \infty$. By the Sunflower Lemma this can only happen if there are less than $\mathrm{SF}(n, k)$ substructures of $\mathcal{M}$ of size at most k. Therefore, it makes sense to restrict attention to structures where $\mathrm{Sub}_k(\mathcal{M})$ is infinite for any sufficiently large k. In particular we restrict our attention to structures that are ultrahomogeneous with (SAP).

Note that for any $\mathcal{M}$ and any $n, k \in \omega$ we have $\mathrm{SF}_{\mathcal{M}}(n, k) \geq n$. The following straightforward proposition shows that for any fixed k this lower bound can be achieved.

Proposition 2.3. *Let $\mathscr{L} = \{f\}$ where f is a unary function. Then for every $n, k \in \omega$ there is an $\mathscr{L}$-structure $\mathcal{M}_k$ where*

- *$\mathcal{M}_k$ is ultrahomogeneous,*
- *$\mathcal{M}_k$ has (SAP),*
- *$\mathcal{M}_k$ is totally categorical,*
- *$\mathrm{SF}_{\mathcal{M}_k}(n, k) = n$.*

Proof. For $k \in \omega$ let C_k be the $\mathscr{L}$-structure with k-elements and where f is a bijection such that for every $x, y \in C_k$ there is an $i \in [k]$ such that $f^i(x) = y$. Let $\mathcal{M}_k$ be the union of ω-many disjoint copies of C_k.

Now let F be any collection of minimal non-empty substructures of $\mathcal{M}_k$. Each element of $\mathcal{F}$ must be isomorphic to C_k, and hence no two elements of $\mathcal{F}$ can have non-empty intersection. Hence, $\mathcal{F}$ is a sunflower of size $|\mathcal{F}|$. □

Proposition 2.3 showed us that, if we fix the size of the subsets we are considering, we can obtain the minimal bound on the size of sets needed to guarantee a large sunflower. We did this by, intuitively, enlarging each point to be a substructure of the desired size. The problem of finding a sunflower among sets of size k in $\mathcal{M}_k$ then reduced to the problem of finding a sunflower of size 1 among ordinary sets.

Next, look at what happens if, instead of fixing the size of the sets we want to consider, we fix the size of the sunflower we want to find. We then have to consider substructures of arbitrary sizes, and so we need a more complicated construction. However, the underlying idea is similar. Our goal is to construct a structure whose substructures act like ordinary sets, but where the size of the substructure corresponding to a set of a given size grows arbitrarily fast. This will then let us use the Sunflower Lemma to ensure that for any fixed n we can find structures $\mathcal{M}$ where $\mathrm{SF}_{\mathcal{M}}(n, k)$ grows arbitrarily slowly as a function of k.

Theorem 2.4. *There is a finite language $\mathscr{L}$, consisting of only unary and binary functions such that, whenever $\alpha\colon \omega \to \omega$ is non-decreasing with $\alpha(0) \geq 3$ and $\lim_{n\to\infty} \alpha(n) = \infty$, there is a locally finite countable $\mathscr{L}$-structure $\mathcal{M}_\alpha$ such that*

(a) *$\mathcal{M}_\alpha$ is ultrahomogeneous, has (SAP), and is totally categorical;*

(b) *if $\mathcal{A}_0, \mathcal{A}_1 \in \mathrm{Sub}(\mathcal{M}_\alpha)$ with $|\mathcal{A}_0| = |\mathcal{A}_1|$, then $\mathcal{A}_0 \cong \mathcal{A}_1$;*

(c) *if* $\mathcal{F} \subseteq \mathrm{Sub}(\mathcal{M}_\alpha)$ *is finite, then there is an* $\mathcal{F}^* \subseteq \mathrm{Sub}(\mathcal{M}_\alpha)$ *and a bijection* $i\colon \mathcal{F} \to \mathcal{F}^*$ *such that*
- *for all* $\mathcal{A}_0, \mathcal{A}_1 \in \mathcal{F}^*$, $\mathcal{A}_0 \cong \mathcal{A}_1$,
- *for all* $\mathcal{A} \in \mathcal{F}$, $\mathcal{A} \subseteq i(\mathcal{A})$,
- *if* $\mathcal{F}_0 \subseteq \mathcal{F}$, *then* $\mathcal{F}_0$ *is a sunflower if and only if* $i[\mathcal{F}_0]$ *is a sunflower;*

(d) *all uniform sunflowers are strongly uniform;*

(e) $(\forall n, k \in \omega) SF_{\mathcal{M}_\alpha}(n,k) \leq \alpha(k) \cdot (n-1)^{\alpha(k)}$.

Proof. Let $\mathscr{L}_0 = \{c, s\}$ where c, s are unary functions. Let $\beta\colon \omega \to \omega$ be an increasing function. Let $\mathcal{N}_\beta^-$ be the unique countable infinite $\mathscr{L}_0$-structure such that

- s is a bijection;
- for each $x \in \mathcal{N}_\beta^-$ there is a $k \in \omega$ such that $s^k(x) = x$; let $i(x)$ be the least such k;
- for all $m \in \omega$ there are infinitely many x with $i(x) = \beta(m)$;
- for each x there is an m such that $i(x) = \beta(m)$; let $m = j(x)$;
- $c^2 = c$;
- for all $x \in \mathcal{N}_\beta^-$, there is a k such that $s^k(x) = c(x)$.

In other words, $\mathcal{N}_\beta$ consists of infinitely many s-cycles of each length in $\{\beta(m) : m \in \omega\}$. Furthermore, each s-cycle has a unique distinguished element that is a fixed point of c.

Let $\mathscr{L} = \mathscr{L}_0 \cup \{p_0, p_1, a\}$ where p_0, p_1 are unary and a is binary. Let $\mathcal{N}_\beta$ be the $\mathscr{L}$-structure where

- $\mathcal{N}_\beta \upharpoonright_{\mathscr{L}_0} = \mathcal{N}_\beta^-$;
- if $c(x) \neq x$, then $p_0(x) = p_1(x) = x$, and for all $y \in \mathcal{N}_\beta$, $a(x,y) = x$ and $a(y,x) = y$;
- if $j(x) = 1$, then $p_0(x) = p_1(x) = x$;
- if $c(x) = x$ and $j(x) = n+1$, then
  - $j(p_0(x)) = 1$,
  - $j(p_1(x)) = n$,
  - $c(p_0(x)) = p_0(x)$ and $c(p_1(x)) = p_1(x)$,
  - $\overline{p_0}(x)$ has size $n+1$;

- if $c(x) = x$, $c(y) = y$, $j(x) = k$ and $j(y) = n$, then
  - if $k \neq 1$ then $a(x, y) = x$,
  - if $k = 1$ and $x \in \overline{p_0}(y)$ then $a(x, y) = x$,
  - if $k = 1$ and $x \notin \overline{p_0}(y)$ then $j(a(x, y)) = n + 1$, $p_0(a(x, y)) = x$, and $p_1(a(x, y)) = y$,
  - if $\overline{p}(x) = \overline{p}(y)$ then $x = y$.

Intuitively, the structure $\mathcal{N}_\beta$ is constructed as follows. First, we divide $\mathcal{N}_\beta$ into infinitely many equivalence classes, where all elements in the same equivalence class are some power of s of each other. Further, the size of the equivalence classes are precisely the range of β (and there are infinitely many equivalence classes of each size).

Inside each equivalence class we have a distinguished element, which is a fixed point of c. These are the only elements for which the functions p_0, p_1, a may be non-trivial.

Now on these distinguished elements we have three functions, p_0, p_1 and a. If x is an element in the equivalence class with $\beta(n)$ many elements, then we want to think of x as a tuple of length $n + 1$, without repetitions, of equivalence classes with $\beta(0)$ many elements. We then think of p_0 as the projection onto the first element of the tuple and p_1 as the tuple that results from removing the first element of the tuple x. We then think of $a(x, y)$ as a pairing function that can add a single element x to the front of the tuple y, but only if x is not already in the tuple y.

For any substructure $\mathcal{A} \subseteq \mathcal{N}_\beta$ let $b(\mathcal{A}) = \{x \in \mathcal{A} : c(x) = x \text{ and } p_0(x) = x\}$, i.e., the collection of distinguished elements of equivalence classes with $\beta(0)$ many elements. The following are then immediate for (possibly infinite) substructures $\mathcal{A}_0, \mathcal{A}_1$ of $\mathcal{N}_\beta$:

- if $b(\mathcal{A}_0) = b(\mathcal{A}_1)$, then $\mathcal{A}_0 = \mathcal{A}_1$;
- if $i\colon b(\mathcal{A}_0) \to b(\mathcal{A}_1)$ is an bijection, then i extends uniquely to an isomorphism $i^*\colon \mathcal{A}_0 \cong \mathcal{A}_1$;
- $b(\mathcal{A}_0) \cap b(\mathcal{A}_1) = b(\mathcal{A}_0 \cap \mathcal{A}_1)$.

Further, for any subset $A \subseteq b(\mathcal{N}_\beta)$ there is a substructure $\mathcal{A} \subseteq \mathcal{N}_\beta$ with $b(\mathcal{A}) = A$. Put together these imply $\mathcal{N}_\beta$ satisfies (a)–(d). All that is left is to choose a β such that $\mathcal{N}_\beta$ satisfies (e) with respect to α for all n.

Let $\gamma_\beta\colon \omega \to \omega$ be the function where

$$\gamma_\beta(n) = \sum_{k \leq n} n^{(k)} \cdot \beta(k).$$

Note γ_β is increasing with $\lim_{n\to\infty} \gamma_\beta(n) = \infty$. We then have for any substructure $\mathcal{A}$ that

$$|\mathcal{A}| = \gamma_\beta\big(\big|s(\mathcal{A})\big|\big).$$

We also have from [3] that

$$\mathrm{SF}_{\mathcal{N}_\beta}\big(n, \gamma_\beta(k)\big) \le k!(n-1)^k \le k!(n-1)^{k!}$$

and so

$$\mathrm{SF}_{\mathcal{N}_\beta}(n,k) \le \big(\gamma^\circ_\beta(k)\big)!(n-1)^{\gamma^\circ_\beta(k)!}$$

(where γ°_β is as defined in Section 1.1). But as β goes to infinity and γ°_β is constant on the interval from $\big[\beta(k), \beta(k+1)\big)$, we can find a β_α such that, for all $k \in \omega$, $\gamma^\circ_{\beta_\alpha}(k)! \le \alpha(k)$. So if we let $\mathcal{M}_\alpha = \mathcal{N}_{\beta_\alpha}$, we have for all $n, k \in \omega$

$$\mathrm{SF}_{\mathcal{M}_\alpha}(n,k) \le \alpha(k) \cdot (n-1)^{\alpha(k)}$$

as desired. □

Remark 2.5. Note that when considering the function $\mathrm{SF}_{\mathcal{M}}(n,k)$ we are considering finite substructures of $\mathcal{M}$ whose size is bounded by k. One might instead wish to consider finite substructures that have a generating set whose size is bounded by k. However, the structure constructed in the proof of Theorem 2.4 shows that this approach would be problematic as every finitely generated substructure is generated by a single element. Specifically, if $(A_i)_{i\in\omega}$ is an increasing collection of finite substructures of $\mathcal{M}_\alpha$, then $\{A_i\}_{i\in\omega}$ is an infinite set system, each of which is generated by a single element, but which contains no sunflower of size 3.

The following corollary, which shows that for any fixed n the map $k \mapsto \mathrm{SF}_{\mathcal{M}}(n,k)$ can be chosen to grow arbitrarily slowly, is then immediate from Theorem 2.4.

Corollary 2.6. *There is a finite language $\mathscr{L}$ such that whenever*

- $n \in \omega$,
- $\alpha\colon \omega \to \omega$ *is non-decreasing with* $\alpha(0) \ge 3\cdot(n-1)^3$ *and* $\lim_{n\to\infty} \alpha(n) = \infty$,

there is a locally finite countable $\mathscr{L}$-structure $\mathcal{M}_{n,\alpha}$ such that

- $\mathcal{M}_{n,\alpha}$ *is ultrahomogeneous, has (SAP), and is totally categorical;*
- *for all k,* $\mathrm{SF}_{\mathcal{M}_{n,\alpha}}(n,k) \le \alpha(k)$.

We have considered the growth rate of the function $\mathrm{SF}_{\mathcal{M}}(n,k)$ where we fix n or k. We end with two conjectures about this function when neither coordinate is fixed.

Conjecture 2.7. Suppose $f\colon \omega \to \omega$ is a non-decreasing function that goes to infinity. Then there is a constant $c_f \in \omega$ such that, whenever $\alpha\colon \omega \to \omega$ is a non-decreasing function that goes to infinity and with $\alpha(0) \geq c_f$, there is a finite functional language $\mathscr{L}$ and a locally finite $\mathscr{L}$-structure $\mathcal{M}_f$ such that for all $k \in \omega$

$$\mathrm{SF}_{\mathcal{M}_f}\big(f(k),k\big) \leq \alpha(k).$$

Conjecture 2.8. Suppose $g\colon \omega \to \omega$ is a non-decreasing function that goes to infinity. There is a constant $d_g \in \omega$ such that, whenever $\alpha\colon \omega \to \omega$ is a non-decreasing function that goes to infinity and with $\alpha(0) \geq d_g$, there is a finite functional language $\mathscr{L}$ and a locally finite $\mathscr{L}$-structure $\mathcal{N}_g$ such that for all $n \in \omega$

$$\mathrm{SF}_{\mathcal{N}_g}(n, g(n)) \leq \alpha(n).$$

References

[1] N. Ackerman, M. Golshani, and M. Mirabi, Cohen generic structures with functions, *arXiv e-prints* (2023), `https://arxiv.org/abs/2310.11582`.

[2] R. Alweiss, S. Lovett, K. Wu, and J. Zhang, Improved bounds for the sunflower lemma, *Ann. of Math. (2)* **194**(3) (2021), 795–815, `https://doi.org/10.4007/annals.2021.194.3.5`.

[3] P. Erdős and R. Rado, Intersection theorems for systems of sets, *Journal of the London Mathematical Society* **35**(1) (1960), 85–90.

[4] J. Fox, J. Pach, and A. Suk, Sunflowers in set systems of bounded dimension, *Combinatorica* **43**(1) (2023), 187–202, `https://doi.org/10.1007/s00493-023-00012-z`.

[5] M. Golshani, Fraïssé limit via forcing, *arXiv e-prints* (2019), `https://arxiv.org/abs/1902.06206`.

[6] Z. Kostana, Cohen-like first order structures, *Ann. Pure Appl. Logic* **174**(1) (2023), Paper No. 103172, 17, `https://doi-org.ezp-prod1.hul.harvard.edu/10.1016/j.apal.2022.103172`.

Nathanael Ackerman
Harvard University, Cambridge, MA 02138, USA
nate@aleph0.net

Mostafa Mirabi
Department of Mathematics and Computer Science, The Taft School, Watertown, CT 06795, USA
mmirabi@wesleyan.edu