ON RESTRICTED ARITHMETIC PROGRESSIONS OVER FINITE FIELDS

BRIAN COOK ÁKOS MAGYAR

Abstract. Let A be a subset of $\mathbb{F}_p^n$, the n-dimensional linear space over the prime field $\mathbb{F}_p$ of size at least δN ($N = p^n$), and let $S_v = P^{-1}(v)$ be the level set of a homogeneous polynomial map $P : \mathbb{F}_p^n \to \mathbb{F}_p^R$ of degree d, for $v \in \mathbb{F}_p^R$. We show, that under appropriate conditions, the set A contains at least $c\,N|S|$ arithmetic progressions of length $l \leq d$ with common difference in S_v, where c is a positive constant depending on δ, l and P. We also show that the conditions are generic for a class of sparse algebraic sets of density $\approx N^{-\gamma}$.

1. Introduction.

1.1. **Background.** A celebrated result of Szemerédi [9] states that a set A of positive upper density of the integers contains arbitrarily long arithmetic progressions $x, x+d, \ldots, x+ld$. There has been various generalizations and extensions, a natural question to ask is whether one may add restrictions on the common difference d. A well-known result of this type is a theorem due to Sárközy [8] and Furstenberg, stating that A contains two distinct elements whose difference is a square. More recently Green has shown [3] that A contains a 3-term arithmetic progression, whose common difference is a sum of two squares. Far reaching results of this type for longer progressions have been obtained recently by Green and Tao [5] and by Wooley and Ziegler [12] where the gap is of the form $p-1$ and $P(p-1)$, p being a prime and P an integral polynomial such that $P(0) = 0$.

The aim of this note is to provide a simple extension of this type in the finite field settings, where $A \subseteq \mathbb{F}_p^n$ is a set of density $\delta > 0$ and one is counting arithmetic progressions in A with gaps in algebraic sets S given as level sets of a family of homogeneous polynomials.

1.2. **Main results.** Let $\mathbb{F}_p^n$ be the n-dimensional linear space above the prime field $\mathbb{F}_p$, and for a fixed $\delta > 0$ let $A \subseteq \mathbb{F}_p^n$ be a set of at least $|A| \geq \delta p^n$ elements. The finite field version of Szemerédi's theorem states that such sets will contain genuine arithmetic progressions of length $l \leq p$ as long as n is large enough. Note that the condition $l \leq p$ is natural as it ensures that progressions in the form $\{x, x+y, \ldots, x+(l-1)y\}$ consist of distinct points for $y \neq 0$. We will need the following quantitative version

Theorem A. *Let $\delta > 0$ and $l \in \mathbb{N}$. For a function $f : \mathbb{F}_p^n \to [0,1]$ satisfying $\mathbb{E}(f(x) : x \in \mathbb{F}_p^n) \geq \delta$ one has that*

$$\mathbb{E}(f(x)f(x+y)\ldots f(x+(l-1)y) : \ x, y \in \mathbb{F}_p^n) \geq c(\delta, l, p), \tag{1.1}$$

where $c(\delta, l, p)$ is a positive constant depending only on δ, l and p.
Here $\mathbb{E}(f(x) : \ x \in S) = \frac{1}{|S|}\sum_{x\in S} f(x)$ denotes the average of a function f over a set S.

Note that, because of the existence of trivial progressions with $y = 0$ the above statement is non-trivial only in large enough dimensions $n \geq n_\delta$. So we will assume that n is sufficiently large from now on, while keeping the characteristics of the field fixed.

Let $S_v = P^{-1}(v)$ be an algebraic set defined as the level set of a family of homogeneous polynomials $P = (P_1, \ldots, P_R) : \mathbb{F}_p^n \to \mathbb{F}_p^R$ of degree d, $v \in \mathbb{F}_p^R$ being a given vector. We will be interested in counting l-term arithmetic progressions in A with common difference y in S_v. It is clear that in

The second author was supported by NSERC Grant 22R44824.

order to make this problem well-defined one needs to make a few assumptions. First S_v needs to be nonempty, and in order to have a more precise formula for the size of S_v, a natural assumption is that the associated set of singular points

$$S_P^* := \{x \in \mathbb{F}_p^n : \ rank\,(Jac_P(x)) < R\}, \tag{1.2}$$

is small. Here $Jac_P(x)$ is the $R \times n$ matrix with entries $\partial_{x_j} P_i(x)$. This will be done by requiring that

$$K := codim(S_P^*) \tag{1.3}$$

is sufficiently large. Note that the dimension of an algebraic set is defined above the algebraic closure $\bar{\mathbb{F}}_p$ of the prime field $\mathbb{F}_p$. To avoid degeneracies like the identical vanishing of certain derivatives, we'll also assume that $d < p$. Let us introduce the parameters $0 < \alpha, \beta, \gamma < 1$ by

$$\gamma n = R, \ \beta n = K, \ p^\alpha = d \tag{1.4}$$

Our main result is the following.

Theorem 1. *Let $\delta > 0$, $\varepsilon > 0$ be given, and let $A \subseteq \mathbb{F}_p^n$ be a set of size $|A| \geq \delta p^n$. Let the polynomial map P and the parameters $0 < \alpha, \beta, \gamma < 1$ be defined as above. Then for $l \in \mathbb{N}$, $l \leq d$ one has uniformly in $v \in \mathbb{F}_p^R$ that*

$$\mathbb{E}(\mathbf{1}_A(x)\mathbf{1}_A(x+y)\ldots\mathbf{1}_A(x+(l-1)y) : \ x \in \mathbb{F}_p^n, \, y \in S_v) \geq c(\varepsilon, \delta, l, p), \tag{1.5}$$

provided

$$\beta - \alpha - (2^d + 1)\gamma \geq \varepsilon, \tag{1.6}$$

with a constant $c(\varepsilon, \delta, l, p) > 0$ depending only on ε, δ, l and p. Here $\mathbf{1}_A$ stands for the indicator function of the set A and $S_v = P^{-1}(v)$.

Remarks:

- If n is large enough, it follows that A contains $\approx |\mathbb{F}_p^n||S_v|$ non-trivial progressions of length l with common difference $y \in S_v$. In particular for every $v \in \mathbb{F}_p^R$ there is a progression with common difference y such that $P(y) = v$. As a byproduct of the proof we also obtain that $|S_v| \approx p^{n-R}$, uniformly in v, thus is a "sparse" set of density of $\approx p^{-\gamma n}$.
- The condition $l \leq d$ seems necessary, as it can be seen from the following example in [2] adapted to the finite field settings. Let $d = 2$, $R = 1$ and $P(x) = x_1^2 + \ldots + x_n^2$. Fix $p > 2$ (say $p = 5$) and let $A = \{x \in \mathbb{F}_p^n : \ P(x) = 0\}$, then A has density $\approx p^{-1}$. By the parallelogram identity: $P(x) - 2P(x+y) + P(x+2y) = 2P(y)$, thus if A contains a 3-term arithmetic progression $\{x, x+y, x+2y\}$ then necessarily $P(y) = 0$. One can construct similar examples for the polynomials $Q(x) = \sum_j x_j^d$ of degree d for all $d \geq 2$. These examples also show the necessity of a condition on the singular set. Indeed (1.5) does not hold for the level sets of $P(x) = (x_1^2 + \ldots, +x_n^2)^{d/2}$ ($d > 2$ even) for $l > 2$, while it dose hold for the level sets of $Q(x) = x_1^d + \ldots + x_n^d$ for $l = d$. The difference is that S_P^* is $n-1$-dimensional while $S_Q^* = \{0\}$.
- For the special case, when $v = 0$ the set A contains non-trivial arithmetic progressions with gap $y \in S = P^{-1}(0)$ of length $l > d$, under the more restrictive conditions that $\gamma \leq c(\delta, l, p, d)$. This is based on the fact that zero set of homogeneous polynomial maps contain a large linear subspace, which follows from a theorem of Chevalley and Warning [4], see Thm. 6.11. This will be discussed in Section 4.

We will also study polynomial maps P for which the conditions of Theorem 1 hold. The point being that with the exception of a sparse set of polynomial maps P Theorem 1 holds, with an absolute constant $\gamma > 0$.

Consider first diagonal forms, when $P_i(x) = \sum_{j=1}^n a_{ij}x_j^d$, $A = \{a_{ij}\}_{1\le i\le R, 1\le j\le n}$ being an $R \times n$ matrix. We call the map $P = (P_1, \dots, P_R)$ *diagonal* associated with the matrix A. We have

Proposition 1. *Let $p \ge 5$ and set $\gamma_0 := \frac{1}{2} - \frac{\log 2}{\log 5}$. If $R = \gamma\, n$ with $0 < \gamma < \gamma_0$, then the probability that*

$$codim(S_P^*) \ge \frac{n}{2},$$

for a random diagonal polynomial map $P = (P_1, \dots, P_R)$, is at least $1 - p^{-(\gamma_0-\gamma)n}$.

This implies that condition (1.6) holds for diagonal maps as long as $\alpha + (2^d + 1)\gamma \le 1/2 - \varepsilon$, in particular for fixed d and $p > d^2$, one may choose γ to be an absolute constant.

Next for quadratic maps $P = (P_1, \dots, P_R)$ with $P_i(x) = \langle A_i x, x\rangle$, A_i being a symmetric $n \times n$ matrix with entries in $\mathbb{F}_p$, we have

Proposition 2. *Let $p \ge 5$, and let $\gamma_0 := \frac{1}{2} - \frac{\log 2}{\log 5}$. Then the probability that*

$$codim(S_P^*) \ge n - R - 2\sqrt{n},$$

for a random quadratic map $P = (P_1, \dots, P_R)$, is at least $1 - p^{-(\gamma_0-\gamma)n}$.

It follows that $\beta \ge 1 - \gamma - \frac{2}{\sqrt{n}}$ for generic quadratic maps. Thus condition (1.6) holds, if $6\gamma < 1 - \log_2 p - \frac{2}{\sqrt{n}}$, in particular with an absolute constant for $p \ge 5$, $n \ge 12$.

In general, for higher degree non-diagonal forms, let $\mathcal{P}(n, d)$ be the space of homogeneous polynomials $P : \mathbb{F}_p^n \to \mathbb{F}_p$ of degree d, which is an $N = \binom{n+d-1}{d}$ dimensional linear space over $\mathbb{F}_p$. In Section 3 we will show that

Proposition 3. *Let $1 \le R, L \le n$ be given such that $L \ge 2R - 1$. Then the locus of polynomial maps $P \in \mathcal{P}(n, d)^R$ such that $dim\, S_P^* \ge L$ is contained in an algebraic set of codimension at least $L - 2R + 2$.*

Since the dimension of an algebraic set is defined over the algebraically closed field $\bar{\mathbb{F}}_p$, it follows that for generic polynomial maps $P : \bar{\mathbb{F}}_p^n \to \bar{\mathbb{F}}_p^R$ the dimension of the singular variety satisfies the bound: $dim\, S_P^* \le 2R - 2$, thus (1.6) holds if: $\alpha + (2^d + 3)\gamma \le 1 - \varepsilon$. We remark however, that this does not in itself imply that condition (1.6) holds for most polynomial maps $P : \mathbb{F}_p^n \to \mathbb{F}_p^R$.

1.3. **Outline of the Proof.** To begin we shall need a generalized von Neumann inequality for restricted progressions. To state it, let us briefly recall the Gowers uniformity norms. The multiplicative derivative of a function f is given by $\Delta_h f(x) = f(x + h)\overline{f(x)}$, and higher derivatives as $\Delta_{h_1,\dots,h_l} = \Delta_{h_l}(\Delta_{h_1,\dots,h_{l-1}})$. The U^l-Gowers norm is then given by

$$||f||_{U^l}^{2^l} = \mathbb{E}(\Delta_{h_1,\dots,h_l} f(x) : x, h_1, \dots, h_l \in \mathbb{F}_p^n).$$

This norm represents the average of f over 'cubes' in $\mathbb{F}_p^n$, sets of the form

$$\{x + \omega_1 h_1 + \dots + \omega_l h_l : \omega_1, \dots, \omega_l \in \{0, 1\}^l\},$$

and is indeed a norm for integers $l > 1$ ($l = 1$ provides a semi-norm). On fact we shall need is the monotonicity formula

$$||f||_{U^{l-1}} \le ||f||_{U^l}. \tag{1.7}$$

For the above definitions and facts on may consult [10]. For a pair of functions $f, g : \mathbb{F}_p^n \to \mathbb{C}$ define the form

$$\widetilde{\Lambda}_l(f, g) = \mathbb{E}(f(x)f(x + r)\dots f(x + (l - 1)r)g(r) : x, r \in \mathbb{F}_p^n).$$

Lemma 1. *(Generalized von Neumann inequality)*
For functions f, g bounded in absolute value by one, one has

$$\widetilde{\Lambda}_l(f,g) \leq ||g||_{U^l}.$$

To apply this result in combination with Theorem A, one finds an appropriate balanced function of the level set $S = P^{-1}(v)$, say $g = \mathbf{1}_S - \rho$ for an appropriate constant ρ, and writes

$$\widetilde{\Lambda}_l(f, \mathbf{1}_S) = \rho \Lambda_l f + \widetilde{\Lambda}_l(f,g). \tag{1.8}$$

The first term applies to Theorem A, while the second can be bounded by $||g||_{U^l}$ by Lemma 1. Then it remains to show that ρ can be chosen properly as to give $||g||_{U^l}$ small. Thus the crucial step is to obtain the following bound, which may be of interest on its own.

Proposition 4. *Let $v \in \mathbb{F}_p^R$ and let $S = P^{-1}(v)$, where $P = (P_1, \ldots, P_R) : \mathbb{F}_p^n \to \mathbb{F}_p^R$ is a homogeneous polynomial map of degree d. Then one has*

$$||\mathbf{1}_S - p^{-R}||_{U^d} \leq (d-1)^{2^{-d}n}\ p^{2^{-d}(R-K)}, \tag{1.9}$$

where $K = codim(S_P^*)$, S_P^* being the singular variety associated to P.

We remark that non-trivial estimates on the U^1 norm of the function $\mathbf{1}_S - p^{-R}$ yield asymptotic formulas for the number of points on the level surfaces S_v. In fact the proof of Proposition 4 is based on the observation that the analytical methods of Birch-Davenport [1], developed to count points on varieties, are easily adapted to estimate higher Gowers norms.

<u>*Proof of Theorem 1*</u>. Let the parameters $0 < \alpha, \beta, \gamma < 1$ be defined as in (1.4) and assume that condition (1.6) holds. First, note that by (1.9) and (1.6) we have that

$$\left|p^{-n}|S| - p^{-R}\right| = \|\mathbf{1}_S - p^{-R}\|_{U^1} \leq p^{n(\alpha+\gamma-\beta)2^{-d}} \leq p^{-\varepsilon n} p^{-R}$$

thus in particular $|S| = p^{n-R}(1 + O(p^{-\varepsilon n}))$.

Let $f = \mathbf{1}_A$, $\rho = p^{-R}$, $g = \mathbf{1}_S - p^{-R}$, then by (1.8)

$$\mathbb{E}\left(\mathbf{1}_A(x)\mathbf{1}_A(x+y)\ldots\mathbf{1}_A(x+(l-1)y) :\ x \in \mathbb{F}_p^n,\, y \in S\right) =$$

$$= p^n|S|^{-1}\left(p^{-R}\Lambda_l f + \widetilde{\Lambda}_l(f,g)\right) = (1 + O(p^{-\varepsilon n}))\left(\Lambda_l f + p^R \widetilde{\Lambda}_l(f,g)\right)$$

By Theorem A the first term satisfies

$$\Lambda_l f \geq c(\delta, l, p)$$

while by (1.9) and (1.6) the second term is at most

$$p^R \widetilde{\Lambda}_l(f,g)) \leq p^R p^{n(\alpha+\gamma-\beta)2^{-d}} \leq p^{-\varepsilon n}$$

This implies that the left side of (1.5) is at least $c(\delta, l, p)/2$ for $n \geq n(\varepsilon, \delta, l, p)$, while it is trivially at least p^{-n} for all n's. This proves Theorem 1. □

It remains to prove Proposition 4. The starting point is the identity

$$\mathbf{1}_S(x) = \mathbb{E}(e(\alpha \cdot (P(x) - v) : \alpha \in \mathbb{F}_p^R) = p^{-R} + p^{-R} \sum_{\alpha \in \mathbb{F}_p^R,\, \alpha \neq 0} e(-\alpha \cdot v)\, e(\alpha \cdot P(x)).$$

The triangle inequality for the Gowers norms then gives

$$||g||_{U^d} \leq p^{-R} \sum_{\alpha \in \mathbb{F}_p^R,\, \alpha \neq 0} (||e(\alpha \cdot P(x))||_{U^d},$$

reducing our task to bounding $||e(\alpha \cdot P(x))||_{U^d}$ for a nonzero α.

We apply the method of Birch [1] to achieve a bound in terms of an explicitly given algebraic set $W^* \subseteq \mathbb{F}_p^{(d-1)n}$, which will be discussed in detail in the next section.

Lemma 2. *If $\alpha \neq 0$ in $\mathbb{F}_p^n$, then we have*

$$||e(\alpha \cdot P(\cdot)||_{U^d}^{2^d} \leq \frac{|W^*|}{p^{(d-1)n}}. \tag{1.10}$$

The set W^* appears in the work of Birch on exponential sums [1] and is closely related to the singular set S_P^*, see (2.6) below. In particular if one embeds $\mathbb{F}_p^n$ into the diagonal $\Delta \subseteq \mathbb{F}_p^{(d-1)n}$ via the map $\Phi(h) = (h, \ldots, h)$, then $\Phi(S_P^*) = W^* \bigcap \Delta$. As Δ is a linear subspace of codimension $(d-2)n$, it follows that $dim(W^*) \leq (d-2)n + dim(S_P^*)$. Thus, $codim(W^*) \geq codim(S_P^*) = K$. Also W^* can be partitioned into algebraic sets W_λ^*, ($\lambda \in \mathbb{F}_p^R \backslash \{0\}$) such that W_λ^* is defined by n equations of degree $d-1$. Then basic facts from algebraic geometry give

Lemma 3. *With W^* as above, we have $|W^*| \leq (d-1)^n p^R p^{(d-1)n-K}$.*

Thus Proposition 4 follows save for the proofs of Lemma 2 and Lemma 3.

2. Exponential sum estimates.

Let $P = (P_1, \ldots, P_R) : \mathbb{F}_p^n \to \mathbb{F}_p^R$ be a polynomial map, P_i being a homogeneous polynomial of degree d, written in the symmetric form

$$P_i(x) = \sum_{1 \leq j_1, \ldots, j_d \leq n} a^i_{j_1 \ldots j_d} x_{j_1} \ldots x_{j_d}, \qquad x = (x_1, \ldots, x_n) \tag{2.1}$$

where $a^i_{j_1 \ldots j_d} = a^i_{\pi(j_1) \ldots \pi(j_d)}$ for any permutation $\pi : \{1, \ldots, d\} \to \{1, \ldots, d\}$. Note that this is possible as $(d!, p) = 1$. For $h \in \mathbb{F}_p^n$ define the differencing operator

$$D_h P_i(x) = P_i(x+h) - P_i(x), \tag{2.2}$$

and note that $deg\,(D_h P) = deg\,(P) - 1$. After applying the differencing operators $d-1$ times one obtains a linear function of the form

$$D_{h^1} \ldots D_{h^{d-1}} P_i(x) = \sum_{j=1}^{n} \Phi^i_j(h^1, \ldots, h^{d-1}) x_j, \tag{2.3}$$

where $\Phi^i_j(h^1, \ldots, h^{d-1})$ is the multilinear form

$$\Phi^i_j(h^1, \ldots, h^{d-1}) = d! \sum_{1 \leq j_1, \ldots, j_d \leq n} a^i_{j_1 \ldots j_d, j}\ h^1_{j_1} \ldots h^{d-1}_{j_{d-1}}, \tag{2.4}$$

for any $(d-1)$-tuple of vectors $(h^1, \ldots, h^{d-1}) \in \mathbb{F}_p^{(d-1)n}$. Note that on the diagonal

$$\Phi^i_j(h, \ldots, h) = (d-1)!\ \partial_{x_j} P_i(h) \tag{2.5}$$

Define the set W^* associated to the polynomial map P by

$$W^* = \{(h^1, \ldots, h^{d-1}) \in \mathbb{F}_p^{(d-1)n} :\ rank\ (\Phi(h^1, \ldots, h^{d-1})) < R\}, \tag{2.6}$$

where $\Phi(h^1, \ldots, h^{d-1})$ is the $R \times n$ matrix with entries $\Phi^i_j(h^1, \ldots, h^{d-1})$ for $1 \leq i \leq R$, $1 \leq j \leq n$.

Proof of Lemma 2. Using the definition of U^d norm:

$$\begin{aligned} ||e(\alpha \cdot P)||_{U^d} &= \mathbb{E}(\Delta_{h_1, \ldots, h_d} e(\alpha \cdot P(x) :\ x, h_1, \ldots, h_d \in \mathbb{F}_p^n) \\ &= \mathbb{E}(e(\alpha \cdot D_{h_1, \ldots, h_d} P(x)) : x, h_1, \ldots, h_d \in \mathbb{F}_p^n). \end{aligned} \tag{2.7}$$

From (2.2) (2.3) and the definition of the matrix $\Phi(h^1,\dots,h^{d-1})$ it is clear that

$$\Delta_{h_1,\dots,h_d} e(\alpha\cdot P(x)) = e(\Phi^T(h^1,\dots,h^{d-1})\alpha\cdot h^d) \tag{2.8}$$

where Φ^T is the transpose of the matrix Φ and "$\cdot$" is the dot product. If $rank\,(\Phi(h^1,\dots,h^{d-1})) = R$ then $\Phi^T(h^1,\dots,h^{d-1})\,\alpha \neq 0$ hence summing (2.8) in the h^d variable vanishes. Thus only the tuples $(h^1,\dots,h^{d-1}) \in W^*$ contribute to $||e(\alpha\cdot P)||_{U^d}$ and Lemma 2 follows. □

If $rank\,(\Phi(h^1,\dots,h^{d-1})) < R$ then its rows $\Phi^1,\dots,\Phi^R$ are linearly dependent, thus one may write

$$W^* = \bigcup_{(\lambda_1,\dots,\lambda_R)\neq 0} W^*(\lambda_1,\dots,\lambda_R),$$

where for $\lambda = (\lambda_1,\dots,\lambda_R) \neq 0$

$$W^*(\lambda) = \{(h^1,\dots,h^{d-1}) \in \mathbb{F}_p^{(d-1)n} :\ \lambda_1\Phi^1(h^1,\dots,h^{d-1}) + \dots + \lambda_R\Phi^R(h^1,\dots,h^{d-1}) = 0\}. \tag{2.9}$$

Note that $W^*(\lambda)$ is a homogeneous algebraic set defined by n equations of degree $d-1$. To estimate the size these sets we need the following basic facts from algebraic geometry.

Lemma 4. [6] *For a homogeneous (affine) algebraic set $U \subseteq \mathbb{F}_p^m$ of degree r and dimension s one has that*

$$|U| \leq r\,p^s \tag{2.10}$$

The degree of the set U is defined as the degree of its image U^0 in the $m-1$-dimensional projective space above $\mathbb{F}_p$. For projective algebraic sets it is shown in [6], Prop. 12.1, that $|U^0| \leq r\,\pi_s(\mathbb{F}_p)$ where $\pi_s(\mathbb{F}_p) = (p^s-1)/(p-1)$ is the size of the $s-1$-dimensional projective space. This implies (2.10) as $|U| = (p-1)|U^0| + 1$.

Lemma 5. *If a homogeneous algebraic set U is defined by n equations of degrees $d_1,\dots,d_n$ then its degree is bounded by*

$$deg\,(U) \leq d_1 d_2 \dots d_n \tag{2.11}$$

The degree of a (projective) algebraic set is defined as the sum of degrees of its irreducible components, and for a projective algebraic variety it can be defined geometrically as the number of intersection points with a generic subspace of complementary dimension or algebraically, see [7] Prop. 7.6 and the definition preceding it. Note that the degree of a hypersurface is the degree of its defining polynomial. Lemma 5 may be viewed as a generalization of Bezout's theorem on the number of intersection of plane algebraic curves, written as an inequality ignoring the multiplicities of the intersection points. It follows easily from the following basic inequality which is an immediate corollary of Thm. 7.7 in [7].

Theorem B. *Let Y be a (projective) variety of dimension at least 1 and let H be a hypersurface not containing Y. Let $Z_1,\dots,Z_s$ be the irreducible components of the intersection $Y\cap H$. Then*

$$\sum_{i=1}^{s} deg\,(Z_i) \leq deg\,(Y)\,deg\,(H) \tag{2.12}$$

Proof of Lemma 5. One may write (2.12) as

$$deg\,(Y\cap H) \leq deg\,(Y)\,deg\,(H) \tag{2.13}$$

as long as $Y \not\subseteq H$ and $dim\,Y \geq 1$. However for $Y \subseteq H$ or when $dim\,Y = 0$ (Y being a point and $deg\,(Y) = 1$) inequality (2.13) holds trivially. It also extends to algebraic sets V by writing them as union of their irreducible components Y and using (2.13) for each Y together with the fact that

each irreducible component of $V \cap H$ is contained in some $Y \cap H$ for an irreducible $Y \subseteq V$. Then (2.11) follows immediately by induction on n. □

Proof of Lemma 3. For a given $\lambda \in \mathbb{F}_p^R \backslash \{0\}$ one has that $dim\, W^*(\lambda) \leq dim\, W^* \leq (d-1)n - K$, and since it is defined by n equations of degree $d-1$ its degree satisfies the bound $deg\,(W^*(\lambda) \leq (d-1)^n$. Hence by (2.10) and (2.11) one has for all $\lambda \in \mathbb{F}_p^R \backslash \{0\}$

$$|W^*(\lambda)| \leq (d-1)^n p^{(d-1)n-K}$$

and Lemma 3 follows from the decomposition $W^* = \cup_\lambda W^*(\lambda)$. □

3. Generic polynomial maps.

Let $P = (P_1, \ldots, P_R)$ be a polynomial map, consisting of diagonal forms $P_i(x) = \sum_j a_{ij} x_j^d$ associated to a matrix A. We call the matrix A *non-degenerate* if $rank\, A' = R$ for every $R \times n/2$ submatrix of A.

Claim 1. *If A is non-degenerate then $dim\, S_P^* \leq n/2$.*

Proof. For $1 \leq j \leq n$ let $\mathbf{a}_j$ be the j-th column of the matrix A. Then the jth column of the Jacobian $Jac_P(x)$ is $dx_j^{d-1}\mathbf{a}_j$ at $x = (x_1, \ldots, x_n)$. If x has at least $n/2$ nonzero coordinates, say $x_{j_1}, \ldots, x_{j_m}$ then the corresponding columns of $Jac_P(x)$ span $\mathbb{F}_p^R$, hence $rank\,(Jac_P(x)) = R$. Thus S_P^* is contained in the union of the $n/2$-dimensional coordinate hyperplanes. □

It is easy to see that most $R \times n$ matrices are non-degenerate. Let A be a random matrix obtained by choosing each of its column vector $\mathbf{a}_i$ independently with probability p^{-R}.

Claim 2. *Let $p \geq 5$, and set $\gamma_0 = \frac{1}{2} - \frac{\log 2}{\log 5}$. If $R = \gamma n$ with $0 < \gamma < \gamma_0$, the the probability that a random matrix is non-degenerate is at least $1 - p^{-(\gamma_0 - \gamma)}$.*

Proof. For a given subspace $M \leq \mathbb{F}_p^R$ of codimension 1, the probability that the columns $\mathbf{a}_{i_1}, \mathbf{a}_{i_2}, \ldots, \mathbf{a}_{i_m}$ of the random matrix A are all contained in M is p^{-m}. Thus the probability that M contains at least $n/2$ columns of A is less than $2^n p^{-n/2}$. Since there are p^R distinct $R-1$ dimensional subspaces, the probability that none of them will contain at least $n/2$ columns of A is at least

$$1 - 2^n p^{R - \frac{n}{2}} \geq 1 - p^{-(\gamma_0 - \gamma)n}.$$

In that case A is non-degenerate. □

Proposition 1 follows immediately from the above Claims.

Next, we consider families of quadratic forms $P = (P_1, \ldots, P_R\,,\ \ P_i(x) = \langle A_i x, x \rangle$ where A_i is a symmetric $n \times n$ matrix with entries in $\mathbb{F}_p$.

Proof of Proposition 2. Let M_r denote the set of $n \times n$ symmetric matrices of rank at most$r$. If $B \in M_r$, then there is an $r' \times r'$ minor B of non-zero determinant, but every minor which properly contains B has determinant 0. If B' is the complementary $n - r' \times n - r'$ minor, then by expanding the corresponding determinant one obtains a linear equation for every entry $a_{ij} \in B'$, moreover each equation involves only one such entry. This gives $\binom{n-r'+1}{2}$ independent equations, and since the number of all possible minors of size $r' \leq r$ is less than 4^n, we have that

$$|M_r| < 4^n p^{\binom{n+1}{2} - \binom{n-r+1}{2}}.$$

If $x \in S_P^*$, then the $R \times n$ matrix with rows $A_1 x, \ldots, A_R x$ has rank less than R, thus there exists scalars λ_i (not all 0), so that

$$x \in Ker\,(\lambda_1 A_1 + \ldots + \lambda_R A_R).$$

Thus if $K = codim\ (S_P^*)$, then there is an r-tuple $(\lambda_1, \ldots, \lambda_R) \neq 0$ for which $dim\,(Ker B_\lambda) \geq n - K - R$, where $B_\lambda = \sum_i \lambda_i A_i$. Hence $B_\lambda \in M_{K+R}$.

For a fixed r-tuple $\lambda \neq 0$ and for a fixed matrix $B \in M_{K+R}$, the number of r-tuple of matrixes such that $\sum_i \lambda_i A_i = B$ is bounded by $p^{(R-1)\binom{n+1}{2}}$, thus the total number of r-tuples of matrixes is at most

$$p^{(R-1)\binom{n+1}{2}} p^{\binom{n+1}{2}-\binom{n-K+R+1}{2}} p^{(1-2\gamma_0)n} \leq p^{R\binom{n+1}{2}} p^{-2\gamma_0 n},$$

assuming that $n-K-R \geq 2\sqrt{n}$, that is $K \leq n-R-2\sqrt{n}$. Proposition 2 follows. □

We will work above $\bar{\mathbb{F}}_p$ the algebraically closed field of characteristic p. Though it is best to view the arguments below by identifying homogeneous algebraic sets with their image in the projective space, we will keep to the affine settings. Let $\mathcal{P}(n,d)$ be the space of homogeneous polynomials $P : \bar{\mathbb{F}}_p^n \to \bar{\mathbb{F}}_p$ of degree d, which is an $N = \binom{n+d-1}{d}$-dimensional linear space over $\bar{\mathbb{F}}_p$. The critical locus $\mathcal{C} \subseteq \mathcal{P}(n,d) \times \bar{\mathbb{F}}_p^n$ is defined as the set of pairs (P,x) such that $x \in S_P^*$, the singular variety of $S_P = P^{-1}(0)$; that is given by the equations $\partial_{x_1} P(x) = \ldots = \partial_{x_n} P(x) = 0$. Its image under the natural projection $\pi : \mathcal{P}(n,d) \times \bar{\mathbb{F}}_p^n \to \mathcal{P}(n,d)$ is the discriminant $\mathcal{D} = \mathcal{D}(n,d)$, which is the locus of polynomials P so that S_P is singular. It is known that $\mathcal{C}$ is a connected smooth variety of dimension equal to N, and $\mathcal{D}$, is a hypersurface in $\mathcal{P}(n,d)$, see Smith-Varley ([11], Cor.2 and Cor.5).

For given $K \geq 1$ let $\mathcal{D}_K = \mathcal{D}_K(n,d)$ be the locus of polynomials P such that $dim\, S_P^* \geq K$, and let $\mathcal{C}_K = \pi^{-1}(D_K)$ be its pre-image. The set $\mathcal{C}_K$ is a Zarisky closed subset of $\mathcal{C}$, moreover the pre-image $\pi^{-1}(P) = S_P^*$ of every point $P \in \mathcal{D}_K$ has dimension at least K. This implies that

$$dim\, \mathcal{D}_K \leq dim\, \mathcal{C}_K - K \leq N - K. \tag{3.1}$$

Thus $\mathcal{D}_K$ has codimension at least K in $\mathcal{P}(n,d)$.

Proof of Proposition 3. For given R let $\mathcal{P}(n,d)^R$ the space of homogeneous polynomial maps $P = (P_1, \ldots, P_R) : \bar{\mathbb{F}}_p^n \to \bar{\mathbb{F}}_p^R$ of degree d. If $x \in S_P^*$ then by definition (1.2), these exists a nonzero $\mu = (\mu_1, \ldots, \mu_R) \in \bar{\mathbb{F}}_p^R$, such that $x \in S_{P_\mu}^*$, where $P_\mu = \mu_1 P_1 + \ldots + \mu_R P_R$. Note that $S_{P_\mu} = S_{P_{\mu'}}$ if $\mu = \lambda \mu'$ for a scalar $\lambda \neq 0$. Thus

$$S_P^* \subseteq \bigcup_{\mu \in \Pi_p^{R-1}} S_{P_\mu}^*, \tag{3.2}$$

where Π_p^{R-1} is the $R-1$ dimensional projective space above $\bar{\mathbb{F}}_p$, considered as an algebraic set in $\bar{\mathbb{F}}_p^R$. This implies that for a given $L \geq 2R-1$, if $dim\, S_P^* \geq L$ then there must exist a $\mu \neq 0$ such that $dim\, S_{P_\mu} \geq L - R + 1$.

Let $\Phi : \Pi_p^{R-1} \times \mathcal{P}(n,d)^R \to \mathcal{P}(n,d)$ be the map defined by $\Phi(\mu, P) = P_\mu$, and let $\pi : \Pi_p^{R-1} \times \mathcal{P}(n,d)^R \to \mathcal{P}(n,d)^R$ be the natural projection. Then the locus of polynomial maps P for which $dim\, S_P^* \geq L$ is contained in

$$\{P \in \mathcal{P}(n,d)^R : \ dim\, S_P^* \geq L\} \subseteq \pi(\Phi^{-1} D_K) \tag{3.3}$$

with $K = L - R + 1$. The tangent map $d\Phi_{(\mu,P)}$ is onto at every point (μ, P) where $\mu \neq 0$, thus the codimension of the algebraic set $\Phi^{-1} D_K$ is at least K. The projection π cannot increase the dimension, hence the codimension of $\pi(\Phi^{-1} D_K) \subseteq \mathcal{P}(n,d)^R$ is at least $K - R + 1$. If $L \geq 2R - 1$ then the set of polynomial maps P for which $dim\ S_P^* \geq L$ has codimension at least $L - 2R + 2 \geq 1$, thus is contained in a proper algebraic set. This proves Proposition 3. □

4. Linear subspaces in homogeneous varieties.

We will show that a homogeneous variety $S = P^{-1}(0)$ contains a large linear subspace M, as an easy corollary of the following result due to Chevalley and Warning ([4], Thm. 6.11)

Theorem C. *Let $Q_1, \ldots, Q_t : \mathbb{F}_p^n \to \mathbb{F}_p$ be polynomials of degree $d_1, \ldots, d_t$ such that $D = d_1 + \ldots + d_t < n$, and $Q_i(0) = 0$ for all i. If S_Q is the common zero set of the polynomials Q_i then*

$$|S_Q| \geq p^{n-D} \tag{4.1}$$

Proposition 5. *Let $S = P^{-1}(0)$, where $P : \mathbb{F}_p^n \to F_p^R$ a homogeneous polynomial map of degree $d < p$. Then S contains a linear subspace M such that*

$$dim\, M \geq c_d (n/R)^{\frac{1}{d}} \tag{4.2}$$

with a constant $c_d > 0$ depending only on d.

Proof. Let M be a maximal subspace, such that $M \subseteq S$. Let $h_1, \ldots, h_m$ be a basis of M. One may write $P_i(x) = Q_i(x, \ldots, x)$ where $Q_i(x_1, \ldots, x_d)$ is a symmetric multi-linear form, as in (2.1). If h is such that $Q_i(h, \ldots, h, h_{i_{k+1}}, \ldots, h_{i_d}) = 0$ for all $1 \leq k \leq d$, $1 \leq i \leq R$, and $1 \leq i_{k+1} \leq \ldots, \leq i_d \leq m$, then $M' = M + \mathbb{F}_p h \subseteq S$ as well. For fixed k this gives $R \binom{m}{k}$ homogeneous equations of degree k. The sum of degrees D of all these equations is bounded by

$$D \leq C_d R\, m^d \tag{4.3}$$

By the Chevalley-Warning's Theorem, the number of such h is at least p^{n-D}. If $m < c_d\, (m/R)^{\frac{1}{d}}$, then $p^{n-D} > p^m = |M|$. Thus one may choose $h \notin M$ such that $M + \mathbb{F}_p h \subseteq S$ contradicting our assumption. This proves the Proposition. □

Corollary 1. *Let $A \subseteq \mathbb{F}_p^n$ of density $\delta > 0$, and let $S = P^{-1}(0)$, where $P : \mathbb{F}_p^n \to F_p^R$ a homogeneous polynomial map of degree $d < p$. If*

$$R < c(\delta, l, d, p)\, n \tag{4.4}$$

then A contains an arithmetic progression $\{x, x + y, \ldots, x + (l-1)y\}$ with common difference $y \in S \backslash \{0\}$.

Proof. Let M be a maximal subspace contained in S. Let $M + x_i$ be a translate of M such that the relative density $\delta_i = |(A \cap (M + x_i)|/|M|$ of A on $m + x_i$ is at least δ. If the dimension m of M is large enough: $m \geq m(\delta, l, p)$ then $A \cap (M + x_i)$ contains a non-trivial arithmetic progression of length l, whose gap y is then in $M \subseteq S$. By (4.2) this happens if $(n/R)^{1/d} > m(\delta, l, d, p)$ for which it is enough to assume (4.4). □

References

[1] J. Birch *Forms in many variables*, Proc. Roy. Soc. Ser. A 265 (1962), 245-263

[2] J. Bourgain *A Szemerédi type theorem for sets of positive density in $\mathbb{R}^k$*, Israeli J. Math, 54 (1986), 307-316

[3] B. Green *On arithmetic structures in dense sets of integers*, Duke Math. J. 114 (2) (2002), 215-238

[4] R. Lidl and H. Niederreiter Finite Fields, Cambridge University Press (1997)

[5] B. Green and T. Tao, *Linear equations in the primes*, Annals of Math. (to appear)

[6] S. Ghorpade and G. Lachaud *Étale cohomology, Lefshetz theorems and number of points of singular verieties over finite fields*, Moscow Math. J. (2) (2002), 589-631

[7] R. Hartshorne, *Algebraic Geometry*, Graduate texts in mathematics: 52, Springer-Verlag (1977)

[8] A. Sárközy, *On difference sets of sequences of integers III*, Acte Math. Acad. Sci. Hungar. 31 (1978), 355-386

[9] E. Szemerédi, *On sets of integers containing no k elements in arithmetic progression*, Acta Arith. 27 (1975), 299-345

[10] T. Tao and V. Vu *Additive combinatorics*, Cambridge University Press (2004)

[11] R. Smith and R. Varley *The tangent cone to the discriminant*, Proc. Conf. in Alg. Geom. Vancouver (1984)

[12] T. Wooley and T. Ziegler, *Multiple recurrence and convergence along the primes*, preprint (2010)

Department of Mathematics, The University of British Columbia, Vancouver, BC, V6T1Z2, Canada
E-mail address: bcook@math.ubc.ca

Department of Mathematics, University of British Columbia, Vancouver, B.C. V6T 1Z2, Canada
E-mail address: magyar@math.ubc.ca