Weil Bounds: A Survey of Decoding Applications

Soham Chatterjee
May 2026

School of Technology and Computer Science
Tata Institute of Fundamental Research

Goal: Learning polynomials from noisy evaluations

Secret Polynomial

$f \in \mathbb{F}[X]$

Degree of $f < d$

Evaluation Points

$\alpha_1, \ldots, \alpha_n \in \mathbb{F}$

All distinct

Evaluation Vector

n-tuple β, such that

All positions $f(\alpha_i) = \beta_i$

For our setting some locations of β got corrupted. r is the noisy evaluation vector.

Let f agrees with r in t many positions.

Question: 1. Can we learn f from the noisy evaluations ?

2. If t is small recover all such f which agrees with r in t many locations

We understand this question very well. The work of Berlekamp-Welch, Sudan (1997) and Guruswami-Sudan (1999), C-Harsha-Kumar solves both of these questions efficiently.

What if: The evaluation vector has small alphabet

Problem: Earlier the evaluation vector has large range of possible values, $|\mathbb{F}|$.

f is the secret polynomial over $\mathbb{F}$ with degree less than d.

Want: To have some function $\chi : \mathbb{F} \to S$ where S is very small subset of $\mathbb{F}$.

Evaluation vector is $\chi \circ f(\alpha_i)$ for each evaluation point. Even the noise is from S

Question: 1. Can we still recover f?

2. Are there examples of such constructions

Idea: Use the power of finite fields (Odd Character Case)

Let $|\mathbb{F}| = p$ where p is a prime. Take p odd.

For all $\alpha \in \mathbb{F}$, if α is non-zero then $\alpha^{p-1} = 1$. So $\alpha^{\frac{p-1}{2}} = \pm 1$

Fact: α is called a square if $\alpha^{\frac{p-1}{2}} = 1$

So take $\chi(\alpha) = \alpha^{\frac{p-1}{2}}$ (Quadratic functions). Then for any α,

$f(\alpha)$ is a square	$f(\alpha)$ is a non-square	$f(\alpha) = 0$
$\chi \circ f(\alpha) = 1$	$\chi \circ f(\alpha) = -1$	$\chi \circ f(\alpha) = 0$

Take the evaluation vector to be $\chi \circ f(\alpha)$ for all α in the finite field.

Every entry of the evaluation vector is from $\{-1,0,1\}$.

Some assumptions on the polynomials

Observation: Since the map $\chi(\alpha) = \alpha^{\frac{p-1}{2}}$ is a multiplicative map for any β

$$\chi(\beta \cdot f(\alpha)) = \chi(\beta) \cdot \chi(f(\alpha))$$

For multiple β we can get the same evaluation vector.

Observation: For any polynomial g, and point α which is not a zero of g

$$\chi(g^2(\alpha)) = \chi^2(g(\alpha)) = 1$$

So f and $f \cdot g^2$ give the same evaluation vector.

Same evaluation vector even in zero-noise case can give multiple solutions.

We always refer to the “minimal” polynomial. The monic and square free one.

How many locations do two evaluation vectors differ

Fact: For any non-zero α such that $\alpha^{\frac{p-1}{2}} = 1$ there exists a β such that $\beta^2 = \alpha$.

Consider evaluation of f at any point α then if $f(\alpha)$ is a square

$$\implies \text{we get a solution for } Y^2 = f(X)$$

f, g are two distinct polynomials of degree less than d

- At α both f, g give same evaluation.
- Evaluation vector of $f \cdot g$ gives 1 or 0 at α

If degree of f is small the bound on number of solutions of $Y^2 = f(X)$

$\implies$ How many points at least the evaluation vectors of two distinct polynomials differ.

Weil Bounds Theorem (Andrè Weil)

Theorem: If $Q(X, Y)$ is an absolutely irreducible polynomial of total degree d over $\mathbb{F}_p$ then number of solutions of Q is

$$|\{(\alpha, \beta) : Q(\alpha, \beta) = 0\}| = p + O(d\sqrt{p})$$

Weil Bounds

For our use case $Q(X, Y) = Y^2 - f(X)$.

Fact: Such curves are called elliptic curves

Stepanov gave an elementary proof of above bound.

Let S_0 is the set of zeros of f and S_1 is the set of all points such that $f^{\frac{p-1}{2}} = 1$.

Number of solutions of $Y^2 = f(X)$ is $|S_0| + 2|S_1|$

Weil Bounds

Condition: degree of f is less than $\sqrt{p}$

- Step 1. Find $P(X) = A(X, X^p) + f^{\frac{p-1}{2}}(X) \cdot B(X, X^p)$ that vanishes at S_1 with high multiplicity M
- Step 2: Then $R(X) = P(X) \cdot f^M(X)$ vanishes at every point of S_0 or S_1 with multiplicity at least M
- Step 3. Total size of S_0 and S_1 is at most the degree of R divided by M

With proper analysis of the degree of $R \implies$ bound the number of zeros of $Y^2 = f(X)$ is $p + O(d\sqrt{p})$

When can we recover from a noisy evaluation vector

Let f, g are two polynomials of degree less than $d << \sqrt{p}$. Monic and square-free.

Weil Bounds $\implies \chi \circ f$ and $\chi \circ g$ differ in at least $O(p)$ many locations.

Recall we obtain the noisy evaluation vector from $\chi \circ f$ where f is the secret polynomial

If noisy evaluation vector agrees with f in more than $p/2$ many locations $\implies f$ is unique.

With at less than $p/2$ many error locations

$\implies$ we should get the secret polynomial of degree less than $\sqrt{p}$.

State of art for recovery of secret polynomial

Bad News: Till last year there has been no known algorithms to recover.

Even the zero-noise case was not known to be efficiently solvable

Recently in the paper by [Kopparty, 2026] he gave an algorithm:

- Recovers from $1/8^{th}$ fraction of the evaluation points being corrupted
- Gave a polynomial time deterministic algorithm

How to recover polynomial ?

Given: Noisy evaluation vector $r \in \{-1,0,1\}^p$ that agrees with f in more than $p/2$ locations.

We'll approach the recovery procedure in the steps of Berlekamp-Welch.

Let S is the set of points α such that $\chi \circ f(\alpha) \neq v(\alpha)$.

Want: A polynomial which vanishes with high multiplicity at any point iff that point is in S.

So consider the polynomial

$$Z_S(X) = \prod_{\alpha \in S} (X - \alpha)^M.$$

So a potential polynomial to work with is $Z_S(X) \cdot f^{\frac{p-1}{2}}(X)$.

Better polynomials to work with

For all α we can write $(Z_S \cdot f^{\frac{p-1}{2}})(\alpha) = r(\alpha) \cdot g(\alpha)$ where g is a polynomial of degree less than p.

Problem: But we want to say the same similarly for all it's M derivatives.

Idea: 1. Take $F(X) = f^{\frac{p-1}{2}+M}(X)$ instead

2. Take a multiple of $Z_S(X)$ of the form:

$$E(X) = \sum_i E_i(X) \cdot (X^p - X)^i$$

Where each E_i is of low degree

Better polynomials to work with

So we have

$E(X)$ multiple of $Z_S(X)$ And $F(X) = f^{\frac{p-1}{2}+M}(X)$

The polynomial we want to work with $E(X) \cdot F(X)$.

Fact: For every $0 \leq \ell < M$,

$$(E \cdot F)^{(\ell)}(\alpha) = r(\alpha) \cdot g_\ell(\alpha)$$

where g_ℓ are of low degree

Setting parameters for the new polynomial

So we want to work with $E(X) \cdot F(X)$ where

$$E(X) = \sum_{i=0}^{c-1} E_i(X) \cdot (X^p - X)^i$$

Such that E_i of degree at most k and E vanishes at all error locations with multiplicity M

Now if number of error locations is $e < (1/8 - \epsilon)\, p$ and degree of f, $d \leq \epsilon\sqrt{p}$ then

There exists some good setting of parameters such that

Then such E exists and for all $0 \leq \ell < M$ we have the relation:

$$(E \cdot F)^{(\ell)}(\alpha) = r(\alpha) \cdot g_\ell(\alpha)$$ with degree of g_ℓ is at most $k + dM$

Notion of Pseudopolynomials

The polynomials of the form

$$Q(X) = \sum_{i=0}^{d} Q_i(X) \cdot (X^p - X)^i$$

Where each Q_i has degree at most k are called $k-$pseudoolynomial

Observe: All derivatives of Q are of degree at most k

Theorem: If H is an irreducible factor of Q with multiplicity m then,

$$m \bmod q \in [0, d+k]$$

So if we factorize $E \cdot F \implies$ factors of E does not have high multiplicity but

Factors of F have very high multiplicity.

There is a rich theory of pseudo polynomials and codes constructed from them in [Kopparty, 2026]

Algorithm

With the setting of parameters as earlier we can design the algorithm

Given:
- A noisy evaluation vector with at most $e \leq (1/8 - \epsilon)p$ many corrupted locations
- Degree $d \leq O(\epsilon\sqrt{p})$

Output: The secret polynomial f with degree at most $d \leq O(\epsilon\sqrt{p})$

Runtime: poly (e, d, p)

Algorithm

Step 1: Set the parapets accordingly.

Step 2: Solve system of linear equations to find $G, g_0, \ldots, g_{M-1}$ such that

$$G^{(\ell)}(\alpha) = r(\alpha) \cdot g_\ell(\alpha)$$

Step 3: Completely factorize G into distinct monic irreducible factors

Step 4: Take the product of factors with multiplicity in the range $[3p/8,\ 7p/8] \bmod p$

How the algorithm finds G related to F ?

The algorithm found G but we want to find $F = f^{\frac{p-1}{2}+M}$

Idea: We can relate F, G using pseudo polynomials

We can find polynomials A, B such that:

- A is a $3M/8-$pseudopolynomial
- B is a $e + 4dM-$pseudopolynomial

$$A(X) \cdot G(X) \equiv B(X) \cdot F(X)$$

Compare Multiplicities

Recall Fact: If H is an irreducible factor of a $k-$pseudopolynomial Q with degree $d \cdot p$ then,

$$\textbf{Mult}\,(H, Q) = \text{Multiplicity of } H \bmod p \in [0,\ d + k] \quad \text{(low)}$$

Found pseudopolynomials A, B such that

$$A(X) \cdot G(X) \equiv B(X) \cdot F(X)$$

Recall:

$G(X)$ found by algorithm $\qquad$ $F(X) = f^{\frac{p-1}{2}}(X)$ where f is secret polynomial

For any factor H of F:

$$\textbf{Mult}\,(H, G) = \textbf{Mult}\,(H, F) + \textbf{Mult}\,(H, B) - \textbf{Mult}\,(H, A)$$

Generalization of the result

In general for maps of the form $\chi_m(\alpha) = \alpha^{\frac{p-1}{m}}$

Take the evaluation vector $\chi_m \circ f(\alpha)$

The corresponding polynomial equation to look

$$Y^m = f(X)$$

The Weil bound for this equation gives the same bound $p + O(d\sqrt{p})$.

$\Longrightarrow$ Less than $p/2$ many errors uniquely determines secret polynomial.

Can generalize the algorithm to recover from $1/8$ fraction of errors

What about even characteristic field ?

For such finite field there is no distinction between 1 and -1

Need: to use a different map to compose on the secret polynomials

Idea: Use the Trace function

Suppose we are working in the field $\mathbb{F}_q$ of size q where $q = 2^n$

So $\mathbb{F}_q$ is a finite extension of the field $\mathbb{F}_2$

The trace function Tr: $\mathbb{F}_q \to \mathbb{F}_2$ is the map

$$\mathrm{Tr}(X) = X + X^2 + X^{2^2} + \cdots + X^{2^{n-1}}$$

Fact: For all α in $\mathbb{F}_q$, $\mathrm{Tr}(\alpha) \in \{0,1\} = \mathbb{F}_2$

Fact: These are called Dual BCH codes

So take the evaluation vector $\mathrm{Tr} \circ f(\alpha)$ for all α in the field. Every entry of evaluation vector is from $\{0,1\}$.

Some assumptions on the polynomials

The trace function, $\mathrm{Tr}(X) = X + X^2 + X^{2^2} + \cdots + X^{2^{n-1}}$ is additive

f is a polynomial of degree less than d.

Write f as: $f(X) = a + g(X) + h(X) + h^2(X)$ where

- $g(X)$ has only odd degree monomials
- $h(X)$ has constant term 0

$$f(X) = 1 + X^2 + X^3 + X^4$$
$$a = 1$$
$$g(X) = X + X^3$$
$$h(X) = X$$

For any α,

$$\mathrm{Tr} \circ h(\alpha) = \mathrm{Tr} \circ h^2(\alpha) \implies \mathrm{Tr} \circ f(\alpha) = \mathrm{Tr}\,(a) + \mathrm{Tr} \circ g(\alpha)$$

So we assume f has only odd degree monomials

How many locations do two evaluation vectors differ

Fact: For any α, $\mathrm{Tr}(\alpha) = 0$ if and only if there exists a β so that $\alpha = \beta^2 - \beta$

Consider evaluation of f at any point α then if $\mathrm{Tr} \circ f(\alpha)$ is zero

$\implies$ we get a solution for $Y^2 - Y = f(X)$

Fact: Such curves are called *Artin-Schreier* equations

f, g are two distinct polynomials of degree less than d

- At α both f, g give same evaluation.
- Evaluation vector of $f + g$ gives 0 at α

If degree of f is small the bound on number of solutions of $Y^2 - Y = f(X)$

$\implies$ How many points at least the evaluation vectors of two distinct polynomials differ.

Weil Bounds

Recall the Weil Bound theorem

Andrè Weil: If $Q(X, Y)$ is an absolutely irreducible polynomial over $\mathbb{F}_q$

$$|\{(\alpha, \beta) : Q(\alpha, \beta) = 0\}| = q + O(\sqrt{q})$$

For our use case $Q(X, Y) = Y^2 - Y - f(X)$.

Bombieri gave an elementary proof of above bound for this case.

The proof follows similar steps like as in the case of $Y^2 - f(X)$

$\implies$ Number of solutions of $Y^2 - Y = f(X)$ is $q + O(d\sqrt{q})$

$\implies$ with less than $q/2$ many error locations f can be uniquely found

Current Situation

Like the previous case until last year we didn't know even to solve for the zero-noisy case

In the recent paper by [Kopparty, 2026]:

They were able to recover from $q/8$ many errors.

The algorithm follows the similar steps as in the previous case:

- They use the power of pseudo polynomials
- They finds the secret polynomial monomial by monomial

Open Questions

- The algorithms so far recovers from only $1/8^{th}$ of the allowed number of errors.

  We don't know how to extend this to $1/2$ of total number of evaluation points.

- Recovering set of all close enough polynomials with more than $1/2$ fraction of errors is open (List decoding).

- Codes using pseudopolynomials and their derivatives evaluated on distinct points.

  They give good distance.

  There is no efficient deterministic algorithm to interpolate from noisy evaluation.

- There are other polynomials whose bounds on number of zeros are known.

  We can ask same questions for them. And it's not been studies in detail.

Thank You