Where are the zeroes of a random p-adic polynomial?

Xavier Caruso

October 30, 2018

1 Training: over finite fields

Throughout this talk, p is a fixed prime number.

Expected number of roots over $\mathbb{F}_p$

Let $P \in \mathbb{F}_p[X]$ be a random polynomial of degree $d \geq 1$. (By "random", we mean that each coefficient is chosen uniformly and independanty in the finite set $\mathbb{F}_p$.)

We are interested in the expected number of roots of P in $\mathbb{F}_p$. In order to compute it, for each $a \in \mathbb{F}_p$, we define the Bernoulli variable B_a by:

$$\begin{aligned} B_a(P) &= 1 \qquad \text{if } a \text{ is a root of } P \\ &= 0 \qquad \text{otherwise.} \end{aligned}$$

The fact that the mapping $\varepsilon_a : \mathbb{F}_p[X] \to \mathbb{F}_p$, $P \mapsto P(a)$ is a surjective ring homomorphism implies that $\mathbb{P}(P(a) = b) = \frac{1}{p}$ for all $b \in \mathbb{F}_p$. Thus $\mathbb{P}(B_a = 1) = \frac{1}{p}$ as well.

Now we observe that the number of zeroes of P in $\mathbb{F}_p$ is

$$Z = \sum_{a \in \mathbb{F}_p} B_a. \tag{1}$$

Therefore:

$$\mathbb{E}[Z] = \sum_{a \in \mathbb{F}_p} \mathbb{E}[B_a] = p \times \frac{1}{p} = 1.$$

Observe that this not does depend on p nor or d (when $d \geq 1$).

Expected number of roots over $\mathbb{F}_q$

Let $q = p^n$ be a power a p. We can count the number of roots of P over $\mathbb{F}_q$ using the same techniques. For each $a \in \mathbb{F}_q$, we define the Bernoulli variable B_a as above and consider the ring homomorphism $\varepsilon_a : \mathbb{F}_p[X] \to \mathbb{F}_q$, $P \mapsto P(a)$. If $n \geq d$, the image of ε_a is $\mathbb{F}_p[a]$ (almost by definition). Therefore $\mathbb{P}(B_a = 1) = \frac{1}{\#\,\mathbb{F}_p[a]}$. Consequently, if Z_n is the random variable counting the number of roots of P in $\mathbb{F}_q$, we have:

$$\mathbb{E}[Z_n] = \sum_{a \in \mathbb{F}_q} \frac{1}{\#\,\mathbb{F}_p[a]} = \sum_{a \in \mathbb{F}_q} p^{-\deg(a)} \tag{2}$$

where $\deg(a)$ is the degree of a, that is the degree of the extension generated by a. Observe that it does not depend on d (as soon as $d \geq n$).

Repartition of roots in $\bar{\mathbb{F}}_p$

Let N_n be the number of elements of $\bar{\mathbb{F}}_p$ with degree n, that is the number of generators of $\mathbb{F}_{p^n}$. From the relation $\sum_{m|n} N_m = p^m$, we derive the formula:

$$N_n = \sum_{m|n} \mu\Big(\frac{n}{m}\Big) p^m = p^n + O(p^{n/2})$$

where μ is the Moebius function. Let $Z'_n(P)$ be the number of roots of P lying in $\mathbb{F}_{p^n}$ but not in a strict subextension. By what we have achieved before, the expected value of Z'_n is:

$$\mathbb{E}[Z'_n] = N_n p^{-n} = 1 + O(p^{-n/2}).$$

Of course we can be more precise for the smallest values of n. Here is what we find:

$$\mathbb{E}[Z'_1] = 1 \quad ; \quad \mathbb{E}[Z'_2] = 1 - \tfrac{1}{p} \quad ; \quad \mathbb{E}[Z'_3] = 1 - \tfrac{1}{p^2}$$
$$\mathbb{E}[Z'_4] = 1 - \tfrac{1}{p^2} \quad ; \quad \mathbb{E}[Z'_5] = 1 - \tfrac{1}{p^4} \quad ; \quad \mathbb{E}[Z'_6] = 1 - \tfrac{1}{p^3} - \tfrac{1}{p^4} + \tfrac{1}{p^5}$$

Roughly speaking, a random polynomial of degree d has, on average, one root in $\mathbb{F}_p$, one more root in $\mathbb{F}_{p^2}$, one more root in $\mathbb{F}_{p^3}$, etc. until $\mathbb{F}_{p^d}$. This gives d roots as expected. Trying to be more precise is actually quite interesting. Indeed, we observe that $\mathbb{E}[Z'_n] < 1$ as soon as $n > 1$. Thus the $\mathbb{E}[Z'_n]'s$ for $1 \leq d \leq n$ cannot sum up to d exactly. This is due to the existence of multiple roots, of course. When d goes to infinity, one can prove that:

$$\sum_{n=1}^{d} \mathbb{E}[Z'_n] = d - \sum_{i=2}^{\infty} \frac{\mu(i)}{p^{i-1}-1} + o(1)$$

The default $\sum_{i=2}^{\infty} \frac{\mu(i)}{p^{i-1}-1}$ (which is a constant, not depending on d) is then the number of "redundant" roots; it is also the expected value for the degree of $\gcd(P, P')$ when the degree of P grows up.

2 Expected number of roots over $\mathbb{Q}_p$

We now move to p-adic polynomials.

We recall that $\mathbb{Z}_p$ is a compact additive group. It is then equipped with a canonical measure of probability, namely its Haar measure. We denote it by μ and extend it to $\mathbb{Q}_p$. We let $|\cdot|$ denote the norm over $\mathbb{Q}_p$, normalized by $|p| = p^{-1}$. With this normalization, we have $\mu(aE) = |a|\mu(E)$ for all $a \in \mathbb{Q}_p$ and all measurable subset E of $\mathbb{Q}_p$.

In what follows, we will denote by Ω_d the set of polynomials with coefficients in $\mathbb{Z}_p$ and degree at most d and equip Ω_d with the Haar measure. Choosing at random an element in Ω_d then amounts at choosing independantly each coefficient with respect to the Haar measure on $\mathbb{Z}_p$.

Kac formula

Kac formula is a continuous equivalent of the trivial summation formula (1). In the p-adic case, it reads:

$$Z_H(P) = \lim_{s\to\infty} p^s \int_H |P'(a)| \cdot \mathbb{1}_{\{|P(a)|\leq p^{-s}\}}\, da \tag{3}$$

where H is a measurable subset of $\mathbb{Q}_p$ and $Z_H(P)$ counts the number of zeroes of P in H. From Kac formula, it is easy to derive the expected value of Z_H. Indeed, one writes:

$$\begin{aligned}\int_{\Omega_d} Z_H(P) dP &= \int_{\Omega_d} \lim_{s\to\infty} p^s \int_H |P'(a)| \cdot \mathbb{1}_{\{|P(a)|\leq p^{-s}\}}\, da\, dP \\ &= \int_H \lim_{s\to\infty} p^s \int_{\Omega_d} |P'(a)| \cdot \mathbb{1}_{\{|P(a)|\leq p^{-s}\}}\, dP\, da = \int_H \rho(a) da\end{aligned}$$

where the function ρ is the "density function" defined by:

$$\rho(a) = \lim_{s\to\infty} p^s \int_{\Omega_d} |P'(a)| \cdot \mathbb{1}_{\{|P(a)|\leq p^{-s}\}}\, dP. \tag{4}$$

Moreover, it turns out that ρ has a simple expression. Indeed first observe that for $a = 0$ and $P = a_0 + a_1 X + \cdots + a_d X^d$, we have:

$$\rho(0) = \lim_{s\to\infty} p^s \int_{\mathbb{Z}_p^2} |a_1| \cdot \mathbb{1}_{\{|a_0|\leq p^{-s}\}}\, da_0\, da_1 = \int_{\mathbb{Z}_p} |a_1| da_1 = \frac{p}{p+1}.$$

More generally, for $a \in \mathbb{Z}_p$, the above result is still valid since the morphism $\Omega_d \to \Omega_d$, $P(X) \mapsto P(X-a)$ preserves the measure. For $a \in \mathbb{Q}_p$, $a \notin \mathbb{Z}_p$, we perform the change of variables $P(X) \mapsto X^d P(\frac{1}{X})$ and get this way $\rho(a) = \rho(\frac{1}{a}) \cdot |a|^{-2}$. Putting everything together, we end up with the simple formula:

$$\begin{aligned} \rho(a) &= \frac{p}{p+1} && \text{if } a \in \mathbb{Z}_p \\ &= \frac{p}{p+1} \cdot \frac{1}{|a|^2} && \text{if } a \notin \mathbb{Z}_p. \end{aligned}$$

From this, we derive $\mathbb{E}[Z_H] = \mu(H) \cdot \frac{p}{p+1}$ if $H \subset \mathbb{Z}_p$ and $\mathbb{E}[Z_{\mathbb{Q}_p}] = 1$ (independently from p and d).

3 Expected number of roots over a finite extension

We now want to count the roots of P, not only in $\mathbb{Q}_p$ but more generally in any given finite extension of $\mathbb{Q}_p$. We fix such an extension K. We denote by n its degree and by $\mathcal{O}_K$ its ring of integers. We endow K with its Haar measure μ_K, normalized by $\mu_K(\mathcal{O}_K) = 1$. We recall that the norm $|\cdot|$ extends uniquely to $\mathcal{O}_K$ and we continue to use to notation $|\cdot|$ to refer to the norm on K. The relation between the measure and the norm now reads $\mu_K(aE) = |a|^n \mu_K(E)$ for $a \in K$ and $E \subset K$. Kac formula is still valid in this extended framework. It now reads:

$$Z_H(P) = \lim_{s\to\infty} p^s \int_H |P'(a)|^n \cdot \mathbb{1}_{\{|P(a)|^n \le p^{-s}\}}\, da \tag{5}$$

where H is any measurable subset of K. We are now tempted to write

$$\mathbb{E}[Z_H] = \int_H \rho_K(a) da \quad \text{where} \quad \rho_K(a) = \lim_{s\to\infty} p^s \int_{\Omega_d} |P'(a)|^n \cdot \mathbb{1}_{\{|P(a)|^n \le p^{-s}\}}\, dP.$$

However we have to be careful since (1) the latter value is not always finite (more precisely, it is infinite as soon as if $\mathbb{Q}_p[a] \neq K$) and (2) it may depend on d. Instead we define:

$$\begin{aligned} \rho_{K,d}(a) &= \lim_{s\to\infty} p^s \int_{\Omega_d} |P'(a)|^n \cdot \mathbb{1}_{\{|P(a)|^n \le p^{-s}\}}\, dP && \text{if } \mathbb{Q}_p[a] = K \\ &= 0 && \text{otherwise} \end{aligned}$$

and we extend this definition to any subextension L of K. We can then prove that the $\rho_{K,d}$'s are well defined and take finite values over K (*cf* Theorem 2 below for a more precise statement). Moreover, we have the following result.

Theorem 1 *For any measurable subset H of K, we have:*

$$\int_{\Omega_d} Z_H(P) dP = \sum_L \int_{H\cap L} \rho_{L,d}(a) da$$

where the sum is extended to all fields L with $\mathbb{Q}_p \subset L \subset K$.

In other words, the above theorem says that $\int_K \rho_{K,d}(a)da$ counts the average number of roots of a random p-adic polynomial that lie in K but not in any proper subfield. Besides, in many cases, we can give alternative and easier formulae for $\rho_{L,d}$.

Theorem 2 *Let L be a field with $\mathbb{Q}_p \subset L \subset K$. For $a \in \mathcal{O}_L$ such that $\mathbb{Q}_p[a] = L$, we have:*

$$\begin{aligned} &\rho_{L,d}(a) = 0 && \text{if } d < [L:\mathbb{Q}_p] \\ &\rho_{L,d}(a) = \frac{|D_L|}{[\mathcal{O}_L : \mathbb{Z}_p[a]]} \cdot \frac{p^d}{p^d + p^{d-1} + \cdots + 1} && \text{if } d = [L:\mathbb{Q}_p] \\ &\rho_{L,d}(a) < \rho_{L,d+1}(a) && \text{if } [L:\mathbb{Q}_p] \le d < 2[L:\mathbb{Q}_p] - 1 \\ &\rho_{L,d}(a) = |D_L| \cdot \int_{\mathbb{Z}_p[a]} |x|^{[L:\mathbb{Q}_p]} dx && \text{if } d \ge 2[L:\mathbb{Q}_p] - 1 \end{aligned}$$

where D_L is discriminant of the extension $L/\mathbb{Q}_p$.

Remark 3 *When $a \notin \mathcal{O}_L$, we have the relation $\rho_{L,d}(a) = \rho_{L,d}(\frac{1}{a}) \cdot |a|^{-2n}$ which gives the value of $\rho_{L,d}(a)$ when combined with Theorem 2 (noting that $\frac{1}{a} \in \mathcal{O}_L$).*

Remark 4 *Theorem 2 implies that $\rho_{L,d}(a) < \frac{|D_L|}{[\mathcal{O}_L:\mathbb{Z}_p[a]]}$ for all d and $a \in \mathcal{O}_L$.*
Moreover, when $\mathbb{Z}_p[a] = L$, one can be more precise and compute the maximal value of $\rho_{L,d}(a)$; we find $\rho_{L,d}(a) = |D_L| \cdot \frac{p^n}{p^n+1}$ for $d \geq 2[L:\mathbb{Q}_p] - 1$.

Furthermore, as a corollary of Theorem 2, we obtain:

Corollary 5 *Let L be a field with $\mathbb{Q}_p \subset L \subset K$ with $L \neq \mathbb{Q}_p$. Then:*

$$\int_L \rho_{L,d}(a)da < |D_L|.$$

Moreover, if f is the residual degree of $L/\mathbb{Q}_p$, we have the estimation:

$$\int_L \rho_{L,d}(a)da = |D_L| \cdot \left(\frac{N_f}{p^f} + O\left(\frac{1}{p^f}\right)\right)$$

where N_f is the number of generators of $\mathbb{F}_{p^f}$ as defined in §1 and the constant hidden in the $O(\cdot)$ is absolute.

4 Repartition of roots in an algebraic closure

Roots in unramified extensions

For all positive integers n, let $\mathbb{Q}_{p^n}$ denote the unique unramified extension of $\mathbb{Q}_p$ of degree n. Let also $\mathbb{Q}_p^{\text{ur}}$ be the union of all $\mathbb{Q}_{p^n}$; this is the maximal unramified extension of $\mathbb{Q}_p$.

By Corollary 5, a random p-adic polynomial of large degree has $N_n p^{-n} + O(p^{-n})$ "new" roots in $\mathbb{Q}_{p^n}$. We observe that this is very close to the number of "new" roots in $\mathbb{F}_{p^n}$ of a random polynomial over $\mathbb{F}_p$ (which is exactly $N_n p^{-n}$). However, the reason for this coincidence is not clear (at least to me).

For polynomials of degree d, the repartition is the following:

$$\begin{aligned}
\mathbb{E}[Z_{\mathbb{Q}_p}] &= 1 \\
\mathbb{E}[Z_{\mathbb{Q}_{p^2}}] &= 1 - \frac{1}{p} + O\Big(\frac{1}{p^2}\Big) \\
\mathbb{E}[Z_{\mathbb{Q}_{p^3}}] &= 1 \qquad + O\Big(\frac{1}{p^2}\Big) \\
&\vdots \\
\mathbb{E}[Z_{\mathbb{Q}_{p^{d-1}}}] &= 1 \qquad + O\Big(\frac{1}{p^2}\Big) \\
\mathbb{E}[Z_{\mathbb{Q}_{p^d}}] &= 1 - \frac{1}{p} + O\Big(\frac{1}{p^2}\Big)
\end{aligned}$$

The total number of roots in $\mathbb{Q}_p^{\text{ur}}$ is then $d - \frac{2}{p} + O(\frac{1}{p^2})$. Observing that a p-adic random polynomial has no multiple roots almost surely, we deduce that it has $\frac{2}{p} + O(\frac{1}{p^2})$ roots outside $\mathbb{Q}_p^{\text{ur}}$ on average.

Roots outside $\mathbb{Q}_p^{\text{ur}}$

It is actually possible to better locate the $\frac{2}{p} + O(\frac{1}{p^2})$ "missing roots". We will do it when $p > 2$. With this assumption, $\mathbb{Q}_p$ has exactly two totally ramified of degree 2, namely $K_1 = \mathbb{Q}_p[\sqrt{p}]$ and $K_2 = \mathbb{Q}_p[\sqrt{up}]$ where $u \in \mathbb{Z}_p^\times$ is not a square. It is easy to check that $|D_{K_1}| = |D_{K_2}| = \frac{1}{p}$. Corollary 5 gives:

$$\mathbb{E}[Z_{K_1}] = \mathbb{E}[Z_{K_2}] = \frac{1}{p} + O\Big(\frac{1}{p^2}\Big)$$

and we have found (a large part of) the missing roots: $\frac{1}{p}$ of them is in K_1, $\frac{1}{p}$ of them is in K_2.

Playing a similar game, one can count the expected number of roots outside the maximal tamely ramified extension of $\mathbb{Q}_p$. Here is what we find.

Theorem 6 *A random p-adic polynomial of degree at least $2p-1$ has $p^{-p+2} + O(p^{-p+1})$ outside the maximal tamely ramified extension of $\mathbb{Q}_p$.*

More precisely, a theorem of Krasner asserts that there exist exactly p^2 totally ramified extensions of $\mathbb{Q}_p$ of degree p and discriminant p^p. By Corollary 5, a p-adic random polynomial has $p^{-p} + O(p^{-p-1})$ roots in each such extension. Summing up the corresponding contributions, we find the $p^{-p+2} + O(p^{-p+1})$ roots promised by Theorem 6.

5 Numerical experimentations

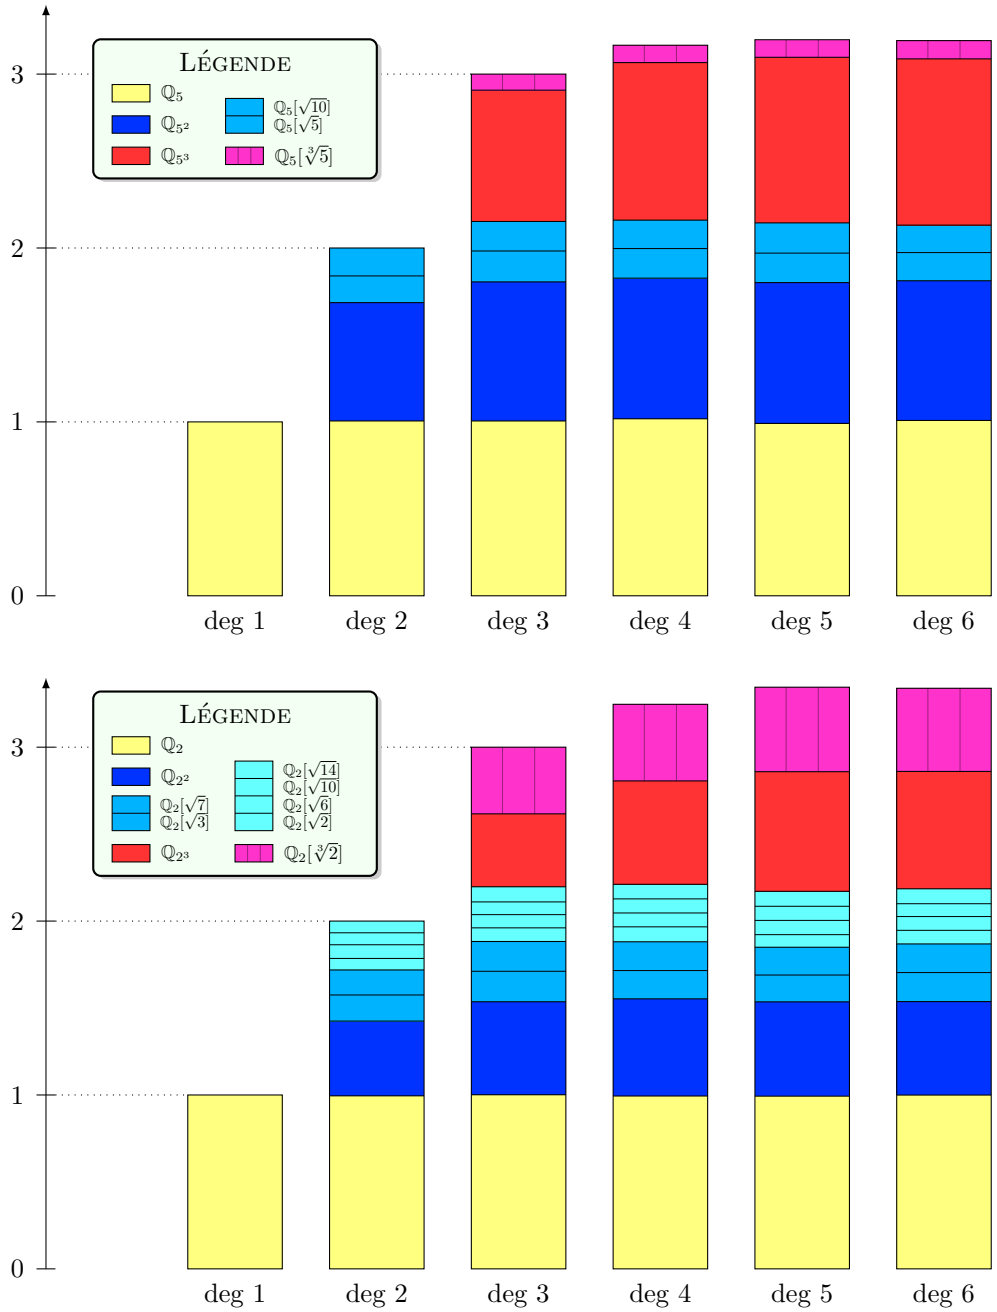

6 Further questions

We can imagine several natural generalizations of the results discussed above.

First of all, we may vary the distribution on the input polynomials: instead of taking uniformly polynomials in Ω_d, one can for instance consider polynomials of the form $X^d + pR(x)$ where R is picked uniformly in Ω_{d-1}. This particular case would lead to very different results (almost all roots would be in totally ramified extensions) and then promise to be interesting. Similarly, we can restrict ourselves to polynomials having a given Newton polygon or, more generally, we can choose a lattice $\mathcal{L}_d$ inside $\mathbb{Q}_p[X]_{\leq d}$ and pick polynomials uniformly in $\mathcal{L}_d$.

Another thing we can do is to consider multivariate polynomials. More precisely, if $P_1, \ldots, P_m$ are polynomials in m variables (with the same m), the set of common zeroes of the P_i's is finite almost surely and we can study its expected cardinality. Using a multidimensional analogue of Kac formula, one can prove that a system of m polynomials in m variables has again 1 root in $\mathbb{Q}_p^m$ on average. Using similar techniques, it should be feasible to study the number of roots in algebraic extensions.

Instead of working with a system of m polynomials exactly, we can also pick a family of c polynomials with $c < m$. The variety it defines has then codimension c, so that its number of points is expected to be infinite. However, it should has a finite $(m-c)$-dimensional measure[1] and we can study its average. It is likely that Kac formula extends to this more general framework and that one can derive from it interesting results.

[1] The d-dimensional measure of the subset A of $\mathbb{Q}_p^m$ is defined as the limit of the quantity $p^{-sd} \cdot \#\left(A \mod p^s\right)$ when s goes to infinity.