NB: throughout this section, $|\cdot|$ will denote the p-adic absolute value.

5.1 Functions and Continuity

We have now built up $\mathbb{Q}_p$ as an analogue of $\mathbb{R}$ (in particular, as another completion of $\mathbb{Q}$). We want to develop a theory of functions on $\mathbb{Q}_p$.

Since we have an absolute value on $\mathbb{Q}_p$, we can define continuity the same way we do in $\mathbb{R}$:

Definition 5.1

Let $U \subset \mathbb{Q}_p$ be an open set. A function $f : U \to \mathbb{Q}_p$ is **continuous** at $x_0 \in U$ if for all $\delta > 0$ there exists $\epsilon > 0$ such that

$$|x - x_0| < \delta \implies |f(x) - f(x_0)| < \epsilon$$

For example, polynomials are continuous everywhere (same proof as in $\mathbb{R}$). However, the function defined by $f(x) = 1/x$ for $x \neq 0$ and $f(0) = 0$ is not continuous at 0, since $\lim_n p^n = 0$ but $1/p^N \to \infty$.

We can also define derivatives similarly!

Definition 5.2

Let $U \subset \mathbb{Q}_p$ be an open set. A function $f : U \to \mathbb{Q}_p$ is **differentiable** at $x_0 \in U$ if the limit

$$f'(x_0) := \lim_{h \to 0} \frac{f(x_0 + h) - f(x_0)}{h}$$

exists. If $f'(x)$ exists for every $x \in U$ we say f is differentiable in U.

For example, polynomials are differentiable everywhere (same proof as in $\mathbb{R}$), and the derivative is what you'd expect.

However, we run into trouble attempting to continue along the real path, since analogues of key theorems needed for calculus and analysis in $\mathbb{R}$ are false. We can state a version of the mean value theorem for $\mathbb{Q}_p$, but it's false! Also, there are functions on $\mathbb{Q}_p$ which are not locally constant but have derivative 0 (for example, consider $f : \mathbb{Z}_p \to \mathbb{Q}_p$ defined by $f(\sum_{i=0}^{\infty} a_i p^i) = \sum_{i=0}^{\infty} a_i p^{2i}$).

Since we are missing such key theorems, we can't develop calculus and analysis for differentiable functions like we do in $\mathbb{R}$. But all is not lost.

5.2 A series of fortunate events

We restrict our attention to functions defined by power series. This is pretty natural since many important functions in $\mathbb{R}$ arise from power series, like e^X and $\sin X$.

Given a formal power series, we want to determine where it defines a function, i.e. where it converges.

Theorem 5.3

Let $f(X) = \sum_{n=0}^{\infty} a_n X^n \in \mathbb{Q}_p[[X]]$ and define

$$\rho = \frac{1}{\limsup \sqrt[n]{|a_n|}}.$$

1. If $\rho = 0$, then $f(x)$ converges only when $x = 0$.
2. If $\rho = \infty$, then $f(x)$ converges for every $x \in \mathbb{Q}_p$.
3. If $0 < \rho < \infty$ and $\lim_{n\to\infty} |a_n|\rho^n = 0$, then $f(x)$ converges if and only if $|x| \leq \rho$.
4. If $0 < \rho < \infty$ and $|a_n|\rho^n$ does not converge to 0, then $f(x)$ converges if and only if $|x| < \rho$.
5. Let $D_f = \{x \in \mathbb{Q}_p : f(x) \text{ converges}\}$. The function $f : D_f \to \mathbb{Q}_p$, $x \mapsto f(x)$ is continuous.

Proof: this theorem follows from the fact that a series converges in $\mathbb{Q}_p$ if and only if the terms of the series converge to 0, so $f(x) = \sum_{n=0}^{\infty} a_n x^n$ converges if and only if $\lim_n |a_n||x|^n = 0$. □

So, for example, for $f(X) = \sum p^n X^n$, $\rho = \infty$ so f converges everywhere. For $g(X) = \sum X^n$, $\rho = 1$ and since the coefficients don't converge to 0, the region of convergence for g is $B(0,1) = p\mathbb{Z}_p$.

Given formal power series

$$f(X) = \sum_{n=0} a_n X^n \text{ and } g(X) = \sum_{n=0} b_n X^n$$

we can define their sum and product series as

$$(f+g)(X) := \sum_{n=0} (a_n + b_n) X^n \text{ and } (fg)(X) = \sum_{n=0} \sum_{k=0}^{n} a_k b_{n-k} X^n.$$

You can check that these series behave how we would expect, that is, that if f, g converge at $x \in \mathbb{Q}_p$, then $f+g$ and fg converge at x, and $(f+g)(x) = f(x)+g(x)$ and $fg(x) = f(x)g(x)$.

We will also want to compose functions; can a composition of functions defined by power series be written as a power series, and if so, how? We can solve recursively for what the coefficients of such a series would be, and we call that series their formal composition.

As it turns out, the formal composition is not the composition as a function unless we have some particular conditions.

Theorem 5.4

Let $f(X) = \sum_{n=0} a_n X^n$ and $g(X) = \sum_{n=0} b_n X^n$, and let $h(X)$ be the formal composition $(f \circ g)(X)$. Let $x \in \mathbb{Q}_p$ and suppose that

1. $g(x)$ converges,
2. f converges on the value $g(x)$, and
3. for all n, we have $|b_n x^n| \leq |g(x)|$

Then $h(x)$ also converges, and $f(g(x)) = h(x)$.

This is a result one would hope for in general, but, alarmingly, you can find series f, g and a value $x \in \mathbb{Q}_p$ such that h does converge, but not to f evaluated at $g(x)$ if the above conditions are not satisfied. We omit the proof of the theorem here, but you can find it in Fernando Gouvea's p-adic Numbers (Theorem 5.3.3).

Given a power series and a point α in its region of convergence, we can recenter the power series around α, writing it as a power series in $X - \alpha$. We can then ask where the new series converges.

Theorem 5.5

Let $f(X) = \sum a_n X^n \in \mathbb{Q}_p[[X]]$, and let $\alpha \in D_f$ (so f converges at α). For each $m \geq 0$, define

$$b_m = \sum_{n \geq m} \binom{n}{m} a_n \alpha^{n-m} \text{ and } g(X) = \sum_{m=0}^{\infty} b_m (X - \alpha)^m.$$

1. The series defining b_m converges for all m
2. $D_f = D_g$ (same region of convergence)
3. For any $x \in D_f$, $f(x) = g(x)$.

We omit the proof (see Gouvea Proposition 5.4.2) but note that it's enough to show that f and g have the same radius of convergence, since $\alpha \in D_f \cap D_f$, and p-adic disks "are either concentric or disjoint (like drops of mercury)"–Yves Andrès.

This is a very cool fact, but it does mean that we can't do analytic continuation the same way we do in $\mathbb{C}$.

We now describe some ways of determining when power series are equal, and some properties of their derivatives.

Theorem 5.6

Let $f, g \in \mathbb{Q}_p[[x]]$, and suppose there is a non-stationary (i.e. not eventually constant) sequence $x_m \in \mathbb{Q}_p$ with $\lim x_m = 0$ such that $f(x_m) = g(x_m)$ for all m. Then $f(X) = g(X)$ (i.e. f, g have the same coefficients).

Proof sketch: this is the same proof as in $\mathbb{R}$. We look at the formal power series of the difference $f - g$, noting that for a power series h, $h(x_m)$ converges to the constant term of h as $x_m \to 0$. □

Theorem 5.7

Let $f(X) = \sum a_n X^n \in \mathbb{Q}_p[[X]]$ and let f' be the formal derivative of f. Let $x \in \mathbb{Q}_p$. If $x \in D_f$, then $x \in D_{f'}$, and

$$f'(x) = \lim_{h \to 0} \frac{f(x_0 + h) - f(x_0)}{h}$$

Proof: First, we note that for $x \neq 0$,

$$|na_n x^{n-1}| \leq |a_n x^{n-1}| = \frac{1}{|x|}|a_n x^n| \to 0$$

and so $f'(x)$ converges (series in $\mathbb{Q}_p$). Next, let $r \in \mathbb{Q}$ such that $D_f = B_{cl}(0,r)$. Suppose $x \neq 0$ and suppose $|h| < |x| \leq r$. Then

$$\frac{f(x+h) - f(x)}{h} = \sum_{n=1}^{\infty} \sum_{m=1}^{n} a_n \binom{n}{m} x^{n-m} h^{m-1}.$$

Then

$$|a_n \binom{n}{m} x^{n-m} h^{m-1}| \leq |a_n| r^{n-1}$$

where the right quantity converges to 0 and does not depend on h, so we can set $h = 0$ to conclude

$$f'(x) = \sum_{n=1}^{\infty} n a_n x^{n-1}.$$

□

Now we can see a compelling reason to focus on power series: we do not have the disturbing phenomenon of non-locally constant functions with derivative 0.

Theorem 5.8

Suppose $f(X), g(X) \in \mathbb{Q}_p[[x]]$, and that f, g both converge for $|x| < \rho$. If $f'(x) = g'(x)$ for all $|x| < \rho$, then there exists a constant $C \in \mathbb{Q}_p$ such that $f(X) = g(X) + C$.

Proof: from Theorem 5.6 and Theorem 5.7, f' and g' have the same coefficients, so f and g have the same coefficients aside from possibly the constant term. □

5.3 Rooting around (because pigs root around)

We'll now explore the zeros of functions coming from power series. There are a lot of wonderful results!

Theorem 5.9

$\mathbb{Z}_p$ is compact.

Proof: $\mathbb{Z}_p$ is a closed subset of $\mathbb{Q}_p$, so it is complete. And for any $\epsilon > 0$, one can find $N \in \mathbb{N}$ such that $p^{-N} < \epsilon$, and

$$\mathbb{Z}_p = \bigcup_{i=0}^{p^N - 1} i + p^N \mathbb{Z}_p$$

is a covering of $\mathbb{Z}_p$ by finitely many balls of radius less than ϵ. So $\mathbb{Z}_p$ is complete and totally bounded, so it is compact.

Theorem 5.10 Strassman's Theorem

Let $f(X) = \sum_{n=0}^{\infty} a_n X^n$ be a nonzero element of $\mathbb{Q}_p[[X]]$. Suppose that $\lim_{n\to\infty} a_n = 0$. Let N be the integer such that

$$|a_N| = \max_n |a_n| \text{ and } |a_n| < |a_N| \text{ for } n > N.$$

Then $f : \mathbb{Z}_p \to \mathbb{Q}_p$ defined by $x \mapsto f(x)$ has at most N zeros. Also, if $\{\alpha_1, ..., \alpha_m\}$ are the zeros of f, then $g \in \mathbb{Q}_p[[X]]$ such that

$$f(X) = (X - \alpha_1)\cdots(X - \alpha_m)g(X)$$

such that g converges on $\mathbb{Z}_p$ and has no zeros in $\mathbb{Z}_p$.

Proof sketch: induct on N and rearrange series to factor out $X - \alpha$ for roots α.

Next we want to consider roots that aren't even necessarily in $\mathbb{Q}_p$. That's right, we want to look in an algebraically closed field. We could take the algebraic closure of $\mathbb{Q}_p$, but it turns out that that's not complete, so we complete that, and thankfully the result is algebraically closed (phew!) We will take the preceding statement as a black box, calling the resulting field $\mathbb{C}_p$. This is summarized in the following theorem:

Theorem 5.11 Complex numbers but make it p-adic

There exists a field $\mathbb{C}_p$ and a valuation function $v_p(\cdot)$ on $\mathbb{C}_p$ (and hence a non-archemidean absolute value $|\cdot| = p^{-v_p(\cdot)}$) on $\mathbb{C}_p$ such that

1. $\mathbb{C}_p$ contains $\overline{\mathbb{Q}_p}$, and the restriction of $|\cdot|$ to $\mathbb{Q}_p$ coincides with the p-adic absolute value
2. $\mathbb{C}_p$ is complete with respect to $|\cdot|$
3. $\mathbb{C}_p$ is algebraically closed
4. $\overline{\mathbb{Q}_p}$ is dense in $\mathbb{C}_p$
5. $\{v_p(x) : x \in \mathbb{C}_p\} = \mathbb{Q}$. In particular, if $x \in \overline{\mathbb{Q}_p}$ has minimal polynomial of degree d, then $v_p(x) \in \frac{1}{d}\mathbb{Z}$.

Now that we are assured that there is a nice field in which we can find all our roots, we explore this bucolic idyll with the following tool:

Definition 5.12

Let $f = a_0 + a_1X + a_2X^2 + \cdots + a_nX^n$ be a polynomial in $K[X]$. Then the **Newton polygon** of f, denoted $NP_p(f)$, is the lower convex hull in $\mathbb{R}^2$ of the points $\{(i, v_p(a_i)) : i = 0, 1, ..., n \text{ and } a_i \neq 0\}$.

One can think of the lower convex hull as being formed by the following procedure: hammer a nail into the plane at each point $(i, v_p(a_i))$, let a rope hang below all the nails, and then pull the rope straight up above the points $(0, v_p(a_0))$ and $(n, v_p(a_n))$ until it is taut.

We illustrate with an example:

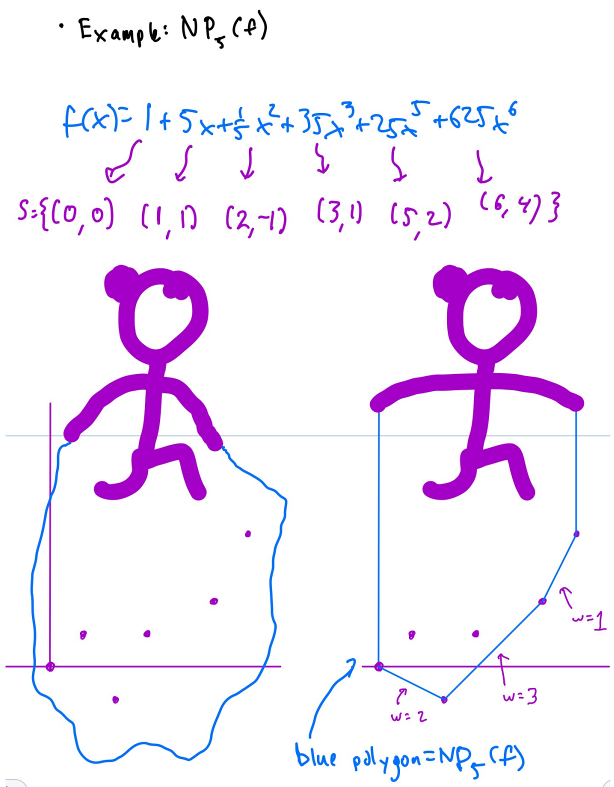

The boundary edges of the Newton polygon of f convey a lot of information about its roots! Define the width of a segment to be its length along the x dimension.

Theorem 5.13

Let K be either $\mathbb{C}_p$ or a finite extension of $\mathbb{Q}_p$. Let $f(X) = a_0 + a_1X + a_2X^2 + ... + a_nX^n \in K[X]$. Let $m_1, ..., m_r$ be the slopes of the boundary edges of $NP_p(f)$, with corresponding widths $w_1, ..., w_r$. Then for each $k : 1 \leq k \leq r$, $f(X)$ has exactly w_k roots (in $\mathbb{C}_p$, counting multiplicities) of absolute value p^{m_k} (that is, of valuation $-m_k$).

Proof: We omit the proof of the number of roots with a given valuation, but we prove that, given a root $\alpha \in \mathbb{C}_p$ with $f(\alpha) = 0$, then $-v_p(\alpha)$ is a slope of a boundary edge of $NP_p(f)$.

Let S denote the set $\{i, v_p(a_i) : 0 \leq i \leq n, a_i \neq 0\}$, whose lower convex hull is $NP_p(f)$. We have:

$$\infty = v_p(0) = v_p(f(\alpha)) = v_p(\sum_{i=0}^{n} a_i\alpha^i) \geq \min_i\{v_p(a_i(\alpha^i)\}$$
$$= \min_i\{v_p(\alpha) \cdot i + v_p(a_i)\} = \min\{v_p(\alpha) \cdot x + y : (x, y) \in S\}$$

If the minimum were uniquely attained, then the inequality would be an equality,which is a contradiction. Hence there must be some $i \neq j$ such that $v_p(\alpha) \cdot i + v_p(a_i) = v_p(\alpha) \cdot j + v_p(a_j)$. Thus, the points $(i, v_p(i))$ and $(j, v_p(j))$ minimize the linear function $v_p(\alpha) \cdot x + y$ along the set S.

Note in general, given a set S of points whose lower convex hull is H, any linear function $l(x, y) = mx + y$ attains its minimum on H at an extremal point, or extremal edge. Thus its minimum on S equals its minimum on the entire convex hull, and is attained at an extremal point or a set of points lying along an extremal edge. One can see this intuitively by varying the line $l(x, y) = c$ for different values of c and noting that, if the line intersects H at some interior point then c can be decreased with the line $l(x, y) = c$ still intersecting H.

In our case, we are minimizing the linear function $l(x, y) = v_p(\alpha) \cdot x + y$ over our set S. As it is minimized at the two points $(i, v_p(i))$ and $(j, v_p(j))$, the edge between these two points is an extremal edge of $NP_p(f)$, whose slope is $-v_p(\alpha)$, the slope of the line $l(x, y) = c$. □

One corollary is Eisenstein's classic criterion for irreducibility. Eisenstein's criterion states that, given a monic polynomial $f(x) = a_0 + a_1X + \cdots + a_{n-1}X^{n-1} + X^n \in \mathbb{Z}[X]$, where $n > 1$, such that p divides a_i for every i but p^2 does not divide a_0, then f is irreducible over $\mathbb{Q}$. To see this using Newton polygons, note that $NP_p(f)$ will have a boundary edge from $(0, 1)$ to $(n, 0)$, whose slope is $-\frac{1}{n}$, so all roots of f have valuation $\frac{1}{n}$.

Now let α be a root of f. If its minimal polynomial over $\mathbb{Q}$ has degree d, then $v_p(\alpha) \in \frac{1}{d}\mathbb{Z}$. But $\frac{1}{n} \notin \frac{1}{d}\mathbb{Z}$, so $d = n$. Thus f is irreducible.

5.4 Connecting the dots (another way)

We will now step back and talk about how to construct p-adic functions via *interpolation*. We will be interested in functions that are uniformly continuous. Recall:

Definition 5.14

Given a field K with absolute value, and a set $S \subset K$, a function $f : S \to K$ is *uniformly continuous* if for every $\epsilon > 0$ there exists $\delta > 0$ such that for all $x, y \in S$,

$$|x - y| < \delta \text{ implies } |f(x) - f(y)| < \epsilon$$

Importantly, the same δ works for a given ϵ, regardless of the choice of x, y. The following Theorem explains the importance of uniform continuity.

Theorem 5.15

Let S be a dense subset of $\mathbb{Z}_p$ and $f : S \to \mathbb{Q}_p$ be a function. Then there exists a continuous function $\tilde{f} : \mathbb{Z}_p \to \mathbb{Q}_p$ such that $\tilde{f}(s) = f(x)$ for all $x \in S$ if and only if f is bounded and uniformly continuous. If such an extension $\tilde{f}$ exists, then it is unique.

Proof: Uniqueness of the extension follows from S being dense in $\mathbb{Z}_p$. Now suppose that a continuous extension $\tilde{f}$ exists. Then it is bounded and uniformly continuous since $\mathbb{Z}_p$ is compact.

Conversely, suppose f is bounded and uniformly continuous. Let $x \in \mathbb{Z}_p$. Since S is dense in $\mathbb{Z}_p$, we can write $x = \lim x_n$ for $x_n \in S$. Since f is bounded and uniformly continuous, you can show that the sequence $f(x_n)$ is Cauchy, hence converges to a limit $\tilde{f} := \lim f(x_n)$. $\square$

For example, we can take $S = \mathbb{Z}$ or even $S = \mathbb{N}$.

Note that in the p-adic setting, we can rephrase uniform continuity as follows.
A function f is uniformly continuous if for all $m \in \mathbb{N}$ there exists some $N \in \mathbb{N}$ such that if

$$a \equiv b \pmod{p^N}$$

then

$$f(a) \equiv f(b) \pmod{p^m}$$