CLASS NOTES FOR APRIL 14, 2000

Announcement: Section 1.2, Questions 3,5 have been deferred from Assignment 1 to Assignment 2. Section 1.4, Question 5 has been dropped entirely.

1. Review of Wednesday class

Let (X, d) be a metric space. A set $E \subset X$ is said to be *dense* if $\overline{E} = X$.

Theorem 1.1. *(Baire category theorem) If X is complete and $V_1, V_2, \ldots$ are open dense sets, then $\bigcap_{n=1}^{\infty} V_n$ is dense.*

A set $E \subset X$ is said to be *nowhere dense* if $\overline{E}$ has no interior. Examples: $\mathbf{Z}$ is nowhere dense in $\mathbf{R}$. The set

$$\{1, 1/2, 1/4, 1/8, 1/16, \ldots\}$$

is also nowhere dense in $\mathbf{R}$. The set $\mathbf{Q} \cap [0, 1]$ is not nowhere dense in $\mathbf{R}$ (it's closure has interior $(0, 1)$).

2. Conclusion of Wednesday class

Another example of nowhere dense sets: any line or circle in $\mathbf{R}^2$ is nowhere dense.

Nowhere dense sets are in some sense the opposite of dense sets. A precise connection is:

Exercise 2.1. *A set E is nowhere dense if and only if $\overline{E}^c$ is open and dense.*

Proof (Optional) We first prove the "only if" part. Suppose E is nowhere dense. Then $\overline{E}$ has no interior. Since closures are closed, $\overline{E}$ is closed, which means that $\overline{E}^c$ are open. Now we have to show that $\overline{E}^c$ is dense.

Let B be an open ball. We have to show that $\overline{E}^c$ intersects the ball B. Suppose for contradiction that $\overline{E}^c$ did not intersect B. Then B would be contained in $\overline{E}$, contradicting the fact that $\overline{E}$ has no interior.

Now suppose that $\overline{E}^c$ is open and dense. Then every ball must contain at least one point in $\overline{E}^c$, which means that no ball can be completely contained inside $\overline{E}$. This means that $\overline{E}$ has no interior, which means that E is nowhere dense. ■

Example: the set $\mathbf{R}$ is nowhere dense in $\mathbf{R}^2$. $\mathbf{R}$ is closed (the only points adherent to $\mathbf{R}$ are the points that are already on $\mathbf{R}$, so the above Exercise implies that $\mathbf{R}^2 \backslash \mathbf{R}$ is open and dense in $\mathbf{R}^2$.

We can use the Baire category theorem to say something about nowhere dense sets:

Theorem 2.2. *Let X be a complete metric space, and let $E_1, E_2, \ldots$ be a sequence of nowhere dense sets in X. Then*

$$X \neq \bigcup_{n=1}^{\infty} E_n. \tag{1}$$

In other words, a complete metric space cannot be covered by a countable number of nowhere dense sets.

Proof The sets E_n are nowhere dense. By the above exercise, this means that the sets $\overline{E_n}^c$ are open and dense. By the Baire category theorem, this implies that the set

$$\bigcap_{n=1}^{\infty} \overline{E_n}^c$$

is dense. In particular, it is non-empty (since the empty set is not dense). Thus we can find a point $x \in X$ such that $x \in \overline{E_n}^c$ for all n. This implies that $x \notin E_n$ for any n. Thus we have found a point in X which is not in $\bigcup_{n=1}^{\infty} E_n$, which proves (1). ∎

This theorem is quite powerful. As just one example, it shows that if you take a totally arbitrary sequence $C_1, C_2, \ldots$ of circles in $\mathbf{R}^2$, then the circles do not completely fill out $\mathbf{R}^2$, i.e. there is always at least one point in $\mathbf{R}^2$ which is not covered by any one of the circles. This is a typical example of how the theorem is used: to show that given any sequence of nowhere dense sets, that there is always at least one point which avoids all of them. We'll see some more applications in a couple weeks.

(Optional): Of course, if you take the set of *all* circles in $\mathbf{R}^2$, of arbitrary center and radius, then these circles will definitely fill out $\mathbf{R}^2$ (every point in $\mathbf{R}^2$ is certainly contained in at least one circle). However, this does not contradict Theorem 2.2 because the set of all circles in $\mathbf{R}^2$ is *uncountable* - it cannot be organized as a sequence.

By the way, René-Louis Baire (1874-1932) divided sets into two categories. He defined a set of the *first category* to be any set which could be written as the countable union of nowhere dense sets. For instance, any set which is made up of a countable number of circles is of the first category. Anything which is not of the first category was placed in the *second category*. Baire's theorem then says that every complete space is in the second category. (Nowadays, the concept of first and second category are not used very much, but this is the historical reason for the name "Baire Category Theorem".

3. Compact sets

We now turn to a trickier notion, that of compactness.

Our motivation shall be the Extremal Value Theorem from lower-division math:

Theorem 3.1. *Let $f : [a, b] \to \mathbf{R}$ be a continuous function on an interval $[a, b]$. Then there are points x_{max}, x_{min} in $[a, b]$ such that $f(x_{max})$ is the maximum of f and $f(x_{min})$ is the minimum of f on $[a, b]$.*

This theorem only works for closed intervals; for open intervals it is false (e.g. $f(x) = 1/x$ on $(0, 1)$) and it is also false for unbounded sets (e.g. $f(x) = x$ on $[0, \infty)$). The topological explanation for this is that closed intervals have a special property, called *compactness*, which makes the Extreme Value Theorem (and many, many other results) work.

Compactness is quite a tricky and unintuitive notion to pin down. Here is the formal definitions:

Definition 3.2. Let (X, d) be a metric space. An *open cover* of X is a (possibly infinite) collection $\{V_\alpha\}_{\alpha \in A}$ of open sets in X such that

$$X \subseteq \bigcup_{\alpha \in A} V_\alpha.$$

We say that a metric space is *compact* if every open cover has a finite subcover.

For instance, consider $\mathbf{R}$ with the usual metric. This space has many open covers; for instance, one can use the balls $B(x, 1)$, as x ranges over all real numbers, to cover $\mathbf{R}$. Some of these balls are redundant (e.g. because $B(1.9, 1)$ and $B(2.1, 1)$ are in the open cover, the ball $B(2, 1)$ is unnecessary and could be removed), but no matter how many redundant balls you remove, you still need an infinite number of balls in this collection to cover $\mathbf{R}$, because one could never cover $\mathbf{R}$ with a finite number of balls. Thus, $\mathbf{R}$ is not compact.

On the other hand, let's consider the set $\{1, 2, 3\}$. This set is compact for the following reason. Let $\{V_\alpha\}_{\alpha \in A}$ be an open cover of $\{1, 2, 3\}$. Since 1 is an element of $\{1, 2, 3\}$, 1 must be contained in one of the V_α; let's say $1 \in V_{\alpha_1}$. Similarly we can find an $\alpha_2 \in A$ such that $2 \in V_{\alpha_2}$, and an $\alpha_3 \in A$ such that $3 \in V_{\alpha_3}$. The sub-collection $\{V_{\alpha_1}, V_{\alpha_2}, V_{\alpha_3}\}$ is now a finite sub-cover of the original open cover which covers $\{1, 2, 3\}$. Since every open cover has a finite sub-cover, $\{1, 2, 3\}$ is compact.

As you can see already, this definition is very unwieldy. Fortunately, there are other ways to characterize compactness for metric spaces which are easier to deal with. To do this we first need some more definitions.

Definition 3.3. A space X is said to be *bounded* if there is some ball $B(x, r)$ which contains X. A space is said to be *totally bounded* if, for every $\varepsilon > 0$, one can cover X by a finite number of open balls of radius ε.

For instance, **R** is not bounded (it can't be enclosed inside a ball) and it is not totally bounded either. The set $[0,1]$ is both bounded and totally bounded (for any $\varepsilon > 0$ one can cover $[0,1]$ by a finite number of ε-balls).

Now consider the integers **Z** with the discrete (or "teleport") metric

$$d(x,y) = \begin{cases} 1 & \text{if } x \neq y \\ 0 & \text{if } x = y \end{cases}$$

In other words, all the integers are distance 1 apart from each other under this metric. In this case, **Z** is actually bounded (e.g. it is contained in $B(0,2)$) but it is not totally bounded (it is not possible to cover **Z** by balls of radius 1/2).

Exercise 3.4. *Show that every totally bounded set is bounded.*

Proof (Optional) Let X be a totally bounded set. Then it can be covered by a finite number of balls of radius 1 (for instance). Let's call these balls $B(x_1,1),\ldots,B(x_n,1)$. Let R denote the radius

$$R = 1 + \max_{1 \leq i \leq n} d(x_i, x_1).$$

From the triangle inequality we see that

$$B(x_i, 1) \subset B(x_1, R)$$

for all $1 \leq i \leq n$. Since the $B(x_i,1)$ cover X, we thus see that

$$X \subset B(x_1, R)$$

and hence X is bounded. ■

Theorem 3.5. *Let (X,d) be a metric space The following statements are equivalent:*

- *(i) X is compact.*
- *(ii) Every sequence in X has at least subsequence which converges in X.*
- *(iii) X is complete and totally bounded.*

The properties (ii) and (iii) are a bit more intuitive to work with than (i). For instance, from (ii) we can see why $(0,1)$ is not compact, and it is plausible that $[0,1]$ is compact.

The proof of this theorem is somewhat complicated. There are three steps: showing that (i) implies (ii); showing that (ii) implies (iii); and showing that (iii) implies (i). Of these, the first two are fairly straightforward. The last one is nastier and will be left to the written notes.

4. Proof of $(i) \implies (ii)$.

Let X be a compact set, and suppose that $x_1, x_2, \ldots$ is a sequence in X. We have to find a subsequence of the x_i which converges in X.

There are two cases:

Case 1: There exists a point $x \in X$ such that every open ball centered at x contained infinitely many elements of the sequence $x_1, x_2, \ldots$.

Then we would be done, for the following reason. Since the ball $B(x, 1)$ contains infinitely many elements of the sequence, we can pick an element x_{n_1} in $B(x, 1)$. Then, since $B(x, 1/2)$ contains infinitely many elements of the sequence, it must contain at least one element x_{n_2} with $n_2 > n_1$. Repeating this, we can find an element $x_{n_3} \in B(x, 1/3)$ with $n_3 > n_2$. More generally, we can choose a subsequence $x_{n_1}, x_{n_2}, \ldots$ such that $n_{j+1} > n_j$ and $x_{n_j} \in B(x, 1/j)$ for all $j = 1, 2, 3, \ldots$. Then it is clear that this sequence is convergent to x, and we are done.

Case 2: For every point $x \in X$ there exists an open ball centered at x which contains at most finitely many elements of the sequence $x_1, x_2, \ldots$.

For every $x \in X$, let B_x be an open ball centered at x which contains at most finitely many elements of the sequence x_n. Then the collection of all these balls $\{B_x\}_{x \in X}$ is an open cover of X, because every $x \in X$ is contained in at least one ball in the collection, namely the ball B_x centered at x. However, this open cover cannot have a finite sub-cover, because every ball only contains a finite number of elements of the sequence, and there are an infinite number of elements of the sequence to cover. This contradicts compactness.

5. Proof of $(ii) \implies (iii)$.

Let X satisfy property (ii), i.e. every sequence in X has a convergent subsequence. We need to show that X is complete and totally bounded.

Let's first show that X is complete. Let $x_1, x_2, \ldots$ be a Cauchy sequence in X. We need to show that this sequence converges.

By (ii), we can find a subsequence $x_{n_1}, x_{n_2}, x_{n_3}, \ldots$ of $x_1, x_2, \ldots$ which converges to some point in X, say x. This is almost what we want, but we need to show that the entire sequence converges to x, not just a subsequence. I'll leave this as an exercise:

Exercise 5.1. *If x_n is a Cauchy sequence, and a sub-sequence x_{n_j} of x_n converges to x, then x_n itself converges to x.*

Proof (Optional) We need to show that for every $\varepsilon > 0$ there exists $N > 0$ such that $d(x_n, x) < \varepsilon$ for all $n > N$. Well, because x_{n_j} converges to x, we can find a $J > 0$ such that

$$d(x_{n_j}, x) < \varepsilon/2$$

for all $j > J$. Also, since x_n is Cauchy, we can find an $N > 0$ such that

$$d(x_n, x_m) < \varepsilon/2$$

for all $n, m > N$.

Since n_j is an increasing sequence, there must exist a $j > J$ such that $n_j > N$. Then for any $n > N$ we have

$$d(x_n, x) \leq d(x_n, x_{n_j}) + d(x_{n_j}, x) < \varepsilon/2 + \varepsilon/2 = \varepsilon$$

which is what we want. ∎

Now we show that X is totally bounded. Pick any $\varepsilon > 0$. We want to show that X can be covered by a finite number of ε-balls.

We will do this in the most naive way possible - simply by drawing ε-balls one by one until we exhaust the space X.

More precisely, pick any $x_1 \in X$. If $B(x_1, \varepsilon)$ already covers X, then we are done. Otherwise, let x_2 be a point not in $B(x_1, \varepsilon)$. If $B(x_1, \varepsilon)$ and $B(x_2, \varepsilon)$ cover X, then we are done. Otherwise, let x_3 be a point which is not contained in either $B(x_1, \varepsilon)$ or $B(x_2, \varepsilon)$. We repeat this process indefinitely.

If this process halts after a finite number of steps, we are done. Now suppose the process never halts, so we find an infinite sequence of points $x_1, x_2, \ldots$ such that each x_n is not contained in any preceding ball $B(x_1, \varepsilon), \ldots, B(x_{n-1}, \varepsilon)$. In particular, we have

$$d(x_n, x_m) \geq \varepsilon \text{ for all } n \neq m.$$

But this implies that this sequence can never have a Cauchy subsequence (because Cauchy sequences need to eventually get within ε of each other). Since all convergent sequences are Cauchy, this implies that the sequence $x_1, x_2, \ldots$ does not have any convergent subsequences, contradicting (ii). Hence the process described above must terminate in a finite number of steps, and we are done.

6. Proof of $(iii) \implies (i)$. (Optional)

We prove the following standard fact:

Theorem 6.1. *Every complete totally bounded metric space is compact.*

Proof Let (X, d) be a complete totally bounded metric space. We need to show that X is compact.

Suppose that $\{V_\alpha\}_{\alpha \in A}$ is an open cover of X. Our objective is to show that one only needs a finite number of these V_α in order to cover X.

This will be a proof by contradiction. Let's assume that the open cover $\{V_\alpha\}_{\alpha \in A}$ has no finite subcover. In other words, any finite collection of V_α will fail to cover X. We will obtain a contradiction from this.

We know that X is totally bounded. In particular, we can cover X by a finite number of balls of radius 1.

Suppose every one of these balls of radius 1 could be covered by a finite number of V_α. Then the whole space X could also be covered by a finite number of V_α, which would be a contradiction. Thus there must be at least one ball, let's call it $B(x_0, 1)$, which cannot be covered by a finite number of V_α.

Since X is totally bounded, and $B(x_0, 1)$ is a subset of X, $B(x_0, 1)$ is totally bounded. In particular, we can cover $B(x_0, 1)$ by a finite number of balls of radius $1/2$. We may assume that the centers of these balls are at a distance at most $1 + \frac{1}{2}$ from x_0, since otherwise the ball would not intersect $B(x_0, 1)$ (why?), and so would play no role in the cover.

Suppose every one of these balls of radius $1/2$ could be covered by a finite number of V_α. Then $B(x_0, 1)$ could also be covered by a finite number of V_α, which would be a contradiction. So there must be at least one ball, let's call it $B(x_1, 1/2)$, which cannot be covered by a finite number of V_α. Also, from the previous discussion we have

$$d(x_0, x_1) \leq 1 + \frac{1}{2}.$$

Since X is totally bounded, and $B(x_1, 1/2)$ is a subset of X, $B(x_1, 1/2)$ is totally bounded. In particular, we can cover $B(x_1, 1/2)$ by a finite number of balls of radius $1/4$. We may assume that the centers of these balls are at a distance at most $\frac{1}{2} + \frac{1}{4}$ from x_1, by similar arguments to before.

By similar arguments to before, we can find a ball, let's call it $B(x_2, 1/4)$, which cannot be covered by a finite number of V_α, and satisfies

$$d(x_1, x_2) \leq \frac{1}{2} + \frac{1}{4}.$$

Continuing in this vein, we can find a sequence $x_0, x_1, x_2, \ldots$ of points in X such that each ball $B(x_n, 2^{-n})$ cannot be covered by a finite number of V_α, and such that

$$d(x_n, x_{n+1}) \leq 2^{-n} + 2^{-n-1}$$

for all $n = 0, 1, 2, \ldots$. In particular, this means that the x_n are a fast Cauchy sequence:

$$\sum_{n=0}^{\infty} d(x_n, x_{n+1}) < \infty.$$

Thus the x_n are a Cauchy sequence. Since X is complete, this means that x_n is a convergent sequence, and has some limit x.

Now, we are given that $\{V_\alpha\}_{\alpha \in A}$ covers X. Since $x \in X$, this means that there must exist at least one $\alpha \in A$ such that V_α contains x. Pick one such α. Since V_α is open, this means that x is in the interior of V_α. This implies that there exists some radius $r > 0$ such that $B(x, r)$ is contained in V_α.

The sequence x_n converges to x, and the sequence 2^{-n} converges to 0. Since $r > 0$, this implies that we can find an integer n such that

$$d(x_n, x) < r/2$$

and

$$2^{-n} < r/2.$$

In particular, this implies that

$$B(x_n, 2^{-n}) \subset B(x, r).$$

But since $B(x, r)$ is contained in V_α, we have

$$B(x_n, 2^{-n}) \subset V_\alpha.$$

This implies that $B(x_n, 2^{-n})$ can be covered by a finite number (in fact, just one) of the V_α, which is a contradiction. The proof is complete. ■