Math 55a: Honors Advanced Calculus and Linear Algebra

Metric topology V: Compactness

So far in our development of metric topology we have been mostly formalizing and generalizing familiar notions. Compactness is more subtle; it is not even easy to give an *a priori* motivation for the importance of compactness — I'll have to ask you to take my word for the claim that it is a key concept in the rigorous development of analysis (and indeed much else in modern mathematics). We shall see that studying compactness in the general context of metric spaces will even clarify some arguments that arise in the familiar contexts of functions of one or several real veriables.

The definition of compactness (see Rudin, 2.31–32 on p.36) requires the notion of an *open cover* of a set E in a metric space X. By definition, a *cover* of E is a collection $\{G_\alpha : \alpha \in I\}$ of subsets of X whose union $\cup_{\alpha\in I} G_\alpha$ contains E; an *open cover* is a cover $\{G_\alpha\}$ with each G_α open. A *subcover* is a subcollection $\{G_\alpha : \alpha \in I'\}$ (some $I' \subseteq I$) such that $\cup_{\alpha\in I'} G_\alpha$ still contains E. The subcover is *finite* if I' is finite; i.e., a finite subcover is a finite list $G_{\alpha_1}, G_{\alpha_2}, \ldots, G_{\alpha_n}$, with each $\alpha_i \in I$, such that $\cup_{i=1}^n G_{\alpha_i} \supseteq E$.

Definition. A subset K of a metric space X is said to be *compact* if every open cover of K has a finite subcover.

For instance, every finite set is compact; if K has the discrete metric, then K is compact if and only if it is finite (why?). Indeed one way to come close to visualizing compactness is that a compact set behaves like set of finite but unspecified size.

Here is an easy consequence of the definition which illustrates the idea and use of open covers and finite subcovers:

Theorem. [cf. Rudin, p.37, 2.34] *If a subset K of a metric space X is compact, then K is bounded and closed.*

Proof: We may assume X is nonempty (else $K = \emptyset$, and we already observed that finite sets are compact). Fix $p \in X$, and let $G_m = B_m(r)$ $(m = 1, 2, 3, \ldots)$. Then $\{G_m\}$ is an open cover of K (indeed it is an open cover of X). Let $\{G_{m_i}\}$ be a finite subcover. Then K is contained in G_M where $M = \max_i m_i$. Thus K is bounded.

Now let p be a point of X not in K. Let G_m be the complement in X of $\overline{B}_{1/m}(p)$. Then $\cup_m G_m = X \setminus \{p\} \supseteq E$, so again $\{G_m\}$ is an open cover of K. Let $\{G_{m_i}\}$ be a finite subcover. Then K is disjoint from the $(1/M)$-neighborhood of p, where again $M := \max_i m_i$. Thus K is closed. □

This used only countable open covers; here's a typical use of open covers of unrestricted cardinality. Fix $\epsilon > 0$, let $I = K$, and let $G_p = B_\epsilon(p)$ for each $p \in K$. Then $\{G_p : p \in K\}$ is an open cover of K. So, if K is compact, we have

a finite subcover — i.e., a finite collection of ϵ-balls that cover K. Equivalently, we have a finite subset $\{p_1, \ldots, p_n\}$ of K such that each point of K is at distance $< \epsilon$ from some p_i. Such a set $\{p_1, \ldots, p_n\}$ is called an *ϵ-net*; a subset of a metric space which has an ϵ-net for each $\epsilon > 0$ is said to be *totally bounded*. So we have shown:

Theorem. *A compact set is totally bounded.*

We shall drop the other shoe next week when we cover completeness and its connection with compactness.

Meanwhile we observe that compactness is a topological property (so you now know what it means for a subset of an arbitrary topological space to be compact). Moreover, like boundedness and unlike the notions of being open or closed, compactness is "intrinsic": it does not depend on the ambient space. That is, we have:

Theorem. [Rudin, p.37, 2.33] *Let X be a metric (or even a topological) space, and $K \subseteq Y \subseteq X$. Then K is compact as a subset of X if and only if it is compact as a subset of Y.*

Proof: Compactness relative to Y is obtained by replacing "open set" by "relatively open subset of Y" — which we have seen already is the same as "$G \cap Y$ for some open subset G of X". (In the general topological setting, that's what we adopted as the *definition* of an open subset of Y.) Suppose K is compact, and $\{V_\alpha\}$ is a cover of K by relatively open subsets of Y. Write each V_α as $G_\alpha \cap Y$ with each G_α open in X. Then $\{G_\alpha\}$ is an open cover of K, so has an finite subcover $\{G_{\alpha_i}\}$; $\{V_{\alpha_i}\}$ is then the desired finite subcover of $\{V_\alpha\}$. Thus, as claimed K is compact as a subset of Y. Conversely, if K is compact as a subset of Y, and $\{G_\alpha\}$ is a cover of K by open subsets of X, then $\{G_\alpha \cap Y\}$ is a cover of K by relatively open subsets of Y, a finite subcover of which yields a finite subcover of $\{G_\alpha\}$. $\square$

So, we may sensibly speak of K as being compact without reference to the ambient space; in particular this is always equivalent to K being compact as a metric space in its own right (as a subset of itself; take $Y = K$ in the above theorem).

Given one compact space, we may obtain many others via the next two results:

Theorem. [Rudin, p.37, 2.35] *A closed subset of a compact set is compact.*

Proof: Let K be a compact metric space and F a closed subset. Then its complement F^c is open. Thus if $\{V_\alpha\}$ is an open cover of F we obtain an open cover Ω of K by adjoining F^c. Since K is compact, Ω has a finite subcover; removing F^c if necessary, we obtain a finite subcollection of $\{V_\alpha\}$ which covers F. This is the desired open cover. $\square$

[Rudin presents this result with K a compact subset of an arbitrary metric space X. But we have seen that such K is necessarily closed, so $F \subseteq K$ is closed relative to K if and only if it already closed in K. Thus Rudin's and our formulations are equivalent.]

Theorem. [Rudin, p.89, 4.14] *The continuous image compact set is compact. That is, if $f : K \to Y$ is a continuous function, and K is compact, then so is $f(K)$.*

Proof: Let $\{V_\alpha\}$ be an open cover of Y. Then $\{f^{-1}(V_\alpha)\}$ is a cover of K, which is open because f is continuous. Thus it has a finite subcover $\{f^{-1}(V_{\alpha_i})\}$. Then $\{V_{\alpha_i}\}$ is a finite subcover of $\{V_\alpha\}$.

Sequential compactness. [See Rudin, pages 51–52.] Compactness can be formulated in several ways which are equivalent but not obviously so; this partly accounts for the concept's power. We next develop one of these equivalents, in terms of sequences and convergence.

A *subsequence* of a sequence $\{p_n\}$ is a sequence of the form $\{p_{n_i}\}$ for any choice of n_i $(i = 1, 2, 3, \ldots)$ such that $n_1 < n_2 < n_3 < \ldots$. If p_n and p are in a metric space X then $\{p_n\}$ has a subsequence converging to p if and only if for each $\epsilon > 0$ there are infinitely many n such that $d(p_n, p) < \epsilon$ (why?). A metric space K is said to be *sequentially compact* if every sequence in K has a convergent subsequence. (So $K \subseteq X$ is sequentially compact if and only if every sequence in K has a subsequence convergent in K.) Then:

Theorem. [cf. Rudin, p.51, 3.6, and the exercises on p.45] *A metric space is compact if and only if it is sequentially compact.*

Proof: ($\Rightarrow$) Let $\{p_n\}$ be a sequence in a compact space K. We shall prove that it has a convergent subsequence by contradiction. If $\{p_n\}$ had no such subsequence, then for any $p \in K$ there would exist $\epsilon_p > 0$ such that $B_{\epsilon_p}(p) \ni p_n$ for only finitely many n. But then $\{B_{\epsilon_p}(p) : p \in K\}$ would be an open cover of K all of whose subcovers are infinite.

($\Leftarrow$) This is trickier. We first show that a *countable* open cover $\{G_n\}$ of a sequentially compact space K has a finite subcover. Equivalently, we claim that $K = \cup_{i=1}^{n} G_i$ for some n. Assume not. Let p_n be a point in the complement of $\cup_{i=1}^{n} G_i$, and let p be the limit of a subsequence in $\{p_n\}$. Since $K = \cup_{i=1}^{\infty} G_i$, this p is in some G_N. Since G_N is open, some neighborhood of p is also contained in G_N. Since p is the limit of a subsequence in $\{p_n\}$, that neighborhood contains p_n for infinitely many n. But $p_n \notin G_N$ once $n > N$. This contradiction proves that $K = \cup_{i=1}^{n} G_i$ for some n, as claimed.

To reduce the general case to this, we shall need:

Proposition. *A sequentially compact space contains a countable dense set.*

[In general a metric space with a countable dense subset is said to be *separable*;

$\mathbf{R}$ is an example of a non-compact separable space, since $\mathbf{Q}$ is a countable dense subset.]

Proof of Proposition: Let K be a sequentially compact space. We may assume that K is infinite, else the statement is trivial. Note first that K is necessarily bounded, else there exists $p, p_n \in K$ with $d(p, p_n) > n$, and $\{p_n\}$ has no convergent subsequence. We inductively construct a dense sequence $\{p_n\}$ in K. The first point p_1 is arbitrary. Having chosen $p_1, \ldots, p_n$, let $\delta_n := \sup_{p \in K} \min_{i \leq n} d(p, p_n)$, and let p_{n+1} be a point such that $d(p_{n+1}, p_i) \geq \delta/2$ for each $i = 1, \ldots, n$. Note that $\delta_n > 0$, thus in particular no two p_n coincide. But $\{p_n\}$ has a convergent subsequence, so we may find for each $\epsilon > 0$ integers m, n with $m < n$ but $d(p_m, p_n) < \epsilon$ (the subsequence eventually gets within $\epsilon/2$ of its limit). Thus $\delta_{n-1} < 2\epsilon$. So, each $p \in K$ is within 2ϵ of p_i for some $i < n$. Since ϵ is an arbitrary positive number, we conclude that $\{p_n\}$ is dense in K as claimed. $\square$

Now let $\{G_\alpha : \alpha \in I\}$ be an open cover of K. Let $\{p_n\}$ be a dense sequence, and consider the family $\mathcal{F}$ of neighborhoods $B_r(p_n)$ with $r \in \mathbf{Q}$ that are contained in some G_α. By a familiar argument, $\mathcal{F}$ is countable. We claim that it is also an open cover. Indeed let $p \in K$ and $G_\alpha \ni p$. Then $G_\alpha \supseteq B_s(p)$ for some $s > 0$. Since $\{p_n\}$ is dense, $d(p, p_n) < s/2$ for some n. Then, for every rational r with $d(p, p_n) < r < s - d(p, p_n)$, we have $p \in B_r(p_n) \subseteq B_s(p) \subseteq G_\alpha$. Thus $p \in B_r(p_n) \in \mathcal{F}$. We can then find a finite subcover of $\mathcal{F}$. Replacing each set in the subcover by a G_α containing it we at last obtain a finite subcover of $\{G_\alpha\}$, and theorem is proved! $\square$

So, sequential compactness is the same as compactness, but we may still write "by sequential compactness" to indicate that we are using the equivalent sequential definition. For instance, that definition yields an easy proof of the next construction of new compact sets from old:

Theorem. *The product of two compact metric spaces is compact.*

Proof: Let K, L be compact spaces and $\{(p_n, q_n)\}$ any sequence in $K \times L$. Extract from $\{p_n\}$ a convergent subsequence $p_{n_i} \to p$. Regard $\{q_{n_i}\}$ as a sequence in L, and extract a convergent subsequence $q_{n_{i_j}} \to q$. Since also $p_{n_{i_j}} \to p$, we conclude that $(p_{n_{i_j}}, q_{n_{i_j}}) \to (p, q)$. Thus $K \times L$ is sequentially compact, and we are done. $\square$ [Apologies for triple $\mathrm{sub}_{\mathrm{sc_{ri_{pts}}}}$!]

After all this, we still do not know what compact sets in $\mathbf{R}$ look like; that, and a third equivalent definition of compactness involving ϵ-nets, will have to await the next handout, concerning Cauchy sequences and completeness.