updated 17/10/19

Lecture 14 : Banach fixed point theorem

We now know that if X is compact and (Y, d_Y) is a complete metric space then $Cts(X, Y)$ is a complete metric space with the d_∞ metric. In today's lecture we examine an important theorem about complete metric spaces, the Banach fixed point theorem, which will be applied next lecture to prove the existence and uniqueness of solutions to (ordinary) differential equations.

Def^n A fixed point of a function $f: X \longrightarrow X$ is $x \in X$ such that $f(x) = x$.

A fixed point problem is the problem of proving the existence and/or uniqueness of a fixed point for a given $f: X \longrightarrow X$. Many problems in mathematics (e.g. root finding, convex optimisation, or solving ODEs) can (through varying degrees of chicanery) be phrased as fixed point problems. Hence, general theorems about fixed points tend to have widespread application.

Example L14-1 (Newton's algorithm) Suppose $g: \mathbb{R} \longrightarrow \mathbb{R}$ is differentiable and we wish to find a solution a of $g(a) = 0$. Newton's method is to iterate, beginning with any $x_0 \in \mathbb{R}$, the formula

$$x_{n+1} := x_n - \frac{g(x_n)}{g'(x_n)}.$$

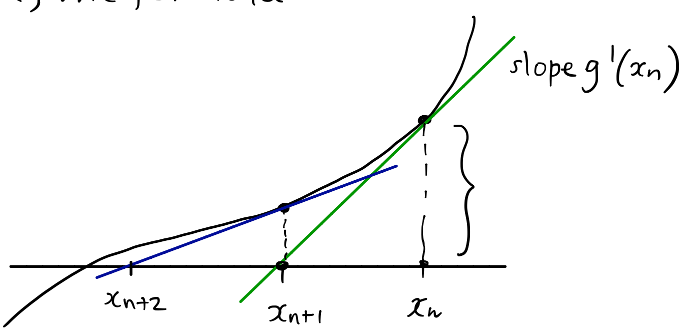

Observe that a fixed point of $f(x) = x - g(x)/g'(x)$ is precisely a root of g where $g'(a) \neq 0$ since

$$a = a - g(a)/g'(a) \iff g(a) = 0$$

<u>Example L14-2</u> Let V be a vector space, $g: V \to V$ a function (perhaps non-linear) and suppose we are given $w \in V$ and wish to solve the equation $g(v) = w$ for v. This is equivalent to finding a fixed point of $f(v) = v + g(v) - w$.

<u>Why bother</u> rephrasing problems as fixed point problems? Well, because finding a fixed point of a function $f: X \to X$ is (at least conceptually) trivial: just iterate f infinitely many times! Choose $x_0 \in X$ and set

$$x_1 = f(x_0),\ x_2 = f(x_1) = f^2(x_0),\ \dots,\ x_{n+1} = f(x_n).$$

In the "limit" (whatever that means) we have $x^* = \lim_{n\to\infty} f^n(x_0)$, and provided f is continuous we find that x^* is a fixed point

$$f(x^*) = f\left(\lim_{n\to\infty} f^n(x_0)\right) = \lim_{n\to\infty} f^{n+1}(x_0) = \lim_{n\to\infty} f^n(x_0) = x^*.$$

This is an informal argument, but it can be made rigorous provided the sequence $(f^n(x_0))_{n=0}^{\infty}$ converges: one way to ensure this is to assume f is a contraction (so the sequence is Cauchy) and that X is complete (so that Cauchy sequence converges). The fact that x^* is a fixed point follows by precisely the above argument once we know it exists.

<u>Def^n</u> Let (X, d) be a metric space. A function $f: X \to X$ is called a <u>contraction</u> if there exists $\lambda \in (0,1)$ such that

$$d(fx, fx') \leq \lambda\, d(x, x') \qquad \forall x, x' \in X.$$

Clearly a contraction is continuous. (We call f a λ-contraction) from his PhD thesis! No pressure.

<u>Theorem L14-1</u> (Banach fixed pt. theorem) Suppose (X,d) is a complete metric space and $f : X \longrightarrow X$ is a contraction. Then f has a unique fixed point. For any $x \in X$ the sequence $(f^n x)_{n=0}^{\infty}$ converges to this fixed point.

<u>Proof</u> Let λ be the contraction factor of f as above. If $f(p) = p$, $f(q) = q$ then

$$d(p,q) = d(fp, fq) \leq \lambda\, d(p,q)$$

which is a contradiction unless $d(p,q) = 0$. Hence a fixed point, if it exists, is unique. It remains to prove existence. Given $x \in X$ set $a_n := f^n(x)$. Note that $d(f^2x, f^2y) \leq \lambda d(fx, fy) \leq \lambda^2 d(x,y)$ and by induction also $d(f^k x, f^k y) \leq \lambda^k d(x,y)$ for $k \geq 1$. We have for $m > n$

$$\begin{aligned}
d(a_m, a_n) &\leq d(a_m, a_{m-1}) + \cdots + d(a_{n+2}, a_{n+1}) + d(a_{n+1}, a_n) \\
&= d(f^m x, f^{m-1} x) + \cdots + d(f^{n+2} x, f^{n+1} x) + d(f^{n+1} x, f^n x) \\
&\leq \lambda^{m-1} d(fx, x) + \cdots + \lambda^{n+1} d(fx, x) + \lambda^n d(fx, x) \\
&= \left[\lambda^{m-1} + \cdots + \lambda^n \right] d(fx, x) \\
&= \lambda^n \Big(\sum_{i=0}^{m-n-1} \lambda^i \Big) d(fx, x) \\
&\leq \lambda^n \Big(\sum_{i=0}^{\infty} \lambda^i \Big) d(fx, x) \\
&= \lambda^n \cdot \frac{1}{1-\lambda} \cdot d(fx, x) \qquad \text{(since } 0 < \lambda < 1\text{)}
\end{aligned}$$

Since $\lambda < 1$ we may make the RHS arbitrarily small by making n sufficiently large, and so it follows that $(a_n)_{n=0}^{\infty}$ is Cauchy. Since X is complete this Cauchy sequence converges, say to $x^* \in X$, that is

$$x^* = \lim_{n\to\infty} a_n = \lim_{n\to\infty} f^n(x).$$

But by Lemma L8-4 we have

$$f(x^*) = \lim_{n\to\infty} f(f^n(x)) = \lim_{n\to\infty} f^{n+1}(x) = x^*$$

so x^* is a fixed point. $\square$

While Example L14-1, L14-2 hint at some of the problems that may be phrased as fixed point problems, it may be still some work to show the function g is a contraction (for example Example L14-2 can be used to prove the Implicit Function Theorem, the hypotheses for which create the circumstances for that particular g to be a contraction). Here is an exercise that walks you through the linear case:

<u>Exercise L14-1</u> Let $A \in M_n(\mathbb{R})$ and let $g: \mathbb{R}^n \longrightarrow \mathbb{R}^n$ be $g(v) = Av$. Define $f: \mathbb{R}^n \longrightarrow \mathbb{R}^n$ by $f(v) = v - Av + w$, where w is a fixed vector. Prove that if there exists $\lambda \in (0,1)$ with

$$\sum_{j=1}^{n} |\delta_{ij} - A_{ij}| \leq \lambda \quad \text{for each } 1 \leq i \leq n$$

then $Av = w$ has a unique solution v, using the Banach fixed point theorem applied to $f: \mathbb{R}^n \longrightarrow \mathbb{R}^n$. (Hint: choose your metric on $\mathbb{R}^n$ wisely).

<u>Exercise L14-2</u> Suppose (X,d) is a compact metric space, and for $\lambda \in (0,1)$ let

$$\mathrm{Cts}_\lambda(X,X) \subseteq \mathrm{Cts}(X,X)$$

be the subspace of λ-contraction mappings, with the subspace topology (as usual $\mathrm{Cts}(X,X)$ has the compact-open topology). Prove that the function

$$\mathrm{fix} : \mathrm{Cts}_\lambda(X,X) \longrightarrow X$$

sending a contraction mapping to its unique fixed point is continuous.

Example L14-3 Applications of the Banach fixed point theorem include:

(1) The Picard theorem on existence of sol^{n}s to ODEs (see Lecture 15).

(2) There is a proof of the Implicit Function Theorem (fundamental to higher calculus aka differential geometry) via the fixed point theorem (see for example http://www.math.jhu.edu/~jmb/note/invfnthm.pdf). This follows the idea outlined in Example L14-2.

(3) The Bellman equation is a functional equation (i.e. an equation in which the = sign means equality of functions, as in a DE) which is foundational in optimal control, dynamic programming and reinforcement learning. The functions involved are value functions $U(s)$ assigning to each possible state s (say of an agent playing an Atari game) its utility. Such a value function determines the agent's behaviour. You may be familiar with one application from DeepMind:

V. Mnih, K. Kavukcuoglu, D. Silver "Playing Atari with deep reinforcement learning" arXiv: 1312.5602 (published in Nature in 2015).

The paper starts by explaining the Bellman equation, how to converge to an optimal policy by iteration (more on that in a moment) and how that's too slow, so instead they use a (deep) neural network as a substitute. The theoretical convergence is based on the Banach fixed pt. thm:

- S. Russell, P. Norvig "Artificial intelligence: a modern approach" 3rd ed. §17.2.3
- R.S. Sutton, A.G. Barto "Reinforcement learning" §4.1.

The following is Exercise 17-6 of Stuart & Russell's AI book:

<u>Exercise L14-3</u> Consider an agent acting in an environment in order to achieve some objective, the degree of attainment of which is measured by scalar rewards. At any given time (which is discrete) the agent can be in one of a number of <u>states</u> S, and if they are in state $s \in S$ then they choose from a finite set of <u>actions</u> $A(s)$. This action is an interaction with the environment, which causes the agent to transition to state s' with probability denoted $P(s' \mid a, s)$. At each time step the agent receives a <u>reward</u> $R(s)$ depending on their state, with $R(s) \in \mathbb{R}$. We assume the set $\{ R(s) \mid s \in S \} \subseteq \mathbb{R}$ is bounded, and that given fixed a, s the probability $P(s' \mid a, s)$ is nonzero only for finitely many s'.

⌜ e.g. $S = \mathbb{Z}$, $A(s) = \{ \text{left}, \text{right} \}$ for all s, and

$$P(s' \mid a, s) = \begin{cases} 1 & \text{if } s' = s+1, a = \text{right} \\ 1 & \text{if } s' = s-1, a = \text{left} \\ 0 & \text{otherwise} \end{cases}$$

and say $R(s) = \min\{1000, e^s\}$ Move right to win! Note the probabilistic aspect is there because sometimes you try and <u>fail</u>, i.e. we would take instead

$$P(s' \mid a, s) = \begin{cases} 1/2 & \text{if } s' = s+1, a = \text{right} \\ 1/2 & \text{if } s' = s, a = \text{right} \\ 1 & \text{if } s' = s-1, a = \text{left} \\ 0 & \text{otherwise} \end{cases}$$ ⌟

The <u>discounted reward</u> of a sequence of states $\underline{s} = (s_0)_{n=0}^{\infty}$ is

$$R(\underline{s}, \gamma) := \sum_{t \geq 0} \gamma^t R(s_t) \qquad \text{(why is this finite?)}$$

where $0 < \gamma < 1$ is a fixed discount factor.

The *optimal control problem* (or reinforcement learning problem, or dynamic programming problem, or cybernetic feedback problem, ...) is to determine how the agent should behave so that its sequence of states $s_0, s_1, \ldots$ maximise the expected discounted reward. Here by "behaviour" we mean the *choice* of action $a \in A(s)$ given a current state s. Let $A = \bigcup_{s \in S} A(s)$ and define a *policy* to be a complete set of such choices, i.e. a function $\pi : S \longrightarrow A$ such that $\pi(s) \in A(s)$ for all $s \in S$.

Given a starting state s, a policy π, and the "transition model" (meaning all the probabilities $P(s' \mid a, s)$) we obtain a probability distribution over state sequences $\underline{s}$, with $P(\underline{s})$ being the probability an agent initially in state s, following π, and subject to the transition model, experiences $\underline{s}$ as its sequence of states.

MORASS: Please avoid

The *expected discounted reward* in this case is

$$U^{\pi}(s) := \mathbb{E}\big(R(\gamma, \underline{s}) \big) = \sum_{\underline{s}} P(\underline{s}) R(\gamma, \underline{s}).$$

The *optimal policy* π^*_s beginning in state s is the one that maximises $U^{\pi}(s)$ over all π, and it turns out this is independent of s, call it π^*. The *true utility* of a state s is then $U^{\pi^*_s}(s)$, which we denote $U(s)$. This all seems unsatisfying: how would you even find such a π^*_s? Recall S may be infinite.

Now here's the brilliant trick! To get around the morass in the box, we can *define* the value function $U(s)$ as the solution (among functions $U: S \longrightarrow \mathbb{R}$) of an equation (the *Bellman equation*)

$$U(s) = R(s) + \gamma \cdot \max_{a \in A(s)} \sum_{s' \in S} \underbrace{P(s' \mid a, s) U(s')}_{\text{only nonzero for finitely many } s'!}$$

finite

The idea is that *if* $U : S \to \mathbb{R}$ is a solution of the Bellman equation *then* the optimal policy of the agent is derived from it via

$$\pi^*(s) := \underset{a \in A(s)}{\operatorname{argmax}} \sum_{s' \in S} P(s' \mid a, s)\, U(s').$$

So to solve the optimal control problem we need only solve the Bellman equation. But this is obviously a *fixed point problem*!

(i) Let X be the set of bounded functions $U : S \to \mathbb{R}$ and prove that $d(U, U') = \sup\{ |U(s) - U'(s)| \mid s \in S \}$ makes X a *complete* metric space.

(ii) Prove that $f : X \longrightarrow X$ defined by

$$f(U)(s) = R(s) + \gamma \cdot \max_{a \in A(s)} \sum_{s' \in S} P(s' \mid a, s)\, U(s')$$

is a contraction with contraction factor γ.

Conclude that by the Banach fixed point theorem f has a unique fixed point (hence the Bellman eq^{n} has a unique sol^{n}) and that given any initial value function U_0 the sequence $U_i := f^i(U_0)$ converges to that sol^{n}, and hence

$$\pi_i^*(s) := \underset{a \in A(s)}{\operatorname{argmax}} \sum_{s' \in S} P(s' \mid a, s)\, U_i(s')$$

"converges" to an optimal policy as $i \to \infty$.

The set of policies does not have any reasonable topology as A is discrete.
We can however consider probabilistic policies and thereby make sense of this.