21 conjugate gradients; penalties; feasible directions

Tuesday, November 24, 2020 8:12 AM

For an elliptic functional $J : \mathbb{R}^n \to \mathbb{R}$, the optimal step-size method can be given as follows:

Note $J(v) = \frac{1}{2}\langle Av, v\rangle - \langle b, v\rangle$ and $\nabla J_v = Av - b$, so

$$0 = \langle \nabla J_{u_{k+1}}, \nabla J_{u_k}\rangle = \langle A(u_k - \rho_k(Au_k - b)) - b,\ Au_k - b\rangle$$

$$\Rightarrow \rho_k = \frac{\|w_k\|^2}{\langle Aw_k, w_k\rangle}, \quad w_k = Au_k - b = \nabla J_{u_k}.$$

So we get steps:

(1) $w_k = Au_k - b$

(2) $\rho_k = \dfrac{\|w_k\|^2}{\langle Aw_k, w_k\rangle}$

(3) $u_{k+1} = u_k - \rho_k w_k$.

fast especially if Aw is cheap to compute e.g. if A is sparse.

Convergence proofs for other variants of gradient descent are available, though they may require more or fewer conditions. e.g. see 13.7 for convergence of gradient descent with variable step-sizes in (infinite-dim) Hilbert spaces where J is elliptic and the gradient of J is Lipschitz. (Prop. 13.6)

Alternately, can also do steepest descent w.r.t. another norm (instead of Euclidean 2-norm)

13.10 Conjugate gradients for unconstrained optimization

One more method for unconstrained optimizations.
Intuition is that setting descent $d_k = -\nabla J_{u_k}$ is not always optimal.

Consider ordinary gradient descent. We have $\langle \nabla J_{u_k}, \nabla J_{u_{k+1}}\rangle = 0$ for consecutive directions. But you might "repeat" directions during a long run. Can we ensure that we never repeat directions?

Normally, given u_k, $u_k = u_k - \rho_k \nabla J_{u_k}$, where $\rho_k = \inf_{\rho\in\mathbb{R}} J(u_k - \rho\nabla J_{u_k})$

Instead, let's try minimizing J over $u_k + \mathcal{G}_k$, where $\mathcal{G}_k = \mathrm{span}\{\nabla J_{u_0}, \dots, \nabla J_{u_k}\} \subseteq \mathbb{R}^n$.

i.e. Find $u_{k+1} \in u_k + \mathcal{G}_k$ and $J(u_{k+1}) = \inf_{v\in u_k + \mathcal{G}_k} J(v)$.

Has several nice properties, assuming J is elliptic.

$v \in u_k + \mathcal{G}_k$

Has several nice properties, assuming J is elliptic.

(1) The gradients ∇J_{u_i} and ∇J_{u_j} are orthogonal $\forall i, j$ with $0 \le i < j \le k$.

Thus, if $\nabla J_{u_i} \neq 0$ for $i = 0, \dots, k$, then $\{\nabla J_{u_i}\}_{i=0,\dots,k}$ are lin. ind., so the method terminates in at most n steps.

(2) Let $\Delta_\ell = u_{\ell+1} - u_\ell = -\rho_\ell d_\ell$. Then $\langle A\Delta_\ell, \Delta_i \rangle = 0 \quad 0 \le i < \ell \le k$.

i.e. Δ_ℓ and Δ_i are "A-conjugate." Thus, if $\Delta_\ell \neq 0$ for $\ell = 0, \dots, k$, then Δ_ℓ are lin. ind.

(3) There is a simple formula to compute d_{k+1} from d_k and to compute ρ_k.

If $\nabla J_{u_i} \neq 0$ for $i = 0, \dots, k$, then we can write

$$d_\ell = \sum_{i=0}^{\ell-1} \lambda_i^\ell \nabla J_{u_i} + \nabla J_{u_\ell}, \quad 0 \le \ell \le k.$$

Then
$$\begin{cases} \lambda_i^k = \dfrac{\|\nabla J_{u_k}\|^2}{\|\nabla J_{u_i}\|^2}, & 0 \le i \le k-1 \\ d_0 = \nabla J_{u_0} \\ d_\ell = \nabla J_{u_\ell} + \dfrac{\|\nabla J_{u_\ell}\|^2}{\|\nabla J_{u_{\ell-1}}\|^2} d_{\ell-1}, & 1 \le \ell \le k. \end{cases}$$

We will not prove these properties in class due to time constraints.

Penalty methods for constrained optimization

So, turning to constrained optimization, we can sometimes just project gradients onto the convex set and use that, but often we can't easily compute $p_U(v)$. Instead, people often use penalty methods instead.

Def. 13.6/13.10 Given a nonempty, closed convex subset $U \subseteq \mathbb{R}^n$, a function $\psi: \mathbb{R}^n \to \mathbb{R}$ is called a *penalty function* for U iff ψ is convex and continuous and if the following conditions hold:

$$\psi(v) \ge 0 \quad \forall \; v \in \mathbb{R}^n, \quad \text{and} \quad \psi(v) = 0 \text{ iff } v \in U.$$

Prop. 13.11/13.19 Let $J: \mathbb{R}^n \to \mathbb{R}$ be a continuous, coercive, strictly convex function; $U \subseteq \mathbb{R}^n$ nonempty, closed, and convex; $\psi: \mathbb{R}^n \to \mathbb{R}$ a penalty function for U, and let $J_\epsilon: \mathbb{R}^n \to \mathbb{R}$ be the penalized function; $U \subseteq \mathbb{R}^n$ nonempty, closed, and convex; $\Psi: \mathbb{R}^n \to \mathbb{R}$ a penalty function for U, and let $J_\varepsilon : \mathbb{R}^n \to \mathbb{R}$ be the penalized objective function given by

$$J_\varepsilon(v) = J(v) + \frac{1}{\varepsilon}\Psi(v) \quad \text{for all } v \in \mathbb{R}^n.$$

Then $\forall \varepsilon > 0$, $\exists$ a unique $u_\varepsilon \in \mathbb{R}^n$ s.t. $J_\varepsilon(u_\varepsilon) = \inf_{v \in \mathbb{R}^n} J_\varepsilon(v)$.

Furthermore, if $u \in U$ is the unique minimizer of J over U, so that $J(u) = \inf_{u \in U} J(v)$, then $\lim_{\varepsilon \to 0} u_\varepsilon = u$.

proof. J is coercive, and $\Psi(v) \geq 0$, so $J_\varepsilon(v) \geq J(v)$ is also coercive.
J is strictly convex, and Ψ is convex, so $J_\varepsilon = J + \frac{1}{\varepsilon}\Psi$ is strictly convex.
Thus, J and J_ε both have unique minimizers $u \in U$, $u_\varepsilon \in \mathbb{R}^n$.

$\forall v \in U$, $J_\varepsilon(v) = J(v)$, so $J_\varepsilon(u) = J(u)$.

Since u_ε is the min of J_ε, $J_\varepsilon(u_\varepsilon) \leq J_\varepsilon(u)$.

Thus, $J(u_\varepsilon) \leq J(u_\varepsilon) + \frac{1}{\varepsilon}\Psi(u_\varepsilon) = J_\varepsilon(u_\varepsilon) \leq J_\varepsilon(u) = J(u)$

positivity of $\Psi(u_\varepsilon)$ — def. — u_ε is a min of J_ε — because $u \in U$ and $\Psi(u) = 0$.

Importantly, $J_\varepsilon(u_\varepsilon) \leq J(u)$.

Since J is coercive, the family $(u_\varepsilon)_{\varepsilon > 0}$ is bounded
(or else $J_\varepsilon(u_\varepsilon) \to +\infty$, and $J_\varepsilon(u_\varepsilon) \leq J(u)$).

By compactness (since we are in $\mathbb{R}^n$), $\exists$ subsequence $(u_{\varepsilon(i)})_{i \geq 0}$ with $\lim_{i \to \infty} \varepsilon(i) = 0$, and some $u' \in \mathbb{R}^n$ s.t. $\lim_{i \to \infty} u_{\varepsilon(i)} = u'$.

But $J(u_\varepsilon) \leq J(u)$, so $J(u') = \lim_{i \to \infty} J(u_{\varepsilon(i)}) \leq J(u)$.

Also $0 \leq \Psi(u_{\varepsilon(i)}) = \varepsilon(i)\left(J_\varepsilon(u_{\varepsilon(i)}) - J(u_{\varepsilon(i)})\right) \leq \varepsilon(i)\left(J(u) - J(u_{\varepsilon(i)})\right)$,

and since $(u_{\varepsilon(i)})_{i \geq 0}$ converges, $J(u) - J(u_{\varepsilon(i)})$ are bounded independently of i. Thus, since $\lim_{i \to \infty} \varepsilon(i) = 0$ and because Ψ is continuous,

$$0 = \lim_{i \to \infty} \Psi(u_{\varepsilon(i)}) = \Psi(u')$$

$\Rightarrow u' \in U$.

But $J(u') \leq J(u)$ and u is the unique minimizer of J over U, $u = u'$.

But $J(u') \leq J(u)$ and u is the unique minimizer of J over U, $u = u'$.
This argument works for every convergent subsequence of $(u_\varepsilon)_{\varepsilon>0}$, so the whole bounded sequence $(u_\varepsilon)_{\varepsilon>0}$ converges to u.

Ex. If $U = \{v \in \mathbb{R}^n \mid \varphi_i(v) \leq 0,\ i=1,\dots,m\}$, $\varphi_i : \mathbb{R}^n \to \mathbb{R}$ convex, then we can define

$$\psi(v) = \sum_{i=1}^{m} \max\{\varphi_i(v), 0\}^2 \quad \text{as a penalty function.}$$

useful to be differentiable.

Cone of feasible directions

Def. 14.1 Given a (real) vector space V, a nonempty subset $C \subseteq V$ is a cone with apex 0 (for short, a cone), if $\forall v \in V$, if $v \in C$, then $\lambda v \in C$ $\forall \lambda > 0$. For any $u \in V$, a cone with apex u is any subset of the form $u + C = \{u + v \mid v \in C\}$, where C is a cone with apex 0.

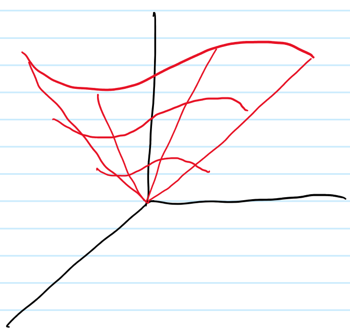

Note that cones are not necessarily convex, and 0 may not be in C.

Contrast this with the polyhedral cones we covered earlier.

Def. 14.2 Let V be a normed vector space and let U be a nonempty subset of V. For any point $u \in U$, the cone $C(u)$ of feasible directions at u is the union of $\{0\}$ and the set of nonzero vectors $w \in V$ for which there exists some convergent sequence $(u_k)_{k\geq 0}$ of vectors s.t.

(1) $u_k \in U$ and $u_k \neq u$ for all $k \geq 0$, and $\lim_{k\to\infty} u_k = u$.

(2) $\displaystyle\lim_{k\to\infty} \frac{u_k - u}{\|u_k - u\|} = \frac{w}{\|w\|}$, with $w \neq 0$.

i.e. $\exists (\delta_k)_{k\geq 0}$, $\delta_k \in V$ s.t.

$$u_k = u + \|u_k - u\| \frac{w}{\|w\|} + \|u_k - u\| \delta_k, \quad \lim_{k\to\infty} \delta_k = 0, \quad w \neq 0$$

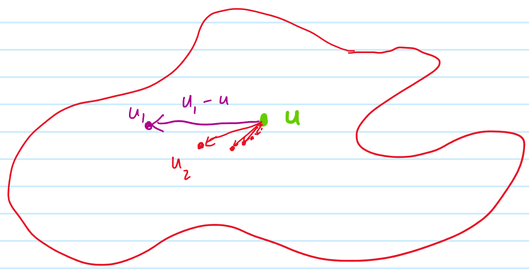

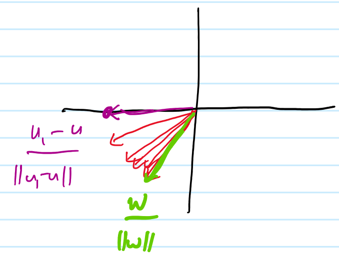

Intuitively, the cone of feasible directions at u is the set of all directions you can go from u while remaining in U.

While it is not generically convex, later we will find conditions where it is a convex cone.

Prop. 14.1 Let U be any nonempty subset of a normed space V.

(1) For any $u \in U$, the cone $C(u)$ of feasible directions at u is closed.

(2) Let $J : \Omega \to \mathbb{R}$ be a function defined on an open subset Ω containing U. If J has a local minimum w.r.t. U at $u \in U$, and if J_u' exists at u, then

$$J_u'(v-u) \geq 0 \qquad \forall\ v \in u + C(u).$$

proof. (1) Let $(w_n)_{n\geq 0}$ be a sequence $w_n \in C(u)$ converging to $w \in V$. May assume $w \neq 0$, since $0 \in C(u)$ by definition, so WLOG may also assume $w_n \neq 0$.

Then $\forall n \geq 0$, $\exists\, (u_k^n)_{k\geq 0}$ in V and some $w_n \neq 0$ s.t.

(1) $u_k^n \in U$ and $u_k^n \neq u$ $\forall k \geq 0$, and $\lim\limits_{k\to\infty} u_k^n = u$.

(2) $\exists\, (\delta_k^n)_{k\geq 0}$ in V s.t.

$$u_k^n = u + \|u_k^n - u\| \frac{w_n}{\|w_n\|} + \|u_k^n - u\|\delta_k^n, \quad \lim_{k\to\infty} \delta_k^n = 0, \quad w_n \neq 0.$$

Let $(\varepsilon_n)_{n\geq 0}$ be a sequence of real numbers $\varepsilon_n > 0$ s.t. $\lim\limits_{n\to\infty} \varepsilon_n = 0$ (e.g. $\frac{1}{n+1}$)

Then for every fixed n, $\exists\ k(n) \in \mathbb{Z}$ s.t.

$$\|u_{k(n)}^n - u\| \leq \varepsilon_n, \qquad \|\delta_{k(n)}^n\| \leq \varepsilon_n.$$

Then
$$u_{k(n)}^n = u + \|u_{k(n)}^n - u\| \frac{w}{\|w\|} + \|u_{k(n)}^n - u\| \left[\delta_{k(n)}^n + \left(\frac{w_n}{\|w_n\|} - \frac{w}{\|w\|} \right) \right]$$

Since $\lim\limits_{n\to\infty} \frac{w_n}{\|w_n\|} = \frac{w}{\|w\|}$, $\lim\limits_{n\to\infty} \frac{u_{k(n)}^n - u}{\ldots} = \frac{w}{\ldots}$ ... $w \in C(u)$

Since $\lim_{n\to\infty} \frac{w_n}{\|w_n\|} = \frac{w}{\|w\|}$, $\lim_{n\to\infty} \frac{u^n_{k(n)} - u}{\|u^n_{k(n)} - u\|} = \frac{w}{\|w\|}$, so $w \in C(u)$.

(2) Let $w = v - u$ be any non zero vector in the cone $C(u)$, and let $(u_k)_{k\geq 0}$ be the seq in $U - \{u\}$ s.t.

(1) $\lim_{k\to\infty} u_k = u$

(2) There is a sequence $(\delta_k)_{k\geq 0}$ of vectors $\delta_k \in V$ s.t.

$$u_k - u = \|u_k - u\| \frac{w}{\|w\|} + \|u_k - u\| \delta_k, \quad \lim_{k\to\infty} \delta_k = 0, \quad w \neq 0.$$

(3) $J(u) \leq J(u_k)$ for all $k \geq 0$.

Since J is differentiable at u,

$$0 \leq J(u_k) - J(u) = J'_u(u_k - u) + \|u_k - u\| \varepsilon_k.$$

for some sequence $(\varepsilon_k)_{k\geq 0}$ s.t. $\lim_{k\to\infty} \varepsilon_k = 0$.

Since J'_u is linear and continuous,

$$0 \leq J'_u(u_k - u) + \|u_k - u\| \varepsilon_k$$

$$= \frac{\|u_k - u\|}{\|w\|}\left[J'_u(w) + \eta_k\right], \quad \text{where} \quad \eta_k = \|w\|\left(J'_u(\delta_k) + \varepsilon_k\right)$$

Since J'_u is continuous, $\lim_{k\to\infty} \eta_k = 0$.

$$\Rightarrow J'_u(w) \geq 0.$$