

Optimization

Classical problems:

- find x st $f(x)$ is minimized

- interesting cases

$$x \in \mathbb{R}^n \quad \left[\begin{array}{l} f \in C^2 \quad (\text{cont. and 1, 2 deriv}) \\ \text{cont} \\ f \in C^1 \\ f \in C^0, \quad f \text{ convex} \\ f \in C^0 \end{array} \right.$$

Constrained:

$$x \in \{u \mid g(u) = 0\}$$

Discrete:

$$x \in \{0, 1\}^n$$

General methods.:Search:

construct a seq x_i
 $st \quad x_N \rightarrow$ right answer

Qns:

- How?
- When to stop?
- Converge how fast?

Closed form:

- uncommon, but sometimes useful
- Write conditions that identify a solution, then find it

→ Very valuable in itself!

Cases and terminology:Continuous optimization

- usually, f is at least continuous
- ~~usually~~, $x \in \mathbb{R}^n$
for the moment

if for $x \in B^n(x^*)$

↑
open ball centered at x^*

we have

$$f(x^*) \leq f(x)$$

then x^* is a local minimum


if we also have for $x \in \mathbb{R}^n$

$$f(x^*) \leq f(x)$$

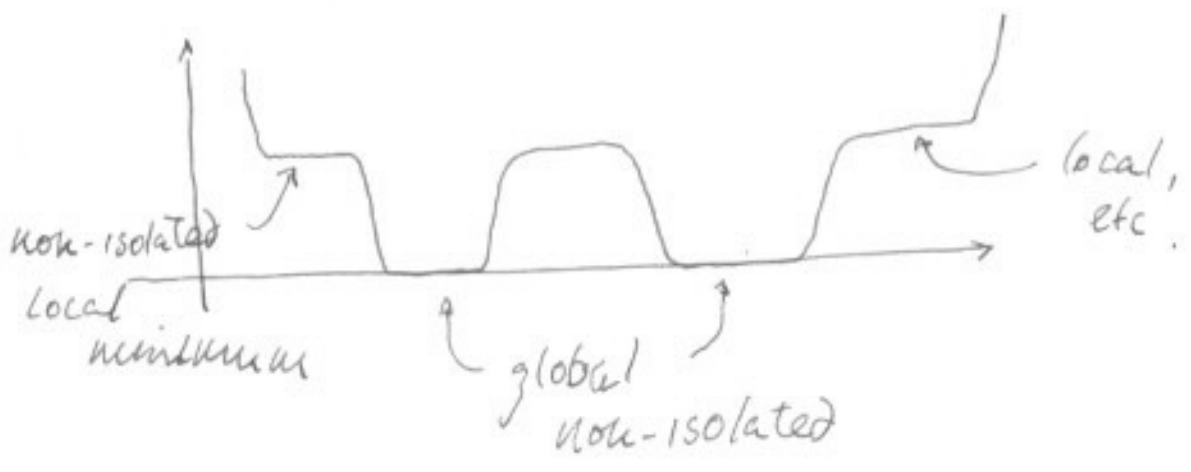
x^* is a global minimum

Some examples suggest

- There can be many x^* that are a global minimum
(But they all have the same $f(x^*)$!)



- Minima don't have to be isolated!





Some closed form examples

① a - 4

General idea: if ~~at~~ any step in any direction takes you uphill you must be at a local min.

Most helpful when f is differentiable
(but more to come)



this means

$$f(x^* + \delta x) \geq f(x^*)$$

for any δx st.

$\|\delta x\|$ is sufficiently small.

↓
this matters!



① a - 5.

Now assume f is continuously differentiable
(i.e. f is continuous; derivative of f exists and
is continuous)

$$f(x^* + \delta x) = f(x^*) + \nabla f^T \delta x + O(\delta x^2).$$

But at x^* , any direction is uphill

$$\text{So } \nabla f = 0$$

for small enough
 $\|\delta x\|^2$

and this is useful.

Recipe 1:

$$f(x) = \frac{x^T A x}{2} + b^T x + c$$

$$\nabla f = Ax + b$$

→ Where and when is there a local minimum?

$$x \mid Ax + b = 0$$

↳ These are our suspects.

But these might not be
local minima

① a-6

change coordinates:

$$u = x + w$$

where

$$Aw = b$$

does this w exist?
what if it doesn't?

then

$$Au = 0$$

is our condition

Cases

A has full rank

$$\rightarrow u = 0$$

Q: ~~are~~ is $\delta u^T A \delta u > 0$ for
all small enough δu ?

i) If $A \succ 0$, yes.

positive definite;
means $w^T A w > 0$ for all $w \neq 0$.

ii) Not isolated local min!

A does not have full rank
 \rightarrow not isolated local min.

Remember this

$f(x) = \frac{x^T A x}{2} + b^T x + c$
 has isolated local min at
 x st. $Ax + b = 0$

if $A \succ 0$
 has isolated local max at
 x st. $Ax + b = 0$

if $-A \succ 0$

Variational:

choose f such that

$$\int F(x, f(x)) dx \text{ is minimized}$$

How do we know we are at a local min?

- Any step in any direction makes cost get bigger
- Easy test if f is ~~differentiable~~ C^1


$$\nabla f = 0$$

- life is harder if f is C^0 or worse

- local test might be difficult

- eg.





we want to fill in Ω

reasonable criterion:

$$\min_{\Omega} \int \|\nabla f\|^2 dA$$

subject to: $f = I$ on $\partial\Omega$

i.e.

- don't create derivatives unnecessarily
- agree w/ boundary.

• How could we solve this?

a)

- discretize, work with discretized derivative and integral
- we are now minimizing a function of a (big!) vector

Alternative:

• What properties does f have?

→ assume that \hat{f} is the soln.

→ now, for ANY test function ϕ ,
such that $\phi = 0$ on $\partial\Omega$,
we have

$$\int_{\Omega} \|\nabla(f + \varepsilon\phi)\|^2 dA \geq \int_{\Omega} \|\nabla f\|^2 dA$$

for small enough $\varepsilon > 0$

(i.e. if you make a small move in any
direction, the value goes UP)

① - 6 - 3

Now, this means

$$\frac{d}{d\varepsilon} \int_{\Omega} \|\nabla f + \varepsilon \phi\|^2 dA = 0 \quad \text{for any } \phi$$

now this is

$$2 \int_{\Omega} \nabla f \cdot \nabla \phi dA = 0 \quad \left(\begin{array}{l} \text{doesn't seem} \\ \text{helpful} \end{array} \right)$$

But recall

$$\nabla \cdot [g \underline{v}] = (\nabla g) \cdot \underline{v} + g (\nabla \cdot \underline{v})$$

$$\text{i.e.} \quad \int_{\Omega} \nabla \cdot [\phi \nabla f] dA = \int_{\Omega} \nabla \phi \cdot \nabla f dA + \int_{\Omega} \phi (\nabla^2 f) dA$$

now

$$\int_{\Omega} \nabla \cdot [\phi \nabla f] dA = \int_{\partial \Omega} (\phi \nabla f) \cdot ds$$

(divergence theorem - remember!)

but $\phi = 0$ on $\partial \Omega$.

$$\text{so } \int_{\Omega} \nabla \phi \cdot \nabla f dA = - \int_{\Omega} \phi (\nabla^2 f) dA = 0$$

But this is true for any ϕ .

$$\text{so } \nabla^2 f = 0$$

and this offers other ways to solve.

① c

- Gives a nice criterion for variational case

$$\frac{d}{d\varepsilon} \left[\int F(u, g(u) + \varepsilon \phi(u)) du \right] \Bigg|_{\varepsilon=0} = 0$$

Variational example



$$\min \int_0^a \sqrt{1 + \left(\frac{dg}{du}\right)^2} du$$

Subject to:

$$g(0) = 0$$
$$g(a) = b$$

i.e. for any test function φ , $\varphi(0) = 0$, $\varphi(a) = 0$ ① d
 we know g is right if

$$\frac{d}{d\varepsilon} \left[\int \sqrt{1 + \left[\frac{d}{dx} (g + \varepsilon \varphi) \right]^2} dx \right] \Big|_{\varepsilon=0} = 0$$

now write $\frac{dg}{du} = g'$, $\frac{d\varphi}{du} = \varphi'$

i.e.

Criterion is:

$$\int \frac{g' \varphi'}{\sqrt{1 + g'^2}} du = 0$$

Not promising; but use integration by parts

(i.e. $\int u v' dx = uv \Big|_{x=0}^{x=a} - \int v u' dx$)

① e

to get

$$\int \frac{d}{du} \left[\frac{g'}{(1+g'^2)^{1/2}} \right] \cdot \varphi \, du = 0$$

for any φ ($\varphi(0) = \varphi(a) = 0$)

this gives $\frac{d}{du} \left[\frac{g'}{(1+g'^2)^{1/2}} \right] = 0$

which means

$$g'' \left[\frac{1 - \frac{g'^2}{(1+g'^2)^{1/2}}}{(1+g'^2)} \right] = 0$$

and $[\]$ is always positive so $g'' = 0$

(The shortest distance between 2 points is a line!)

Equations obtained like this are

Euler Lagrange equations

14

Variational problems can be quite delicate


- a solution could not exist in reasonable function spaces.

example

$$\min_f \int_0^1 f(u)^2 + \left[\left(\frac{df}{du} \right)^2 - 1 \right]^2 du = J[f]$$

$$\text{st. } f(0) = 0 ; f(1) = 0$$

Solu looks like



but as ϵ gets smaller, we have

that J gets smaller, too

→ but there can't be a

limit

So no solution.

Failure of a solution to exist
occurs in practical problems

19

• Simple issues:

$$\begin{array}{ccc} \operatorname{argmin}_u & u^2 & u \in (0, 1] \\ & & \uparrow \\ & & \text{open.} \end{array}$$

(This doesn't happen all that often,
but is worth keeping in mind)

• Harder issues

$$\operatorname{argmin}_f \int_0^1 \left[\frac{1}{\sqrt{1+(f')^2}} - \frac{1}{\sqrt{2}} \right]^2 dx$$

$$\text{s.t. } f(0) = 0$$

$$f(1) = 0$$

(h)

This sort of thing turns up in shape from shading problems rather often.

Notice I can get a min of the objective if $f'^2 = 1$.

→ So, if $f \in C^\infty$ no solution (there can't be a C^∞ function st $f'^2 = 1$, $f(0) = 0$, $f(1) = 0$)

→ if $f \in C^0$, too many solutions!

^, ~, ~~~~~ etc.

~~Now assume that~~

①i

In practice, we usually turn variational problems into continuous optimization problems by writing

$$f = \sum_i a_i g_i$$

↑ basis functions

then solving for a_i

BUT

- bad stuff can happen if original problem is poorly set
- The reasoning comes in useful later

Crucial, Take Home point:

You are at a minimum if every available step is uphill

Generally, we will construct a series, ①-j

\underline{x}_n of points

$$\underline{x}_n = \underline{x}_{n-1} + \alpha \underline{d}$$

↑ search direction
↑ step length, dist, etc.

we construct \underline{d} from some model of f at \underline{x}_{n-1} ; then choose α .

we stop if there is no \underline{d} which will ~~take~~ get a reduction in f .

Qn's

- a) what is \underline{d} ?
- b) what is α ?
- c) how to stop?.

a descent direction \underline{d} at \underline{x}_n

has the property that

$$f(\underline{x}_n + \varepsilon \underline{d}) < f(\underline{x}_n) \quad \text{for } \varepsilon < \Gamma_d$$

some + number

I: What if there aren't any?
(You're at a min!)

II: There are many ways to get
a descent dir'n because there are
many descent dir'n's.

- assume f at least C^2
- random (works about 1/2 the time!)

$$\underline{d} = -\nabla f \Big|_{\underline{x}_n}$$

(this is gradient descent).

gd works (-ish!) because

$$f(\underline{x}_n + \varepsilon \underline{d}) \approx f(\underline{x}_n) + \varepsilon \nabla f^T \underline{d} + O(\varepsilon |\underline{d}|^2)$$

(sufficiently small ε , etc)

$$\text{but } \varepsilon \nabla f^T \underline{d} = -\varepsilon \nabla f^T \nabla f < 0$$

so you can get descent unless $\nabla f^T \nabla f = 0$

This suggests a more general construction.

$$\underline{d} = -B \nabla f$$

↑ some $\begin{pmatrix} \text{p.d.} \\ \text{p.s.d.} \end{pmatrix}$ matrix.

$$f(\underline{x}_n + \varepsilon \underline{d}) \approx f(\underline{x}_n) - \varepsilon \nabla f^T B \nabla f + O(\varepsilon |\underline{d}|^2)$$

and if B p.d., you get descent
 B p.s.d., you never go up.

Common case

$$B = \text{diag}[0 \dots 1 \dots 0]$$

↑ update some vars, fix others
B p.s.d. obvious - never gets worse)

- This is often called coordinate descent
- Often seen in homemade opt. strategies (fix one; update others; etc).

Both coordinate descent and gradient descent can be

Wildly Inefficient


(though sometimes they're hard to beat!)

(2)

Important example:


Gradient descent being naughty.

$$f(x, y) = \frac{1}{2} [x^2 + \sigma y^2] \quad \sigma \gg 1$$



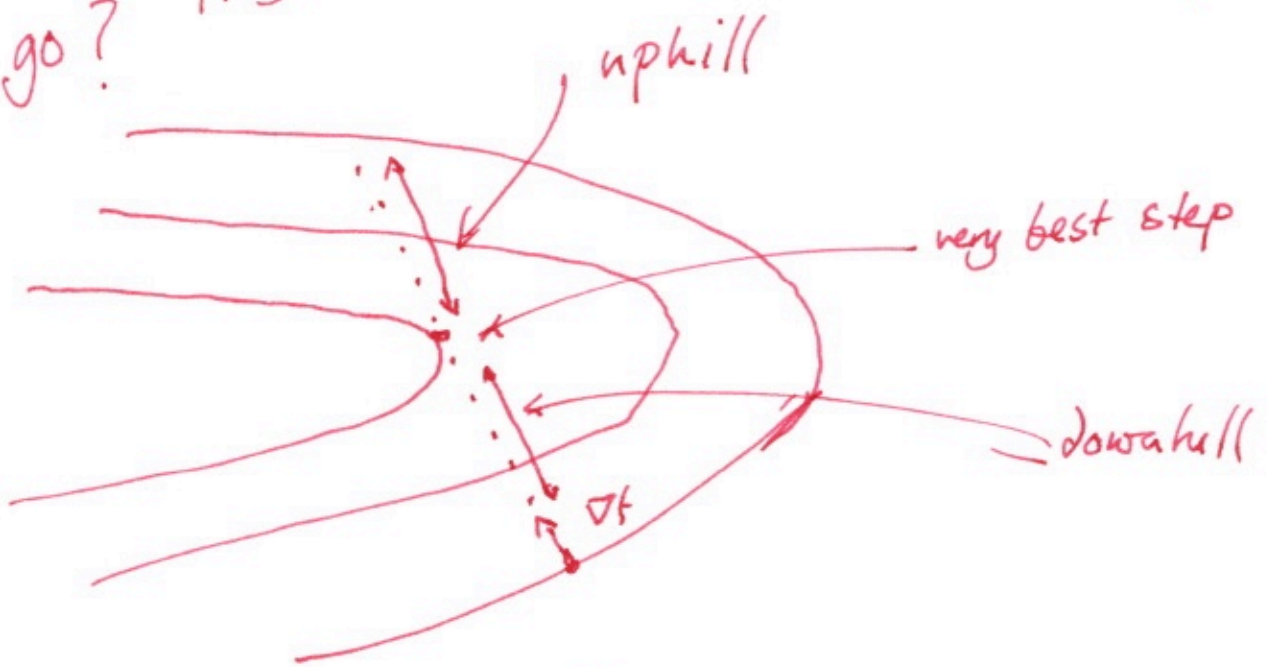
Notice:

$\nabla f|_{x_w}$ must be perpendicular to level curve at x_w

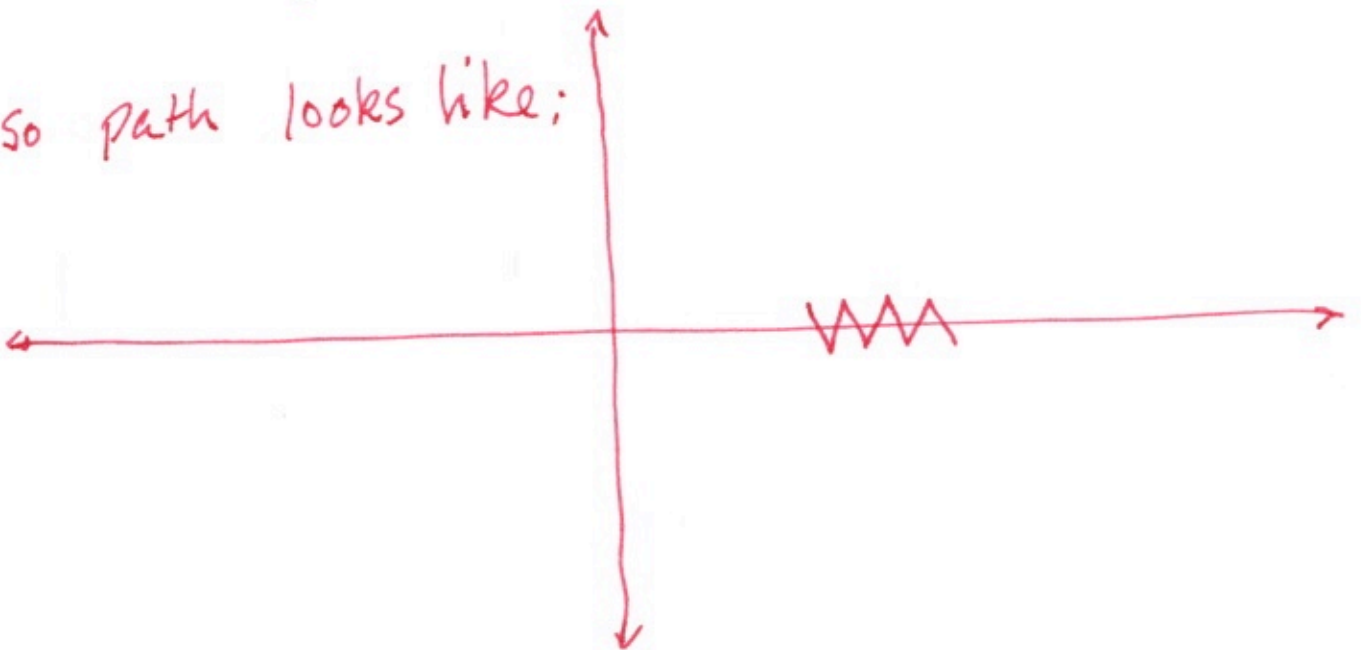


(easy argument: ∇f must be dir'n in which you leave $f(x) = c$ fastest \therefore Normal)

but apply to drawing - how far do you go? ^{(2) - a}



so path looks like:



Can demonstrate this algebraically (3)

$$f(x) = \frac{1}{2} x^T A x \quad \text{where } A \text{ p.d.}$$

• so best step is always: $-x$

• $-$ grad dir'n is $-Ax$.

• now consider $\cosine \left(\begin{array}{l} \text{angle between best} \\ \text{and grad. dir'n} \end{array} \right)$

this is

$$\frac{x^T A x}{(x^T x)^{1/2} (x^T A^T A x)^{1/2}}$$

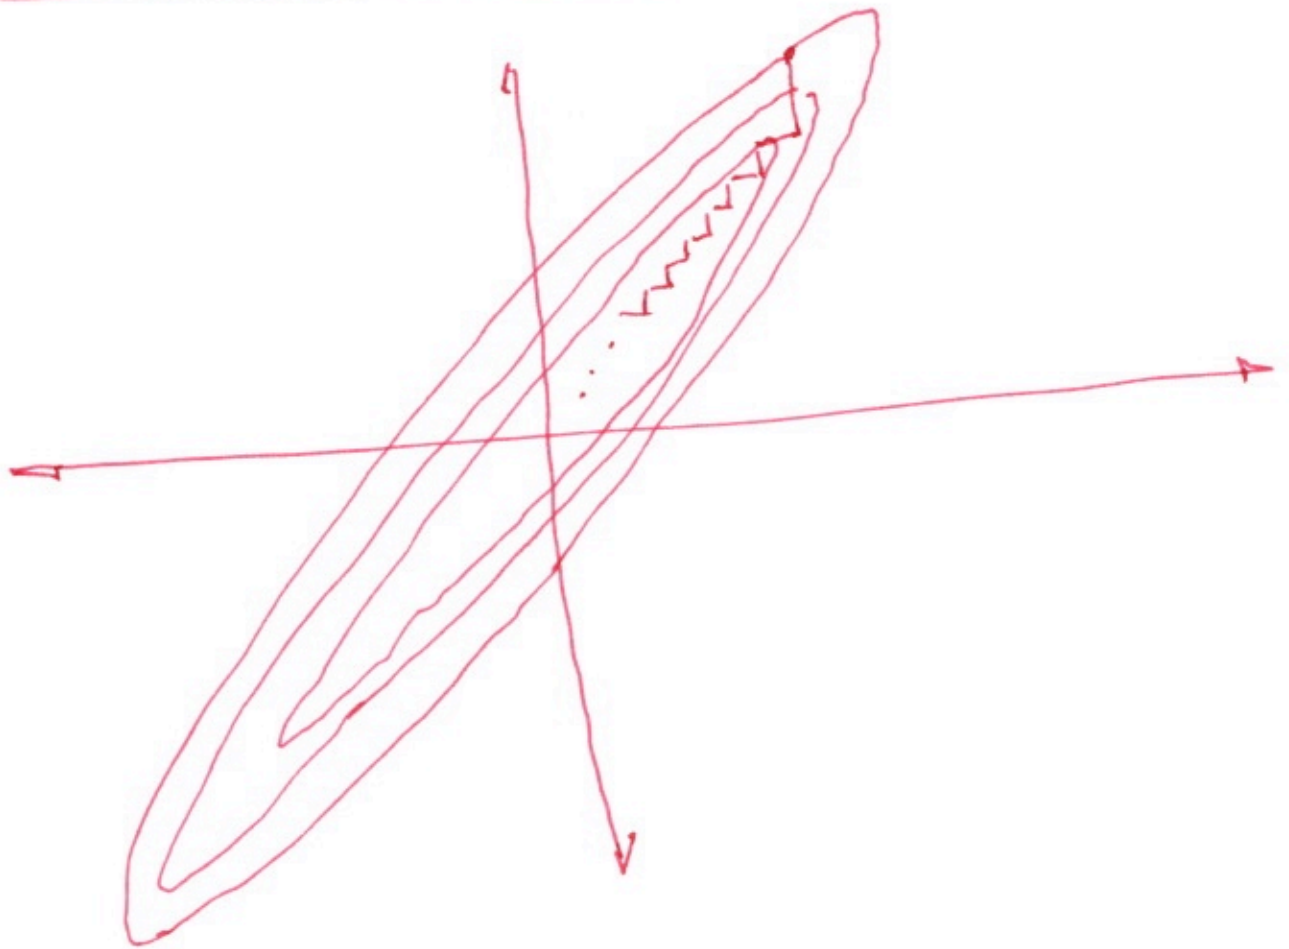
simplify by noticing x could be unit
(u, v) and $A = \begin{pmatrix} 1 & 0 \\ 0 & \sigma \end{pmatrix}$

to get

$$\frac{(\sigma - 1)}{(\sigma^2 - 1)} \rightarrow 0 \quad \text{as } \sigma \rightarrow \infty$$

i.e. If f grows v. fast in one dir, slowly in another, the gradient points nearly at right angles to good dir !! (3)-a

Coordinate Descent has a similar problem



Newton's method:

$$f(x_n + \delta) \approx f(x_n) + \nabla f^T \delta + \frac{\delta^T H_f \delta}{2} + O(\|\delta\|^3)$$

δ is the Hessian:

$$H_f = \begin{bmatrix} \frac{\partial^2 f}{\partial x_0 \partial x_0} & \frac{\partial^2 f}{\partial x_0 \partial x_1} & \dots \\ \frac{\partial^2 f}{\partial x_1 \partial x_0} & \dots & \dots \\ \vdots & \vdots & \ddots \end{bmatrix}$$

Symmetric matrix of second partial derivatives

minimize wrt δ :

$$H_f \delta + \nabla f = 0$$

$$\text{so } H_f \delta = -\nabla f$$

③ - c

example:

$$f(x) = \frac{1}{2} x^T A x + b^T x$$

A p.d

(min is at
 $Ax = -b$)

Start at x_0

$$A\delta = -b - Ax_0$$

so $A(\delta + x_0) = -b$

so $\delta + x_0 = \text{true soln, in one step}$

example:

traditional alg for square roots
is actually Newton's method

desired: r st $r^2 = c$

alg: $r_{n+1} = \frac{1}{2} \left(r_n + \frac{c}{r_n} \right)$

experiment shows \checkmark fast convergence
to a good soln.

(3) - d

consider minimizing

$$f(r) = \frac{1}{2} (r^2 - c)^2.$$

(min where $r^2 = c$).

newton's method:

$$r_{n+1} = r_n - \frac{(r_n^2 - c)r_n}{(3r_n^2 - c)}$$

write $\varepsilon = r_n^2 - c$, and assume small wrt r_n

$$\text{then } r_{n+1} = r_n - \frac{\varepsilon r_n}{2r_n^2 + \varepsilon}$$

$$\approx r_n - \frac{\varepsilon}{2r_n} = r_n - \left[\frac{r_n}{2} - \frac{c}{2r_n} \right]$$

$$= \frac{1}{2} \left[r_n + \frac{c}{r_n} \right]$$

Q. Does Newton's method always give a descent direction?

$$f(x+d) = f(x) + \nabla f d + \frac{1}{2} d^T H_f d + O(d)^3$$

$$d = -H_f^{-1} \nabla f$$

$$f(x+d) - f(x) \approx -\frac{1}{2} \nabla f^T H_f^{-1} \nabla f$$

- so we're ok if H_f is positive definite

Q: is p a descent direction?

A: if $p^T \nabla f < 0$

Notice we can obtain descent dir's from

$$p_k = -B_k^{-1} \nabla f_k$$

↑
iteration.

Q: how do we obtain descent dir? (5)

A:

$$\underline{d} = -B \nabla f$$

↑
_____ where this is p.d. / p.s.d

gradient descent:

$$B = I$$

coordinate descent:

$$B = [\text{Some projection}]$$

Newton's method:

$$B = H_s^{-1}$$

(but this will only
work if H is p.d.!))

A:

Various Quasi-Newton, etc methods (6)

$$B = \left[\begin{array}{l} \text{matrix that is:} \\ \text{(a) rather close to } H^{-1} \\ \text{(b) p.d.} \\ \text{(c) easy to work with} \end{array} \right]$$

we will see a number of ways of obtaining these.

Next Q:

- given a descent dir, how far do we go along it?

We have p_k and wish to choose a step length α .

consider $f(x_k + \alpha p_k)$ $\alpha > 0$


What α are acceptable?

- ideally, α is global minimizer
- sufficient decrease

$$f(x_k + \alpha p_k) \leq f(x_k) + c_1 \alpha \nabla f_k^T p_k$$

$$0 < c_1 < 1$$

for some constant $(\text{typically } 1e^{-4})$ } [Armijo condition]
} Wolfe



Sufficient decrease is not enough

→ very small α are OK.

$$\nabla f(x_k + \alpha p_k)^T p_k \geq c_2 \nabla f_k^T p_k$$

$$c_1 \leq c_2 \leq 1$$


notice:

write $\phi(\alpha) = f(x_k + \alpha p_k)$

then $\frac{d\phi}{d\alpha} = \nabla f(x_k + \alpha p_k)^T p_k$

so condition is:

$$\frac{d\phi}{d\alpha} \geq c_2 \nabla f_k^T p_k = c_2 \left. \frac{d\phi}{d\alpha} \right|_{\alpha=0}$$



Notice sign of slope!

c_2 is usually 0.4 (Newton)
0.1 (conj. grad.)

Wolfe conditions

$$f(x_k + \alpha_k p_k) \leq f(x_k) + c_1 \alpha_k \nabla f_k^T p_k$$

$$\nabla f(x_k + \alpha_k p_k)^T p_k \geq c_2 \nabla f_k^T p_k$$

Notice: for f continuously diff,
 f bounded below along $x_k + \alpha p_k$, $\alpha > 0$
there exist intervals satisfying
these conds.

Alg: for $\tilde{\alpha} > 0$, $\rho \in (0, 1)$
 $\alpha = \tilde{\alpha}$
repeat until [sufficient descent]
 $\alpha = \rho \alpha$
end

OK for Newton; not as good for others.

Step length selection:

$$\varphi(\alpha) = f(x_0 + \alpha p_k)$$

Sufficient decrease is then,

$$\varphi(\alpha) \leq \varphi(0) + c_1 \alpha \varphi'(0).$$

• guess α_0

→ OK ; stop

→ Not OK ; there is an OK step in interval.

• we know $\varphi(0)$, $\varphi(\alpha_0)$, $\varphi'(0)$

• build quadratic interpolate

$$\begin{aligned} \varphi(\alpha) &= \left(\frac{\varphi(\alpha_0) - \varphi(0) - \alpha_0 \varphi'(0)}{\alpha_0^2} \right) \alpha^2 \\ &\quad + \varphi'(0) \alpha \\ &\quad + \varphi(0) \end{aligned}$$

• minimize in α to get α_1

→ α_1 OK ; stop

→ else construct cubic

interpolate of $\varphi(0)$ $\varphi'(0)$ $\varphi(\alpha_0)$

$\varphi(\alpha_1)$

minimize ; α_2

→ α_2 OK stop

→ else cubic with $\varphi(0)$, $\varphi'(0)$,

two most recent α

It can be shown that if $x_k \rightarrow x^*$ superlinearly, then the ratio in this expression converges to 1. If we adjust the choice (3.60) by setting

$$\alpha_0 \leftarrow \min(1, 1.01\alpha_0),$$

we find that the unit step length $\alpha_0 = 1$ will eventually always be tried and accepted, and the superlinear convergence properties of Newton and quasi-Newton methods will be observed.

A LINE SEARCH ALGORITHM FOR THE WOLFE CONDITIONS

The Wolfe (or strong Wolfe) conditions are among the most widely applicable and useful termination conditions. We now describe in some detail a one-dimensional search procedure that is guaranteed to find a step length satisfying the *strong* Wolfe conditions (3.7) for any parameters c_1 and c_2 satisfying $0 < c_1 < c_2 < 1$. As before, we assume that p is a descent direction and that f is bounded below along the direction p .

The algorithm has two stages. This first stage begins with a trial estimate α_1 , and keeps increasing it until it finds either an acceptable step length or an interval that brackets the desired step lengths. In the latter case, the second stage is invoked by calling a function called **zoom** (Algorithm 3.6, below), which successively decreases the size of the interval until an acceptable step length is identified.

A formal specification of the line search algorithm follows. We refer to (3.7a) as the *sufficient decrease condition* and to (3.7b) as the *curvature condition*. The parameter α_{\max} is a user-supplied bound on the maximum step length allowed. The line search algorithm terminates with α_* set to a step length that satisfies the strong Wolfe conditions.

Algorithm 3.5 (Line Search Algorithm).

Set $\alpha_0 \leftarrow 0$, choose $\alpha_{\max} > 0$ and $\alpha_1 \in (0, \alpha_{\max})$;

$i \leftarrow 1$;

repeat

Evaluate $\phi(\alpha_i)$;

if $\phi(\alpha_i) > \phi(0) + c_1\alpha_i\phi'(0)$ or $[\phi(\alpha_i) \geq \phi(\alpha_{i-1})$ and $i > 1]$

$\alpha_* \leftarrow \mathbf{zoom}(\alpha_{i-1}, \alpha_i)$ and **stop**;

Evaluate $\phi'(\alpha_i)$;

if $|\phi'(\alpha_i)| \leq -c_2\phi'(0)$

set $\alpha_* \leftarrow \alpha_i$ and **stop**;

if $\phi'(\alpha_i) \geq 0$

set $\alpha_* \leftarrow \mathbf{zoom}(\alpha_i, \alpha_{i-1})$ and **stop**;

Choose $\alpha_{i+1} \in (\alpha_i, \alpha_{\max})$;

$i \leftarrow i + 1$;

end (repeat)

Note that the se
the order of the argu
the knowledge that th
conditions if one of th

(i) α_i violates the s

(ii) $\phi(\alpha_i) \geq \phi(\alpha_{i-1})$

(iii) $\phi'(\alpha_i) \geq 0$.

The last step of the al
implement this step v
we can simply set α_{i+1}
important that the suc
a finite number of iter

We now specify
its input arguments is

(a) the interval bound
conditions;

(b) α_{lo} is, among all
condition, the on

(c) α_{hi} is chosen so th

Each iteration of **zoom**
of these endpoints by c

Algorithm 3.6 (zoom)

repeat

Interpolate (u

a trial st

Evaluate $\phi(\alpha_j)$;

if $\phi(\alpha_j) > \phi(0)$

$\alpha_{hi} \leftarrow \alpha_j$

else

Evaluate

if $|\phi'(\alpha_j)| \leq -c_2\phi'(0)$

set $\alpha_* \leftarrow \alpha_j$

if $\phi'(\alpha_j) \geq 0$

$\alpha_{lo} \leftarrow \alpha_j$

$\alpha_{hi} \leftarrow \alpha_j$

end (repeat)

Note that the sequence of trial step lengths $\{\alpha_i\}$ is monotonically increasing, but that the order of the arguments supplied to the **zoom** function may vary. The procedure uses the knowledge that the interval (α_{i-1}, α_i) contains step lengths satisfying the strong Wolfe conditions if one of the following three conditions is satisfied:

- (i) α_i violates the sufficient decrease condition;
- (ii) $\phi(\alpha_i) \geq \phi(\alpha_{i-1})$;
- (iii) $\phi'(\alpha_i) \geq 0$.

The last step of the algorithm performs extrapolation to find the next trial value α_{i+1} . To implement this step we can use approaches like the interpolation procedures above, or we can simply set α_{i+1} to some constant multiple of α_i . Whichever strategy we use, it is important that the successive steps increase quickly enough to reach the upper limit α_{\max} in a finite number of iterations.

We now specify the function **zoom**, which requires a little explanation. The order of its input arguments is such that each call has the form **zoom** $(\alpha_{lo}, \alpha_{hi})$, where

- (a) the interval bounded by α_{lo} and α_{hi} contains step lengths that satisfy the strong Wolfe conditions;
- (b) α_{lo} is, among all step lengths generated so far and satisfying the sufficient decrease condition, the one giving the smallest function value; and
- (c) α_{hi} is chosen so that $\phi'(\alpha_{lo})(\alpha_{hi} - \alpha_{lo}) < 0$.

Each iteration of **zoom** generates an iterate α_j between α_{lo} and α_{hi} , and then replaces one of these endpoints by α_j in such a way that the properties (a), (b), and (c) continue to hold.

Algorithm 3.6 (**zoom**).

```

repeat
  Interpolate (using quadratic, cubic, or bisection) to find
    a trial step length  $\alpha_j$  between  $\alpha_{lo}$  and  $\alpha_{hi}$ ;
  Evaluate  $\phi(\alpha_j)$ ;
  if  $\phi(\alpha_j) > \phi(0) + c_1\alpha_j\phi'(0)$  or  $\phi(\alpha_j) \geq \phi(\alpha_{lo})$ 
     $\alpha_{hi} \leftarrow \alpha_j$ ;
  else
    Evaluate  $\phi'(\alpha_j)$ ;
    if  $|\phi'(\alpha_j)| \leq -c_2\phi'(0)$ 
      Set  $\alpha_* \leftarrow \alpha_j$  and stop;
    if  $\phi'(\alpha_j)(\alpha_{hi} - \alpha_{lo}) \geq 0$ 
       $\alpha_{hi} \leftarrow \alpha_{lo}$ ;
     $\alpha_{lo} \leftarrow \alpha_j$ ;
end (repeat)

```

If the new estimate α_j happens to satisfy the strong Wolfe conditions, then **zoom** has served its purpose of identifying such a point, so it terminates with $\alpha_* = \alpha_j$. Otherwise, if α_j satisfies the sufficient decrease condition and has a lower function value than x_{l_0} , then we set $\alpha_{l_0} \leftarrow \alpha_j$ to maintain condition (b). If this setting results in a violation of condition (c), we remedy the situation by setting α_{hi} to the old value of α_{l_0} . Readers should sketch some graphs to see for themselves how **zoom** works!

As mentioned earlier, the interpolation step that determines α_j should be safeguarded to ensure that the new step length is not too close to the endpoints of the interval. Practical line search algorithms also make use of the properties of the interpolating polynomials to make educated guesses of where the next step length should lie; see [39, 216]. A problem that can arise is that as the optimization algorithm approaches the solution, two consecutive function values $f(x_k)$ and $f(x_{k-1})$ may be indistinguishable in finite-precision arithmetic. Therefore, the line search must include a stopping test if it cannot attain a lower function value after a certain number (typically, ten) of trial step lengths. Some procedures also stop if the relative change in x is close to machine precision, or to some user-specified threshold.

A line search algorithm that incorporates all these features is difficult to code. We advocate the use of one of the several good software implementations available in the public domain. See Dennis and Schnabel [92], Lemaréchal [189], Fletcher [101], Moré and Thuente [216] (in particular), and Hager and Zhang [161].

One may ask how much more expensive it is to require the strong Wolfe conditions instead of the regular Wolfe conditions. Our experience suggests that for a “loose” line search (with parameters such as $c_1 = 10^{-4}$ and $c_2 = 0.9$), both strategies require a similar amount of work. The strong Wolfe conditions have the advantage that by decreasing c_2 we can directly control the quality of the search, by forcing the accepted value of α to lie closer to a local minimum. This feature is important in steepest descent or nonlinear conjugate gradient methods, and therefore a step selection routine that enforces the strong Wolfe conditions has wide applicability.

NOTES AND REFERENCES

For an extensive discussion of line search termination conditions see Ortega and Rheinboldt [230]. Akaike [2] presents a probabilistic analysis of the steepest descent method with exact line searches on quadratic functions. He shows that when $n > 2$, the worst-case bound (3.29) can be expected to hold for most starting points. The case $n = 2$ can be studied in closed form; see Bazaraa, Sherali, and Shetty [14]. Theorem 3.6 is due to Dennis and Moré.

Some line search methods (see Goldfarb [132] and Moré and Sorensen [213]) compute a direction of negative curvature, whenever it exists, to prevent the iteration from converging to nonminimizing stationary points. A direction of negative curvature p_- is one that satisfies $p_-^T \nabla^2 f(x_k) p_- < 0$. These algorithms generate a search direction by combining p_- with the steepest descent direction $-\nabla f_k$, often performing a curvilinear backtracking line search.

It is difficult to determine the relative contributions of the steepest descent and negative curvature directions. Because of this fact, the approach fell out of favor after the introduction of trust-region methods.

For a more thorough treatment of the modified Cholesky factorization see Gill, Murray, and Wright [130] or Dennis and Schnabel [92]. A modified Cholesky factorization based on Gershgorin disk estimates is described in Schnabel and Eskow [276]. The modified indefinite factorization is from Cheng and Higham [58].

Another strategy for implementing a line search Newton method when the Hessian contains negative eigenvalues is to compute a direction of negative curvature and use it to define the search direction (see Moré and Sorensen [213] and Goldfarb [132]).

Derivative-free line search algorithms include golden section and Fibonacci search. They share some of the features with the line search method given in this chapter. They typically store three trial points that determine an interval containing a one-dimensional minimizer. Golden section and Fibonacci differ in the way in which the trial step lengths are generated; see, for example, [79, 39].

Our discussion of interpolation follows Dennis and Schnabel [92], and the algorithm for finding a step length satisfying the strong Wolfe conditions can be found in Fletcher [101].

EXERCISES

- 3.1 Program the steepest descent and Newton algorithms using the backtracking line search, Algorithm 3.1. Use them to minimize the Rosenbrock function (2.22). Set the initial step length $\alpha_0 = 1$ and print the step length used by each method at each iteration. First try the initial point $x_0 = (1.2, 1.2)^T$ and then the more difficult starting point $x_0 = (-1.2, 1)^T$.
- 3.2 Show that if $0 < c_2 < c_1 < 1$, there may be no step lengths that satisfy the Wolfe conditions.
- 3.3 Show that the one-dimensional minimizer of a strongly convex quadratic function is given by (3.55).
- 3.4 Show that the one-dimensional minimizer of a strongly convex quadratic function always satisfies the Goldstein conditions (3.11).
- 3.5 Prove that $\|Bx\| \geq \|x\|/\|B^{-1}\|$ for any nonsingular matrix B . Use this fact to establish (3.19).
- 3.6 Consider the steepest descent method with exact line searches applied to the convex quadratic function (3.24). Using the properties given in this chapter, show that if the initial point is such that $x_0 - x^*$ is parallel to an eigenvector of Q , then the steepest descent method will find the solution in one step.

(12a)

Now we are generating a sequence $\{x_i\}$ by finding P_k, α_k and accepting

$$x_{k+1} = x_k + \alpha_k P_k$$

Q: How does this seq behave?

A:

Q: To what does it converge?

Q: How fast?

Some answers by defining

$$\cos \theta_k = \frac{-\nabla f_k^T P_k}{\|\nabla f_k\| \|P_k\|}$$

Thm: (Zoutendijk)

Consider iteration of given form, α_k satisfying Wolfe cond, f bounded below continuously diff in an open set Λ containing $L = \{x : f(x) \leq f(x_0)\}$. Assume ∇f is Lipschitz. Then

$$\sum_k \cos^2 \theta_k \|\nabla f_k\|^2 < \infty$$

(12c)

Rates of Convergence:

- We have, for $f \in C^2$, exact line search (i.e. best α_k), x^* the min, Gradient descent behaves like

$$f(x_{k+1}) - f(x^*) \leq \rho^2 [f(x_k) - f(x^*)]$$

↑
related to Hessian

- For Newton, if x_0 sufficiently close to x^*

$$\|x_{k+1} - x^*\| \leq L \|x_k - x^*\|^2$$

↑
related to Hessian

Newton's method with Hessian modification

Problem: H may not be P.D.,

so

$H p_k = -\nabla f$ may not give

a descent direction

Strategy: modify H to be P.D.

$$B_k = H_f + E_k$$

↑
chosen to make B_k P.D

This will converge globally if

~~k~~ $\{H_f(x_k)\}$ is bounded

\Rightarrow

$$\|B_k\| \|B_k^{-1}\| \leq C > 0$$

(1A)

Generally would like K_k small
(so as to preserve Hessian info)

1: Add a multiple of Identity:

choose $\beta > 0$
if $\min_i h_{ii} > 0$

$$\tilde{\tau}_0 = 0$$

else

$$\tilde{\tau}_0 = -\min(h_{ii}) + \beta;$$

end.

for $k = \dots$

attempt cholesky factorization of $H + \tau_k I$
if OK return factor

else

$$\tau_{k+1} = \max(2\tau_k, \beta)$$

end

end.

we are searching for τI to
make H p.d.

Cholesky:

$$A = L L^T$$

lower triang

- works if A PD, otherwise

get a $\sqrt{0}$, $\sqrt{-ve}$

Modified Cholesky

$$A = L D L^T$$

where

L is lower triang, 1's

on Diag

D is Diag, +ve Diag

Note: if A pd, then D elements are +ve

Cholesky:

for $j = 1 \dots n$

$$c_{jj} = a_{jj} - \sum_{s=1}^{j-1} d_s l_{js}^2$$

$$d_j = c_{jj}$$

for $i = j+1 \dots n$

$$c_{ij} = a_{ij} - \sum_{s=1}^{j-1} d_s l_{is} l_{js}$$

$$l_{ij} = c_{ij} / d_j$$

end

end

Now: d_{jj} all positive if A PD.

Modify alg so that

$$d_{jj} = \max \left(|c_{jj}|, \left(\frac{\theta_j}{\beta} \right)^2, \delta \right)$$

$$\theta_j = \max_{j < i \leq n} |c_{ij}|$$

and this gives a "factorization"
where

$$d_j \geq \delta$$
$$|m_{ij} = l_{ij} \sqrt{d_j}| \leq \beta$$

Desirable for
error control.

Improvements

- Permute rows and columns to reduce the size of the modification.
- This will give guaranteed bounds \Rightarrow global convergence.