

opt.grad Gradient descent

Consider a multivariate function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ that represents some cost or value. This is called an objective function, and we often want to find an $X \in \mathbb{R}^n$ that yields f 's extremum: minimum or maximum, depending on whichever is desirable.

objective function
extremum

It is important to note however that some functions have no finite extremum. Other functions have multiple. Finding a global extremum is generally difficult; however, many good methods exist for finding a local extremum: an extremum for some region $R \subset \mathbb{R}^n$.

global extremum

local extremum

The method explored here is called gradient descent. It will soon become apparent why it has this name.

gradient descent

Stationary points

Recall from basic calculus that a function f of a single variable had potential local extrema where $df(x)/dx = 0$. The multivariate version of this, for multivariate function f , is

$$\text{grad } f = 0. \tag{1}$$

A value X for which Eq. 1 holds is called a stationary point. However, as in the univariate case, a stationary point may not be a local extremum; in these cases, it called a saddle point.

stationary point

saddle point

Consider the hessian matrix H with values, for independent variables x_i ,

hessian matrix

$$H_{ij} = \partial_{x_i x_j}^2 f. \tag{2}$$

For a stationary point X , the second partial derivative test tells us if it is a local maximum, local minimum, or saddle point:

second partial derivative test

minimum If $H(X)$ is positive definite (all its eigenvalues are positive),

positive definite

X is a local minimum.

maximum If $H(X)$ is negative definite (all its eigenvalues are negative),

negative definite

X is a local maximum.

saddle If $H(X)$ is indefinite (it has both positive and negative eigenvalues),

indefinite

X is a saddle point.

These are sometimes called tests for concavity: minima occur where f is convex and maxima where f is concave (i.e. where $-f$ is convex).

convex

concave

It turns out, however, that solving Eq. 1 directly for stationary points is generally hard.

Therefore, we will typically use an iterative technique for estimating them.

The gradient points the way

Although Eq. 1 isn't usually directly useful for computing stationary points, it suggests iterative techniques that are. Several techniques rely on the insight that the gradient points toward stationary points. Recall from

the gradient points toward stationary points

Lec. [vecs.grad](#) that $\text{grad } f$ is a vector field that points in the direction of greatest increase in f .

Consider starting at some point x_0 and wanting to move iteratively closer to a stationary point.

So, if one is seeking a maximum of f , then choose x_1 to be in the direction of $\text{grad } f$. If one is seeking a minimum of f , then choose x_1 to be opposite the direction of $\text{grad } f$.

The question becomes: how far α should we go in (or opposite) the direction of the gradient?

Surely too-small α will require more iteration and too-large α will lead to poor convergence or missing minima altogether. This framing of the problem is called line search. There are a few common methods for choosing α , called the step size, some more computationally efficient than others.

line search

step size α

Two methods for choosing the step size are described below. Both are framed as

minimization methods, but changing the sign of the step turns them into maximization methods.

The classical method

Let

$$\mathbf{g}_k = \text{grad } f(\mathbf{x}_k), \quad (3)$$

the gradient at the algorithm's current estimate \mathbf{x}_k of the minimum. The classical method of choosing α is to attempt to solve analytically for

$$\alpha_k = \underset{\alpha}{\text{argmin}} f(\mathbf{x}_k - \alpha \mathbf{g}_k). \quad (4)$$

This solution approximates the function f as one varies α . It is approximate because as α varies, so should \mathbf{x} . But even with α as the only variable, Eq. 4 may be difficult or impossible to solve. However, this is sometimes called the "optimal" choice for α . Here "optimality" refers not to practicality but to ideality. This method is rarely used to solve practical problems.

The algorithm of the classical gradient descent method can be summarized in the pseudocode of Algorithm grad.1. It is described further in Kreyszig.¹

1. Kreyszig, *Advanced Engineering Mathematics*, § 22.1.

Algorithm grad.1 Classical gradient descent

```

1: procedure classical_minimizer( $f, \mathbf{x}_0, T$ )
2:   while  $\delta \mathbf{x} > T$  do ▷ until threshold  $T$  is met
3:      $\mathbf{g}_k \leftarrow \text{grad } f(\mathbf{x}_k)$ 
4:      $\alpha_k \leftarrow \underset{\alpha}{\text{argmin}} f(\mathbf{x}_k - \alpha \mathbf{g}_k)$ 
5:      $\mathbf{x}_{k+1} \leftarrow \mathbf{x}_k - \alpha_k \mathbf{g}_k$ 
6:      $\delta \mathbf{x} \leftarrow \|\mathbf{x}_{k+1} - \mathbf{x}_k\|$ 
7:      $k \leftarrow k + 1$ 
8:   end while
9:   return  $\mathbf{x}_k$  ▷ the threshold was reached
10: end procedure
```

The Barzilai and Borwein method

In practice, several non-classical methods are used for choosing step size α . Most of these

construct criteria for step sizes that are too small and too large and prescribe choosing some α that (at least in certain cases) must be in the sweet-spot in between. Barzilai and Borwein² developed such a prescription, which we now present.

Let $\Delta \mathbf{x}_k = \mathbf{x}_k - \mathbf{x}_{k-1}$ and $\Delta \mathbf{g}_k = \mathbf{g}_k - \mathbf{g}_{k-1}$. This method minimizes $\|\Delta \mathbf{x} - \alpha \Delta \mathbf{g}\|^2$ by choosing

$$\alpha_k = \frac{\Delta \mathbf{x}_k \cdot \Delta \mathbf{g}_k}{\Delta \mathbf{g}_k \cdot \Delta \mathbf{g}_k}. \tag{5}$$

The algorithm of this gradient descent method can be summarized in the pseudocode of [Algorithm grad.2](#). It is described further in Barzilai and Borwein.³

2. Jonathan Barzilai and Jonathan M. Borwein. ?Two-Point Step Size Gradient Methods? inIMA Journal of Numerical Analysis: 8.1 (january 1988), pages 141–148. issn: 0272-4979. doi: 10.1093/imanum/8.1.141. This includes an innovative line search method.

3. *ibidem*.

Algorithm grad.2 Barzilai and Borwein gradient descent

```

1: procedure barzilai_minimizer(f,x0,T)
2:   while  $\delta \mathbf{x} > T$  do  $\triangleright$  until threshold T is met
3:      $\mathbf{g}_k \leftarrow \text{grad } f(\mathbf{x}_k)$ 
4:      $\Delta \mathbf{g}_k \leftarrow \mathbf{g}_k - \mathbf{g}_{k-1}$ 
5:      $\Delta \mathbf{x}_k \leftarrow \mathbf{x}_k - \mathbf{x}_{k-1}$ 
6:      $\alpha_k \leftarrow \frac{\Delta \mathbf{x}_k \cdot \Delta \mathbf{g}_k}{\Delta \mathbf{g}_k \cdot \Delta \mathbf{g}_k}$ 
7:      $\mathbf{x}_{k+1} \leftarrow \mathbf{x}_k - \alpha_k \mathbf{g}_k$ 
8:      $\delta \mathbf{x} \leftarrow \|\mathbf{x}_{k+1} - \mathbf{x}_k\|$ 
9:      $k \leftarrow k + 1$ 
10:  end while
11:  return  $\mathbf{x}_k$   $\triangleright$  the threshold was reached
12: end procedure

```

Example opt.grad-1

Consider the functions (a) $f_1 : \mathbb{R}^2 \rightarrow \mathbb{R}$ and (b) $f_2 : \mathbb{R}^2 \rightarrow \mathbb{R}$ defined as

$$f_1(\mathbf{x}) = (x_1 - 25)^2 + 13(x_2 + 10)^2 \tag{6}$$

$$f_2(\mathbf{x}) = \frac{1}{2} \mathbf{x} \cdot \mathbf{A} \mathbf{x} - \mathbf{b} \cdot \mathbf{x} \tag{7}$$

where

$$\mathbf{A} = \begin{bmatrix} 20 & 0 \\ 0 & 10 \end{bmatrix} \quad \text{and} \tag{8a}$$

$$\mathbf{b} = \begin{bmatrix} 1 & 1 \end{bmatrix}^T. \tag{8b}$$

re: Barzilai and Borwein gradient descent



Use the method of Barzilai and Borwein^a starting at some x_0 to find a minimum of each function.

a. Barzilai and Borwein, [Two-Point Step Size Gradient Methods?](#)

First, load some Python packages.

```
from sympy import *
import numpy as np
import matplotlib.pyplot as plt
from IPython.display import display, Markdown, Latex
from tabulate import tabulate
```

```
-----
ModuleNotFoundError                                Traceback
↳ (most recent call last)

/tmp/ipykernel_1136564/820890629.py in <module>
      3 import matplotlib.pyplot as plt
      4 from IPython.display import display,
      ↳ Markdown, Latex
----> 5 from tabulate import tabulate

ModuleNotFoundError: No module named 'tabulate'
```

We begin by writing a class `gradient_descent_min` to perform the gradient descent. This is not optimized for speed.

```
class gradient_descent_min():
    """ A Barzilai and Borwein gradient descent class.

    Inputs:
    * f: Python function of x variables
    * x: list of symbolic variables (eg [x1, x2])
    * x0: list of numeric initial guess of a min of f
    * T: step size threshold for stopping the descent

    To execute the gradient descent call descend method.

    nb: This is only for gradients in cartesian
        coordinates! Further work would be to implement
        this in multiple or generalized coordinates.
        See the grad method below for implementation.
    """

    def __init__(self, f, x, x0, T):
        self.f = f
```

```

self.x = Array(x)
self.x0 = np.array(x0)
self.T = T
self.n = len(x0) # size of x
self.g = lambda f,x,self.grad(f,x), 'numpy')
self.xk = np.array(x0)
self.table = {}

def descend(self):
    # unpack variables
    f = self.f
    x = self.x
    x0 = self.x0
    T = self.T
    g = self.g
    # initialize variables
    N = 0
    x_k = x0
    dx = 2*T # can't be zero
    x_km1 = .9*x0-.1 # can't equal x0
    g_km1 = np.array(g(*x_km1))
    N_max = 100 # max iterations
    table_data = [[N,x0,np.array(g(*x0)),0]]
    while (dx > T and N < N_max) or N < 1:
        N += 1 # increment index
        g_k = np.array(g(*x_k))
        dg_k = g_k - g_km1
        dx_k = x_k - x_km1
        alpha_k = abs(dx_k.dot(dg_k)/dg_k.dot(dg_k))
        x_km1 = x_k # store
        x_k = x_k - alpha_k*g_k
        # save
        table_data.append([N,x_k,g_k,alpha_k])
        self.xk = np.vstack((self.xk,x_k))
        # store other variables
        g_km1 = g_k
        dx = np.linalg.norm(x_k - x_km1) # check
    self.tabulater(table_data)

def tabulater(self,table_data):
    np.set_printoptions(precision=2)
    tabulate.LATEX_ESCAPE_RULES={}
    self.table['python'] = tabulate(
        table_data,
        headers=["N","x_k","g_k","alpha"],
    )
    self.table['latex'] = tabulate(
        table_data,
        headers=[
            "$N$", "$\\bm{x}_k$", "$\\bm{g}_k$", "$\\alpha$"
        ],
        tablefmt="latex_raw",
    )

def grad(self,f,x): # cartesian coord's gradient

```



```
/tmp/ipykernel_1136564/1459049274.py in <module>
----> 1 Latex(gd.table['latex'])
```

```
KeyError: 'latex'
```

Now let's lambdify the function f_1 so we can plot.

```
f1_lambda = lambdify((x1,x2),f1(x), 'numpy')
```

Now let's plot a contour plot with the gradient descent overlaid.

```
fig, ax = plt.subplots()
# contour plot
X1 = np.linspace(-100,100,100)
X2 = np.linspace(-50,50,100)
X1, X2 = np.meshgrid(X1,X2)
F1 = f1_lambda(X1,X2)
plt.contourf(X1,X2,F1)
plt.colorbar()
# gradient descent plot
from mpl_toolkits.mplot3d import Axes3D
from matplotlib.collections import LineCollection
xX1 = gd.xk[:,0]
xX2 = gd.xk[:,1]
points = np.array([xX1, xX2]).T.reshape(-1, 1, 2)
segments = np.concatenate(
    [points[:-1], points[1:]], axis=1
)
lc = LineCollection(
    segments,
    cmap=plt.get_cmap('Reds')
)
lc.set_array(np.linspace(0,1,len(xX1))) # color segs
lc.set_linewidth(3)
ax.autoscale(False) # avoid the scatter changing limits
ax.add_collection(lc)
ax.scatter(
    xX1,xX2,
    zorder=1,
    marker="o",
    color=plt.cm.Red(np.linspace(0,1,len(xX1))),
    edgecolor='none'
)
plt.show()
```

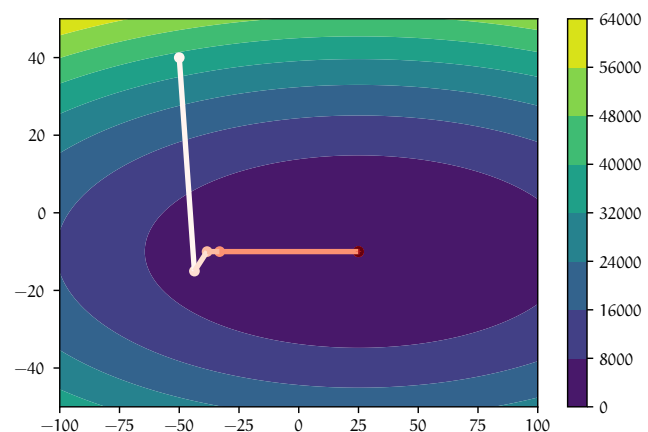


Figure grad.1:

Now consider f_2 .


```
A = Matrix([[10,0],[0,20]])
b = Matrix([[1,1]])
def f2(x):
    X = Array([x]).tomatrix().T
    return 1/2*X.dot(A*X) - b.dot(X)
gd = gradient_descent_min(f=f2,x=x,x0=[50,-40],T=1e-8)
```

Perform the gradient descent.

```
gd.descend()
```

```
-----
NameError                                Traceback
↳ (most recent call last)

/tmp/ipykernel_1136564/2845911865.py in <module>
----> 1 gd.descend()

/tmp/ipykernel_1136564/2142203784.py in
↳ descend(self)
     55     g_km1 = g_k
     56     dx = np.linalg.norm(x_k - x_km1) #
     ↳ check
----> 57     self.tabulater(table_data)
     58
     59     def tabulater(self,table_data):

/tmp/ipykernel_1136564/2142203784.py in
↳ tabulater(self, table_data)
     59     def tabulater(self,table_data):
     60     np.set_printoptions(precision=2)
----> 61     tabulate.LATEX_ESCAPE_RULES={}
     62     self.table['python'] = tabulate(
     63     table_data,

NameError: name 'tabulate' is not defined
```

Print the interesting variables.

```
print(gd.table['python'])
```

```
-----
KeyError                                Traceback
↳ (most recent call last)

/tmp/ipykernel_1136564/1459049274.py in <module>
----> 1 Latex(gd.table['latex'])
```

KeyError: 'latex'

Now let's lambdify the function f_2 so we can plot.

```
f2_lambda = lambdify((x1,x2),f2(x),'numpy')
```

Now let's plot a contour plot with the gradient descent overlaid.

```
fig, ax = plt.subplots()
# contour plot
X1 = np.linspace(-100,100,100)
X2 = np.linspace(-50,50,100)
X1, X2 = np.meshgrid(X1,X2)
F2 = f2_lambda(X1,X2)
plt.contourf(X2,X1,F2)
plt.colorbar()
# gradient descent plot
from mpl_toolkits.mplot3d import Axes3D
from matplotlib.collections import LineCollection
xX1 = gd.xk[:,0]
xX2 = gd.xk[:,1]
points = np.array([xX1, xX2]).T.reshape(-1, 1, 2)
segments = np.concatenate(
    [points[:-1], points[1:]], axis=1
)
lc = LineCollection(
    segments,
    cmap=plt.get_cmap('Reds')
)
lc.set_array(np.linspace(0,1,len(xX1))) # color segs
lc.set_linewidth(3)
ax.autoscale(False) # avoid the scatter changing limits
ax.add_collection(lc)
ax.scatter(
    xX1,xX2,
    zorder=1,
    marker="o",
    color=plt.cm.Red(np.linspace(0,1,len(xX1))),
    edgecolor='none'
)
plt.show()
```

. Python code in this section was generated from a Jupyter notebook named `gradient_descent.ipynb` with a python3 kernel.