

Optimization Algorithms: Theory and Software Implementation

Prof. Thirumulanathan D

Department of Mathematics

Institute of IIT Kanpur

Lecture: 24

Hello everyone, this is the fourth lecture of week 5. Recall that we were learning about Newton's method. In the previous three lectures, we learned the general form of Newton's algorithm and argued that this method has faster convergence than the methods discussed before. The other methods have linear convergence, while Newton's method has quadratic convergence. We also showed this at the end of the previous lecture

Other than this, we also mentioned three issues. One is no global convergence; there is only local convergence. We also proved local convergence at the end of the previous lecture. The direction we consider may not be a descent direction if the Hessian is not positive definite. The other issue was regarding the inversion of the Hessian matrix. We saw a lot of examples citing where those issues come from.

In this lecture, we will start looking at ways of getting around these issues. We do not have global convergence, but can we twist the algorithm a little bit to get global convergence? Similarly, for the other issues, we may not have H_k to be positive definite, and the inversion issues will be there. Is there a small tweak we can do to the algorithm so that we actually nullify these issues and get on track? That is what we will be discussing in this lecture.

Let me recall the issues: no global convergence, $-H_k^{-1}g_k$ need not be a descent direction, and matrix inversion is a computationally expensive process. We will first consider the global convergence issue.

I am going to introduce to you a method called the damped Newton's method. This is a small tweak of the standard Newton's method. What is damped Newton's method? I will just write down the algorithm; it is exactly Newton's method except for a change in one of the lines.

First, initialize as usual: x_0 , tolerance, and $k = 0$. While $\|g_k\| > \text{tolerance}$, do the following:

1. Fix $d_k = -H_k^{-1}g_k$.
2. Choose α_k by backtracking line search.
3. Set $x_{k+1} = x_k + \alpha_k d_k$.
4. Set $k = k + 1$.

Finally, output $x^* = x_k$.

(Refer Slide Time 5:56)

Lecture 2

$$\min_x f(x)$$

$$\text{s.t. } x \in S$$

Let $S = \mathbb{R}$.

- If $f(x) = x$, no global/local solutions.
- If $f(x) = (x-1)x(x+1)$, no global solution.
A local solution at $x^* = 0.5$
- If $f(x) = x^2$, a unique local solution at $x^* = 0$.
This is also a global solution.
- If $f(x) = (x-2)(x-1)x(x+2)$, a local solution at $x^* = 1.5$.
A global solution at some $x^* \in [-2, 0]$.

Weierstrass theorem: Consider f to be a continuous function, and the constraint set S be non-empty, closed and bounded. Then there exists a global solution x^* such that $f(x^*) \leq f(y) \forall y \in S$.

0:07:25 0:19:55

Lec-02 Dr.Thiru

Upcoming Earnings Search ENG IN 4:01 PM 5/20/2025

The only difference between Newton's method and damped Newton's method is that you choose α_k by backtracking line search in the damped version, whereas in the standard method you blindly choose $\alpha_k = 1$. I claim that this would actually give us global convergence.

Let us verify it for the example where we had the issue of no global convergence. The example was

$$f(x) = \sqrt{(x_1^2 + 1)} + \sqrt{(x_2^2 + 1)}.$$

If you had $x_0 = (0.5, 0.5)$, you had convergence. If you had $(1, 1)$, it was shuttling between $(1, 1)$ and $(-1, -1)$ repeatedly. When you had $(2, 2)$, it blew up to infinity. We will use damped Newton's method to see if the issue is rectified.

We will implement this. We set the descent direction d as $-H^{-1}g$. We choose an initial $\alpha = 1$ and then, while the Armijo condition is not satisfied, we keep shrinking α by a factor of ρ . Let's set $c_1 = 0.75$ and $\rho = 0.8$. Then we update $x = x + \alpha * d$. Everything else is the same.

For $(2, 2)$, when we ran the standard Newton's method, we got infinity in six steps. Let's see what happens here. In three steps, you can see that you are at $(0, 0)$, which is the right answer. Let's try a bigger number, say $(20, 20)$, to see if that is also taken care of. Fortunately, in four steps, you actually converge to 0.

This means that if you choose the damped Newton's algorithm instead of the standard algorithm, it slows down a bit—that is correct—because you are not having $\alpha = 1$. The step size is being pulled down. But by doing that, the damped Newton's method actually achieves global convergence.

To understand this physically: suppose you have someone who runs very fast to the destination when you give them the direction, but they sometimes overstep. What do you do? You try to pull their speed down, telling them to go slow. That is exactly what the damped Newton's method does. Newton's method is going too fast, so the damped version damps the speed in each step by using a smaller step size. This takes care of the fact that it does not overstep in each iteration. That is basically the damped Newton's method. It is a very straightforward method; there is nothing very sophisticated happening compared to the standard method.

This takes care of the no global convergence issue. Suppose you start at a particular point; you do not know x^* , so you will obviously start at an arbitrary point. Because of the no global convergence property, the answer might blow up. What would you do next? You don't know whether to increase or decrease the value; you might have to choose a point like (1, 3) based on the problem, but you don't know all these things. It is actually hard when the answer depends on the initial point. The damped Newton's method comes to our rescue. If the answer gets wrong because of the no global convergence property, you can just move to the damped Newton's method and you will get the answer.

One more thing about the damped Newton's method: you might be thinking whether it would take care of the other issues as well. For example, if $-H_k^{-1}g_k$ was not a descent direction (when H_k is not positive definite), would the damped Newton's method take care of that? Unfortunately, that is not possible because the issue there is not the step size; it is actually the direction. Even if you choose an α that is less than 1, the direction is still incorrect. The damped Newton's method does not work if the issue occurs because $-H_k^{-1}g_k$ is not a descent direction.

(Refer Slide Time 12:43)

```

return np.array([x[0]/np.sqrt(x[0]**2+1),\
                 x[1]/np.sqrt(x[1]**2+1)])

def H(x):
    return np.array([[x[0]**2+1]**(-1.5),0],
                    [0,(x[1]**2+1)**(-1.5)])

x=np.array([20,20])
k,c1,rho=0,0.75,0.8
print("%d,(%.3e,%.3e),%.3e"%(k,x[0],x[1],np.linalg.norm(grad(x))))
while(np.linalg.norm(grad(x))>1e-6):
    d=-np.linalg.inv(H(x))@grad(x)
    alpha=1
    while(f(x+alpha*d)-f(x)>c1*alpha*grad(x).dot(d)):
        alpha=rho
        x=x+alpha*d
        k+=1
    print("%d,(%.3e,%.3e),%.3e"%(k,x[0],x[1],np.linalg.norm(grad(x))))

```

0,(2.000e+01,2.000e+01),1.412e+00
1,(6.088e-01,6.088e-01),7.354e-01
2,(-5.878e-02,-5.878e-02),8.298e-02
3,(2.030e-04,2.030e-04),2.872e-04
4,(-8.371e-12,-8.371e-12),1.184e-11

Example for Singular Hessian
12:43 / 31:50
0s completed at 15:32

Let me illustrate that. Recall the example

$$f(x) = x_1^2 e^{x_2} + x_2^2 e^{x_1}.$$

When you start with $(-\sqrt{2}, -\sqrt{2})$, we had an increase of the f value, and the convergence was to a saddle point. Let's try the damped method here to see if it works. You can see that there is no change; the issue persists. So, the damped Newton's method works for no global convergence, but it does not rectify the other property. It rectifies no global convergence but does not rectify $-H_k^{-1}g_k$ being an ascent direction (the opposite of a descent direction).

(Refer Slide Time 14:31)

```

[2*x[0]*np.exp(x[1])+2*x[1]*np.exp(x[0]),
 2*np.exp(x[0])+x[0]**2*np.exp(x[1])]]

x=np.array([-np.sqrt(2),-np.sqrt(2)])
c1,rho=0,0.75,0.8
print("%d,(%.3e,%.3e),%.3e,%.3e"%(k,x[0],x[1],f(x),np.linalg.norm(grad(x))))
while(np.linalg.norm(grad(x))>1e-6):
    d=-np.linalg.inv(H(x))@grad(x)
    alpha=1
    while(f(x+alpha*d)-f(x)>c1*alpha*grad(x).dot(d)):
        alpha*=rho
    x=x+alpha*d
    k+=1
    print("%d,(%.3e,%.3e),%.3e,%.3e"%(k,x[0],x[1],f(x),np.linalg.norm(grad(x))))

0,(-1.414e+00,-1.414e+00),9.725e-01,2.848e-01
1,(-1.914e+00,-1.914e+00),1.081e+00,3.424e-02
2,(-1.997e+00,-1.997e+00),1.083e+00,1.295e-03
3,(-2.000e+00,-2.000e+00),1.083e+00,2.175e-05
4,(-2.000e+00,-2.000e+00),1.083e+00,6.177e-12

```

Now the question is, how do we rectify $-H_k^{-1}g_k$ not being a descent direction? What we do is we actually modify H_k to a slightly different matrix. I will call that the modified Newton's method. Before that, I will explain the theory behind it.

The issue, as you can see, is that if H_k is not positive definite, then there exists an eigenvalue λ_i that is less than or equal to 0. A way of characterizing a positive definite matrix is to say that every eigenvalue of that matrix is positive. For a negative definite matrix, every eigenvalue is negative. For a positive semi-definite matrix, every eigenvalue is non-negative. For a negative semi-definite matrix, every eigenvalue is non-positive. For an indefinite matrix, there exists at least one eigenvalue that is positive and one that is negative.

Since it is not positive definite, that means there should be at least one eigenvalue which is non-positive. The idea is to modify this matrix so that every eigenvalue is positive. Construct $\hat{H}_k = H_k + \lambda I$, where $\lambda = 0.5 - \min(\lambda_i)$. If λ_i 's are the eigenvalues of H_k , find the minimum of those λ_i 's. Since there exists at least one eigenvalue which is non-positive, you will have $\min(\lambda_i)$ to be either 0 or some negative number, and λ will be something greater than 0.5. I claim that \hat{H}_k is a positive definite matrix.

Why is that? The eigenvalues of \hat{H}_k are $\lambda +$ eigenvalues of H_k . The proof is simple: if λ_i are the eigenvalues of H_k , then $|H_k - \lambda_i I|$ will be singular for all i . Now,

$|\hat{H}_k - (\lambda + \lambda_i)I| = 0$, implies that $|(H_k + \lambda I) - (\lambda + \lambda_i)I| = 0$, implies that $|H_k - \lambda_i I| = 0$. This proves that the eigenvalues of \hat{H}_k are $\lambda + \lambda_i$.

(Refer Slide Time 20:22)

Damped Newton's method

- * Initialize $x^0, tol, k=0$.
- * while $(\|g^k\| > tol)$:
 - * $d^k = -(H^k)^{-1} g^k$
 - * choose α^k by backtracking line search
 - * $x^{k+1} = x^k + \alpha^k d^k$
 - * $k = k+1$
- * Output $x^* = x^k$.

The method rectifies "no global convergence" but does not rectify " $-H^k d^k$ being an ascent direction".

If $H^k \neq 0$, then \exists an eigen value $\lambda_i \leq 0$. So, to modify H^k to be a matrix with all positive eigenvalues,
let $\hat{H}_k = H_k + \lambda I$
where $\lambda = \frac{1}{2} - \min_i(\lambda_i)$
Eigenvalues of $\hat{H}_k = \lambda + (\text{Eigenvalues of } H_k)$
Since $|H_k - \lambda_i I| = 0 \Rightarrow |H_k + \lambda I - (\lambda + \lambda_i)I| = 0 \Rightarrow |\hat{H}_k - (\lambda + \lambda_i)I| = 0$.

Since I have written $\lambda = 0.5 - \min(\lambda_i)$, you can see that the minimum eigenvalue of \hat{H}_k is 0.5, which is positive. You might ask, what is the significance of 0.5? You can put any positive number. One thing is, if it is very close to zero, then inversion will make it too high. If the lowest eigenvalue is not 0 but is 0.00001 or 10^{-8} , then H^{-1} will blow up, which is not what you want. You also do not want λ to be too high, because then H_k and \hat{H}_k differ by too much. We want to retain the properties of H_k as far as possible.

So, we choose $\lambda = 0.5 - \min_i(\lambda_i)$.

Instead of the descent direction being $-H_k^{-1}g_k$, I am going to make it $-\hat{H}_k^{-1}g_k$. That is the change. The **modified Newton's method** is as follows. Initialize as usual.

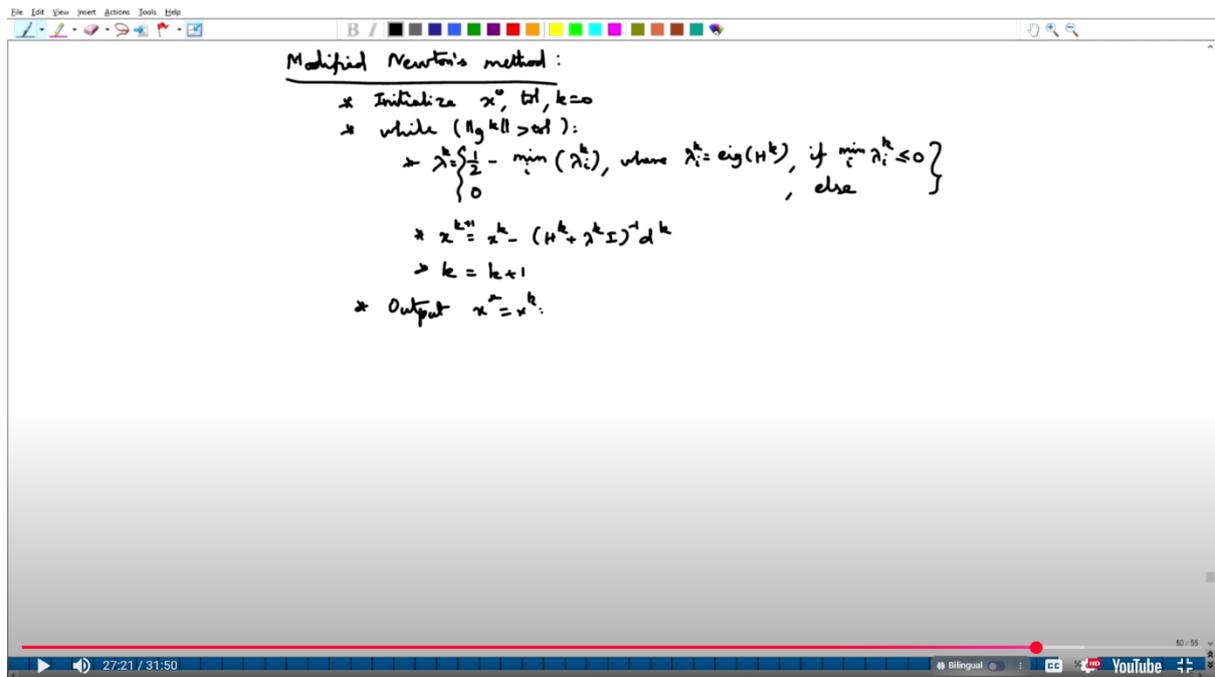
Then, define $\lambda = 0.5 - \min(\lambda_i)$, where λ_i are the eigenvalues of H_k . But this is only if $\min(\lambda_i) \leq 0$. Otherwise, if H_k is already positive definite, we set $\lambda = 0$ so we don't disturb the algorithm.

Then, $x_{k+1} = x_k - (H_k + \lambda I)^{-1}g_k$,

$k = k + 1$, and finally output

$x^* = x_k$.

(Refer Slide Time 27:20)



I claim that this will fix the issue of $-H_k^{-1}g_k$ not being a descent direction. The issue came because H_k was not a positive definite matrix.

Now $H_k + \lambda_k I$ is a positive definite matrix, so it should help us.

Let's verify if this particular method works in the example where we were failing. Let's implement this. We need to find λ first.

We can use `np.linalg.eigvals(H(x))` to get the eigenvalues.

Let's call this array L . Then, $l = \min(L)$. If $l \leq 0$, we set $\lambda_val = 0.5 - l$. Otherwise, we set $\lambda_val = 0$. Then, the descent direction becomes

$$d = -\text{inv}(H(x) + \lambda_val * I) * g.$$

Let's check if this corrects itself. I started with $(-\sqrt{2}, -\sqrt{2})$. You can see that we have converged to $(0, 0)$. Not just for this example, but even for the example $f(x) = x_1^2 - x_2^2$. Without this modification, you can recall it converged to $(0, 0)$, which is a saddle point. But if you implement this modified method, you can see that it actually goes to $(0, -\infty)$, correctly identifying the unbounded descent. This makes it clear how the modified Newton's method gets over the issue of $-H_k^{-1}g_k$ not being a descent direction.

(Refer Slide Time 30:50)

```

4, (1.000e+00, 1.000e+00), 3.942e-07

import numpy as np
def f(x):
    return x[0]**2-x[1]**2

def grad(x):
    return np.array([2*x[0], -2*x[1]])

def H(x):
    return np.array([[2, 0], [0, -2]])

x=np.array([-np.sqrt(2), -np.sqrt(2)])
k=0
print("%d, (%1.3e, %1.3e), %1.3e, %1.3e"%(k, x[0], x[1], f(x), np.linalg.norm(grad(x))))
while(np.linalg.norm(grad(x))>1e-6):
    l=np.min(np.linalg.eigvals(H(x))) #eigvals prints only the eigenvalue
    x=x-np.linalg.inv(H(x)+(1<=0)*(0.5-1)*np.eye(2))@grad(x)
    k+=1
    print("%d, (%1.3e, %1.3e), %1.3e, %1.3e"%(k, x[0], x[1], f(x), np.linalg.norm(grad(x))))

0, (-1.414e+00, -1.414e+00), 0.000e+00, 4.000e+00
1, (0.000e+00, 0.000e+00), 0.000e+00, 0.000e+00

```

```

424, (-8.222e-109, -3.264e+296), -inf, inf
425, (-4.568e-109, -1.632e+297), -inf, inf
426, (-2.538e-109, -8.161e+297), -inf, inf
427, (-1.410e-109, -4.080e+298), -inf, inf
428, (-7.832e-110, -2.040e+299), -inf, inf
429, (-4.351e-110, -1.020e+300), -inf, inf
430, (-2.417e-110, -5.101e+300), -inf, inf
431, (-1.343e-110, -2.550e+301), -inf, inf
432, (-7.461e-111, -1.275e+302), -inf, inf
433, (-4.145e-111, -6.376e+302), -inf, inf
434, (-2.303e-111, -3.188e+303), -inf, inf
435, (-1.279e-111, -1.594e+304), -inf, inf
436, (-7.107e-112, -7.970e+304), -inf, inf
437, (-3.948e-112, -3.985e+305), -inf, inf
438, (-2.194e-112, -1.992e+306), -inf, inf
439, (-1.219e-112, -9.962e+306), -inf, inf
440, (-6.770e-113, -4.981e+307), -inf, inf
441, (-3.761e-113, -inf), -inf, inf
442, (nan, inf), nan, nan
<ipython-input-9-343f82e64f87>:3: RuntimeWarning: overflow encountered in scalar power
return x[0]**2-x[1]**2
<ipython-input-9-343f82e64f87>:16: RuntimeWarning: overflow encountered in matmul
x=x-np.linalg.inv(H(x)+(1<=0)*(0.5-1)*np.eye(2))@grad(x)
<ipython-input-9-343f82e64f87>:16: RuntimeWarning: invalid value encountered in matmul
x=x-np.linalg.inv(H(x)+(1<=0)*(0.5-1)*np.eye(2))@grad(x)

```

It also gets across the issue that H_k could be singular. If H_k is singular, it means at least one of the eigenvalues is 0. If that is the minimum eigenvalue, then you can use this method to modify it to have all positive eigenvalues (i.e., non-singular) and get over that issue as well.

We will end this lecture. In the next lecture, we will look at some more properties of Newton's method. Thank you.