
Feasible Direction Methods

Reference: N&S Chapter 15

Idea:

- We have a collection of methods that work for unconstrained problems:
 - Newton's method
 - steepest descent
 - quasi-Newton
 - ...
- Can we modify them to work when we have constraints?

One advantage of these methods over Penalty and Barrier methods (to be discussed later): **We only need to evaluate the function $f(x)$ for points in the feasible region.**

This can be important if the function is undefined outside the feasible region or if we need to stop early.

Example: If $f(x)$ involves taking the \log of one of the variables, then the function is undefined for zero and negative values of this variable. []

The plan:

- Handle linear equality constraints: generating feasible directions.
- Handle linear inequality constraints: active set strategies.
- Handle more general constraints:
 - sequential quadratic programming
 - reduced-gradient methods

Handling linear equality constraints

Problem:

$$\begin{aligned} & \min_x f(x) \\ \text{subject to} & \\ & Ax = b \end{aligned}$$

We'll assume that $A \in \mathcal{R}^{m \times n}$.

The nature of the constraints

- **Case 1: $m \geq n$.** If we have this many constraints, then either there is exactly one feasible point, or there are no feasible points, or some constraint is redundant and should be removed.

Example:

$$\begin{aligned} & \min_x f(x) \\ \text{subject to} & \\ & \begin{bmatrix} 1 & 2 \\ 4 & 3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 3 \\ 7 \end{bmatrix} \end{aligned}$$

has the solution $x_1 = 1, x_2 = 1$.

If the constraints are

$$\begin{bmatrix} 1 & 2 \\ 4 & 3 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 3 \\ 7 \\ 5 \end{bmatrix}$$

then there is no solution.

If we change the right-hand side to $[3, 7, 3]^T$, then the third constraint is redundant and should be removed.

□

- **Case 2: $m < n$ and A has rank $k < m$.** In this case, some of the constraints are linearly dependent. Again, there may be **no** solution; otherwise, at least one constraint is redundant and should be removed. (**Easier said than done!**)

Example: If the constraints are

$$\begin{bmatrix} 1 & 2 & 0 \\ 2 & 4 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 3 \\ 5 \end{bmatrix}$$

then the constraint matrix has rank $1 < 2$ and there is no solution.

If we change the second entry of the right-hand side to 6, then the second constraint is redundant. If we remove it, we are in Case 3. []

- **Case 3: A has rank $m < n$.** This is the interesting case! We can satisfy the constraints, but we still have some freedom to minimize.

Example: If the constraints are

$$\begin{bmatrix} 1 & 2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 3 \end{bmatrix}$$

then we can express any solution to this linear system as a particular solution + a linear combination of vectors that form a basis for the **nullspace** of A .

For this example, one such representation is

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix} + v_1 \begin{bmatrix} 2 \\ -1 \end{bmatrix}$$

since

$$\begin{bmatrix} 1 & 2 \end{bmatrix} \begin{bmatrix} 2 \\ -1 \end{bmatrix} = 0.$$

[]

So from now on, we will assume $\text{rank}(A) = m < n$. This is a **constraint qualification**.

Clearly, in order to solve this problem, we will have to pull out some linear algebra machinery.

The nullspace of a matrix

Recall that

- z is in the nullspace of A if $Az = 0$.
- The nullspace of a matrix is a **subspace** of \mathcal{R}^n : if z and q are in the nullspace, so is $v_1z + v_2q$ for any scalars v_1 and v_2 .

Computing a basis for the nullspace of a matrix

There are two popular ways to compute a basis for the nullspace:

- Method 1: Basic variables
- Method 2: QR

Method 1: Basic variables

Partition A as

$$A = [B \quad N]$$

where

- $B \in \mathcal{R}^{m \times m}$,
- $N \in \mathcal{R}^{m \times (n-m)}$,
- $\text{rank}(B) = m$.

Note: We can rearrange the columns of A if necessary, and rearrange the variables, too, in order to achieve the full-rank condition.

Let

$$Z = \begin{bmatrix} -B^{-1}N \\ I \end{bmatrix}.$$

Then

$$AZ = 0,$$

so any linear combination of the columns of Z is in the nullspace of A .

Since Z is full rank (by inspection of its bottom block) and has rank equal to the dimension of the nullspace of A , its columns must be a basis for it!

Method 2: QR

Suppose we have factored A as

$$A^T = QR$$

where $Q \in \mathcal{R}^{n \times n}$ and $R \in \mathcal{R}^{n \times m}$ is upper trapezoidal. Then the last $n - m$ columns of Q form a basis for the null space of A .

How does this help us?

Once we have a basis for the nullspace of A , we can express any feasible point as

$$x = \hat{x} + Zv$$

where $v \in \mathcal{R}^{n-m}$ and \hat{x} solves the equations:

$$A\hat{x} = b.$$

So now our minimization problem becomes

$$\min_v f(\hat{x} + Zv)$$

with no constraints at all!

So we can solve this using our old favorite methods.

Unquiz: What is the Newton direction for this problem?

Computational issues

- Computing Z using Method 1 is easier, but it can lead to unnecessary ill-conditioning in the reduced Hessian matrix $H_w \equiv Z^T H(\hat{x} + Zv)Z$: the **condition number** satisfies

$$\kappa(H_w) \leq \kappa(H)\kappa(Z^T Z).$$

- The QR method guarantees that the condition number of the reduced Hessian matrix is no worse than the condition number of the Hessian matrix, since the columns of Z are orthonormal.

Now we have the basics of these feasible direction algorithms for linear equality constraints. Let's see what happens for linear inequality constraints.

Handling linear inequality constraints

Problem:

$$\begin{array}{ll} \min_x & f(x) \\ \text{subject to} & Ax \geq b \end{array}$$

Suppose someone told you that at the optimal solution x^* , the first k constraints held as equalities

$$A_w x = b_w$$

but the other constraints were inactive:

$$A_{\bar{w}} x > b_{\bar{w}}.$$

Then you could replace this problem by the **equality constrained problem**

$$\begin{array}{ll} \min_x & f(x) \\ \text{subject to} & A_w x = b_w. \end{array}$$

Notation: We will let W be the **index set** for the set of active constraints and \bar{W} be its complement. In the example we just did,

$$W = \{1, 2, \dots, k\}, \quad \bar{W} = \{k + 1, \dots, m\}.$$

Our strategy

We can't expect someone to tell us what constraints are active at the solution point, but **we can take a guess**, reduce the function value, and repeat until we can't do any better.

Active set algorithms

The general scheme: Given an initial feasible point $x^{(0)}$, set $k = 0$. Until optimality,

1. Choose W = the set of active constraint indices.

2. Using $x^{(k)}$ as an initial guess, solve

$$\min_x f(x)$$

subject to

$$A_w x = b_w$$

(or, at least step toward a solution to this problem) **but don't allow the algorithm to walk outside the feasible set**. Call your final point $x^{(k+1)}$ and set $k = k + 1$.

Let's think about this in a bit more detail.

Finding a descent direction

We want to compute a descent direction for

$$\min_v f(x + Zv)$$

where x is our current guess at the solution and Z is a basis for the nullspace of the matrix of the constraints that are active at x .

- **The Newton direction:** The search direction for the v variables is

$$(Z^T H(x) Z)p = -Z^T g(x)$$

(or, equivalently, Zp for the x variables)

- **The steepest descent direction:**

$$p = -Z^T g(x)$$

- **Quasi-Newton direction:**

$$B_w p = -Z^T g(x)$$

Complication: Every time we change the active set W , we have the issue of how to modify B_w .

How far to walk?

Step in that direction until we satisfy the Wolfe (or G-A) linesearch conditions **or** until we encounter another constraint.

Unquiz: For Newton's method, we used an initial step length guess of 1. Now we need to reduce this length if another constraint is violated first. How can we compute how far can we step in direction v before encountering another constraint? []

Modifying W

Clearly, if we encounter another constraint on our step, we need to add it to the set of active constraints.

What if we don't encounter a constraint, and we succeed (eventually) in minimizing the function over the current set of active constraints W . **Have we solved our problem?**

Lagrange multipliers to the rescue

We need to check the optimality conditions for constrained optimization: is the gradient a nonnegative linear combination of the columns of A_w ?

So we are trying to solve the (possibly inconsistent) linear system

$$A_w^T \lambda = g(x).$$

- If we have a QR factorization of A_w^T , then the **least squares** solution to this problem is

$$\lambda = R^{-1}Q^T g.$$

- If we used Method 1, then let

$$\lambda = \begin{bmatrix} B^{-T} & 0 \end{bmatrix} g(x).$$

Then

$$\begin{aligned} A_w^T \lambda &= \begin{bmatrix} B^T \\ N^T \end{bmatrix} \begin{bmatrix} B^{-T} & 0 \end{bmatrix} g(x) \\ &= \begin{bmatrix} I & 0 \\ N^T B^{-T} & 0 \end{bmatrix} g(x) \end{aligned}$$

and λ is a solution to the first k equations.

Either of these two methods yields us **estimates** for the current values of the Lagrange multipliers.

So now we know how to test optimality:

- **If no constraints are active**, we need to check that $g(x) = 0$ and $H(x)$ is positive definite.
- **When constraints are active**, we need to check that the Lagrange multipliers satisfy $\lambda \geq 0$ and $A_w^T \lambda = g$ and $Z^T H(x) Z$ is positive definite.

An important special case: Linear Programming

Example: linear programming

Consider this problem:

$$\min_x c^T x$$

subject to

$$Ax = b, \quad x \geq 0.$$

where $A \in \mathcal{R}^{m \times n}$.

What does our algorithm look like in this special case?

We need to see how to

- solve the problem for the current set of active constraints.
- check optimality.
- choose a new set of active constraints.

Solving the problem for the current set of active constraints

Suppose we are currently at a point x . Partition (and rearrange) x so that its nonzero components x_B come first, and rearrange the columns of A in the same way. Call the zero components x_N , suppose it has $n - m$ components, and write the active constraints as

$$\begin{bmatrix} B & N \\ 0 & I \end{bmatrix} \begin{bmatrix} x_B \\ x_N \end{bmatrix} = \begin{bmatrix} b \\ 0 \end{bmatrix}$$

Then this is a system of n equations in n unknowns and has a unique solution as long as B is nonsingular.

Therefore, by solving this system, we determine the optimal solution for this set of active constraints without doing any minimization or line search.

Checking optimality

We need to compute the Lagrange multipliers

$$\begin{bmatrix} B^T & 0 \\ N^T & I \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} = c$$

so

$$\begin{aligned} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} &= \begin{bmatrix} B^{-T} & 0 \\ -N^T B^{-T} & I \end{bmatrix} \begin{bmatrix} c_B \\ c_N \end{bmatrix} \\ &= \begin{bmatrix} B^{-T} c_B \\ c_N - N^T B^{-T} c_B \end{bmatrix} \end{aligned}$$

To be optimal, we need $\lambda_2 \geq 0$.

Jargon: The entries of λ_2 are called the **reduced costs**.

(λ_1 is unrestricted in sign since its constraints are equality constraints.)

Choosing a new set of active constraints

Suppose that λ_2 is not nonnegative.

Then we need to drop a constraint so that we can reduce the function value.

Which constraint can we drop? If we drop any constraint with a negative component in λ_2 , the function value will go down.

Customarily, we choose the constraint corresponding to the most negative multiplier.

This gives us a new set of active constraints, so we are ready for the next iteration.

A variant

Suppose we drop constraint j and let x_j be nonzero.

What is a feasible direction?

We need

$$\begin{bmatrix} B & N \\ 0 & I \end{bmatrix} p = \begin{bmatrix} 0 \\ e_j \end{bmatrix}$$

or

$$p = \begin{bmatrix} B^{-1} & -B^{-1}N \\ 0 & I \end{bmatrix} \begin{bmatrix} 0 \\ e_j \end{bmatrix} = \begin{bmatrix} -B^{-1}N \\ I \end{bmatrix} e_j.$$

Now that we have our direction, we can step along it until we hit a new constraint. Then we have n active constraints again, and we can repeat the process.

The Simplex Method for Linear Programming

The algorithm we just described is the [simplex method for linear programming](#). The original variant is due to George Dantzig in the 1940s, building on work of many others, and there has been considerable effort since then. (See N&S p. 143, 180, 233.)

We've cleaned it up to present it in matrix terms, rather than the traditional tableau form.

Unquiz: Write up the algorithm.

Some interesting questions about the simplex method

- complexity.
- geometry: walk around edge.
- stability issues.
- history

Another important special case: Quadratic programming

In quadratic programming, we minimize a quadratic function subject to linear equality and inequality constraints.

- We encountered such problems in our discussion of **trust regions**.
- We'll see them in our SQP algorithm later in these notes.
- They are important in their own right.

It is a **good exercise** to work through what feasible direction algorithms look like in the special case of quadratic programming.

Handling nonlinear constraints

Two methods:

- sequential quadratic programming
- reduced gradient approach

Sequential quadratic programming

My question: This isn't a "feasible point" method, so why do N&S include it in this chapter?

A good reference: Boggs and Tolle, *Acta Numerica* 1995 p.1

SQP

An **almost** feasible point method.

Idea: For **unconstrained optimization**, we generate a search direction for

$$\min_x f(x)$$

from a quadratic model $q(x+p)$ to f at x .

This is **Newton's method applied to $f(x)$** .

So for **constrained optimization**, we'll try to generate a feasible direction for

$$\min_x f(x)$$

subject to

$$c(x) = 0$$

(linear or nonlinear constraints) from a quadratic model. We'll apply **Newton's method to the Lagrangian**

$$L(x, \lambda) = f(x) - \lambda^T c(x).$$

There are two good ways to think about this.

Derivation 1: Quadratic model

We want a **saddle point** of the Lagrangian, minimizing with respect to x and maximizing with respect to λ . Let's hold λ fixed. Then we have

$$L(x^{(k)} + p, \lambda^{(k)}) \approx L(x^{(k)}, \lambda^{(k)}) + (\nabla_x L(x^{(k)}, \lambda^{(k)}))^T p + \frac{1}{2} p^T \nabla_{xx}^2 L(x^{(k)}, \lambda^{(k)}) p$$

and

$$c(x^{(k)} + p) \approx c(x^{(k)}) + A(x^{(k)})^T p$$

So we choose to determine p by solving the **quadratic programming problem**

$$\min_p (\nabla_x L(x^{(k)}, \lambda^{(k)}))^T p + \frac{1}{2} p^T \nabla_{xx}^2 L(x^{(k)}, \lambda^{(k)}) p$$

subject to

$$c(x^{(k)}) + A(x^{(k)})^T p = 0.$$

Derivation 2: Newton's method

We want the gradient of the Lagrangian to be zero, so we apply Newton's method to

$$\nabla L(x, \lambda) = 0.$$

This gives

$$\begin{bmatrix} x \\ \lambda \end{bmatrix}^{(k+1)} = \begin{bmatrix} x \\ \lambda \end{bmatrix}^{(k)} + \alpha \begin{bmatrix} p \\ v \end{bmatrix}^{(k)}$$

where we define the directions p and v by solving the **Newton equation**

$$\nabla^2 L(x^{(k)}, \lambda^{(k)}) \begin{bmatrix} p \\ v \end{bmatrix}^{(k)} = -\nabla L(x^{(k)}, \lambda^{(k)}).$$

What does this look like? $L(x, \lambda) = f(x) - \lambda^T c(x)$, so

$$\nabla L(x^{(k)}, \lambda^{(k)}) = \begin{bmatrix} g(x^{(k)}) - A^T(x^{(k)})\lambda^{(k)} \\ -c(x^{(k)}) \end{bmatrix}$$

$$\nabla^2 L(x^{(k)}, \lambda^{(k)}) = \begin{bmatrix} H(x^{(k)}) - G(x^{(k)}, \lambda^{(k)}) & -A^T(x^{(k)}) \\ -A(x^{(k)}) & 0 \end{bmatrix}$$

where G is the derivative with respect to x of $A^T \lambda$.

Fact: The **Newton equation** is the first order optimality condition for the quadratic programming problem in Derivation 1.

The SQP idea

Repeat until optimality:

- Generate a direction from solving the **Newton equation**.
- Take a step in that direction.

Advantages of SQP:

- Newton's method gives hope of **quadratic convergence rate**!
- Builds on well-understood technology: solving quadratic programming problems.
- One of the simpler algorithms for nonlinear constraints.

Disadvantages of SQP:

- It does not generate a sequence of **feasible points**, although it ultimately forces feasibility. (See, for example, the last column of Table 15.1 in N&S, p. 513.)
- Each iteration is **very expensive** and ordinarily we must compromise, making approximations:
 - Either neglect or approximate G .
 - Might use Quasi-Newton to approximate the (1,1) block of the matrix in the Newton equation.
 - Don't solve the QP; just insure sufficient progress. (But beware of trouble with feasibility if we stop early.)

Local expert

There are a lot of implementation issues that we have not discussed here.

Andre Tits (ECEE) is an expert on this method and has a very nice software package implementing it.

Reduced gradient methods

I don't think these methods have much of a future, but see N&S Section 15.6 if you are interested.

Final words

- Feasible direction methods and reduced gradient methods are losing popularity.
- SQP is still a good option.