

Introduction to Optimization

Lecture 2: Continuous Optimization I

September 28, 2018
TC2 - Optimisation
Université Paris-Saclay



Dimo Brockhoff
Inria Saclay – Ile-de-France

Course Overview

Date		Topic
Fri, 27.9.2019	DB	Introduction
Fri, 4.10.2019 (4hrs)	AA	Continuous Optimization I: differentiability, gradients, convexity, optimality conditions
Fri, 11.10.2019 (4hrs)	AA	Continuous Optimization II: constrained optimization, gradient-based algorithms, stochastic gradient
Fri, 18.10.2019 (4hrs)	DB	Continuous Optimization III: stochastic algorithms, derivative-free optimization, critical performance assessment [1 st written test]
Wed, 30.10.2019	DB	Discrete Optimization I: graph theory, greedy algorithms
Fri, 15.11.2019	DB	Discrete Optimization II: dynamic programming, heuristics [2 nd written test]
Fri, 22.11.2018		final exam

Details on Continuous Optimization Lectures

Introduction to Continuous Optimization

- examples (from ML / black-box problems)
- typical difficulties in optimization

Mathematical Tools to Characterize Optima

- reminders about differentiability, gradient, Hessian matrix
 - unconstrained optimization
 - first and second order conditions
 - convexity
-

- constraint optimization

Gradient-based Algorithms

- stochastic gradient
 - quasi-Newton method (BFGS)
-


Learning in Optimization / Stochastic Optimization

- CMA-ES (adaptive algorithms / Information Geometry)
- PhD thesis possible on this topic


method strongly related to ML / new promising research area
interesting open questions

What Makes a Function Difficult to Solve?

- dimensionality
(considerably) larger than three
- non-separability
dependencies between the objective variables
- ill-conditioning
- ruggedness
non-smooth, discontinuous, multimodal, and/or noisy function



a narrow ridge



cut from 3D example,
solvable with an
evolution strategy

Curse of Dimensionality

- The term *Curse of dimensionality* (Richard Bellman) refers to problems caused by the **rapid increase in volume** associated with adding extra dimensions to a (mathematical) space.
- Example: Consider placing 100 points onto a real interval, say $[0,1]$. To get **similar coverage**, in terms of distance between adjacent points, of the 10-dimensional space $[0,1]^{10}$ would require $100^{10} = 10^{20}$ points. The original 100 points appear now as isolated points in a vast empty space.
- Consequently, a **search policy** (e.g. exhaustive search) that is valuable in small dimensions **might be useless** in moderate or large dimensional search spaces.

Separable Problems

Definition (Separable Problem)

A function f is separable if

$$\operatorname{argmin}_{(x_1, \dots, x_n)} f(x_1, \dots, x_n) = \left(\operatorname{argmin}_{x_1} f(x_1, \dots), \dots, \operatorname{argmin}_{x_n} f(\dots, x_n) \right)$$


⇒ it follows that f can be optimized in a sequence of n independent 1-D optimization processes

Example:

Additively decomposable functions

$$f(x_1, \dots, x_n) = \sum_{i=1}^n f_i(x_i)$$

Rastrigin function




Non-Separable Problems

Building a non-separable problem from a separable one [1,2]


Rotating the coordinate system

- $f: \mathbf{x} \mapsto f(\mathbf{x})$ separable
- $f: \mathbf{x} \mapsto f(R\mathbf{x})$ non-separable

R rotation matrix



R
→



[1] N. Hansen, A. Ostermeier, A. Gawelczyk (1995). "On the adaptation of arbitrary normal mutation distributions in evolution strategies: The generating set adaptation". Sixth ICGA, pp. 57-64, Morgan Kaufmann


[2] R. Salomon (1996). "Reevaluating Genetic Algorithm Performance under Coordinate Rotation of Benchmark Functions; A survey of some theoretical and practical aspects of genetic algorithms." BioSystems, 39(3):263-278

III-Conditioned Problems: Curvature of Level Sets

Consider the convex-quadratic function

$$f(\mathbf{x}) = \frac{1}{2} (\mathbf{x} - \mathbf{x}^*)^T H (\mathbf{x} - \mathbf{x}^*) = \frac{1}{2} \sum_i h_{i,i} x_i^2 + \frac{1}{2} \sum_{i,j} h_{i,j} x_i x_j$$

H is Hessian matrix of f and symmetric positive definite



gradient direction $-f'(\mathbf{x})^T$

Newton direction $-H^{-1}f'(\mathbf{x})^T$

*Ill-conditioning means **squeezed level sets** (high curvature).
Condition number equals nine here. Condition numbers up to 10^{10}
are not unusual in real-world problems.*

If $H \approx I$ (small condition number of H) first order information (e.g. the gradient) is sufficient. Otherwise **second order information** (estimation of H^{-1}) information necessary.

Reminder: Different Notions of Optimum


Unconstrained case

- local vs. global
 - local minimum \mathbf{x}^* : \exists a neighborhood V of \mathbf{x}^* such that
$$\forall \mathbf{x} \in V: f(\mathbf{x}) \geq f(\mathbf{x}^*)$$
 - global minimum: $\forall \mathbf{x} \in \Omega: f(\mathbf{x}) \geq f(\mathbf{x}^*)$
- strict local minimum if the inequality is strict

Mathematical Characterization of Optima


Objective: Derive general characterization of optima

Example: if $f: \mathbb{R} \rightarrow \mathbb{R}$ differentiable,
 $f'(x) = 0$ at optimal points



- generalization to $f: \mathbb{R}^n \rightarrow \mathbb{R}$?
- generalization to constrained problems?

Remark: notion of optimum independent of notion of derivability





optima of such function can be easily
approached by certain type of methods

Reminder: Continuity of a Function

$f: (V, \| \cdot \|_V) \rightarrow (W, \| \cdot \|_W)$ is continuous in $x \in V$ if

$\forall \epsilon > 0, \exists \eta > 0$ such that $\forall y \in V: \|x - y\|_V \leq \eta; \|f(x) - f(y)\|_W \leq \epsilon$




Reminder: Differentiability in 1D (n=1)

$f: \mathbb{R} \rightarrow \mathbb{R}$ is differentiable in $x \in \mathbb{R}$ if

$$\lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h} \text{ exists, } h \in \mathbb{R}$$

Notation:

$$f'(x) = \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h}$$



The derivative corresponds to the slope of the tangent in x .

Reminder: Differentiability in 1D ($n=1$)

Taylor Formula (Order 1)

If f is differentiable in x then


$$f(x + h) = f(x) + f'(x)h + o(\|h\|)$$

i.e. for h small enough, $h \mapsto f(x + h)$ is approximated by $h \mapsto f(x) + f'(x)h$

$h \mapsto f(x) + f'(x)h$ is called a **first order approximation** of $f(x + h)$

Reminder: Differentiability in 1D ($n=1$)

Geometrically:



The notion of derivative of a function defined on \mathbb{R}^n is generalized via this idea of a linear approximation of $f(x + h)$ for h small enough.

How to generalize this to arbitrary dimension?

Gradient Definition Via Partial Derivatives

- In $(\mathbb{R}^n, \|\cdot\|_2)$ where $\|\mathbf{x}\|_2 = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$ is the Euclidean norm deriving from the scalar product $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^T \mathbf{y}$

$$\nabla f(\mathbf{x}) = \begin{pmatrix} \frac{\partial f}{\partial x_1} \\ \vdots \\ \frac{\partial f}{\partial x_n} \end{pmatrix}$$

- Reminder: partial derivative in x_0

$$f_i: y \rightarrow f(x_0^1, \dots, x_0^{i-1}, y, x_0^{i+1}, \dots, x_0^n)$$

$$\frac{\partial f}{\partial x_i}(x_0) = f_i'(x_0)$$

Exercise: Gradients

Exercise:

Compute the gradients of

a) $f(x) = x_1$ with $x \in \mathbb{R}^n$

b) $f(x) = a^T x$ with $a, x \in \mathbb{R}^n$

c) $f(x) = x^T x (= \|x\|^2)$ with $x \in \mathbb{R}^n$

Exercise: Gradients

Exercise:

Compute the gradients of

- a) $f(x) = x_1$ with $x \in \mathbb{R}^n$
- b) $f(x) = a^T x$ with $a, x \in \mathbb{R}^n$
- c) $f(x) = x^T x (= \|x\|^2)$ with $x \in \mathbb{R}^n$

Some more examples:

- in \mathbb{R}^n , if $f(x) = x^T A x$, then $\nabla f(x) = (A + A^T)x$
- in \mathbb{R} , $\nabla f(x) = f'(x)$

Gradient: Geometrical Interpretation


Exercise:

Let $L_c = \{\mathbf{x} \in \mathbb{R}^n \mid f(\mathbf{x}) = c\}$ be again a level set of a function $f(\mathbf{x})$.
Let $\mathbf{x}_0 \in L_c \neq \emptyset$.

Compute the level sets for $f_1(\mathbf{x}) = \mathbf{a}^T \mathbf{x}$ and $f_2(\mathbf{x}) = \|\mathbf{x}\|^2$ and the gradient in a chosen point \mathbf{x}_0 and observe that $\nabla f(\mathbf{x}_0)$ is **orthogonal** to the level set in \mathbf{x}_0 .

Again: if this seems too difficult, do it for two variables (and a concrete $\mathbf{a} \in \mathbb{R}^2$) and draw the level sets and the gradients.

More generally, the gradient of a differentiable function is orthogonal to its level sets.



Taylor Formula – Order One

$$f(\mathbf{x} + \mathbf{h}) = f(\mathbf{x}) + (\nabla f(\mathbf{x}))^T \mathbf{h} + o(\|\mathbf{h}\|)$$

Reminder: Second Order Derivability in 1D

- Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be a differentiable function and let $f': x \rightarrow f'(x)$ be its derivative.
- If f' is differentiable in x , then we denote its derivative as $f''(x)$
- $f''(x)$ is called the *second order derivative* of f .


Taylor Formula: Second Order Derivative

- If $f: \mathbb{R} \rightarrow \mathbb{R}$ is two times differentiable then

$$f(x+h) = f(x) + f'(x)h + f''(x)h^2 + o(\|h\|^2)$$

i.e. for h small enough, $h \rightarrow f(x) + hf'(x) + h^2f''(x)$ approximates $h \rightarrow f(x+h)$

- $h \rightarrow f(x) + hf'(x) + h^2f''(x)$ is a quadratic approximation (or order 2) of f in a neighborhood of x



- The second derivative of $f: \mathbb{R} \rightarrow \mathbb{R}$ generalizes naturally to larger dimension.

Hessian Matrix

In $(\mathbb{R}^n, \langle x, y \rangle = x^T y)$, $\nabla^2 f(x)$ is represented by a symmetric matrix called the Hessian matrix. It can be computed as

$$\nabla^2(f) = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2} & \frac{\partial^2 f}{\partial x_1 \partial x_2} & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_n} \\ \frac{\partial^2 f}{\partial x_2 \partial x_1} & \frac{\partial^2 f}{\partial x_2^2} & \cdots & \frac{\partial^2 f}{\partial x_2 \partial x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1} & \frac{\partial^2 f}{\partial x_n \partial x_2} & \cdots & \frac{\partial^2 f}{\partial x_n^2} \end{bmatrix}$$

Exercise on Hessian Matrix

Exercise:

Let $f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T A \mathbf{x}$, $\mathbf{x} \in \mathbb{R}^n$, $A \in \mathbb{R}^{n \times n}$.

Compute the Hessian matrix of f .

If it is too complex, consider $f: \begin{cases} \mathbb{R}^2 \rightarrow \mathbb{R} \\ \mathbf{x} \rightarrow \frac{1}{2} \mathbf{x}^T A \mathbf{x} \end{cases}$ with $A = \begin{pmatrix} 9 & 0 \\ 0 & 1 \end{pmatrix}$

Taylor Formula – Order Two


$$f(\mathbf{x} + \mathbf{h}) = f(\mathbf{x}) + (\nabla f(\mathbf{x}))^T \mathbf{h} + \frac{1}{2} \mathbf{h}^T (\nabla^2 f(\mathbf{x})) \mathbf{h} + o(\|\mathbf{h}\|^2)$$

Back to Ill-Conditioned Problems

We have seen that for a convex quadratic function

$$f(x) = \frac{1}{2}(x - x_0)^T A(x - x_0) + b \text{ of } x \in \mathbb{R}^n, A \in \mathbb{R}^{n \times n}, A \text{ SPD}, b \in \mathbb{R}^n:$$

- 1) The level sets are ellipsoids. The eigenvalues of A determine the lengths of the principle axes of the ellipsoid.



For $n = 2$, let λ_1, λ_2 be the eigenvalues of A .

- 2) The Hessian matrix of f equals to A .

Ill-conditioned convex quadratic problems are problems with large ratio between largest and smallest eigenvalue of A which means large ratio between longest and shortest axis of ellipsoid.

This corresponds to having an ill-conditioned Hessian matrix.

Gradient Direction Vs. Newton Direction

Gradient direction: $\nabla f(\mathbf{x})$

Newton direction: $(H(\mathbf{x}))^{-1} \cdot \nabla f(\mathbf{x})$

with $H(\mathbf{x}) = \nabla^2 f(\mathbf{x})$ being the Hessian at \mathbf{x}

Exercise:

Let again $f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T A \mathbf{x}$, $\mathbf{x} \in \mathbb{R}^2$, $A = \begin{pmatrix} 9 & 0 \\ 0 & 1 \end{pmatrix} \in \mathbb{R}^{2 \times 2}$.

Plot the gradient and Newton direction of f in a point $\mathbf{x} \in \mathbb{R}^n$ of your choice (which should not be on a coordinate axis) into the same plot with the level sets, we created before.

Optimality Conditions for Unconstrained Problems

Optimality Conditions: First Order Necessary Cond.

For 1-dimensional optimization problems $f: \mathbb{R} \rightarrow \mathbb{R}$

Assume f is differentiable

- \mathbf{x}^* is a local optimum $\Rightarrow f'(\mathbf{x}^*) = 0$

not a sufficient condition: consider $f(x) = x^3$

proof via Taylor formula: $f(\mathbf{x}^ + \mathbf{h}) = f(\mathbf{x}^*) + f'(\mathbf{x}^*)\mathbf{h} + o(\|\mathbf{h}\|)$*

- points \mathbf{y} such that $f'(\mathbf{y}) = 0$ are called **critical** or **stationary** points

Generalization to n -dimensional functions

If $f: U \subset \mathbb{R}^n \mapsto \mathbb{R}$ is differentiable

- necessary condition: If \mathbf{x}^* is a local optimum of f , then $\nabla f(\mathbf{x}^*) = 0$

proof via Taylor formula

Second Order Necessary and Sufficient Opt. Cond.

If f is twice continuously differentiable

- **Necessary condition:** if \mathbf{x}^* is a local minimum, then $\nabla f(\mathbf{x}^*) = 0$ and $\nabla^2 f(\mathbf{x}^*)$ is positive semi-definite

proof via Taylor formula at order 2

- **Sufficient condition:** if $\nabla f(\mathbf{x}^*) = 0$ and $\nabla^2 f(\mathbf{x}^*)$ is positive definite, then \mathbf{x}^* is a strict local minimum

Proof of Sufficient Condition:

- Let $\lambda > 0$ be the smallest eigenvalue of $\nabla^2 f(\mathbf{x}^*)$, using a second order Taylor expansion, we have for all \mathbf{h} :

- $$f(\mathbf{x}^* + \mathbf{h}) - f(\mathbf{x}^*) = \nabla f(\mathbf{x}^*)^T \mathbf{h} + \frac{1}{2} \mathbf{h}^T \nabla^2 f(\mathbf{x}^*) \mathbf{h} + o(\|\mathbf{h}\|^2)$$
$$> \frac{\lambda}{2} \|\mathbf{h}\|^2 + o(\|\mathbf{h}\|^2) = \left(\frac{\lambda}{2} + \frac{o(\|\mathbf{h}\|^2)}{\|\mathbf{h}\|^2} \right) \|\mathbf{h}\|^2$$

Convex Functions

Let U be a convex open set of \mathbb{R}^n and $f: U \rightarrow \mathbb{R}$. The function f is said to be **convex** if for all $\mathbf{x}, \mathbf{y} \in U$ and for all $t \in [0,1]$

$$f((1-t)\mathbf{x} + t\mathbf{y}) \leq (1-t)f(\mathbf{x}) + tf(\mathbf{y})$$

Theorem

If f is differentiable, then f is convex if and only if for all \mathbf{x}, \mathbf{y}

$$f(\mathbf{y}) - f(\mathbf{x}) \geq (\nabla f(\mathbf{x}))^T (\mathbf{y} - \mathbf{x})$$

if $n = 1$, the curve is on top of the tangent

If f is twice continuously differentiable, then f is convex if and only if $\nabla^2 f(\mathbf{x})$ is positive semi-definite for all \mathbf{x} .

Convex Functions: Why Convexity?

Examples of Convex Functions:

- $f(\mathbf{x}) = a^T \mathbf{x} + b$
- $f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T A \mathbf{x} + a^T \mathbf{x} + b$, A symmetric positive definite
- the negative of the entropy function (i. e. $f(\mathbf{x}) = -\sum_{i=1}^n x_i \ln(x_i)$)

Exercise:

Let $f: U \rightarrow \mathbb{R}$ be a convex and differentiable function on a convex open U .

Show that if $\nabla f(\mathbf{x}^*) = 0$, then \mathbf{x}^* is a global minimum of f

Why is convexity an important concept?

Constrained Optimization

Equality Constraint

Objective:

Generalize the necessary condition of $\nabla f(x) = 0$ at the optima of f when f is in \mathcal{C}^1 , i.e. is differentiable and its differential is continuous

Theorem:

Be U an open set of $(E, \|\cdot\|)$, and $f: U \rightarrow \mathbb{R}$, $g: U \rightarrow \mathbb{R}$ in \mathcal{C}^1 .

Let $a \in E$ satisfy

$$\begin{cases} f(a) = \inf \{f(x) \mid x \in \mathbb{R}^n, g(x) = 0\} \\ g(a) = 0 \end{cases}$$

i.e. a is optimum of the problem

If $\nabla g(a) \neq 0$, then there exists a constant $\lambda \in \mathbb{R}$ called *Lagrange multiplier*, such that

$$\underbrace{\nabla f(a) + \lambda \nabla g(a)} = 0 \quad \text{Euler – Lagrange equation}$$

i.e. gradients of f and g in a are colinear

Geometrical Interpretation Using an Example

Exercise:

Consider the problem

$$\inf \{ f(x, y) \mid (x, y) \in \mathbb{R}^2, g(x, y) = 0 \}$$

$$f(x, y) = y - x^2 \quad g(x, y) = x^2 + y^2 - 1 = 0$$

- 1) Plot the level sets of f , plot $g = 0$
- 2) Compute ∇f and ∇g
- 3) Find the solutions with $\nabla f + \lambda \nabla g = 0$
equation solving with 3 unknowns (x, y, λ)
- 4) Plot the solutions of 3) on top of the level set graph of 1)

Interpretation of Euler-Lagrange Equation

Intuitive way to retrieve the Euler-Lagrange equation:

- In a local minimum a of a constrained problem, the hypersurfaces (or level sets) $f = f(a)$ and $g = 0$ are necessarily tangent (otherwise we could decrease f by moving along $g = 0$).
- Since the gradients $\nabla f(a)$ and $\nabla g(a)$ are orthogonal to the level sets $f = f(a)$ and $g = 0$, it follows that $\nabla f(a)$ and $\nabla g(a)$ are colinear.

Generalization to More than One Constraint

Theorem

- Assume $f: U \rightarrow \mathbb{R}$ and $g_k: U \rightarrow \mathbb{R}$ ($1 \leq k \leq p$) are \mathcal{C}^1 .
- Let a be such that
$$\begin{cases} f(a) = \inf \{f(x) \mid x \in \mathbb{R}^n, & g_k(x) = 0, & 1 \leq k \leq p\} \\ g_k(a) = 0 \text{ for all } 1 \leq k \leq p \end{cases}$$
- If $(\nabla g_k(a))_{1 \leq k \leq p}$ are linearly independent, then there exist p real constants $(\lambda_k)_{1 \leq k \leq p}$ such that

$$\nabla f(a) + \sum_{k=1}^p \lambda_k \nabla g_k(a) = 0$$

↑
Lagrange multiplier

again: a does not need to be global but local minimum

The Lagrangian

- Define the Lagrangian on $\mathbb{R}^n \times \mathbb{R}^p$ as

$$\mathcal{L}(x, \{\lambda_k\}) = f(x) + \sum_{k=1}^p \lambda_k g_k(x)$$

- To find optimal solutions, we can solve the optimality system

$$\left\{ \begin{array}{l} \text{Find } (x, \{\lambda_k\}) \in \mathbb{R}^n \times \mathbb{R}^p \text{ such that } \nabla f(x) + \sum_{k=1}^p \lambda_k \nabla g_k(x) = 0 \\ g_k(x) = 0 \text{ for all } 1 \leq k \leq p \end{array} \right.$$

$$\Leftrightarrow \left\{ \begin{array}{l} \text{Find } (x, \{\lambda_k\}) \in \mathbb{R}^n \times \mathbb{R}^p \text{ such that } \nabla_x \mathcal{L}(x, \{\lambda_k\}) = 0 \\ \nabla_{\lambda_k} \mathcal{L}(x, \{\lambda_k\})(x) = 0 \text{ for all } 1 \leq k \leq p \end{array} \right.$$

Inequality Constraint: Definitions

Let $\mathcal{U} = \{x \in \mathbb{R}^n \mid g_k(x) = 0 \text{ (for } k \in E), g_k(x) \leq 0 \text{ (for } k \in I)\}$.

Definition:

The points in \mathbb{R}^n that satisfy the constraints are also called *feasible* points.

Definition:

Let $a \in \mathcal{U}$, we say that the constraint $g_k(x) \leq 0$ (for $k \in I$) is *active* in a if $g_k(a) = 0$.

Inequality Constraint: Karush-Kuhn-Tucker Theorem

Theorem (Karush-Kuhn-Tucker, KKT):

Let U be an open set of $(\mathbb{R}^n, || \cdot ||)$ and $f: U \rightarrow \mathbb{R}$, $g_k: U \rightarrow \mathbb{R}$, all \mathcal{C}^1

Furthermore, let $a \in U$ satisfy

$$\left\{ \begin{array}{l} f(a) = \inf\{f(x) \mid x \in \mathbb{R}^n, g_k(x) = 0 \text{ (for } k \in E), g_k(x) \leq 0 \text{ (for } k \in I)\} \\ g_k(a) = 0 \text{ (for } k \in E) \\ g_k(a) \leq 0 \text{ (for } k \in I) \end{array} \right. \quad \text{also works again for } a \text{ being a local minimum}$$

Let I_a^0 be the set of constraints that are active in a . Assume that $(\nabla g_k(a))_{k \in E \cup I_a^0}$ are linearly independent.

Then there exist $(\lambda_k)_{1 \leq k \leq p}$ that satisfy

$$\left\{ \begin{array}{l} \nabla f(a) + \sum_{k=1}^p \lambda_k \nabla g_k(a) = 0 \\ g_k(a) = 0 \text{ (for } k \in E) \\ g_k(a) \leq 0 \text{ (for } k \in I) \\ \lambda_k \geq 0 \text{ (for } k \in I_a^0) \\ \lambda_k g_k(a) = 0 \text{ (for } k \in E \cup I) \end{array} \right.$$

Inequality Constraint: Karush-Kuhn-Tucker Theorem

Theorem (Karush-Kuhn-Tucker, KKT):

Let U be an open set of $(E, || ||)$ and $f: U \rightarrow \mathbb{R}$, $g_k: U \rightarrow \mathbb{R}$, all \mathcal{C}^1

Furthermore, let $a \in U$ satisfy

$$\left\{ \begin{array}{l} f(a) = \inf\{f(x) \mid x \in \mathbb{R}^n, g_k(x) = 0 \text{ (for } k \in E), g_k(x) \leq 0 \text{ (for } k \in I)\} \\ g_k(a) = 0 \text{ (for } k \in E) \\ g_k(a) \leq 0 \text{ (for } k \in I) \end{array} \right.$$

Let I_a^0 be the set of constraints that are active in a . Assume that $(\nabla g_k(a))_{k \in E \cup I_a^0}$ are linearly independent.

Then there exist $(\lambda_k)_{1 \leq k \leq p}$ that satisfy

$$\left\{ \begin{array}{l} \nabla f(a) + \sum_{k=1}^p \lambda_k \nabla g_k(a) = 0 \\ g_k(a) = 0 \text{ (for } k \in E) \\ g_k(a) \leq 0 \text{ (for } k \in I) \\ \lambda_k \geq 0 \text{ (for } k \in I_a^0) \\ \lambda_k g_k(a) = 0 \text{ (for } k \in E \cup I) \end{array} \right.$$

either active constraint
or $\lambda_k = 0$

Descent Methods

Descent Methods

General principle

- ① choose an initial point x_0 , set $t = 0$
- ② while not happy
 - choose a **descent direction** $d_t \neq 0$
 - **line search:**
 - choose a step size $\sigma_t > 0$
 - set $x_{t+1} = x_t + \sigma_t d_t$
 - set $t = t + 1$

Remaining questions

- how to choose d_t ?
- how to choose σ_t ?

Gradient Descent

Rationale: $\mathbf{d}_t = -\nabla f(\mathbf{x}_t)$ is a descent direction
indeed for f differentiable

$$\begin{aligned} f(x - \sigma \nabla f(x)) &= f(x) - \sigma \|\nabla f(x)\|^2 + o(\sigma \|\nabla f(x)\|) \\ &< f(x) \text{ for } \sigma \text{ small enough} \end{aligned}$$

Step-size

- optimal step-size: $\sigma_t = \underset{\sigma}{\operatorname{argmin}} f(\mathbf{x}_t - \sigma \nabla f(\mathbf{x}_t))$
- **Line Search:** **total** or partial optimization w.r.t. σ
Total is however often too "expensive" (needs to be performed at each iteration step)
Partial optimization: execute a limited number of trial steps until a loose approximation of the optimum is found. Typical rule for partial optimization: **Armijo rule** (see next slides)

Typical stopping criterium:

norm of gradient smaller than ϵ

The Armijo-Goldstein Rule

Choosing the step size:

- Only to decrease f -value not enough to converge (quickly)
- Want to have a reasonably large decrease in f

Armijo-Goldstein rule:

- also known as backtracking line search
- starts with a (too) large estimate of σ and reduces it until f is reduced enough
- what is enough?
 - assuming a linear f e.g. $m_k(x) = f(x_k) + \nabla f(x_k)^T (x - x_k)$
 - expected decrease if step of σ_k is done in direction \mathbf{d} :
 $\sigma_k \nabla f(x_k)^T \mathbf{d}$
 - actual decrease: $f(x_k) - f(x_k + \sigma_k \mathbf{d})$
 - stop if actual decrease is at least constant times expected decrease (constant typically chosen in $[0, 1]$)

The Armijo-Goldstein Rule

The Actual Algorithm:

Input: descent direction \mathbf{d} , point \mathbf{x} , objective function $f(\mathbf{x})$ and its gradient $\nabla f(\mathbf{x})$, parameters $\sigma_0 = 10$, $\theta \in [0, 1]$ and $\beta \in (0, 1)$

Output: step-size σ

Initialize σ : $\sigma \leftarrow \sigma_0$

while $f(\mathbf{x} + \sigma\mathbf{d}) > f(\mathbf{x}) + \theta\sigma\nabla f(\mathbf{x})^T\mathbf{d}$ **do**

$\sigma \leftarrow \beta\sigma$


end while

Armijo, in his original publication chose $\beta = \theta = 0.5$.

Choosing $\theta = 0$ means the algorithm accepts any decrease.


The Armijo-Goldstein Rule

Graphical Interpretation




The Armijo-Goldstein Rule

Graphical Interpretation



The Armijo-Goldstein Rule

Graphical Interpretation



Newton Algorithm

Newton Method

- descent direction: $-\left[\nabla^2 f(x_k)\right]^{-1} \nabla f(x_k)$ [so-called **Newton direction**]
- The Newton direction:
 - minimizes the best (locally) quadratic approximation of f :
$$\tilde{f}(x + \Delta x) = f(x) + \nabla f(x)^T \Delta x + \frac{1}{2} (\Delta x)^T \nabla^2 f(x) \Delta x$$
 - points towards the optimum on $f(x) = (x - x^*)^T A (x - x^*)$
- however, Hessian matrix is expensive to compute in general and its inversion is also not easy

quadratic convergence

$$\left(\text{i.e. } \lim_{k \rightarrow \infty} \frac{|x_{k+1} - x^*|}{|x_k - x^*|^2} = \mu > 0 \right)$$

Remark: Affine Invariance

Affine Invariance: same behavior on $f(x)$ and $f(Ax + b)$ for $A \in \text{GLn}(\mathbb{R}) = \text{set of all invertible } n \times n \text{ matrices over } \mathbb{R}$

- Newton method is affine invariant

see http://users.ece.utexas.edu/~cmcaram/EE381V_2012F/Lecture_6_Scribe_Notes.final.pdf

- same convergence rate on all convex-quadratic functions
- Gradient method not affine invariant

Quasi-Newton Method: BFGS

$x_{t+1} = x_t - \sigma_t H_t \nabla f(x_t)$ where H_t is an **approximation** of the inverse Hessian

Key idea of Quasi Newton:

successive iterates x_t, x_{t+1} and gradients $\nabla f(x_t), \nabla f(x_{t+1})$ yield second order information

$$q_t \approx \nabla^2 f(x_{t+1}) p_t$$

where $p_t = x_{t+1} - x_t$ and $q_t = \nabla f(x_{t+1}) - \nabla f(x_t)$

Most popular implementation of this idea: **Broyden-Fletcher-Goldfarb-Shanno (BFGS)**

- default in MATLAB's `fminunc` and python's `scipy.optimize.minimize`

Conclusions

I hope it became clear...

...what are the difficulties to cope with when solving numerical optimization problems

in particular dimensionality, non-separability and ill-conditioning

...what are **gradient** and **Hessian**

...what is the difference between **gradient** and **Newton direction**

...and that adapting the step size in descent algorithms is crucial.