

Introduction to Optimization

Lecture 6: Continuous Optimization III (Gradient-based Optimization)

October 30, 2015

TC2 - Optimisation

Université Paris-Saclay

Anne Auger

INRIA Saclay – Ile-de-France



Dimo Brockhoff

INRIA Lille – Nord Europe

Course Overview

Date		Topic
Fri, 18.9.2015	DB	Introduction and Greedy Algorithms
Fri, 25.9.2015	DB	Dynamic programming and Branch and Bound
Fri, 2.10.2015	DB	Approximation Algorithms and Heuristics
Fri, 9.10.2015	AA	Introduction to Continuous Optimization
Fri, 16.10.2015	AA	Introduction to Continuous Optimization II
Fri, 30.10.2015	AA	Gradient-Based Algorithms
Fri, 6.11.2015	AA	Stochastic Algorithms and Derivative-free Optimization
20.11.2015		Exam

all classes + exam are from 14h till 17h15 (incl. a 15min break)
here in PUIO-D101/D103

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

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 $(E, || ||)$ 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

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

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 $t = 1$
- 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

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

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 mid-term exam

Stopping criteria:

norm of gradient smaller than ϵ

Gradient Descent: Simple Theoretical Analysis

Assume f is twice continuously differentiable, convex and that

$\mu I_d \preceq \nabla^2 f(x) \preceq L I_d$ with $\mu > 0$ holds, assume a fixed step-size $\sigma_t = \frac{1}{L}$

Note: $A \preceq B$ means $x^T A x \leq x^T B x$ for all x

$$x_{t+1} - x^* = x_t - x^* - \sigma_t \nabla^2 f(y_t)(x_t - x^*) \text{ for some } y_t \in [x_t, x^*]$$

$$x_{t+1} - x^* = \left(I_d - \frac{1}{L} \nabla^2 f(y_t) \right) (x_t - x^*)$$

$$\begin{aligned} \text{Hence } \|x_{t+1} - x^*\|^2 &\leq \left\| I_d - \frac{1}{L} \nabla^2 f(y_t) \right\|^2 \|x_t - x^*\|^2 \\ &\leq \left(1 - \frac{\mu}{L} \right)^2 \|x_t - x^*\|^2 \end{aligned}$$

$$\text{Linear convergence: } \|x_{t+1} - x^*\| \leq \left(1 - \frac{\mu}{L} \right) \|x_t - x^*\|$$

algorithm slower and slower with increasing condition number

Non-convex setting: convergence towards stationary point

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

Affine Invariance

Affine Invariance: same behavior on $f(x)$ and $f(Ax + b)$ for $A \in \text{GL}_n(\mathbb{R})$

- 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`