Decentralized learning in control and optimization for networks and dynamic games

Alexandre Proutiere

KTH

Objectives

- Cooperative setting
- Provide a survey of recent advances in distributed optimization techniques

$$\min \sum_{i=1}^{m} f_i(x)$$

  - m independent agents **cooperating** towards a single objective
  - How can they reach the desired minimizer?
  - How much should they communicate?
  - How fast can they reach the objective?

Objectives

- Competitive setting
- Provide a survey of recent advances in convergence to Nash Equilibria in games

$$\forall i = 1, \ldots, m, \quad \min_{x_i} f_i(x_i, x_{-i})$$

  - m independent agents **competing** towards different objectives
  - Does the notion of Nash Equilibrium make sense?
  - Are there natural learning algorithms leading to NEs?
  - Can agents / players select socially efficient NEs?
  - How fast can they reach equilibrium?

Applications

- Large networked systems
  - Internet
  - AdHoc networks
  - Data centers
  - Sensor networks
  - Social networks
  - Economic networks
  - …
- New interaction paradigms
  - Resource allocation
  - Coordination
  - Estimation
  - Games over networks
  - …

Decentralized interactions

- We need new tools to understand the way agents interact in these large-scale networked complex systems

- Challenges
  - Lack of central authority
  - Network dynamics
  - Stochastic phenomena
  - Lack of (or partial) local communication among agents
  - …

Concrete examples

- Resource allocation in communication networks
  - Internet Congestion Control
  - Power control in wireless systems
  - Routing
  - Load balancing
- Games
  - Load balancing games
  - Routing games
  - Power control games
  - Marriage problems
  - …

On the way, you'll learn

- How to optimize a function

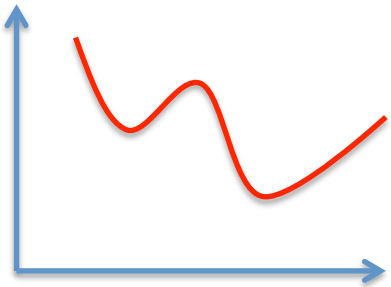

- How to sell your house

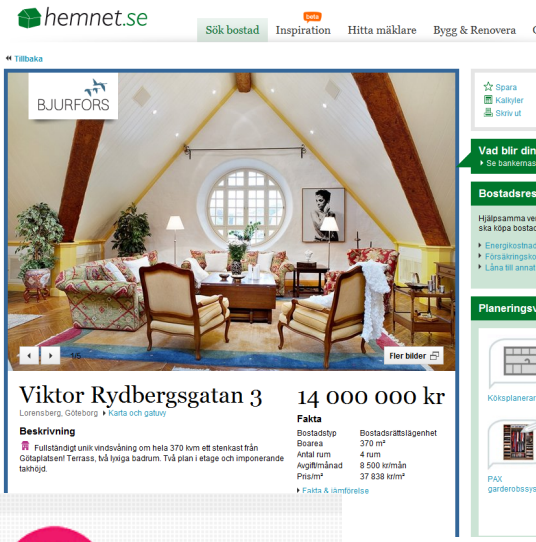

- How to choose your way to work

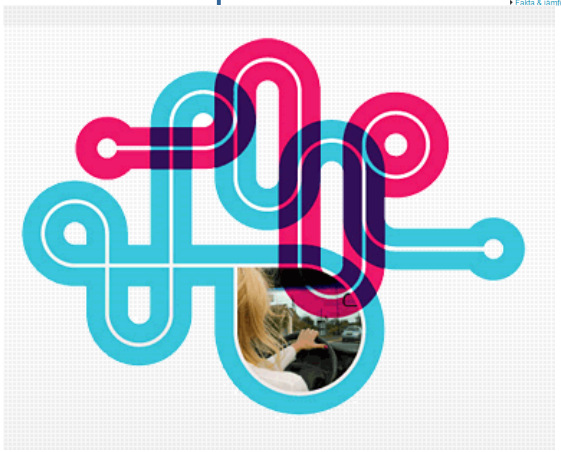

On the way, you'll learn

- How to optimize your portfolio

- How to gamble

- How to select your boy/girl friend

Outline

- Part I: Centralized optimization
- Part II: Distributed optimization
- Part III: Bandits and adversarial optimization
- Part IV: Dynamics in games

Part I

- Gradient-free (or 0^{th} order) methods
  - Hit and run algorithm
  - Oblivious random direct search
  - Random local search
  - Simulated annealing
  - Gradient estimation methods
- Gradient-descent (or 1^{st} order) methods
  - Unconstrained optimization
  - Constrained optimization: duality
- Fixed point iterations

Part II

- Internet congestion control
  - Distributed optimization with separable objective function
- Two miracles in resource allocation in wireless networks
  - Distributed optimization with un-separable objective function, and without message passing
  - Power control
  - Carrier Sensing Multiple Access
- Parallel computations
  - Joint consensus and gradient descent methods
  - Just gradient descent
- Colorings
  - Combinatorial optimization: a sampling approach
- Distributed gradient free optimization

Part III

- Multi-armed bandit problems
  - Notion of regret
  - Stochastic bandits
  - Adversarial bandits
- Stochastic bandit problem
  - IID setting
  - Lower bound on regret
  - UCB policies, finite time analysis
  - Asymptotically optimal policies
- Adversarial bandit problems
  - Models
  - Multiplicative update algorithms

Part IV

- Nash Dynamics
- Fictitious play
- No-regret algorithms
- Trial and error learning

How to reach me?

- By email: alepro@kth.se