Artificial Intelligence Methods for Social Good

M1-1 [Optimization]:

Optimization Problems

08-537 (9-unit) and 08-737 (12-unit)

Instructor: Fei Fang
feifang@cmu.edu
Wean Hall 4126

Outline

- Optimization Problem
- Convex Optimization Problem
- Linear Program (LP)
- (Mixed) Integer Linear Program (MILP)

Learning Objectives

- Understand the concept of
  - Optimization Problem
  - Convex Optimization Problem
    - Convex Function
    - Convex Set
  - Linear Program
  - (Mixed) Integer Linear Program
- Know how to find the method/algorithm/solver/package for different optimization problems
- Briefly describe the following algorithms
  - Gradient Descent Algorithm
  - Simplex Algorithm
- Solve a simple CO/LP/MILP by hand or by calling a solver

Learning Objectives (Alternative View)

- Understand what is highlighted on whiteboard
- Check the slides for reference
- No need to understand all the material in the slides

Optimization Problem: What

- Def 1: Optimization Problem
  - Determine value of optimization variable within feasible region/set to optimize optimization objective
- Def 1.1: Optimal solution
  - Best value of optimization variables
- Def 1.2: Optimal value
  - Best value of optimization objective
- Exp 1

x_i	1.0	2.0	3.5
y_i	2.1	3.98	7.0

$$\min_a \sum_{i=1}^{3} (y_i - a x_i)^2$$
$$\text{s.t. } a \in \mathbb{R}$$

Optimization Problem: How

- No general way to solve
- Many algorithms developed for special classes of optimization problems
  - Convex optimization problem (CO)
  - Linear Program (LP)
  - (Mixed) Integer Linear Program (MILP)
  - Quadratic program (QP)
  - (Mixed) Integer Quadratic program (MIQP)
  - Semidefinite program (SDP)
  - Second-order cone program (SOCP)
- Examples of solvers and code packages for these problems
  - Cplex (LP, MILP, QP)
  - Gurobi (LP, MILP, MIQP)
  - GLPK (LP, MILP)
  - Cvxopt (CO)
  - DSDP5 (SDP)
  - MOSEK (QP, SOCP)
  - Yalmip (SDP)

How to Apply without Going Further

- Lazy mode ☺
- Formulate a problem as an optimization problem
- Identify which class the formulation belongs to
- Call the corresponding solver
- Done!

Then Why Go Further?

- Learn how to identify which class the problem formulation belongs to
- Understand which formulations can be solved more efficiently
- Choose/Convert to the right formulation
- Open the black box to learn key ideas, useful for developing advanced solutions

Convex Optimization Problem: What

- A special case of optimization problem
- Def 2: Convex Optimization Problem
  - An optimization problem whose optimization objective is a convex function and feasible region is a convex set
- Def 3: Convex function
  - Value in the middle point is lower than average value
- Def 4: Convex set
  - Any middle point of two points in the set is also in the set

Convex Optimization Problem: Why

- Why we care about convex optimization problem?
  - Can be solved in a neat way in theory
    - Reliably and efficiently (polynomial time solvable)
  - Easy to solve in practice
    - Commercial and non-commercial solvers / packages
  - Many practical problems are convex
    - Linear regression

Convex Optimization Problem: How

- Can be solved in a neat way in theory
  - Def 5: locally and globally optimal
  - Thm 1: if convex, locally optimal $\rightarrow$ globally optimal
    - Find local optimum is sufficient!
  - Many COs are polynomial time solvable
  - Unconstrained and differentiable
    - Set derivative to be 0
      - Closed form solution
      - Newton's Method (if twice differentiable)
    - Move towards a promising direction
      - Alg 1: Gradient Descent
  - Constrained and differentiable
    - Interior Point Method
  - Non-differentiable
    - ϵ-Subgradient Method
    - Cutting Plane Method

Convex Optimization Problem: How

- Easy to solve in practice
  - Examples: fmincon (MATLAB), cvxpy (Python), cvxopt (Python), cvx (MATLAB)

Convex Optimization Problem: Which

- Which functions are convex?
  - Prove by definition
  - Check secondary derivative
    - Prop 1: if one-dimensional, secondary derivative > 0
    - Prop 2: if multi-dimensional, Hessian matrix is positive semidefinite
  - Check known convex functions
    - Linear, Exponential, L2 norm, Squared L2 norm
  - Use properties
    - Prop 3: Sum of convex functions is convex
    - Prop 4: Convexity is preserved under a linear transformation

Quiz 1: Convex Optimization Problem

- Which of the following functions are convex?
  - A: $f(a) = \sum_{i=1}^{3}(y_i - ax_i)^2$
  - B: $f(x) = x \log x$
  - C: $f(x, y) = -\sqrt{1 - x^2 - y^2}$
  - D: $f(x, y) = x^T A^T y, A \in \mathbb{R}^{m \times n}, x, y \in \mathbb{R}^n$
  - E: $f(x) = x^3$

Convex Optimization Problem: Which

- Which sets are convex?
  - Prove by definition
  - Check known convex sets
    - Region defined by linear inequalities, Unit ball
  - Use properties
    - Prop 5: Union of convex sets is convex

Convex Optimization Problem: Apply

- Model a problem as a convex optimization problem
  - Write it down as an optimization problem
  - Prove it is convex (convex function + convex set)
- Solve the convex optimization problem
  - Build up the model
  - Call the solver
- Map the solution back to the original problem

Learning Objectives

- Understand the concept of
  - Optimization Problem
  - Convex Optimization Problem
    - Convex Function
    - Convex Set
  - Linear Program
  - (Mixed) Integer Linear Program
- Know how to find the method/algorithm/solver/package for different optimization problems
- Briefly describe the following algorithms
  - Gradient Descent Algorithm
  - Simplex Algorithm
- Solve a simple CO/LP/MILP by hand or by calling a solver

Convex Optimization: Additional Resources

- Text book
  - *Convex Optimization, Chapters 1-4*
    Stephen Boyd and Lieven Vandenberghe
    Cambridge University Press
    https://web.stanford.edu/~boyd/cvxbook/
- Online course
  - Stanford University, Convex Optimization I (EE 364A), taught by Stephen Boyd
    - http://ee364a.stanford.edu/courseinfo.html
    - https://youtu.be/McLq1hEq3UY

Outline

- Optimization Problem
- Convex Optimization Problem
- Linear Program (LP)
- (Mixed) Integer Linear Program (MILP)

Linear Program: What

- A special case of convex optimization problem
- Def 6: Linear Program
  - An optimization problem whose optimization objective is a linear function and feasible region is defined by linear constraints
- Exp 2: Maximize Profit

	Price	Labor	Machine
Product 1	$30	0.2 hour	4 hour
Product 2	$30	0.5 hour	2 hour
Total		<=90	<=800

Quiz 2: Linear Program

- In Exp 2, which constraints determine point B?
  - $x = 0$
  - $y = 0$
  - $0.2x + 0.5y = 90$
  - $4x + 2y = 800$

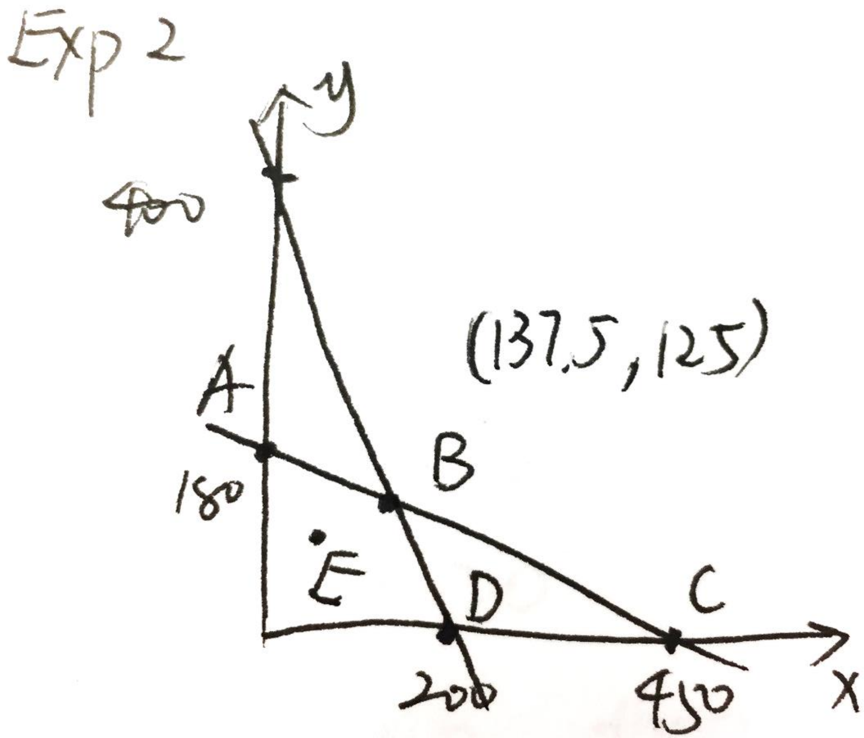

Quiz 3: Linear Program

- In Exp 2, which point is the optimal solution?
  - A
  - B
  - C
  - D
  - E

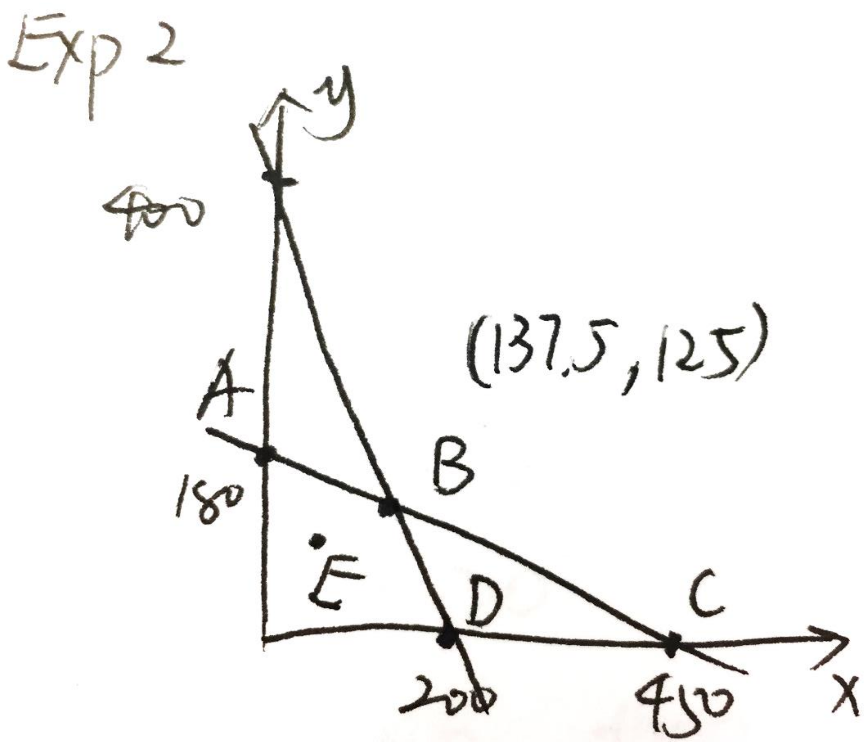

Linear Program: Why

- Polynomial time solvable
- Efficient solvers available
- Many applications

Linear Program: How

- At least one vertex of the polytope is an optimal solution (why?)
- Naïve approach: Enumerate vertices
  - Def 7: Standard form LP
    - An LP with only equality constraints on non-negative variables
  - Select n constraints (why?)
  - Solve linear system
  - Check other constraints
- Ellipsoid algorithm
  - Polynomial-time in theory, poor performance in practice
- Simplex Algorithm
  - Iteratively move to a neighboring vertex with better objective value
  - Alg 2

Linear Program: How

- Solve LPs in practice:
  - All the solvers/algorithms for Convex Optimization problems can be applied
  - Additional solvers/algorithms for LPs
    - linprog (MATLAB), linprog (in Python package SciPy)
    - PuLP (Python)
    - Cplex, Gurobi
    - https://www.informs.org/ORMS-Today/Public-Articles/June-Volume-38-Number-3/Software-Survey-Linear-Programming

Learning Objectives

- Understand the concept of
  - Optimization Problem
  - Convex Optimization Problem
    - Convex Function
    - Convex Set
  - Linear Program
  - (Mixed) Integer Linear Program
- Know how to find the method/algorithm/solver/package for different optimization problems
- Briefly describe the following algorithms
  - Gradient Descent Algorithm
  - Simplex Algorithm
- Solve a simple CO/LP/MILP by hand or by calling a solver

Linear Program: Additional Resources

- Text book
  - *Applied Mathematical Programming, Chapters 2-4*
  - By Bradley, Hax, and Magnanti (Addison-Wesley, 1977)
  - http://web.mit.edu/15.053/www/AMP.htm
- Online course
  - https://ocw.mit.edu/courses/electrical-engineering-and-computer-science/6-251j-introduction-to-mathematical-programming-fall-2009/index.htm

Outline

- Optimization Problem
- Convex Optimization Problem
- Linear Program (LP)
- (Mixed) Integer Linear Program (MILP)

(Mixed) Integer Linear Program: What

- Def 8: Mixed Integer Linear Program (MILP)
  - An optimization problem whose optimization objective is a linear function and feasible region is defined by linear constraints and integer constraints
- Non-convex
- Exp 3 (0-1 Knapsack)

(Mixed) Integer Linear Program: Why

- Super powerful!
- Will be used again and again throughout the course
- Many heuristic algorithms and solvers developed

(Mixed) Integer Linear Program: How

- General case: MILP is NP-Complete
  - Runtime is exponential
- Naïve approach
  - Enumerate values for integer variables, then solve LP
- LP relaxation
  - Remove integer constraint
  - Can often lead to optimal solutions
    - Check total unimodularity
  - Can provide a upper bound (for minimization problem)
    - Larger feasible region
  - Can provide a heuristic solution through proper rounding
- Branch and bound
  - Combine the two: create a binary search tree for integer variables, use bounds to prune the branches (LP relaxation and rounding), expand a non-integer variable in the most promising node (based on LP relaxation)
- Cutting Plane
  - Add constraint to exclude the currently best non-integer solution

(Mixed) Integer Linear Program: How

- Practically efficient solvers: Cplex, Gurobi, intlinprog (MATLAB), SCIP solver

(Mixed) Integer Linear Program: Additional Resources

- Text book
  - *Applied Mathematical Programming, Chapter 9*
  - By Bradley, Hax, and Magnanti (Addison-Wesley, 1977)
  - http://web.mit.edu/15.053/www/AMP.htm
- Online course
    - https://ocw.mit.edu/courses/sloan-school-of-management/15-083j-integer-programming-and-combinatorial-optimization-fall-2009/index.htm

Learning Objectives

- Understand the concept of
  - Optimization Problem
  - Convex Optimization Problem
    - Convex Function
    - Convex Set
  - Linear Program
  - (Mixed) Integer Linear Program
- Know how to find the method/algorithm/solver/package for different optimization problems
- Briefly describe the following algorithms
  - Gradient Descent Algorithm
  - Simplex Algorithm
- Solve a simple CO/LP/MILP by hand or by calling a solver

Summary

- Optimization Problem
- Convex Optimization Problem (CO)
- Linear Program (LP)
- (Mixed) Integer Linear Program (MILP)

- Key take-away:
  - Model a problem as an optimization problem, identify which class it belongs to and call the corresponding solver
  - CO/LP can be efficiently solved in theory and practice
  - MILP can often be efficiently solved in practice

Summary

Optimization Problem	Key Algorithm In Class	In Practice
Convex Optimization Problem (CO)	Gradient descent	Call CO solvers
Linear Program (LP)	Simplex	Call LP solvers
(Mixed) Integer Linear Program (MILP)		Call MILP solvers
Continuous differentiable nonconvex optimization problem		Call CO solver that finds local optimum many times with random restarts