Homework #4
CS675 Fall 2026

Due Friday Oct 30, by 5pm

General Instructions The following assignment is meant to be challenging. Feel free to discuss with fellow students, though please write up your solutions independently and acknowledge everyone you discussed the homework with on your writeup. I also expect that you will not attempt to consult outside sources, including generative AI models, for solutions to any of these homework problems — doing so would be considered cheating. Please see the course website for details on the academic integrity policy.

Several of these problems are drawn from the following texts, each of which is linked on the course website: Luenberger and Ye (4th edition), Korte and Vygen (5th edition), and Boyd and Vendenberghe. Please make sure you are using the correct edition of each of the books by using the links on the course website.

We request that you submit your homework as a pdf file, by email to the TA.

Finally, whenever a question asks you to "show" or "prove" a claim, please provide a formal mathematical proof.

Problem 1. (14 points)
In this problem we will establish some basic properties of polyhedra and linear programs. Let $\mathcal{P} \subseteq \mathbb{R}^n$ be a nonempty polyhedron, let $V \subseteq \mathcal{P}$ be the vertices of $\mathcal{P}$, let $OPT(c) = \max_{x \in \mathcal{P}} c^\intercal x$ for each $c \in \mathbb{R}^n$, and let $OPTSOL(c) = \operatorname{argmax}_{x \in \mathcal{P}} c^\intercal x$ for each $c \in \mathbb{R}^n$.

a (2 points). Fix $x \in V$, and let $C_x = \{c \in \mathbb{R}^n : x \in OPTSOL(c)\}$. Note that C_x is the family of objectives for which x is an optimal solution. Show that C_x is a convex cone.

b (4 points). Fix $x \in V$. Show that there is a linear objective $c \in \mathbb{R}^n$ such that $OPTSOL(c) = \{x\}$.

c (4 points). Let $C_\infty := \{c \in \mathbb{R}^n : OPT(c) = \infty\}$. Note that C_∞ is the family of objectives for which the LP is unbounded. Clearly, the origin $\vec{0}$ is not in C_∞, so C_∞ is not a cone. However, is $C_\infty \cup \left\{\vec{0}\right\}$ a cone? Is it a convex cone? Prove or disprove.

d (4 points). Let $C_V := \{c \in \mathbb{R}^n : OPT(c) < \infty\}$. Note that C_V is the family of objectives for which the LP is bounded. Is C_V a cone? Is it a convex cone? Prove or disprove.

Problem 2. (8 points)
Recall the single-source shortest path problem presented in class, where m denotes the number of edges of a graph and n denotes the number of vertices. Recall that the Bellman-Ford algorithm runs in time $O(mn)$, and computes the shortest path from the given source s to every other vertex in the graph. Moreover, the classical Dijkstra's algorithm solves the same problem on graphs with nonnegative-weighted edges in time $O(m + n \log n)$. Show to combine both algorithms to solve the *all-pairs shortest path* problem, in graphs with arbitrary edge weights but no negative cycles, in time $O(mn + n^2 \log n)$ time. In this problem, you must compute a shortest path between every pair of vertices in the graph.
Hint: Run Bellman-Ford once and Dijkstra $O(n)$ times. Use the dual.

Problem 3. (8 points)
Let $G = (L \cup R, E)$ be a bipartite graph with left hand side nodes L, right hand side nodes R, and $|L| = |R| = n$. As shorthand, for $S \subseteq L$ we use $\mathcal{N}(S) = \{v \in R : \exists u \in S \text{ s.t. } (u, v) \in E\}$ to denote the neighbors of S. Use König's theorem from class to prove *Hall's theorem*: G has a perfect matching (i.e. a matching with n edges) if and only if $|\mathcal{N}(S)| \geq |S|$ for all $S \subseteq L$.

Problem 4. (12 points)
Let $G = (V, E)$ be a bipartite graph, and let $\mathcal{M} \subseteq 2^E$ be the family of perfect matchings. Suppose that $\mathcal{M} \neq \emptyset$. You are given $\theta \in [0, 1]^E$ in the relative interior of the perfect matching polytope — i.e. you are gauranteed that there exists some probability distribution p over matchings with $p_M > 0$ for all $M \in \mathcal{M}$, and moreover $\sum_{M \in \mathcal{M}: e \in M} p_M = \theta_e$ for all $e \in E$. In fact, there are many such distributions, and you wish to find one with the most entropy.

(a) [6 points]. Consider the following convex program for finding the maximum entropy distribution over matchings with marginal probabilities θ. We restrict the domain to $p \succ 0$.

$$\begin{array}{ll} \text{maximize} & \sum_{M \in \mathcal{M}} p_M \log \frac{1}{p_M} \\ \text{subject to} & \sum_{M \in \mathcal{M}: e \in M} p_M = \theta_e, \quad \text{for } e \in E. \\ & \sum_{M \in \mathcal{M}} p_M = 1 \end{array}$$

This program has impractically many variables, so it is instructive to consider its dual instead. Derive the dual of this convex program. Your dual should have many fewer variables, on the order of the number of edges in the graph.

(b) [6 points]. Let p^* be the maximum entropy distribution over matchings with marginal probabilities θ. Show that p^* exhibits the following product structure: There exist weights $w_e > 0$ on the edges $e \in E$ such that p^*_M is proportional to $\prod_{e \in M} w_e$. Show how to obtain w from an optimal solution to your dual program.

Problem 5. (16 points)
We have noted in class that convex objective functions have a desirable property, one which is required by efficient convex programming solvers. Namely, one can implement a separation oracle for any sub-level set of f by simply evaluating the gradient (or, more generally, a subgradient) of f at a suitable point. One might wonder if the same property holds for the more general class of *quasi-convex functions*: functions whose sub-level sets are all convex. It turns out that this is

not the case in general, which explains why quasi-convex optimization problems are not always tractable.

(a) [2 points]. Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex, and assume that you can efficiently evaluate $f(x)$ and $\bigtriangledown f(x)$ at any point $x \in \mathbb{R}^n$. Given a threshold α, show how to implement a separation oracle for the set $S(\alpha) = \{x \in \mathbb{R}^n : f(x) \leq \alpha\}$.

(b) [4 points]. Exhibit a continuous and quasi-convex function f, a threshold α, and a point $x \notin S(\alpha)$ such that $\bigtriangledown f(x) = 0$ (i.e. the gradient of f provides no information on how to separate x from $S(\alpha)$).

(c) [10 points]. Describe a continuous quasi-convex function f for which you can efficiently evaluate f and $\bigtriangledown f$, yet for which minimizing $f(x)$ is computationally hard. You may reduce from the special case of the boolean satisfiability problem SAT, often referred to as UNIQUE-SAT, in which you are promised that there is at most one satisfying assignment for every input formula. Note: UNIQUE-SAT is known to be NP-hard under randomized polynomial-time reductions (Valiant-Vazirani theorem).