ADFOCS 2015

Exercises on Maximizing Submodular Functions

August 17-21, 2015

1. In the *maximum coverage problem*, we have a set of elements E, and m subsets of elements $S_1, \ldots, S_m \subseteq E$, each with a nonnegative weight $w_j \geq 0$. The goal is to choose k elements such that we maximize the weight of the subsets that are covered. We say that a subset is covered if we have chosen some element from it. Thus we want to find $S \subseteq E$ such that $|S| = k$ and that we maximize the total weight of the subsets j such that $S \cap S_j \neq \emptyset$.

   (a) Give a $(1 - \frac{1}{e})$-approximation algorithm for this problem.

   (b) (Not easy) Show that if an approximation algorithm with performance guarantee better than $1 - \frac{1}{e} + \epsilon$ exists for the maximum coverage problem for some constant $\epsilon > 0$, then $P = NP$. (Hint: For this problem, you need to know that there is no $(c \ln n)$-approximation algorithm for $c < 1$ for the set cover problem unless $P = NP$.)

2. A *matroid* $(E, \mathcal{I})$ is a set E of ground elements together with a collection $\mathcal{I}$ of subsets of E; that is, if $S \in \mathcal{I}$, then $S \subseteq E$. A set $S \in \mathcal{I}$ is said to be *independent*. The independent sets of a matroid obey the following two axioms:

   - If S is independent, then any $S' \subseteq S$ is also independent.
   - If S and T are independent, and $|S| < |T|$, then there is some $e \in T - S$ such that $S \cup \{e\}$ is also independent.

   An independent set S is a *base* of the matroid if no set strictly containing it is also independent.

   (a) Given an undirected graph $G = (V, E)$, show that the forests of G form a matroid; that is, show that if E is the ground set, and $\mathcal{I}$ the set of forests of G, then the matroid axioms are obeyed.

   (b) Show that for any matroid, every base of the matroid has the same number of ground elements.

   (c) For any given matroid, suppose that for each $e \in E$, we have a nonnegative weight $w_e \geq 0$. Give a greedy algorithm for the problem of finding a maximum-weight base of a matroid.

3. Let $(E, \mathcal{I})$ be a matroid as defined above, and let f be a monotone, submodular function such that $f(\emptyset) = 0$. Consider the following local search algorithm for finding a maximum-value base of the matroid: First, start with an arbitrary base S. Then consider all pairs $e \in S$ and $e' \notin S$. If $S \cup \{e'\} - \{e\}$ is a base, and $f(S \cup \{e'\} - \{e\}) > f(S)$, then set $S \leftarrow S \cup \{e'\} - \{e\}$. Repeat until a locally optimal solution is reached.

The goal of this problem is to show that a locally optimal solution has value at least half the optimal value.

(a) We begin with a simple case: suppose that the matroid is a *uniform matroid*; that is, $S \subseteq E$ is independent if $|S| \leq k$ for some fixed k. Prove that for a locally optimal solution S, $f(S) \geq \frac{1}{2}\,\mathrm{OPT}$.

(b) To prove the general case, it is useful to know that for any two bases of a matroid, X and Y, there exists a bijection $g : X \to Y$ such that for any $e \in X$, $S - \{e\} \cup \{g(e)\}$ is independent. Use this to prove that for any locally optimal solution S, $f(S) \geq \frac{1}{2}\,\mathrm{OPT}$.

4. We showed that a randomized greedy algorithm could be used to obtain a $\frac{1}{2}$-approximation algorithm for maximizing nonnegative, nonmonotone submodular functions. Here we consider a deterministic variant of that algorithm, shown below. Recall that we defined the function f over the set $\{1, \ldots, n\}$, and defined $\hat{X}_i \equiv X_i \cup \{i+1, \ldots, n\}$. For your proof, you may find it useful again to consider $\mathrm{OPT}_i = X_i \cup (\mathrm{OPT} \cap \{i+1, \ldots, n\})$, where OPT is an optimal set. Recall that for randomized algorithm we showed that

$$E[f(OPT_{i-1}) - f(OPT_i)] \leq \frac{1}{2} E[f(X_i) - f(X_{i-1}) + f(\hat{X}_i) - f(\hat{X}_{i-1})].$$

$X_0 \leftarrow \emptyset$
for $i \leftarrow 1$ to n **do**
  $a_i \leftarrow f(X_{i-1} \cup \{i\}) - f(X_{i-1})$
  $r_i \leftarrow f(\hat{X}_{i-1} - \{i\}) - f(\hat{X}_{i-1})$
  if $a_i \geq r_i$ **then**
    $X_i \leftarrow X_{i-1} \cup \{i\}$
  else
    $X_i \leftarrow X_{i-1}$
return X_n

(a) Prove that it gives a $\frac{1}{3}$-approximation algorithm for the problem, for f a nonnegative, nonmonotone submodular function. What inequality leads to a $\frac{1}{3}$-approximation algorithm?

(b) (Challenge problem) Show that if f is symmetric (that is, $f(X) = f(E - X)$ for all X), then the same algorithm gives a $\frac{1}{2}$-approximation algorithm.