Abstract

Suppose we are given a submodular function f over a set of elements, and we want to maximize its value subject to certain constraints. Good approximation algorithms are known for such problems under both monotone and non-monotone submodular functions. We consider these problems in a stochastic setting, where elements are not all active and we can only get value from active elements. Each element e is active independently with some known probability p_e, but we don't know the element's status *a priori*. We find it out only when we *probe* the element e—probing reveals whether it's active or not, whereafter we can use this information to decide which other elements to probe. Eventually, if we have a probed set S and a subset $active(S)$ of active elements in S, we can pick any $T \subseteq active(S)$ and get value $f(T)$. Moreover, the sequence of elements we probe must satisfy a given *prefix-closed constraint*—e.g., these may be given by a matroid, or an orienteering constraint, or deadline, or precedence constraint, or an arbitrary downward-closed constraint—if we can probe some sequence of elements we can probe any prefix of it. What is a good strategy to probe elements to maximize the expected value? In this paper we study the gap between adaptive and non-adaptive strategies for f being a submodular or a fractionally subadditive (XOS) function. If this gap is small, we can focus on finding good non-adaptive strategies instead, which are easier to find as well as to represent. We show that the adaptivity gap is a constant for monotone and non-monotone submodular functions, and logarithmic for XOS functions of small *width*. These bounds are nearly tight. Our techniques show new ways of arguing about the optimal adaptive decision tree for stochastic problems.