Discrete Applied Mathematics 186 (2015) 112–127

Contents lists available at ScienceDirect

Discrete Applied Mathematics

journal homepage: www.elsevier.com/locate/dam

Informative path planning as a maximum traveling salesman problem with submodular rewards[☆]

Syed Talha Jawaid, Stephen L. Smith [*]

Department of Electrical and Computer Engineering, University of Waterloo, Waterloo ON, N2L 3G1, Canada

ARTICLE INFO

Article history:
Received 31 July 2013
Received in revised form 30 August 2014
Accepted 5 January 2015
Available online 2 February 2015

Keywords:
Submodular function maximization
Traveling salesman problem
Informative path planning

ABSTRACT

In this paper we extend the classic problem of finding the maximum weight Hamiltonian cycle in a graph to the case where the objective is a submodular function of the edges. We consider a greedy algorithm and a 2-matching based algorithm, and we show that they have approximation factors of $\frac{1}{2+\kappa}$ and $\max\{\frac{2}{3(2+\kappa)}, \frac{2}{3}(1-\kappa)\}$ respectively, where κ is the curvature of the submodular function. Both algorithms require a number of calls to the submodular function that is cubic to the number of vertices in the graph. We then present a method to solve a multi-objective optimization consisting of both additive edge costs and submodular edge rewards. We provide simulation results to empirically evaluate the performance of the algorithms. Finally, we demonstrate an application in monitoring an environment using an autonomous mobile sensor, where the sensing reward is related to the entropy reduction of a given set of measurements.

1. Introduction

The maximum weight Hamiltonian cycle is a classic problem in combinatorial optimization. It consists of finding a path in a graph that starts and ends at the same vertex and visits all other vertices exactly once while maximizing the sum of the weights (i.e., the reward) on the edges traversed. Also referred to as the max-TSP, the problem is NP-hard; however, a number of approximation algorithms have been developed. In [10] four simple approximation algorithms are analyzed. The authors show that greedy, best-neighbor, and 2-interchange heuristics all give a $\frac{1}{2}$ approximation to the optimal tour. They also show that a 2-matching based heuristic, which first finds a perfect 2-matching and then converts that to a tour, gives a $\frac{2}{3}$ approximation. The simple and elegant Serdyukov's algorithm [31] – which combines a perfect 2-matching and a 1-matching to compute a tour – gives a $\frac{3}{4}$ approximation. The best known deterministic algorithm is given in [29] and it achieves a $\frac{7}{9}$ approximation in $O(n^3)$ time. A number of randomized algorithms also exist such as the one given in [16] that achieves a $\frac{25}{33}$ approximation ratio. In [37] it was shown that this algorithm can be derandomized while maintaining its approximation factor. In this paper we look at extending the max-TSP problem to the case of submodular rewards.

This extension is motivated, in part, by the application of mobile sensing robots to persistently monitor a large environment [33] such as for obtaining data on natural phenomenon or for security. It is desirable to have a closed walk or a tour over which the sensing robot travels to minimize re-planning during long term operation. Applications include tasks such as monitoring oil spills [4], forest fires [3] or underwater ocean monitoring [34].

☆ This research was partially supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

* Corresponding author.
E-mail addresses: stjawaid@uwaterloo.ca (S.T. Jawaid), stephen.smith@uwaterloo.ca (S.L. Smith).

http://dx.doi.org/10.1016/j.dam.2015.01.004

Informative path planning involves using preexisting information about the environment (such as a probability distribution) to plan a path that maximizes the information gained. It is a topic that has been researched with various different approaches. For example, in [23] the authors look into creating trajectories to best estimate a Gaussian random field by intelligently generating rapidly-exploring random cycles. One common way to measure information content is using mutual information. In [14], this metric is used to place static sensors in a Gaussian field. Other papers investigate maximizing the knowledge at specific points by planning a path for a sensing robot while also taking into account budget constraints [32,1]. The metrics used to determine the quality of the sensing, such as mutual information, are usually submodular in nature. The defining property of a submodular function is that of decreasing marginal value, i.e., adding an element to a set will result in a larger increase in value than adding it to a superset of that set. For example, if a sensor is placed close to another, then the benefit gained by the second sensor will be less than if the first sensor had not already been placed. Other areas where submodular functions arise include viral marketing, active learning [12] and AdWords assignment [13].

Our problem can be stated as maximizing a submodular function over the edges of a graph subject to the constraint that the selected edges form a tour. Generally, one way to represent constraints in a combinatorial optimization problem is through the concept of independence systems or its many specializations, including p-systems and matroids. Although unconstrained minimization of a submodular function can be achieved in polynomial time [30,18], maximizing a non-decreasing submodular function over an independence system constraint is known to be NP-hard. For a monotone submodular function, a number of approximation algorithms exist for optimizing over multiple matroid constraints [11]. Some bounds that include the dependence on curvature are presented in [5]. Local search methods have been found to be particularly useful [36,9] for both monotone and non-monotone functions. Various results exist for non-monotone submodular function maximization for both the unconstrained case [8] and for the case where constraints exist, such as multiple matroid and knapsack constraints [24,25] or p-system constraints [15]. The use of continuous relaxations of the submodular function have also lead to optimal approximation bounds [2] for the case of a single matroid as well as improved bounds for a combination of various constraints [35].

Contributions: The contributions of this paper are to present and analyze two algorithms for constructing a maximum reward tour on a graph. The metric used in maximizing the "reward" of a particular tour is a positive monotone submodular function of the edges. We frame this problem as an optimization over an independence system constraint and present two approximation algorithms. The first method is greedy and gives a $\frac{1}{2+\kappa}$ approximation. The second method creates a 2-matching and then turns it into a tour. This gives a $\max\left\{\frac{2}{3(2+\kappa)}, \frac{2}{3}(1-\kappa)\right\}$ approximation where κ is the curvature of the submodular function. Both techniques require $O(|V|^3)$ value oracle calls to the submodular function. The algorithms are also extended to directed graphs. To obtain these results, we establish a few new properties of submodular maximization, including the generalization of a known bound on the performance of the greedy algorithm to functions with curvature κ (Theorem 2.9), and a result on the minimum possible value of disjoint subsets within a function's base set (Theorem 3.5).

We also present an approach for the case of a multi-objective optimization consisting of submodular (sensing) rewards on the edges along with modular (travel) costs. We incorporate these two objectives into a single function which is no longer monotone nor positive. We provide bounds on the performance of our algorithms in this case, which depend on the relative weight of the rewards.

A preliminary version of this work was presented in [19]. This paper substantially improves upon those results by providing complete proofs of the approximation ratios, a discussion on extending Serdyukov's algorithm, an approach for multi-objective optimization with weights and costs, simulations demonstrating dependence on curvature as well as an application in informative path planning.

Organization: The organization of this paper is as follows. In Section 2 we review independence systems and submodularity and we formalize our problem. In Section 3 we analyze two different strategies for approximating a solution; namely, a greedy approach and a 2-matching based tour. We also discuss extending Serdyukov's algorithm. In Section 4 we extend our algorithms to the case where the graph is directed. In Section 5 we discuss a method to incorporate costs into the optimization. Finally, we provide some simulation results in Section 6 and demonstrate an example application in monitoring.

2. Preliminaries and problem formulation

Here we present preliminary concepts and formulate the maximum submodular TSP problem. We give a brief summary of results on combinatorial optimization over independence systems and we extend a known result on optimization over p-systems to incorporate curvature.

2.1. Independence systems

Combinatorial optimization problems can often be formulated as the maximization or minimization of an objective function $f : \mathcal{F} \to \mathbb{R}$ over a set system $(E, \mathcal{F})$, where E is the base set of all elements and $\mathcal{F} \subseteq 2^E$. An **independence system** is a set system that is closed under subsets (i.e., if $A \in \mathcal{F}$ then $B \subseteq A \implies B \in \mathcal{F}$). Sets in $\mathcal{F}$ are referred to as **independent sets**. Maximal independent sets (i.e., all $A \in \mathcal{F}$ such that $A \cup \{x\} \notin \mathcal{F}, \forall x \in E \setminus A$) are the **bases**.