VEHICLE ROUTING UNDER CORRELATED REWARDS WITH AN APPLICATION TO FISH TRAWLING

Fahrettin ÇAKIR*

Abstract

Routing problems where a vehicle travels through a network to collect rewards are known as orienteering problems. In these types of problems, typically the objective is to maximize the total collected reward. Certain real life problems in this class exhibit uncertainty and spatial and possibly temporal correlation. We develop a methodology that handles correlated rewards in stochastic orienteering problems. We model the problem as a Markov Decision Process and use the Approximate Dynamic Programming framework to generate routing solutions using a variable neighborhood search technique. To model uncertainty, we use various predictive methods ranging from neural networks to Kriging (geospatial Gaussian interpolation). We test our methodology on a fishing dataset for the Baltic Sea. We find that as the predictive performance of a method increases, so does the average reward collected from a fishing trip. However, this relationship is not linear, but it is nonlinear. We suggest that under stochastic rewards, spending time building highly accurate prediction methods is valuable to increase the profit generated from a vehicle routing trip.

Keywords: Stochastic Orienteering, Fish Trawling, Spatiotemporal Correlation

Özet

Kar toplama amacıyla belirlenen güzergahlar, güzergah belirleme problemleri sınıfında önemli bir yer teşkil ediyor. Gerçek hayatta ise bilinmezliklerden kaynaklanan değişkenler mekansal ve zaman olarak korelasyon göstermekte. Bu tür şartlarda kullanılabilecek bir metodoloji sunuyoruz bu araştırmamızda. Problemimizi Markov Karar Süreci olarak modelleyip, Approximate Dinamik Programlama teknikleri ile çözüyoruz. Parametre değişkenliğini ise çeşitli öngörü modelleri kullanarak test ediyoruz. Bulgularımıza göre öngörü modelinin kesinliği ile optimizasyon sonucu bulunan güzergah değeri arasında pozitif ama lineer olmayan bir ilişki mevcut.

Anahtar Kelimeler: Kar Toplama Amacıyla Güzergah Belirleme, Balıkçılık, Mekansal Ve Zamana Dayalı Korelasyon

Introduction

We address the problem of how to route a fishing vessel as it trawls for fish over a number of fishing areas. It can be considered as an orienteering problem in which a vehicle traverses a network to maximize its reward or expected reward. Specifically, a fishing trawler departs from an origin port, visits several fishing grounds, and lands its catch at a destination port. At each period, the fishing trawler needs to decide whether to keep fishing in the current fishing area or move onto a new one. We assume the fishing trawler harvests one or more species over several fishing grounds. The harvest is uncertain but we have historical data with which to estimate it. The per unit value of each species is known, pl. The objective of the fishing trawler is to maximize the expected total value of catch over all species by traversing between the fishing grounds starting from an origin port to a destination port, while respecting time and capacity constraints.

A common solution approach to handle uncertainty is to use an a priori solution approach (fixed-route policy). This approach produces a routing solution that is easier to compute for large scale routing problems. While an a priori solution approach has its benefits, a dynamic solution approach would constitute an improvement over fixed-route policies because it enables modification of the routing plan as observations (catches) are made. There has been recent work (Goodson et al., 2013) that develop dynamic solutions based on fixed-route policies in a rollout

* Dr. Öğr. Üye., Istanbul Sabahattin Zaim University/Department of Business Administration, Istanbul/Turkey, fahrettin.cakir@izu.edu.tr

framework. Rollout procedures are heuristic solution approaches that select the next best action at a current state using an estimate of the rewards-to-go for the set of states approachable from this current state. We employ their solution methodology and apply it to our stochastic orienteering problem.

Rollout algorithms rely on the valuation of the future at every step of the decision process. Further, this valuation depends on how uncertainty is quantified. Modeling uncertainty is key in these types of problems as there will be a discrepancy between how well a routing solution can perform under the assumptions of the model versus how well it performs when faced with real data. Using raw data, various statistical and data mining (data science) methods can be employed to aid the rollout algorithm. We explore several data analysis methods suitable for our data and validate the correct approach that creates higher value.

In our solution approach, we incorporate data by using forecasting models built using historical catch data. We integrate previous observations made during a fishing trip into the decision of where to harvest for fish.

This work makes several contributions including taking into account correlation of rewards in the field of stochas tic orienteering problems. We model the correlation into an MDP formulation by incorporating previous observations as a state variable. We propose short term forecasting models for the South Baltic sea cod harvesting. We shed light on the relationship between the quality of a prediction model and quality of routing obtained as a result.

Orienteering problems are class of problems closely related to vehicle routing problems. Roughly speaking, routing is carried out to maximize total reward in orienteering as opposed to minimizing cost in vehicle routing problems.

The closest study to ours is the work by (Ilhan et al., 2008). The authors consider the problem of finding a tour that visits a subset of nodes on a graph within a prespecified time limit and maximizes the probability of collecting more than a certain threshold level of profit. Each node is associated with a random profit (reward) that is normally distributed. The authors use a bi-objective genetic algorithm to solve their model. In addition to the fact that our problem does not make normality assumptions on the rewards, we make dynamic decision on a graph with correlated rewards between nodes.

There are also other studies that consider orienteering problems with stochastic features. (Campbell et al., 2011) introduce a problem with stochastic travel and service times. The reward function in their work is related to whether or not a node can be visited by a deadline. As a result, stochastic travel and service times lead to uncertainty in the total reward that can be accrued. They present computational results and characterize optimal solutions under certain assumptions.

Various studies have proposed dynamic models of decision making. (Babcock et al., 2000) present a model where bottom trawlers target different species under management-imposed limits on landings for each species. The paper focuses on the effect of trip limits on the choice of fishing strategy in terms of targeting bottom rockfish or deepwater Dover sole. (Dorn, 2001) focuses on the sequential decision making process for a fleet of vessels based in Seattle, WA, operating in the Pacific hake fishery. By modeling the dynamics for the Pacific hake fishery, the author develops a dynamic sequential decision making model in terms of a Markov Decision Process. The author models the sequence of decisions which include what threshold prey density should be used to decide to leave an area, how observations can be combined into the decision to harvest an area, and how much time should be spent searching an area before deciding to leave or not. It is assumed that areas exhibit independent population dynamics unlike a correlated reward structure such as ours. Moreover, the author assumes the vessel moves only to adjacent areas if it decides to fish in a different area. We do not have such a restriction. In order to model the decision to move to a new area, the author updates the mean catch rate of all areas according to a Kalman filter update process from the observed catch rates. The Kalman filter is an update process that combines previous observations with new ones. Since areas exhibit independent dynamics, the Kalman filter is applied to each area independently, hence there is no correlation assumed to exist between mean catch rates between areas. The author uses threshold based decision rules to move to a new area: if the abundance, measured in mean catch rate, is lower than a threshold, the vessel leaves for the adjacent fishing area. The same rule applies to start fishing in an area after searching produces an estimate for an area. These decision rules are justified by way of the marginal value theorem in optimal foraging theory. (Dorn, 2001) carries out numerical search over parameters that govern the decision rules previously discussed to find the combination of parameters that maximize daily revenue using simulation. His results suggest that the reward surface is flat in the parameter neighborhood of the maximum reward. When new observations have a relatively lower weight than previous observations in updating the state, then a high catch rate threshold is a poor strategy because the vessel is always seeking high density areas, but never believes information (observations) that indicate a high mean catch rate. When new observations have a relatively higher weight than previous observations in updating the state, then a low catch rate threshold does not substantially reduce the average reward because the vessel does not act upon information from fishing and searching until it indicates that fish density is very low.

References

Justin C Goodson, Jeffrey W Ohlmann, and Barrett W Thomas. Rollout policies for dynamic solutions to themultivehicle routing problem with stochastic demand and duration limits. Operations Research, 61(1):138–154,2013.

Taylan Ilhan, Seyed M. R. Iravani, and Mark S. Daskin. The orienteering problem with stochastic profits. IIE Transactions, 40(4):406–421, February 2008. 00018.

Ann M Campbell, Michel Gendreau, and Barrett W Thomas. The orienteering problem with stochastic traveland service times. Annals of Operations Research, 186(1):61–81, June 2011

Elizabeth A Babcock and Ellen K Pikitch. A dynamic programming model of fishing strategy choice in amultispecies trawl fishery with trip limits. Canadian Journal of Fisheries and Aquatic Sciences, 57(2):357–370, February 2000.

Martin W Dorn. Fishing behavior of factory trawlers: a hierarchical model of information processing anddecision-making. ICES Journal of Marine Science: Journal du Conseil, 58(1):238–252, January 2001.