Artificial Intelligence 214 (2014) 45–65

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Joint search with self-interested agents and the failure of cooperation enhancers ☆

Igor Rochlin, David Sarne *, Moshe Mash

Department of Computer Science, Bar-Ilan University, Ramat-Gan, Israel

A R T I C L E I N F O

Article history:
Received 30 January 2013
Received in revised form 9 May 2014
Accepted 11 May 2014
Available online 16 May 2014

Keywords:
Multi-agent economic search
Cooperation
Self-interested agents

A B S T R A C T

This paper considers the problem of autonomous agents that need to pick one of several options, all plausible however differ in their value, which is a priori uncertain and can be revealed for a cost. The agents thus need to weigh the benefits of revealing further values against the associated costs. The paper addresses the problem in its multi-agent joint form, such that not a single but rather a group of agents may benefit from the fruits of the search. The paper formally introduces and analyzes the joint search problem, when carried out fully distributedly, and determines the strategies to be used by the agents both when fully cooperative and when self-interested. The analysis is used to demonstrate that elements that can easily be proved to be beneficial with fully cooperative agents' search (e.g., extension of the search horizon, increase in the number of cooperating agents) can actually degrade individual and overall expected utility in the self-interested case. The analysis contributes to the advancement of joint search theories, and offers important insights for system designers, enabling them to determine the mechanisms that should be included in the markets and systems they design.

1. Introduction

Consider the problem of a CS graduate student named Jill, whose paper was accepted to one of the top conferences in the field and now needs to search for a way of traveling to the conference. While Jill knows there are many options to travel to the conference venue (e.g., different flights of diverse airlines to nearby airports and different means of transportation from each airport) she does not know a priori their feasibility of getting her to the conference on time, and more importantly, the cost of each option. Checking an alternative potentially involves several activities (e.g., checking locations on the map and checking the companies' web-sites for routes, timetables, fares and availability) thus incurs some "opportunity cost". Therefore Jill will not necessarily seek the cheapest alternative that can be found, but rather at any time throughout the process she will weigh the benefits of additional search against its costs. The optimal search strategy dictates continued search only insofar as its expected utility is greater than its associated cost.

The above setting is the archetypal setting of *costly search* [12,70,51,30,34,13,26] (which essence is "optimal stopping"). In general, the setting considers a decision maker (a "searcher") that needs to choose one of several available opportunities, any of which is associated with some value to her. The value of each opportunity, e.g., the price, but more generally: expense, reward, utility, is a priori unknown to the searcher however can be obtained for a cost, denoted "search cost" (either monetary or in terms of resources that need to be consumed). The value and search costs are assumed to be measured

☆ Preliminary results of this work appeared in Proceedings of the 2012 IEEE/WIC/ACM International Conference on Intelligent Agent Technology.
* Corresponding author.

http://dx.doi.org/10.1016/j.artint.2014.05.004

on the same scale and the searcher can obtain the value of as many opportunities requested, incurring the search cost of each of them. The goal of the searcher is to maximize her overall expected utility, defined as the value of the opportunity eventually picked minus the accumulated search costs. A *strategy* for this problem is a mapping from prior findings to the next action – which can be either terminated the search, or continue by checking another specific opportunity.[1] The problem as formulated above applies to a variety of real-world situations, including: job search, buying and selling goods, house search, technology R&D, making decisions on a bank to deposit funds, a vacation, where to drill an oil well, or a path to route packets, and many more [47,70,41,64].

At times, not a single but rather a group of agents may benefit from the fruits of the search. In our example, Jill's advisor can ask another student from her research group, Jack, who will be attending the same conference, to join Jill in checking different alternatives for traveling to its location, i.e., execute the search *jointly*. Numerous other examples of joint search can be found in other domains. For example, a drilling company may send multiple agents to explore possible drilling sites, and the best site of all those found by all agents will be chosen. Similarly when looking to fill-in a position, HR managers can interview candidates in parallel and recruit the best candidate found. The benefits of the multi-agent joint search is twofold. First, since each search result can benefit many agents, the relative cost of search is reduced, and the overall welfare increases. Secondly, the search space can be divided according to the expertise of the different agents, if such expertise exists.[2]

While multi-agent search has been considered in the past, it has been limited to the case of a single agent (e.g., "a representative agent") searching on behalf of the group [25,63,11,43,10,62]. More important, most previous work assumes that the agents are fully cooperative, and that their shared goal is to maximize the joint utility. While this may be the case in some situations, it is not so in a growing number of others. Rather, agents are frequently self-interested, i.e., may represent different entities, and attempt to maximize their individual utility rather than the joint utility [58,16,42,37]. A selfish agent will engage in search only if it is individually beneficial. Moreover, if other agents search, it will prefer to take advantage of their search, rather than do the work itself. For example, Jill may find it more beneficial to spend her time working on her research or resting and rely on Jack's findings. The analysis of such settings calls for a strategic, incentive driven approach, seeking stable, equilibrium solutions.

In this paper we supply a comprehensive analysis of a model of joint search, both for the case of fully cooperative and self-interested agents. For the fully-cooperative case the optimal strategy is proved to be based on the reservation-value (threshold). For the self-interested-agents case we prove a specific structure of the strategies used in equilibrium, wherein each agent first determines whether it will engage in search at all, and if so it inevitably uses a reservation-value-based search strategy. The analysis for this case considers Bayesian Nash equilibria, introducing the sets of equations that need to be solved to extract agents' strategies and the conditions that need to be checked for validating the stability of these solutions. The analysis for both settings is extended for the case where communication is enabled throughout the search process such that findings are shared continuously rather than only at the end of the process.

The analysis facilitates demonstrating that methods and instruments (termed "enhancers") that are easily proved to be beneficial in the fully cooperative case, can actually have a negative impact, both on individual and overall performance, in the self-interested case. The explanation for the phenomena is that when the agents are self-interested they might prefer to limit their individual search efforts while counting on potential findings of others. In this case, many potential solutions, in the form of search strategies that are beneficial from the individual and overall expected utility point of view, become unstable, and the resulting stable solutions are such that the agents find it beneficial to search to a lesser extent.

One key aspect of joint search that the paper emphasizes is communication between the agents along the search process. Communication is known to be a critical enhancer of joint search with fully cooperative agents [56,54] and of coordination between agents in general [57]. Alas, as demonstrated in later sections, the use of communication by the self-interested agents in joint search often results in substantial performance degradation (both individually and in total). This stems from a somewhat unique characteristic of the value an individual finds in communication in the model considered – a self-interested agent finds the communication to be beneficial only when it receives a report of a favorable finding from another agent. When the agent is the one to report, the report may lead to a reduction in the extent of search carried out by all other agents, and consequently to an expected loss for the agent.

The study of the effect of different joint search "enhancers", and in particular communication throughout the process, which often turns out to be counter-intuitive, provides market designers and platform owners with a better understanding of the benefit and usefulness of enabling such search "enhancers" in their systems. In the context of the Jack and Jill example the implication for the student's advisor can potentially be that it is better to divide all students in the group who need to attend the conference into sub-groups, each executing the joint search separately or ask them to execute the joint search with no communication between them (e.g., have them work on computers located at different offices so they cannot see each other's findings).

[1] Notice that the concept of "search" in our context is different from state-space search that is common in AI. The latter is an active process in which an agent finds a sequence of actions that will bring it from the initial state to a goal state [22]. In our case all opportunities are plausible goal states, however differ in their value to the searcher, and the goal function takes into consideration both the opportunity value and the search costs.

[2] In case of buyers' cooperation, the agents can also benefit from a volume discount through their cooperation, however this property holds only for that specific domain.