IJCAI-16 (author)

IJCAI-16

Help Log out

Submission 2032 | IJCAI-16 | Support | News | EasyChair

IJCAI-16 Submission 2032

Update authors
Review response

Submission information updates are disabled.

Click "Review response" to **see the reviews sent to you for the review response period and your response**.

For all questions related to processing your submission you should contact the conference organizers. Click here to see information about this conference.

All **reviews sent to you** can be found at the bottom of this page.

Paper 2032	
Title:	A Polynomial Time Optimal Algorithm for Robot-Human Search under Uncertainty
Paper	
Track:	IJCAI-16 Main Track
Author keywords:	Robot-Human search Uncertain Knowledge Human Availability Optimal Search Strategy
EasyChair keyphrases:	reveal cost (160), expected value (140), search rule (130), average utility (130), known item (130), unknown item (100), human availability (100), human robot interaction (95), known collect reward (79), polynomial time (70), robot decision making (63), probability distribution function (63), expected utility (60), human help (60), search strategy (60), zireveal ziask (50), dynamic programming (50), optimal strategy (50), ask index (50), search problem (50), cost creveal (50), pandora problem (50), human robot (50), distribution function fi (47), cost cask (40), closed box (40), time optimal (40), ask action (40), search process (40), unknown location (40)
Topics:	Contribution Type::Integrated contribution involving multiple areas of AI, Evaluation Methodology::Application-driven, Human-Aware AI (IJCAI-16 Theme)::Planning and Decision Support for Human-Machine Teams, Planning and Scheduling::Robot Planning, Robotics and Vision::Human Robot Interaction
Abstract:	This paper studies a search problem involving a robot that is searching for a certain item in an uncertain environment (e.g., searching minerals on Mars) that allows only limited interaction with humans. The uncertainty of the environment comes from the rewards of undiscovered items and the availability of costly human help. The goal of the robot is to maximize the reward of the items found while minimising the search costs. We show that this search problem is polynomially solvable with a novel integration of the human help, which has not been studied in the literature before. Furthermore, we empirically evaluate our solution with simulations and show that it significantly outperforms several benchmark approaches.
Time:	Jan 27, 15:21 GMT
Technical Track	(Main) The LEAD author of this submission is a (graduate or undergraduate)

(Required only for the technical track) Confirm that your submission follows IJCAI-16 requirements	student I confirm that this paper is not currently under review at any other archival venue, and will not be submitted to any such venue until the end of IJCAI reviewing period (4/4/2016)

Authors

first name	last name	email	country	organization	Web page	corresponding?
Shaofei	Chen	chensf005@163.com	China	National University of Defense Technology		✔
Tim	Baarslag	tbaarslag@gmail.com	United Kingdom	University of Southampton		
Dengji	Zhao	dengji.zhao@gmail.com	United Kingdom	University of Southampton		
Jing	Chen	chenjing001@vip.sina.com	China	National University of Defense Technology		
Lincheng	Shen	LCShen@nudt.edu.cn	China	National University of Defense Technology		

Reviews

Review 1

Significance:	**1**: (low (minimal contribution or weak impact))
Soundness:	**3**: (correct)
Scholarship:	**1**: (important related work missing or mischaracterizes prior research)
Clarity:	**3**: (well written)
Breadth of Interest:	**3**: (some interest beyond specialty area)
Summary Rating:	**-2**: (--)
Confidence:	**3**: (highly confident)
Summarize the Main Contribution of the Paper:	The paper presents an optimal algorithmic solution for a restricted setting of a robotic search problem where humans' help is available to some extent. The authors model their proposed setting and by using reduction to Pandora's problem derive an optimal search rule. The proposed algorithm is evaluated in simulation against 4 baseline models.
	I think that the paper is not well positioned with respect to the existing literature. First, state-of-the-art HRI solutions deal with much more than "architecture mechanisms" , a very relevant example is Rosenfeld at el. Intelligent Agent Supporting Human-Multi-Robot Team Collaboration from IJCAI 15. Second, it's unclear how does the paper relate to Hazon et al. physical search notion? How is this paper different/similar to other approaches? This brings me to my main concern - as far as I understand, the paper is disregarding the physical search costs, which is part of the fundamental issues in robotic search. That is, once the robot moved to a location X the costs of moving to location Y has (most likely) changed. The

Comments for the Authors:	formulation as an instance of Pandora's problem induces a very harsh assumption that traveling is "free". In general, the motivation for the paper's assumptions is lacking. I don't have a problem with most of the assumptions (mainly with the one stated above), still I think those should be better motivated. For example, why does the robot can only collect one item? If a single robot is supervised by a single operator, why should the operator not be available to provide help at any given moment? The human is assumed to be able to perfectly answer every query, why? Isn't it possible the human is merely adjusting the probabilities (providing a noisy signal)? Why is p assumed to be constant over time? I assume both revel and ask costs are in terms of time. But, they can be overlapping. While waiting for an answer the robot can start its excavating without any extra cost (in time). More disturbing is the following statement - "...a robot's strategy should maximize the overall benefit resulting from the search process, defined as the value of the option eventually collected, minus the costs accumulated along the process". Aren't we comparing apples and oranges here? Costs are in terms of time and rewards are in terms of money/etc. Is there a bound on the robot's time? Does it get a reward for finishing quickly or with a non-empty battery? In Section 5 the authors say "...we do not have to compare the computational time with other related algorithms with harder complexity such as those proposed in [Rosenthal and Veloso, 2012; Kang and Ouyang, 2011; Hazon et al., 2013]." This is a very problematic statement. The mentioned papers solves different problems, they are incomparable. The experimentation is nicely presented, but given that the Pandora solution is optimal, why is the experimentation even needed? Overall, my main concern with this paper is that it fails to explain why its proposed settings are relevant to HRI.

Review 2

Significance:	**2**: (medium (modest contribution or average impact))
Soundness:	**3**: (correct)
Scholarship:	**2**: (relevant literature cited but could expand)
Clarity:	**3**: (well written)
Breadth of Interest:	**3**: (some interest beyond specialty area)
Summary Rating:	**2**: (++)
Confidence:	**2**: (reasonably confident)
Summarize the Main Contribution of the Paper:	This paper studies a search problem involving a robot that is searching for a certain item in an uncertain environment, with limited allowed interaction with humans. The robot keeps on revealing items with the help from the human, and then selects one of the revealed items to collect. Therefore, the goal of the robot is to maximize the reward of the obtained item while minimizing the sum of search costs. One key point is the fact that is taken into account the fact that the human is not always available, and the robot has to decide whether to try to ask for help (at a given cost, also if there will not be a reply) or revealing the item by himself. The authors focus on a class of the robot search problem where costs of visiting potential locations are independent with each other, proposing a new model of robot-human search (RHS) and a novel optimal algorithm to solve it.

	An approach based on a dynamic programming formulation for robot’s decision making is presented, as well as algorithms to solve it, showing that this search problem is polynomially solvable (with an index based search policy inspired by Pandora's rules) with an integration of the human help. The proposed approach is empirically evaluate with simulations, and is shown that it significantly outperforms several benchmark approaches.
Comments for the Authors:	I found this paper well written, with an interesting balance between theory and practice. The problem addressed looks interesting and realistic. I have a question. This solution is clearly designed for entailing autonomy, in scenario with difficult communications with the human operator, like deep space missions. Hence this has to be run on-board, most likely with very limited computational resources. Did you consider this aspect in your experiments? Did you run them on-board of any rover? You do not compare the computational time with other approaches, but what is actually this computational time? And, how your result compare with other mentioned algorithms, in terms of utility I mean? I think also this aspect should be considered, not only the computational time.

Review 3

Significance:	**2**: (medium (modest contribution or average impact))
Soundness:	**3**: (correct)
Scholarship:	**3**: (excellent coverage of related work)
Clarity:	**3**: (well written)
Breadth of Interest:	**3**: (some interest beyond specialty area)
Summary Rating:	**3**: (+++)
Confidence:	**2**: (reasonably confident)
Summarize the Main Contribution of the Paper:	In this paper, the authors present an algorithm for combined human-robot search under uncertainty. It is assumed that the cost for searching a site is independent of the cost for other sites, but the reward that is obtained at a site is uncertain. A human operator can assist the robot by stating the reward for the site; however, there is a cost for asking the human. The authors develop a polynomial time optimal algorithm that maximizes the reward while using human help for search. It is implicitly assumed that the collection cost at each site is identical and the human always know the right valuation. The paper is clearly written and the proofs are well explained. The limitation of the paper is the fact that authors ignore motion cost of the robots which makes its overall application limited.
Comments for the Authors:	(1) Do you assume that the collection cost at each site is identical? If so, can this assumption be relaxed. (2) When you are defining the reward function it may be good to state explicitly that the cost of revealing is considered as part of the action cost when the robot takes the reveal action whereas the cost of asking is the action cost when the robot asks the human. In the latter case, the cost of revealing is put in as part of the reward function (Equation (1)). I was confused initially with the definition of the reward function as to why the cost of revealing is considered only in the reward for asking the human.

Review 4

Significance:	**2**: (medium (modest contribution or average impact))
Soundness:	**3**: (correct)
Scholarship:	**3**: (excellent coverage of related work)
Clarity:	**2**: (mostly readable with some scope of improvement)
Breadth of Interest:	**3**: (some interest beyond specialty area)
Summary Rating:	**3**: (+++)
Confidence:	**3**: (highly confident)
Summarize the Main Contribution of the Paper:	This paper introduces a twist on a standard search optimization problem. The theoretical advance is minor, but the identification of an important aspect of human-robot interaction in this context is worthwhile. The simulated results are narrow and not particularly realistic, but they are carefully and comprehensively performed and provide solid evidence for the merits of the approach.
Comments for the Authors:	The paper is generally well-written, but substantial disfluencies (subject-verb agreement, number, tense, etc) creep in. These should be corrected. The theoretical reduction of your problem to Pandora is simple and clear and improves the paper a great deal. Of course, many of the interesting aspects of the human-robot advisory relationship are entirely elided in this paper (though not all! I was delighted to see the exploration of human availability in your experimental results). The model will ultimately have to account for a number of different factors: - Search costs are not known up front; at best, the robot has access to a prior over possible search costs. - Humans are not oracles; their advice may improve the search cost estimate, but they can also be wrong. - A very sophisticated robot may even be able to learn a distribution over human advice quality and relate this to particular tasks - The task-independence assumption is a large one, of which the authors are well aware, as they defend their decision on page one. In a robot-on-Mars scenario, the cost of simply moving from exploration site to exploration site is very large, so any real planning must take into account the geographic and route dependence of visiting sites of interest.