

Solving AI's moving-target search problem

August 28 2017



Credit: IBM

As we look to get from point A to point B, people increasingly have options that go beyond the typical bus or taxi. Various flavours of ride and car sharing are now available in cities worldwide, making getting around more convenient than ever. Yet, as anyone that uses such services knows, it can be difficult for a driver to find you using GPS alone. And, of course, you have to stay put until they do—even if plans change.

In the language of artificial intelligence (AI) research, this is a "moving-target search problem." The user is a moving target and the cab is an agent that needs to find the user so they can board the cab. Together with a team of collaborators, I'm working on a solution that could one day allow passengers in need of transportation to book intelligent cabs, shared cars or autonomous vehicles that will be able to find them and pick them up at the right time and precise location.

At IBM Research – Ireland we've developed a new moving target algorithm that allows users to keep changing their location as they wish, instead of being forced to wait at a pre-arranged location. The service provider in this scenario also benefits from the ability to minimize running costs such as fuel.

This solution is illustrated in our research paper entitled "A Scalable Approach to Chasing Multiple Moving Targets with Multiple Agents," which I presented at the International Joint Conference on Artificial Intelligence (IJCAI) in Melbourne, Australia. At the conference, I demonstrated how different agents are assigned to different targets (users) and how to plan the movements of the chaser agents (cabs).

The algorithm is very fast and it easily scales to problems with hundreds of agents—much larger than could previously be addressed. It provides a solid AI foundation that will allow to apply large-scale moving target search to useful real-life applications. Besides intelligent cab systems, this could be applied to a healthcare scenario, such as matching in-home care providers with patients.

A care service company typically manages a large number of care workers under some constraints, such as the availability and expertise of each worker, making it challenging to match them with the right patient. In the AI research community, this is a so-called multi-agent search and planning problem where each agent (care worker) finds a plan

(determining a set of care receivers to visit), under some objectives (e.g., minimizing the delay to service) and constraints (e.g., all receivers need to be treated, duplicate services to the same care receiver are unnecessary).

In another paper we presented at the AI conference IJCAI –"Efficient Optimal Search under Expensive Edge Cost Computation"—we address the problem of minimizing delays of service caused by uncertainty of transportation. In this scenario, the benefits are similar to the taxi scenario in that the care services experience reduced running costs by matching caregivers closer to the patient's location while the patient benefits from reduced waiting time.

This approach builds on top of a journey planning system IBM Research has developed for several years and is very robust in the sense that the maximum delay of the service is theoretically guaranteed even if uncertain events occur such as missed connections of trains/buses. Our experimental results with real transportation networks show the promise of the approach.

Both the taxi and care giver scenarios present a real-time search problem that requires the agent to quickly match a user to a service provider with little information. In another paper we're presenting at IJCAI, "Online Bridged Pruning for Real-Time Search with Arbitrary Lookaheads," we describe an algorithm that improves the behavior of a real-time agent significantly. The algorithm identifies places in the environment that the agent has no reason to revisit, and prevents the agent from returning there. The results show that this work is a solid step ahead in the journey of having intelligent agents make very good decisions in the presence of limited information and limited thinking time.

For example, real-time search enables an agent to quickly find another route to reach its destination, when detecting that the current road is

blocked off due to an unexpected event such as a traffic accident. We clearly do not have to revisit the origin, and just find another route from the current location. Real-time search can efficiently do so based on the acquired knowledge about which roads were passable without any problems. Real-time search can be extended even to [search](#) and rescue operations against disasters such as earthquakes and floods, where autonomous robots have limited knowledge about the current road-map due to many roads that become suddenly blocked off.

More information: Fan Xie et al. A Scalable Approach to Chasing Multiple Moving Targets with Multiple Agents, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence* (2017). [DOI: 10.24963/ijcai.2017/624](https://doi.org/10.24963/ijcai.2017/624)

Masataro Asai et al. Efficient Optimal Search under Expensive Edge Cost Computation, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence* (2017). [DOI: 10.24963/ijcai.2017/596](https://doi.org/10.24963/ijcai.2017/596)

Carlos Hernandez et al. Online Bridged Pruning for Real-Time Search with Arbitrary Lookaheads, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence* (2017). [DOI: 10.24963/ijcai.2017/72](https://doi.org/10.24963/ijcai.2017/72)

Provided by IBM

Citation: Solving AI's moving-target search problem (2017, August 28) retrieved 12 May 2024 from <https://phys.org/news/2017-08-ai-moving-target-problem.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is

provided for information purposes only.