

Bounty Hunting and Human-Agent Group Task Allocation

Drew Wicke, Sean Luke

Department of Computer Science,
George Mason University, Fairfax, VA USA
dwicke@gmu.edu, sean@cs.gmu.edu

Abstract

Much research has been done to apply auctions, markets, and negotiation mechanisms to solve the multiagent task allocation problem. However, there has been very little work on human-agent group task allocation. We believe that the notion of *bounty hunting* has good properties for human-agent group interaction in dynamic task allocation problems. We use previous experimental results comparing bounty hunting with auction-like methods to argue why it would be particularly adept at handling scenarios with unreliable collaborators and unexpectedly hard tasks: scenarios we believe highlight difficulties involved in working with humans collaborators.

1 Introduction

Robots and agents currently operate in limited capacities in our daily lives. In the future, however, as these interactions increase and become more complex, many real world problems will involve more than one human and agent interacting and collaborating in non-obvious ways. Present approaches in Multiagent Systems (MAS) and Multirobot Systems (MRS) often only consider agent-agent interactions or robot-robot interactions, but many real-world situations that robots will encounter will involve humans. Although there has been little research into the ergonomics of MAS and MRS, the knowledge gained from these fields is ripe for translation into the more complex domain of multiple human-agent environments.

We are interested in the interaction between humans and multiple agents. In this paper we will focus on multiagent task allocation, where the problem is to determine how to assign tasks to agents in order to maximize the system utility. Research in this area largely focuses on auction and market based methods where agents bid their valuation of the task and the winner is allocated the task (Korsah, Stentz, and Dias 2013). Recently there has been work done in motivating an alternative mechanism, *bounty hunting* (Wicke, Freelan, and Luke 2015; Wicke, Wei, and Luke 2016). Here we will argue for the benefits of bounty hunting as a model for multiple agents in mixed agent-human scenarios.

In bounty hunting, multiple agents may independently attempt to complete various available tasks: task ownership

is not exclusive. When a task becomes available, it is assigned a gradually-increasing *bounty* which an agent will win only if it completes the task before any other agent does. An agent may commit to work on at most one task at a time, and the fact that it has committed to a task is known to all agents. An agent may also abandon tasks, or lose a task if it was completed by someone else first. Agents quickly learn to divvy up the space of tasks according to which agent is most adept at completing it.

This paper first motivates the need for considering the bounty hunting approach to task allocation by arguing from a human perspective. We further consider approaches to multiagent task allocation and the current state of the art in human-agent group and human-robot coordination. We then give a more formal description of bounty hunting. Finally, we will argue through our previous experiments that certain key features, such as non-exclusivity and intuitive allocation based on committing to tasks for rewards, make bounty hunting well suited for human-agent group systems.

2 Motivation for Bounty Hunting in Mixed Human-Agent Environments

As we argue later, bounty hunting's non-exclusivity is well suited to how humans approach tasks. Indeed bounty hunting is a human notion, used in everything from "have you seen this lost dog" signs, to modern-day bail enforcement agents. Non-exclusivity means there is no need for any special mechanisms to handle human-agent negotiation, nor protocols for performing multiple tasks simultaneously, and this reduces the human-agent communication load. This is in strong contrast to auctions, where tasks are allocated exclusively and the agents will be directly affected by the decisions of the humans. Additionally, bounty hunting's flexibility might be particularly useful when coordinating with humans who often juggle multiple tasks, such as in the healthcare field where nurses must handle multiple patients and tasks simultaneously (Garrett and Caldwell 2006).

We think that bounty hunting is also more intuitive than auctions from a human perspective, since it can be difficult to formulate a bid for a task. In an auction, agents bid their valuation for a task, and this valuation is an estimate for how difficult it will be for them to complete. Humans, however, tend to experience the *planning fallacy*, wherein they under-

estimate how long a task will take them to complete (Buehler, Griffin, and Ross 1994). This may then lead to misallocating tasks that might be better done by another human or agent. Additionally, humans bidding in auctions against other humans exhibit the *winner’s curse*, in which humans end up paying more for an item than it is worth (Malhotra and Bazerman 2008). Other research suggests that the effect of the winner’s curse is reduced when a human bids against only computer agents, rather than only other humans (Van Den Bos et al. 2008). Nevertheless, in a sufficiently complex multiple human-agent scenario, the winner’s curse may come into play. We discuss in Section 5 experiments that we conducted in past work, which exhibit qualities of both the winner’s curse and the planning fallacy, and examine their effects on bounty hunters and auction agents.

3 Prior Task Allocation Research

The problem of task allocation is very common in the real world. Examples include tasking firefighters with putting out fires, or tasking volunteers with organizing a food pantry or delivering food, both of which are subject to research in MAS, but currently lack human-factors research (Santos and Bazzan 2011; Aleksandrov et al. 2015). Integrating humans into the task allocation problem is necessary in order to advance this goal in more realistic scenarios, though it of course can significantly complicate the problem. The following brief overview of current research in multiagent and multirobot task allocation, as well as work in human-agent task allocation problems, will clarify this need.

Multiagent and Multirobot Task Allocation A taxonomy of solutions to the general multiagent task allocation problem has been offered (Gerkey 2004) with three major dimensions: the type of task, the type of agent, and the information available to determine task assignment. This taxonomy has recently been improved to classify more task allocation problems, like those where tasks have dependencies (Korsah, Stentz, and Dias 2013). Current solution methods include centralized approaches using k-armed bandits, integer programming, genetic algorithms, and combinatorial optimization; and decentralized approaches like auctions, markets, and task swapping.

The most common decentralized approach, auctions, were popularized by MURDOCH, an auction method which focused on minimizing resource usage, task completion time, and communication overhead (Gerkey and Matarić 2002). This approach trusted agents to truthfully bid their task valuations, though sensor noise or an unknown environment could cause their bids to be inaccurate. TraderBots, another auction framework, was designed to be more flexible and extensible (Jones et al. 2006). CoMutaR merged multiagent task allocation with coordination through the use of auctions under the assumption of truthful agents (Shiroma and Campos 2009).

Markets have also been used in multi-robot task allocation (Pustowka and Caicedo 2012; Schneider et al. 2005). There are many similarities between auctions and markets as they both rely on the agents bidding on tasks. Another approach is *token-passing* (Farinelli et al. 2005). This method,

like auctions and markets, also assumes task exclusivity. But like some market methods, it allows the tasks to be reassigned by passing the token to other agents. Bounty hunting has not been used as a mechanism for multiagent task allocation except in our past work (Wicke, Freelan, and Luke 2015; Wicke, Wei, and Luke 2016).

Human-Agent Task Allocation There are a number of points of interaction when dealing with human-agent task allocation. Tasks may be allocated to humans and agents for independent completion based off of the human preferences. However, human preferences may have unintended negative team results, such as overall inefficiency (Gombolay, Huang, and Shah 2015). Additionally, the interaction may be at a decision support level where the agent assists the human in selecting which task to complete. Max-Sum, Multiagent Markov Decision Processes (MMDPs), and Partially Observable Markov Decision Processes (POMDPs) have been studied as ways to allocate tasks in human-agent task allocation systems (Delle Fave et al. 2012; Ramchurn et al. 2015; 2016; Roncone, Mangin, and Scassellati 2017). However, these methods use a decentralized optimization method (Max-Sum) or a centralized planner (MMDP or POMDP) rather than a multiagent algorithm to allocate tasks. These studies have also shown from an ergonomics perspective that transparency is critical. Knowing which tasks and roles the robots will take on is important for human operators (Roncone, Mangin, and Scassellati 2017).

4 Bounty Hunting

The bounty hunting framework applied to a dynamic task allocation problem is defined by the tuple $\langle A, S, I, Q, M \rangle$ and a bondsman that assigns to the task a base bounty b_0 and a bounty rate r (see (Wicke, Freelan, and Luke 2015) for details on setting base bounty and bounty rate). We have a set of bounty hunting agents $A : \{a_1, a_2, \dots\}$ whose goal is to maximize the expected bounty per timestep they will receive by completing tasks. We define a set of task classes $S : \{S_1, S_2, \dots\}$, and for each such task class S_i , there exists a set of possible tasks $\{I_{i,1}, I_{i,2}, \dots\}$ of that class which might appear. We will presume that at most one task of a given class will be posted at any time. At a given timestep t there is a set $Q^{(t)}$ of uncompleted tasks available for the bounty hunters to complete. These tasks are indexed by the integers $1, 2, \dots, i, \dots$. We define agent commitment to a task at time t and the bounty the agent will receive when they complete the task by $M_t : A \times \{Q^{(t)} \cup \{\square\}\} \rightarrow \mathbb{R}$, where we define the empty task as $\{\square\}$ which has no bounty.

Committing does not imply exclusivity: other agents may and can commit to work on the same task. Further, if for some reason an agent was not able to complete tasks, their bounty would rapidly rise to the point that other agents would be incentivized to take the tasks from them. In our previous work we have shown that it is rational for agents to pursue tasks and not to wait indefinitely (Wicke, Freelan, and Luke 2015).

Non-exclusivity can be inefficient. In (Wicke, Freelan, and Luke 2015; Wicke, Wei, and Luke 2016) we have studied

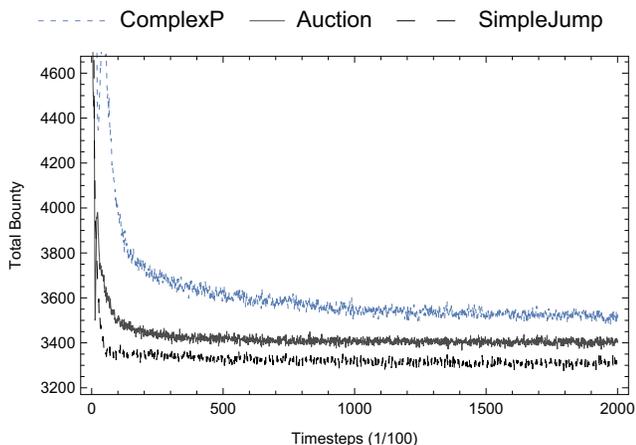


Figure 1: Experiment 1, Static Environment, 200,000 steps. Lower values are better. Legends are read from left to right corresponding to the plots top to bottom.

how the agents may rapidly learn to divvy up the task space according to which agents are most efficient at each task. Our metric is based on the sum total bounty on outstanding tasks at any time: this is minimized by completing as many tasks as rapidly as possible, so their bounty does not climb. We compared bounty hunting against several other methods, including two exclusive methods, one of which was an auction based approach, where agents bid for tasks based on the bounty and were exclusively allocated the tasks by the auctioneer.

5 Bounty Hunting and Human-Agent Teams: Argument from Past Experiments

The non-exclusivity of the bounty hunting model, its dynamic incentive structure, and the agents’ ability to adapt, combine to make bounty hunting effective at dealing with even quite noisy and dynamic scenarios. We argue that these same features may enable bounty hunters to work effectively with inconsistent, unpredictable, or confused human collaborators experiencing the winner’s curse and the planning fallacy.

To make this argument we will draw from some past experiments we have done in bounty hunting in dynamic environments (Wicke, Wei, and Luke 2016). In that test scenario, four robots live on a small soccer field, and as balls appear randomly on the field the robots chase after and retrieve the balls. Balls fell into certain *task classes* according to locations in which they appeared on the field, and the robots rapidly learned to divvy up the field so as not to overly compete for tasks. A task (a ball) was posted with a bounty of 100, and thereafter the bounty increased by 1 each timestep until the task was completed.

We focused on an adaptive bounty hunter algorithm which can abandon tasks, called *SimpleJump*. We compared it to *ComplexP*, a previously-studied high-performing bounty hunter algorithm which could not abandon tasks after committing, and to *Auction*, an auction-like exclusive method.

Experiment 1 was the baseline comparison: while balls

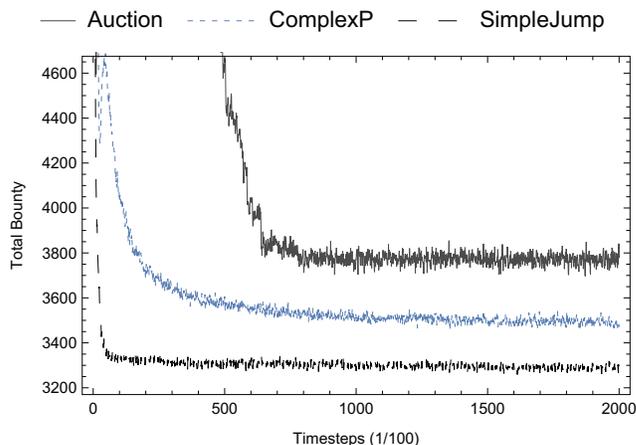


Figure 2: Experiment 4, Unreliable Collaborators, 200,000 steps. Lower values are better. Legends are read from left to right corresponding to the plots top to bottom.

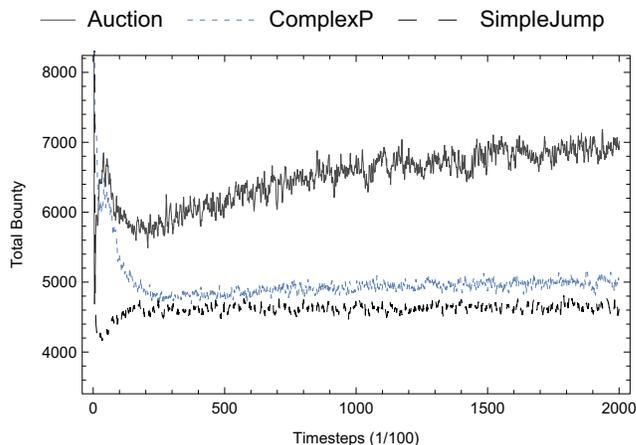


Figure 3: Experiment 5, Unexpectedly Hard Tasks, 200,000 steps. Lower values are better. Legends are read from left to right corresponding to the plots top to bottom.

would appear randomly on the field, the environment did not change dynamically in any significant way, nor did the makeup or nature of collaborators. Here non-exclusivity should show little advantage. Figure 1 shows how all three techniques converged: SimpleJump edged out Auction largely due to its ability to abandon tasks, but the difference between the three was relatively small.

Experiments 2, 4, and 5 tested agents in different dynamically changing environments. We draw two to show here: Experiment 4 involved unreliable collaborating agents, and Experiment 5 involved unexpectedly and suddenly difficult tasks. We think that both experimental scenarios match naturally with human involvement.

We begin with Experiment 4. Here two additional collaborators were added: these collaborators committed to (or bid heavily for) tasks entirely at random and moved 10x slower than the other agents. In many respects these collaborators modeled humans grabbing arbitrary tasks and performing

poorly at them. Also, in some sense this experiment focuses on the winner's curse. In the winner's curse scenario, humans bid amounts unrelated to the true value of the task. Unreliable agents bidding highly at random simulates the noise generated by humans experiencing the winner's curse.

The results in Figure 2 show that even with these unreliable collaborators the SimpleJump bounty hunting method could maintain functionality at a similar level to the baseline (Experiment 1). But the auction mechanism's performance was considerably degraded.

In Experiment 5, with a 10% probability a given task would be unexpectedly 10x harder for a randomly-selected agent to do. This is similar in effect to the planning fallacy, because the agents believe that it will take less time to complete a task than it actually takes. Obviously this would prove problematic for an exclusive method; but task abandonment would also be of benefit.

Exactly this is shown in Figure 3. Here, while SimpleJump is not able to maintain the same efficiency as the baseline, its performance is still statistically better than the auction mechanism, and also better than ComplexP.

6 Conclusion and Future Work

We have motivated the need to further study multiagent task allocation approaches for human-agent groups by pointing to the pitfalls such as the winner's curse and the planning fallacy, which appear due to human behavior. We have reviewed the literature on task allocation methods in multiagent systems and we have examined current research in human-agent task allocation. We have argued that bounty hunting is particularly well suited to the challenge of human-agent task allocation by reviewing previous experimental results in light of human-agent groups. Our findings suggest that since bounty hunting is robust, intuitive, and the tasks are non-exclusive that bounty hunting best accommodates the weaknesses and styles of human collaborators. We propose that further research into implementing bounty hunting for human-agent groups is needed to better understand the effect of non-exclusive tasks and the interaction between adaptive agents and humans.

References

- Aleksandrov, M.; Aziz, H.; Gaspers, S.; and Walsh, T. 2015. Online fair division: Analysing a food bank problem. In *IJCAI*, 2540–2546.
- Buehler, R.; Griffin, D.; and Ross, M. 1994. Exploring the "planning fallacy": Why people underestimate their task completion times. *Journal of personality and social psychology* 67(3):366.
- Delle Fave, F. M.; Rogers, A.; Xu, Z.; Sukkarieh, S.; and Jennings, N. R. 2012. Deploying the max-sum algorithm for decentralised coordination and task allocation of unmanned aerial vehicles for live aerial imagery collection. *ICRA* 469–476.
- Farinelli, A.; Iocchi, L.; Nardi, D.; and Patrizi, F. 2005. Task assignment with dynamic token generation. In *Monitoring, Security, and Rescue Techniques in Multiagent Systems*. Springer. 467–477.
- Garrett, S. K., and Caldwell, B. S. 2006. Task coordination and group foraging in healthcare delivery teams. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 50, 1059–1063. SAGE Publications.
- Gerkey, B. P., and Matarić, M. J. 2002. Sold!: auction methods for multirobot coordination. *IEEE Transactions on Robotics and Automation* 18(5):758–768.
- Gerkey, B. P. 2004. A formal analysis and taxonomy of task allocation in multi-robot systems. *IJRR* 23(9):939–954.
- Gombolay, M. C.; Huang, C.; and Shah, J. A. 2015. Coordination of Human-Robot Teaming with Human Task Preferences. *AAAI Fall Symposium Series on AI-HRI*.
- Jones, E. G.; Browning, B.; Dias, M. B.; Argall, B.; Veloso, M.; and Stentz, A. 2006. Dynamically formed heterogeneous robot teams performing tightly-coordinated tasks. In *ICRA*, 570–575.
- Korsah, G. A.; Stentz, A.; and Dias, M. B. 2013. A comprehensive taxonomy for multi-robot task allocation. *IJRR* 32(12):1495–1512.
- Malhotra, D., and Bazerman, M. H. 2008. Psychological Influence in Negotiation: An Introduction Long Overdue. *Journal of Management* 34(3):509–531.
- Pustowka, A., and Caicedo, E. 2012. Market-based task allocation in a multi-robot surveillance system. In *Robotics Symposium and Latin American Robotics Symposium*, 185–189.
- Ramchurn, S. D.; Fischer, J. E.; Ikuno, Y.; Wu, F.; Flann, J.; and Waldo, A. 2015. A study of human-agent collaboration for multi-UAV task allocation in dynamic environments. *IJCAI* 1184–1192.
- Ramchurn, S. D.; Wu, F.; Jiang, W.; Fischer, J. E.; Reece, S.; Roberts, S.; Rodden, T.; Greenhalgh, C.; and Jennings, N. R. 2016. Human-agent collaboration for disaster response. *AAMAS* 30:82–111.
- Roncone, A.; Mangin, O.; and Scassellati, B. 2017. Transparent Role Assignment and Task Allocation in Human Robot Collaboration. *ICRA*.
- Santos, F. D., and Bazzan, A. L. 2011. Towards efficient multiagent task allocation in the robocup rescue: a biologically-inspired approach. *AAMAS* 22(3):465–486.
- Schneider, J.; Apfelbaum, D.; Bagnell, J. A. D.; and Simmons, R. 2005. Learning opportunity costs in multi-robot market based planners. In *ICRA*, 1151–1156.
- Shiroma, P. M., and Campos, M. F. M. 2009. Comutar: A framework for multi-robot coordination and task allocation. In *IROS*, 4817–4824.
- Van Den Bos, W.; Li, J.; Lau, T.; Maskin, E.; Cohen, J. D.; Montague, P. R.; and McClure, S. M. 2008. The value of victory: social origins of the winner's curse in common value auctions. *Judgment and decision making* 3(7).
- Wicke, D.; Freelan, D.; and Luke, S. 2015. Bounty hunters and multiagent task allocation. In *AAMAS*.
- Wicke, D.; Wei, E.; and Luke, S. 2016. Throwing in the towel: Faithless bounty hunters as a task allocation mechanism. In *IJCAI workshop on Interactions with Mixed Agent Types*.