

Considerations for Task Allocation in Human-Robot Teams

Arsha Ali, Dawn M. Tilbury, Lionel P. Robert Jr.

University of Michigan
Ann Arbor, MI 48109
{arshaali, tilbury, lprobert}@umich.edu

Abstract

In human-robot teams where agents collaborate together, there needs to be a clear allocation of tasks to agents. Task allocation can aid in achieving the presumed benefits of human-robot teams, such as improved team performance. Many task allocation methods have been proposed that include factors such as agent capability, availability, workload, fatigue, and task and domain-specific parameters. In this paper, selected work on task allocation is reviewed. In addition, some areas for continued and further consideration in task allocation are discussed. These areas include level of collaboration, novel tasks, unknown and dynamic agent capabilities, negotiation and fairness, and ethics. Where applicable, we also mention some of our work on task allocation. Through continued efforts and considerations in task allocation, human-robot teaming can be improved.

1 Introduction

The motivation for human-robot teaming is based on the premise that humans and robots are heterogeneous and can have complementary skills. Having complementary strengths and weaknesses can allow humans and robots to work together effectively.

With advanced robots capable of executing certain tasks on their own, one question that arises is how tasks should be allocated in human-robot teams to achieve effective teamwork and performance. Many task allocation methods have been proposed for human-robot teams so the division of tasks is clear, but there are still some areas for continued and further consideration in task allocation methods. This paper presents some of these areas and briefly discusses some of our work on task allocation where applicable. To facilitate our discussion, we refer to an indivisible task simply as a task, where a task is generally allocated to one agent (either a human or a robot).

2 Background

An early report, commonly referred to as Fitts' list, details the strengths of humans versus machines. In Fitts' list, the strengths attributed to humans include detection, perception, judgement, induction, improvisation, and long-term memory, while the strengths attributed to machines include speed,

Presented at the AI-HRI Symposium at AAAI Fall Symposium Series (FSS) 2022

power, computation, replication, simultaneous operations, and short-term memory (Fitts et al. 1951). These descriptions have been used as a basis for function allocation (Hancock and Scallen 1996).

Following the idea that agents have different capabilities, many human-robot and multi-robot task allocation methods are based on the capabilities of agents and the capabilities required for tasks. Ranz, Hummel, and Sihn's method allocates tasks by matching agent capabilities with task requirements, where agent capabilities consider the elements of cost, time, and quality (Ranz, Hummel, and Sihn 2017). Budinská and Havlík's method allocates spatially distributed tasks by considering the binary capabilities and location required for the task with the capabilities and location of robots (Budinská and Havlík 2016). In addition, AL-Buraiki and Payeur's method assigns robots with specialized capabilities to a task based on the fit between the task and agent and the agent's availability (AL-Buraiki and Payeur 2019).

Task allocation methods have also considered minimizing or moderating factors such as human workload and fatigue. For example, Ge et al. developed an allocation method based on Markov decision process (MDP) that aims to minimize the mental load of all operators in a human-machine system (Ge et al. 2018). Frame, Boydston, and Lopez created a method that continuously evaluates task performance and human workload to redistribute tasks between a human and automation for a surveillance task (Frame, Boydston, and Lopez 2020). In addition, Hu and Chen include human fatigue as a continuous-time MDP in their task allocation method (Hu and Chen 2017).

Task allocation methods may also be tailored to specific types of tasks or domains, such as assembly tasks. For example, Malik and Bilberg's method allocates assembly tasks between a human and a robot based on the physical features of components (e.g., size, weight, shape), as well as other attributes such as part presentation and safety considerations (Malik and Bilberg 2019).

3 Areas for Consideration

As described above, there has been great progress and many methods proposed regarding task allocation. In this paper, we identify some areas that could benefit from continued and further consideration in task allocation methods for human-robot teams.

3.1 Level of Collaboration

Just as humans and automation can collaborate at different levels (e.g., see (Sheridan and Verplank 1978; Parasuraman, Sheridan, and Wickens 2000)), humans and robots can collaborate at different levels too. One study proposes four levels of collaboration including no coexistence (humans and robots are physically separated), coexistence (humans and robots share at least some workspace but do not share goals), cooperation (humans and robots share at least some workspace and have a shared goal), and collaboration (humans and robots share the workspace while working simultaneously on a shared object) (Aaltonen, Salmi, and Marstio 2018). Task allocation methods may need to consider the level of collaboration of the human-robot team for a given task for optimal task allocation. Similarly, the type of task allocation problem needs to be well-defined in order to develop an appropriate task allocation method. The level of collaboration may influence the type of task allocation problem. Considering questions such as what constitutes an indivisible task, how many tasks an agent can execute simultaneously, how many agents can be allocated to the same task, what tasks are likely to occur in the future, whether tasks can be re-allocated, and how tasks and agents depend on one another can aid in defining the problem and developing a corresponding task allocation method.

For example, consider a task of lifting a 50 lbs item with a team of one human and one robot. Consider that the human can lift up to 20 lbs and the robot can lift up to 40 lbs. If there is no coexistence between the human and the robot, simple logic based on capabilities only would indicate that the task of lifting the 50 lbs item should be allocated to the agent that is most capable of the task, or to no agent if no agent on the team is capable. In this example, the task could either be allocated to the robot or the task could be ignored, neither of which is ideal. If the task is allocated to the robot, there is a risk of damage to the item and to the robot. If the task is ignored, future tasks that build on having lifted the 50 lbs item may not be possible. However, if the human-robot team is at the higher levels of collaboration, the agents could lift the 50 lbs item together. The joint collaboration and performance from these two agents is better than one agent acting alone. In a shared workspace with a shared goal, the agents can collaborate simultaneously on the same object. Thus, the level of collaboration and type of problem can influence the outcome of the task allocation method. Fortunately, taxonomies for multi-robot task allocation do exist (e.g., see (Gerkey and Mataric 2004; Korsah, Stentz, and Dias 2013)). These can be used as a starting point to guide the development of human-robot task allocation methods.

3.2 Novel Tasks

Novel tasks that the human-robot team has not experienced before may occur, especially in dynamic situations. For example, this could include office, military, and household settings. Thus, task allocation methods would become more robust if they are able to effectively allocate both existing and novel tasks. As a simple example, consider a human-robot team that has previously always been tasked with picking up

and sorting cubes and rectangular prisms. In the future, if the tasks also involve picking up and sorting different objects such as spheres and cones, how should these novel tasks be allocated? As a baseline, a task allocation method may default the allocation of all novel tasks to a subset of agents on the team, randomly amongst all agents on the team, or neglect the tasks altogether. However, this may be a suboptimal allocation of novel tasks and better approaches and solutions may exist.

Part of the challenge in allocating novel tasks has to do with the difficulty in representing and characterizing tasks. Tkach and Amador recently developed a task allocation method for police officers dealing with tasks with unknown locations, arrival times, and importance levels (Tkach and Amador 2021). In our work, we have proposed representing both existing and novel tasks on continuous scales. Tasks are characterized by the levels of different capabilities required for the task (Ali et al. 2021, 2022). Such a representation allows for a task allocation method to handle both existing and novel tasks, since all tasks are represented in the same manner. However, how to transform a concrete task from the physical environment to a set of required capabilities is not addressed fully in many existing task allocation methods. Knowing the correct levels of capabilities to represent a task is a limitation of our work.

3.3 Unknown and Dynamic Agent Capabilities

Agents on a human-robot team may be unfamiliar with the capabilities of their teammates if they have had limited interaction. The question that follows is on how a task allocation method can learn agent capabilities when they are initially unknown. Learning an agent’s capabilities is especially applicable to human-robot teams where agents are being swapped in or out (e.g., a robot malfunctions and is replaced by a different robot, humans change during a shift).

A human can estimate the capabilities of another agent through interactions. Consider a scenario where a human initially thinks another agent can lift 50 lbs. By observing the agent attempting to lift various items, the human updates their estimate of how much that agent can lift. Say the human observed this agent successfully lift 75 lbs, but struggled with lifting 80 lbs. The human can use these observations to approximate the weight the agent can lift to 75 lbs. This updated information can then be used for more informed decision making in the future. When a lifting task arrives, the human has a better idea of how likely it is their teammate could successfully lift the item. If a 25 lbs item arrives, the human would likely feel the agent could lift this item successfully. However, if a 100 lbs item arrives, the human would likely think the agent is incapable of lifting this item successfully. In our work, we use a similar idea for learning an agent’s capabilities when they are initially unknown by observing the agent’s task performance. We propose representing an agent’s unknown capabilities by a belief distribution that gets refined and updated with the history of task outcomes from that agent. The belief in an agent’s capabilities informs how much that agent can be trusted to succeed at a task. Trust in an agent then informs the task allocation method (Ali et al. 2021, 2022).

A related question is how to deal with dynamic agent capabilities. Capabilities may grow through practice or training. Capabilities may also diminish if they are used infrequently or with fatigue. For example, with practice, an agent can improve their capabilities to parallel park a vehicle. However, if these capabilities are not used often, the agent's capabilities to parallel park may be degraded. Observing many task failures or reduced performance from an agent could indicate their capabilities have decreased and need to be reassessed for optimal task allocation. At the expense of failing at a task, an agent could be "tested" to see if they can succeed on a task that was initially thought to be beyond the agent's capabilities. Again, there may be better approaches and solutions than these preliminary ideas. In short, determining when and how to (re)assess an agent's capabilities is important to avoid poor task allocations.

3.4 Negotiation and Fairness

During task allocation, for their own reasons, an agent may disagree with the agent responsible for task allocation. When such disagreements occur, agents will need a way to negotiate the allocation of a task until they reach a consensus. To start, the task allocation method will need to determine whether there are any disagreements among agents. One idea could be to simply request input when an agent disagrees with the allocation of a task. Once it is determined that disagreements between agents are present, how agents will negotiate and whether one agent will have the ultimate authority will have to be considered.

Ronccone, Mangin, and Scassellati propose a negotiation step in a method in which the robot tells a human that it will perform a task or asks the human to perform a task, to which the human can respond "yes" or "no" (Ronccone, Mangin, and Scassellati 2017). In the event that two robots intend to work on the same task, Budinská and Havlík prescribe each robot with either a selfish, altruistic, or neutral behavior. The two robots negotiate the task by comparing their behavior type, analogous to a "rock-paper-scissors" game (Budinská and Havlík 2016). Regardless of how negotiation is done, at least for human-robot teams, some form of bi-directional communication will likely be necessary.

Along the same lines, task allocation methods may also need to consider the fairness of resulting task allocations. For example, it may be a poor choice to allocate more tasks or the same types of tasks to one agent compared to other agents. Consider an example in a hospital setting where one agent always gets tasked with preparing deliveries, another agent always makes the deliveries, and another agent always checks in on the patients. On any given day, the frequency of deliveries to prepare and make may be disproportionate to the frequency with which patients need to be checked on. Not only could this mean some agents are overly busy while others are idle, but agents, especially humans, may become tired from doing the same types of tasks. While a robot may not "care" about the number and types of tasks it is allocated, the allocation of tasks to a robot impacts the allocation of tasks to a human on a human-robot team. A robot could be the agent who typically makes hospital deliveries, but a human may also want to make deliveries to use other

capabilities and for "a change of pace."

Likewise, prioritizing the preferences of one agent over another in tasks they would like to execute and those they would like to avoid could be perceived as unfair. Unfairness could lead to frustration or low satisfaction in agents, which could in turn harm team performance. One study showed that when a robot distributed wooden blocks unequally between two humans in a tower construction task, team members reported a more negative perception of team relationship compared to teams in which the robot distributed blocks equally (Jung et al. 2020). Going back to the example, a frustrated agent may incorrectly prepare a hospital delivery, which could have serious consequences. Considering a robot as a moral regulator, Kim and Phillips hypothesize that increasing fairness can increase robot legitimacy and the willingness to accept and comply with the robot (Kim and Phillips 2021). This hypothesis may extend to a task allocating robot as well. Hence, both negotiation and fairness could help establish and maintain team relationship and team performance.

3.5 Ethics

Many task allocation methods are developed such that humans only need to execute tasks as opposed to both allocating and executing tasks, for reasons such as a reduction in human workload and faster completion time. Essentially, a robot or automation serves as the task allocator. In AL-Buraiki and Payeur's method, however, a human operator is involved by selecting a minimum fitting threshold which must be achieved for an agent to be assigned to a task (AL-Buraiki and Payeur 2019).

An ethical question is whether humans should be allowed to be removed entirely from the role of the task allocator and in what contexts. Even a well-designed task allocation method could inappropriately allocate a task at times (e.g., if the task is a novel task), whereas a human may quickly determine a more appropriate allocation through common sense. In such a case, a human's input could prevent disastrous outcomes. These issues speak directly to how much autonomy should be given to a robot and the inability to hold a robot accountable (Robert et al. 2020). Ultimately, in areas where human lives are at risk (e.g., healthcare settings, military settings), task allocation must not only consider who can do what ability-wise but also consider who ethically should be allowed to do what. Embedding ethics into technology is not easy (Paul et al. 2022). Future research must seek to understand the inherent appropriateness and limitations of any task allocation method. Furthermore, if humans are removed from the task allocation process or infrequently negotiate, they may come to rely on and overtrust the task allocation method. If the robot or automation doing the task allocation suddenly becomes unavailable, the human may be ill-prepared to inherit the role of the task allocator.

These ideas are similar to the role of a semi-autonomous vehicle. A semi-autonomous vehicle may make an inappropriate decision whereas a human may make a better decision using their common sense. Even if temporarily, drivers are removed from the driving task when the vehicle is driving itself and can overtrust the semi-autonomous vehicle in con-

texts where trust is not justified (Azevedo-Sa et al. 2020). Further, several studies have shown that humans may not be well prepared to resume the driving task if automation becomes unavailable (e.g., see (Gold et al. 2013; Mok et al. 2015; Du et al. 2020)). Semi-autonomous vehicles are designed to operate in some specific contexts, meaning there are times when the human is allowed to disengage to a degree from driving and times when the human needs to drive. Similarly, a task allocation method may at times benefit from human input.

Trust in the task allocation method and the task allocator agent is important. Regarding automation, when there is undertrust, trust is below capabilities and automation is disused (Lee and See 2004). When there is overtrust, trust exceeds capabilities and automation is misused (Lee and See 2004). Both disuse and misuse can harm performance (Azevedo-Sa et al. 2020). Calibrated trust is when trust matches the automation’s capabilities (Lee and Moray 1994; Muir 1987). Similarly, in task allocation, if the task allocation method and task allocator agent are undertrusted, they may not be used and the potential benefits (e.g., reduced workload) cannot be realized. If the task allocation method and task allocator agent are overtrusted, resulting task allocations may be followed in cases that are not justified. Calibrated trust would mean having an appropriate level of trust in the task allocation method and task allocator agent for given purposes and contexts. In other words, there would be an understanding of when to use and when not to use the task allocation method and task allocator agent.

If novel tasks are not a concern and the human-robot team is operating in a well-defined and structured environment, it may be plausible to allow a robot to be the task allocator, relieving cognitive effort from a human. However, if the human-robot team is operating in a dynamic environment where tasks are frequently novel or have serious consequences, it may be more appropriate for a robot to recommend and/or explain task allocations to a human or to have a human serve as the sole task allocator. In sum, the decision on whether humans should be involved in task allocation, to what degree, and how frequently should be considered carefully.

4 Conclusion

We present some areas for continued and further consideration in task allocation methods that could improve human-robot teaming and team performance. Depending on the context and what aspects are to be prioritized in a task allocation method, a human may need to be involved in the task allocation decision at some level and work closely with a robot to achieve the best possible allocation of tasks. By improving human-robot teaming, the benefits of introducing robots to society can be realized.

5 Acknowledgements

Research was sponsored by the Army Research Office and was accomplished under Cooperative Agreement Number W911NF-21-2-0168. The views and conclusions contained in this document are those of the authors and should not

be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

- Aaltonen, I.; Salmi, T.; and Marstio, I. 2018. Refining levels of collaboration to support the design and evaluation of human-robot interaction in the manufacturing industry. *Procedia CIRP*, 72: 93–98.
- AL-Buraiki, O.; and Payeur, P. 2019. Probabilistic Task Assignment for Specialized Multi-Agent Robotic Systems. In *2019 IEEE International Symposium on Robot and Sensors Environments*, 1–7. IEEE.
- Ali, A.; Azevedo-Sa, H.; Tilbury, D. M.; and Robert, L. P. 2022. Heterogeneous human-robot task allocation based on artificial trust. *Scientific Reports*, 12(15304): 1–15.
- Ali, A.; Azevedo-Sa, H.; Tilbury, D. M.; and Robert Jr., L. P. 2021. Using Trust for Heterogeneous Human-Robot Team Task Allocation. In *Proceedings of the AI-HRI Symposium at AAAI-FSS 2021*.
- Azevedo-Sa, H.; Kumar Jayaraman, S.; Yang, X. J.; Robert Jr., L. P.; and Tilbury, D. M. 2020. Context-Adaptive Management of Drivers’ Trust in Automated Vehicles. *IEEE Robotics and Automation Letters*, 5: 6908–6915.
- Budinská, I.; and Havlík, Š. 2016. Task allocation within a heterogeneous multi-robot system. In *2016 Cybernetics & Informatics (K&I)*, 1–5. IEEE.
- Du, N.; Kim, J.; Zhou, F.; Pulver, E.; Tilbury, D. M.; Robert, L. P.; Pradhan, A. K.; and Yang, X. J. 2020. Evaluating effects of cognitive load, takeover request lead time, and traffic density on drivers’ takeover performance in conditionally automated driving. In *12th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, 66–73.
- Fitts, P. M.; Viteles, M.; Barr, N.; Brimhall, D.; Finch, G.; Gardner, E.; Grether, W.; Kellum, W.; and Stevens, S. 1951. Human engineering for an effective air-navigation and traffic-control system, and appendixes 1 thru 3.
- Frame, M. E.; Boydston, A. S.; and Lopez, J. S. 2020. Development of an Autonomous Manager for Dynamic Human-Machine Task Allocation in Operational Surveillance. In *IEEE International Conference on Human-Machine Systems*, 1–4. IEEE.
- Ge, M.; Song, Q.; Qiu, H.; Wang, J.; and Fang, W. 2018. An MDP-Based Task Allocation Model for a Class of Uncertain Human-Machine System. In *2018 IEEE 8th Annual International Conference on CYBER Technology in Automation, Control, and Intelligent Systems (CYBER)*, 148–153. IEEE.
- Gerkey, B. P.; and Mataric, M. J. 2004. A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems. *The International Journal of Robotics Research*, 23(9): 939–954.
- Gold, C.; Damböck, D.; Lorenz, L.; and Bengler, K. 2013. “Take over!” How long does it take to get the driver back

- into the loop? In *Proceedings of The Human Factors and Ergonomics Society Annual Meeting*, volume 57, 1938–1942. Sage Publications Sage CA: Los Angeles, CA.
- Hancock, P. A.; and Scallen, S. F. 1996. The Future of Function Allocation. *Ergonomics in Design: The Quarterly of Human Factors Applications*, 4(4): 24–29.
- Hu, B.; and Chen, J. 2017. Optimal Task Allocation for Human–Machine Collaborative Manufacturing Systems. *IEEE Robotics and Automation Letters*, 2(4): 1933–1940.
- Jung, M. F.; Difranzo, D.; Shen, S.; Stoll, B.; Claire, H.; and Lawrence, A. 2020. Robot-Assisted Tower Construction—A Method to Study the Impact of a Robot’s Allocation Behavior on Interpersonal Dynamics and Collaboration in Groups. *ACM Transactions on Human-Robot Interaction*, 10(1): 1–23.
- Kim, B.; and Phillips, E. 2021. Humans’ Assessment of Robots as Moral Regulators: Importance of Perceived Fairness and Legitimacy. In *Proceedings of the AI-HRI Symposium at AAAI-FSS 2021*.
- Korsah, G. A.; Stentz, A.; and Dias, M. B. 2013. A comprehensive taxonomy for multi-robot task allocation. *The International Journal of Robotics Research*, 32(12): 1495–1512.
- Lee, J. D.; and Moray, N. 1994. Trust, self-confidence, and operators’ adaptation to automation. *International Journal of Human-Computer Studies*, 40(1): 153–184.
- Lee, J. D.; and See, K. A. 2004. Trust in Automation: Designing for Appropriate Reliance. *Human Factors*, 46(1): 50–80.
- Malik, A. A.; and Bilberg, A. 2019. Complexity-based task allocation in human-robot collaborative assembly. *Industrial Robot: The International Journal of Robotics Research and Application*, 46: 471–480.
- Mok, B.; Johns, M.; Lee, K. J.; Miller, D.; Sirkin, D.; Ive, P.; and Ju, W. 2015. Emergency, automation off: Unstructured transition timing for distracted drivers of automated vehicles. In *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, 2458–2464. IEEE.
- Muir, B. M. 1987. Trust between humans and machines, and the design of decision aids. *International Journal of Man-Machine Studies*, 27(5-6): 527–539.
- Parasuraman, R.; Sheridan, T. B.; and Wickens, C. D. 2000. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 30(3): 286–297.
- Paul, S.; Yuan, L.; Jain, H. K.; Robert, L. P.; Spohrer, J.; and Lifshitz-Assaf, H. 2022. Intelligence Augmentation: Human Factors in AI and Future of Work. *AIS Transactions on Human-Computer Interaction*, 14(3): 426–445.
- Ranz, F.; Hummel, V.; and Sihn, W. 2017. Capability-based Task Allocation in Human-robot Collaboration. *Procedia Manufacturing*, 9: 182–189.
- Robert, L. P.; Pierce, C.; Marquis, L.; Kim, S.; and Alahmad, R. 2020. Designing fair AI for managing employees in organizations: a review, critique, and design agenda. *Human-Computer Interaction*, 35(5-6): 545–575.
- Roncone, A.; Mangin, O.; and Scassellati, B. 2017. Transparent role assignment and task allocation in human robot collaboration. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 1014–1021. IEEE.
- Sheridan, T. B.; and Verplank, W. L. 1978. Human and computer control of undersea teleoperators. Technical report, Massachusetts Inst of Tech Cambridge Man-Machine Systems Lab.
- Tkach, I.; and Amador, S. 2021. Towards addressing dynamic multi-agent task allocation in law enforcement. *Autonomous Agents and Multi-Agent Systems*, 35(11): 1–18.