Multi-Agent Task Allocation Techniques for Harvest Team Formation

Helen Harman(✉) and Elizabeth I. Sklar

Lincoln Institute for Agri-food Technology, University of Lincoln, Lincoln, UK
{hharman,esklar}@lincoln.ac.uk

Abstract. With increasing demands for soft fruit and shortages of seasonal workers, farms are seeking innovative solutions for efficiently managing their workforce. The harvesting workforce is typically organised by farm managers who assign workers to the fields that are ready to be harvested. They aim to minimise staff time (and costs) and distribute work fairly, whilst still picking all ripe fruit within the fields that need to be harvested. This paper posits that this problem can be addressed using multi-criteria, multi-agent task allocation techniques. The work presented compares the application of Genetic Algorithms (GAs) vs auction-based approaches to the challenge of assigning workers with various skill sets to fields with various estimated yields. These approaches are evaluated alongside a previously suggested method and the teams that were manually created by a farm manager during the 2021 harvesting season. Results indicate that the GA approach produces more efficient team allocations than the alternatives assessed.

Keywords: Multi-agent system · Applied AI · Harvest management

1 Introduction

Seasonal workers are frequently employed on farms to pick ripe soft fruit (e.g. strawberries, raspberries, cherries and blackberries) during the harvesting season. These workers are usually managed by harvest managers, who make daily decisions about which fields should be picked and which workers to assign to each of those fields. This process typically involves an awkward manual process of juggling spreadsheets produced from different software systems and can be quite time consuming, particularly at the height of the season when there are hundreds of workers available, with varying skills and levels of experience. As there is an increasing demand for soft fruits and shortages in seasonal workers [6,20,28], farms are requiring more innovative solutions for managing their workforce. Having sufficient labour and effectively managing the workforce will help ensure all ripe fruit is harvested. If this is not achieved, unharvested ripe produce will rot in the field. This situation results in food waste and a loss of investment for the grower [5].

F. Dignum et al. (Eds.): PAAMS 2022, LNAI 13616, pp. 217–228, 2022.
https://doi.org/10.1007/978-3-031-18192-4_18

A range of strategies to address the labour shortage issue are being explored. This includes the introduction of robotic devices to assist with harvesting and crop-care tasks [3, 20, 21, 34, 35]. In contrast, our research investigates practical applications of *Artificial Intelligence (AI)* for managing the harvesting workforce, drawing on literature from *Multi-Agent Task Allocation* (MATA). Our approach can be implemented *now*, before a newfangled robotic workforce is ready to be deployed, and can also manage a hybrid human-robot workforce in the future.

MATA methods seek to distribute a set of tasks fairly amongst a set of agents. Our previous MATA work [12], showed that variations of Round Robin (RR) could be adapted to the problem of assigning workers (agents) to fields (tasks). This paper investigates applying Genetic Algorithms (GAs), Ordered Single Item (OSI) auctions and Sequential Single Item (SSI) auctions to this problem. GAs take a population-based approach, inspired by natural selection, they progressively adapt the suitability of the individuals for their environment. In contrast, OSI and SSI are auction-based methods, in which bidders bid for items and the bidder that places the best bid (according to the auctioneer) is assigned the item being bid on. To apply these approaches to our application domain, we define a fitness function that incorporates the farm's desire to minimise staff time (due to the shortages of labour) and balance workload amongst workers. If the workload is unbalanced, some workers would have very few fruits to pick while others would spend too long working. If these discrepancies are large, then the workforce can become disgruntled—as pickers are typically paid by the volume they pick, those with few fruits to pick earn less. Since workers are usually free to leave one farm and move to another, the farm managers would like to keep their workers happy so that their workforce remains intact during the season.

The GA, OSI and SSI methods introduced in this paper are evaluated on real-world data provided by a commercial fruit farm during their 2021 harvesting season. This paper also compares these approaches to the method proposed in [12] and the teams that were manually created by farm managers. This paper is organised as follows. Section 2 highlights related work in the literature on MATA. Section 3 describes our approach to allocating human workers to teams. Section 4 explains the experiments we conducted, within a real-world scenario, in order to evaluate the impact of our approach. Section 5 presents and analyses our experimental results. Finally, we close (Sect. 6) with a summary of our contributions and directions for future work.

2 Background

A key challenge in multi-agent and multi-robot systems is to decide which *tasks* should be assigned to which agents so that the overall execution of a *mission* (set of tasks to be executed within a particular overall timeframe) is *efficient*: resources are used effectively, so that time and energy are not wasted and, often, some reward is maximised. Many different types of *task allocation* mechanisms have been explored within the multi-agent systems (MAS) and multi-robot systems (MRS) communities, generally addressing what are referred to as *MATA* or *multi-robot task allocation (MRTA)* problems.

An effective approach to MATA is the use of Genetic Algorithms (GA) [16]. GAs aim to minimise (or maximise) a fitness function by *evolving* a set of possible solutions—the population. A solution is represented as a "chromosome", (usually) a vector of values referred to as genes. To evolve the population, i.e. to create a new generation, the strongest (i.e. fittest) individuals are selected to survive and to be adapted. Various methods for representing a solution exist. For instance, the Traveling Salesman Problem (TSP) has been represented as a binary string, where each city is represented by a binary number and the order they appear in a chromosome is the order in which they are visited. As large problems require lengthy strings and many infeasible routes could be generated, often approaches use integer strings [22,29]. The fitness function for the TSP calculates the duration and distance associated with visiting each city in the sequence a chromosome represents.

A selection strategy is chosen to retain the most promising individuals and/or maintain diversity (by selecting a mixture of low- and high-ranked individuals). A wide range of methods have been proposed and evaluated within literature [2,10,17]. This includes truncation selection, where the best (least/most fit) individuals are selected; tournament selection, where the best individual from a random subset of population is selected, and fitness proportionate selection, where the probability of an individual being selected is based on their fitness. The two most common types of adaptation (also known as "reproduction operators" [15]) are crossover and mutation. Single-point crossover combines two individuals by randomly choosing a gene and swapping all genes that fall after this gene with the genes of the other individual. Researchers have also investigated multi-point crossover in which the crossing over starts/stops at multiple genes. In mutation, a gene's value is changed; for example, a binary gene can be changed from 0 to 1 or vice versa, or an integer gene could be changed to a different, randomly selected, value.

Multi-objective algorithms (MOA) aim to minimise (or maximise) multiple fitness functions (or objectives) [4,8,19]. Often MOA attempt to find the solutions that appear on the *pareto front* (i.e. a set of possible good solutions—depending on which objective is preferred). Nevertheless, as we intend our system to determine automatically which fields pickers will be assigned to, our approach must be able to return a single solution. In our work, this is performed by finding the weighted sum of the two objectives. This approach has been taken in prior work, such as [11], in which a path length is combined with the path smoothness (amount of rotation) to plan a robot's route.

Auction-based approaches to MATA are often controlled by a centralised auction manager. The auction manager announces item(s) to the bidders, the bidders place a value on the item(s), and the auction manager decides which bidders should be rewarded which item(s). Auctions usually repeat in "rounds" until all items have been allocated. Auctions take into account both the self-interests of individual bidders as well as group goal(s) represented by the auction manager—hence their popularity in multi-agent systems, which seek to balance both sets of, potentially conflicting, goals.

An *Sequential Single-Item (SSI)* auction [18] is a particularly popular method. In SSI, several tasks are announced to bidders at one time. Each bidder, responds with a bid representing the value (utility) of the task to them, incorporating cost to execute and potential reward. The auction manager then determines the winner by picking the bidder with the best (lowest/highest) bid for any task. SSI has been a popular choice for multi-robot task allocation, and many variants have been studied (e.g. [14,23–25,32,33]).

GAs have been applied to a wide range of domains. This includes deciding where to apply nitrates to soil [27], allocating jobs to high-performance computer resources [9], assigning search and rescue tasks after natural disasters [26] and finding appropriate learning experiences for students [37]. Likewise, auctions have been applied to many domains, including allocating harvesting tasks to agents within a single field [13], controlling traffic [30], managing ambulance dispatch [31] and planning the routes of heterogeneous robots [38]. The objective of our work is to automate the process of managing the workforce of a large commercial fruit farm. This workforce could consist of humans and/or robotic workers. Specifically, we investigate applying GAs, OSI and SSI to the complex problem of assigning workers to fields.

3 Approach

On a daily basis, farm managers decide which fields should be picked based on yield estimates and customer demand. They assign workers to "teams", each of which will harvest in one or more of the selected fields. The aim of our work is to decide which workers should be assigned to which field(s), saving farm managers from having to undertake this time-consuming job. This section introduces the problem and the fitness function, and then explains how GAs and auction-based approaches can each be applied to the problem.

To help address labour shortages and reduce farm expenditures, *staff time* (the total time worked by all staff each day) must be minimised. However, the solution with the minimum staff time has an unevenly distributed workload: the quickest worker is assigned to the field with the highest estimated yield, the second fastest to the field with the second highest estimated yield and so on. Each field has a single worker, except the field with lowest estimated yield, which all remaining workers are assigned to. This is problematic for two reasons. First, a worker picking alone will not manage to harvest the whole field—they will tire and/or run out of allowable time[1]. Since soft fruit must be harvested within a specific timeframe, unpicked fruit could spoil. Second, workers picking the field with the lowest yield would have relatively little work, and thus lower earnings. To prevent these problems, our approach also aims to minimise *execution time*—a proxy for evenness (and thus fairness), calculated as the duration of time between when the first picker starts work and when the last picker stops work on a given day. Our method assumes that enough pickers are available so that all ripe fruit can be harvested each day.

[1] Labour regulations restrict the number of hours per day a worker can work.

3.1 Problem Description

For each date within the picking season, the farm manager selects a set of fields (F) that require harvesting by the workers (W) who are available on that date. Each of the fields has an *estimated yield*. In our approach, groups of fields that are picked by the same team of workers are treated as one field (i.e. their estimated yields are summed). Our aim is to assign each worker ($w \in W$) to a field ($f \in F$), so that staff time and execution time are minimised. A solution is only valid if each worker is assigned to exactly one field (or group of fields) and no fields have zero workers.

Each of the workers has a set of skills, indicating how quickly they can pick each type of fruit. Their picking speed (in grams per second) is computed using data that is already produced by commercial fruit farms. This data is recorded so that a farm can pay the pickers by piece-rate (i.e. by the amount of fruit they pick). When a worker's picking speed is unknown for a particular type of fruit (e.g. they have not picked that type of fruit before), a picking speed of 1 is used. This is lower than the speed of any picker with experience.

3.2 Solution Fitness

We have developed a single fitness function that reflects and combines the two factors we aim to minimise: staff time and execution time. To calculate this we first need to introduce how to estimate the time it takes to pick a field. This section details the components of our fitness function.

The **estimated picking time** (*ept*) is calculated for each field ($f \in F$) selected for picking on a particular date (d), assuming it is picked by a specific team of workers ($f.W$). It is calculated by dividing the *estimated yield* (for field f on date d) by the sum of the picking speeds ($w.ps$) for the workers assigned to that field (Eq. 1). If a field has no workers, the *ept* of that field is infinity.

$$ept(f) = \begin{cases} \frac{f.estimated\_yield}{\sum_{w \in f.W} w.ps(f.\mathsf{fruit})}, & \text{if } |f.W| \geq 1 \\ \infty, & \text{otherwise} \end{cases} \tag{1}$$

The first factor we aim to minimise is the average *ept* of workers (*eptw*), representing *staff time*. This is calculated by summing the time all workers spend picking and dividing this by the number of workers. This is shown in Eq. 2, in which W is the set of all workers available on the date being scheduled.

$$eptw(F, W) = \frac{\sum_{f \in F}(ept(f) \times |f.W|)}{|W|} \tag{2}$$

Equation 3 calculates the second factor we aim to minimise: the average *ept* of the fields (*eptf*), which represents the *execution time*. The *ept* of the fields is summed and then this is divided by the number of fields.

$$eptf(F) = \frac{\sum_{f \in F} ept(f)}{|F|} \tag{3}$$

The weighted sum of *eptw* and *eptf* is calculated to find the fitness of a solution. Raw staff time (sum of the *ept* of workers) and execution time (maximum of the *ept* of fields) are not combined because staff time is a much greater value than execution time and is likely to undergo larger changes when the workers are rearranged. This would result in staff time being more dominant, whereas, *eptw* and *eptf* are of similar scale and can be balanced against each other. The fitness function is shown in Eq. 4. In Sect. 5.2, we discuss the results for different values of the linear weighting factor (α).

$$\mathit{fitness}(F, W, k) = (\mathit{eptw}(F, W) * \alpha) + (\mathit{eptf}(F) * (1 - \alpha)) \tag{4}$$

3.3 GA Approach

Our approach is implemented using the Jenetics library [39]. After creating the initial population, the following steps are repeated until the termination criteria has been reached (the best fitness remains unaltered for 5 generations):

1. Select survivors.
2. Select individuals for adaptation.
3. Adapt to create offspring (using mutation and single-point crossover).
4. Combine survivors and offspring to form new generation. 60% of the new generation contains the offspring and 40% is formed of the survivors. This represents a 60:40 exploration:exploitation ratio.

For our problem, the size of a chromosome is equal to the number of workers ($|W|$), with each gene representing which field each worker is assigned to. A gene is an integer value in the range 0 to $|F| - 1$. An example, for a problem with 5 workers and 3 fields, is shown in Fig. 1.

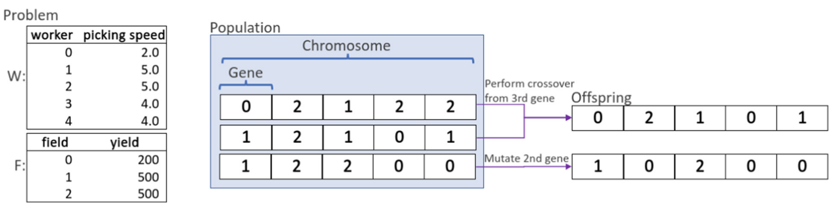

Fig. 1. Example population representing 3 possible solutions for assigning 5 workers to 3 fields. For the first chromosome, the first worker is assigned to field 0; the second, fourth and fifth to field 2, and the third to field 1. The offspring on the top right shows the result when single-point crossover (from 3rd gene) is performed on the top two chromosomes. The fitness of the original chromosomes are (top down) 71.28, 74.54 and 63.33; the offspring have fitness (top down) of 59.26 and 60.51. In this case, crossover and mutation have produced individuals who are stronger (have a lower fitness) than those in the original population. This example assumes that there is a single type of fruit; multiple types of fruit are used in our experiments.

3.4 Auction Approach

Typically, in auction methods, an auction manager advertises items (e.g. tasks) to the agents. The agents bid on these items and the auction manger decides which bid wins, and thus which agent is assigned the items(s) they have bid on. For our problem domain, it is the fields that bid on the workers. Each agent (field) aims to minimise its own *ept* (the individual's goal); whereas, the auction manager aims to minimise the *fitness* (the team's goal). Thus, the cost of a bid is the field's *ept* if it were to be assigned the worker being bid on. The auction manager replaces the *ept* of this field when calculating the fitness (Eqs. 2, 3 and 4) with the bid's cost. We have implemented two auction-based approaches.

- *Ordered Single Item (OSI)* : Workers are sorted slowest first, and one worker at a time is advertised to the fields. The auction manager assigns the field whose bid produces the lowest *fitness* the worker being bid on.
- *Sequential Single Item (SSI)* : All unassigned workers are advertised to all fields and each field responds with a bid for the worker that results in the lowest *ept* (i.e. for the worker that has the quickest picking speed for the type of fruit the field contains). As with OSI, the auction manager assigns the field whose bid produces the lowest *fitness* the worker the field bid on.

To handle tiebreaks, if two bids have an equally low fitness, then the winner is the field that has the fewest agents. If the fields also have the same number of agents, the field with the highest yield wins. This guarantees that each field will be assigned at least one worker (if $|W| \geq |F|$).

4 Experiments

Our experiments are designed to evaluate the effectiveness, for managing a harvest workforce, of our GA approach compared to two auction-based approaches: OSI and SSI. First, we evaluate the impact changing the linear weighting factor (α) has on the execution time and staff time. Then we compare the methods to the RR variant proposed in [12] and the actual teams created by farm mangers.

4.1 Data

A commercial fruit farm provided us with data during their 2021 harvesting season, which involved 182 picking days. This data was for strawberry, raspberry, cherry, and blackberry fields, a total of 30 fields. Note, this data has been used within our previous work [12], and thus processing this real-world data is not novel here. The following was provided incrementally during the 2021 season:

- **Estimated yield :** A spreadsheet containing an approximate volume of ripe fruit for each field the farm plans to pick on the morrow.
- **Worker list :** A list of available workers.
- **Recorded picking data :** The amount of fruit picked so far by each picker. This data is used to calculate the picking speeds of the workers, and to extract the actual teams (that were manually created by farm managers).

4.2 Metrics

To evaluate our proposed approaches, for each picking day, we calculate *staff time*—the sum of the *ept* of all workers across all fields, and *execution time*—the maximum *ept* of the fields. For determining the significance of our results, we applied statistical testing and factor analysis, where appropriate. A Shapiro-Wilk test [36] was performed to check if each sample is normally distributed. For all samples there is 95% chance that the sample is normally distributed, thus ANalysis Of VAriance (ANOVA) tests [1,7] were performed (for which the F test statistics are reported). The significance of results is indicated by p, the probability of the results occurring randomly. All plots show mean; the error bars indicate the standard deviation.

To tune our GA, we ran experiments to choose amongst 8 different selectors, 15 initial population sizes, 10 mutation probabilities and 15 crossover probabilities. The setup we found produced the lowest fitness is used for the experiments presented here: truncation selection, population size of 1500, crossover probability of 0.9 and mutation probability of 0.0 (i.e., no mutation). Our GA approach was ran 5 times. The result of the run that produced the lowest fitness is shown here since this is for the team that would be recommended to farm managers. OSI, SSI and the RR variant [12] are deterministic, and thus were ran once.

5 Results

5.1 Trade-off Between Staff Time and Execution Time

As mentioned in Sect. 3, only taking into account staff time causes the workload to be unbalanced. Our results, shown in Fig. 2, demonstrate this. When $\alpha = 1.0$ (the point at the top left of the plots), the lowest staff time was achieved; however, this also resulted in the highest execution time. As α was decreased, execution time decreased and staff time increased.

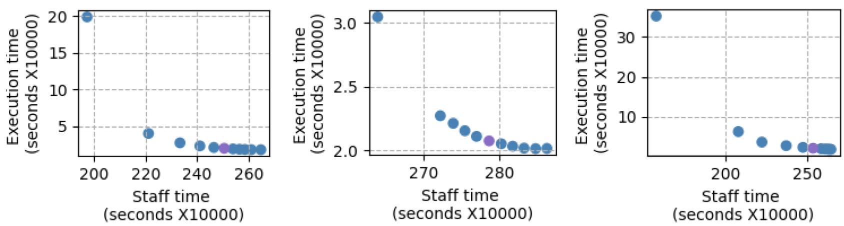

(a) GA **(b)** OSI **(c)** SSI

Fig. 2. Scatter plots showing the trade-off between the two criteria: staff time and execution time. The data points in purple indicate the results used in the comparison shown in Fig. 3. Note the different x and y axis ranges

5.2 Comparison to Alternative Approaches

The results for comparing our GA, OSI and SSI approaches with $\alpha = 0.5$ to the RR variations proposed in [12] (i.e. RR3) and to the teams manually created by farm managers (Actual) are shown in Fig. 3.

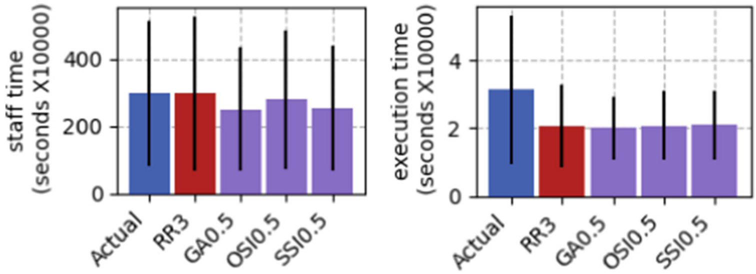

(a) (F=2.24, p=0.063) **(b)** (F=22.58, p=0.000)

Fig. 3. The estimated staff time (left) and execution time (right) for the Actual teams, the RR3 approach proposed in [12], and the GA, OSI and SSI approaches described in this paper for $\alpha = 0.5$. The sample size is equal to the number of picking days. (Note, that Actual uses the actual picking days and actual yield; whereas, the remaining approaches use the estimated data.)

Our GA approach achieved the lowest staff time and execution time. SSI achieved the second lowest staff time and OSI had the second lowest execution time. The difference in staff time is not statistically significant, but the difference in execution time is. Nevertheless, based on the mean staff times, if all workers were paid the UK hourly minimum wage (of £9.50), employing the teams proposed by our GA approach could save the farm an average of £1229.72 per day. The downside of the GA approach is that, for these experiments, it took longer to run than the other approaches. Often there is a trade-off between run-time and solution quality. Within our application domain, when picking has finished for the day, farm managers send us the data required to created the team allocations for the following day. Thus far, all approaches we have tested would produce a result within a timely manner; we therefore, focus on which created the most efficient team allocations. In future experimentation, the farm managers' opinions on the run-time of our approach will be gathered.

6 Summary and Future Work

This paper explored applying Genetic Algorithms (GAs), Ordered Single Item (OSI) auctions and Sequential Single Item (SSI) auctions to the problem of allocating workers to the fields selected for harvesting. We proposed a fitness function that combines minimising the average staff time of workers with minimising the average execution time of the fields. Minimising staff time causes the workload to become unevenly distributed, and this execution time must also be considered. Our experiments compared our GA, OSI and SSI approach to the Round Robin (RR) variant proposed in [12] and the teams manually created by farm managers during the 2021 harvesting season. The resulting staff time and execution time of the GA approach was less than the alternatives.

During future work, we will explore the seasonal variation in our results. For instance, at the start of the season, when there are few fields with ripe fruits and few workers, a lower initial population could be needed than in the height of the season. During the upcoming (2022) harvesting season, we are also planning to perform further trials with our approach. This will hopefully include deploying the teams our system proposes on a real-world farm.

Acknowledgments. This work was supported by Research England [Lincoln Agri-Robotics] as part of the Expanding Excellence in England (E3) Programme and by Ceres Agri-tech.

References

1. Anscombe, F.: The validity of comparative experiments. J. Roy. Stat. Soc. Ser. A (General) **111**(3), 181–211 (1948)
2. Blickle, T., Thiele, L.: A comparison of selection schemes used in evolutionary algorithms. Evol. Comput. **4**(4), 361–394 (1996). https://doi.org/10.1162/evco.1996.4.4.361
3. Das, G., Cielniak, G., From, P., Hanheide, M.: Discrete event simulations for scalability analysis of robotic in-field logistics in agriculture-a case study. In: IEEE International Conference on Robotics and Automation, Workshop on Robotic Vision and Action in Agriculture (2018)
4. Deb, K., Jain, H.: An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: Solving problems with box constraints. IEEE Trans. Evol. Comput. **18**(4), 577–601 (2014). https://doi.org/10.1109/TEVC.2013.2281535
5. Doward, J., Baldassari, V.: Red Alert: UK farmers warn of soft fruit shortage. The Guardian (2018)
6. Duckett, T., Pearson, S., Blackmore, S., Grieve, B., Smith, M.: Agricultural Robotics White Paper: The Future of Robotic Agriculture (2018). www.ukras.org/wp-content/uploads/2018/10/UK_RAS_wp_Agri_web-res_single.pdf. Accessed 10 Mar 2020)
7. Fisher, R.A.: Statistical methods for research workers (1925)
8. Fonseca, C.M., Fleming, P.J.: Genetic algorithms for multiobjective optimization: formulation discussion and generalization. In: Proceedings of the 5th International Conference on Genetic Algorithms, pp. 416–423. Morgan Kaufmann Publishers Inc., San Francisco (1993)
9. Gabaldon, E., Lerida, J.L., Guirado, F., Planes, J.: Multi-criteria genetic algorithm applied to scheduling in multi-cluster environments. J. Simul. **9**(4), 287–295 (2015). https://doi.org/10.1057/jos.2014.41
10. Goldberg, D.E., Deb, K.: A comparative analysis of selection schemes used in genetic algorithms. In: Foundations of Genetic Algorithms, vol. 1, pp. 69–93. Elsevier (1991). https://doi.org/10.1016/B978-0-08-050684-5.50008-2

11. Hao, K., Zhao, J., Yu, K., Li, C., Wang, C.: Path planning of mobile robots based on a multi-population migration genetic algorithm. Sensors **20**(20) (2020). https://doi.org/10.3390/s20205873
12. Harman, H., Sklar, E.: Multi-agent task allocation for fruit picker team formation. In: Proceedings of the 21th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2022, International Foundation for Autonomous Agents and Multiagent Systems (2022)
13. Harman, H., Sklar, E.I.: A practical application of market-based mechanisms for allocating harvesting tasks. In: Dignum, F., Corchado, J.M., De La Prieta, F. (eds.) PAAMS 2021. LNCS (LNAI), vol. 12946, pp. 114–126. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-85739-4_10
14. Heap, B., Pagnucco, M.: Repeated sequential single-cluster auctions with dynamic tasks for multi-robot task allocation with pickup and delivery. In: Klusch, M., Thimm, M., Paprzycki, M. (eds.) MATES 2013. LNCS (LNAI), vol. 8076, pp. 87–100. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40776-5_10
15. Holland, J.H.: Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence. University of Michigan Press (1975)
16. Holland, J.H.: Genetic algorithms. Sci. Am. **267**(1), 66–73 (1992). https://www.jstor.org/stable/24939139
17. Katoch, S., Chauhan, S.S., Kumar, V.: A review on genetic algorithm: past, present, and future. Multimedia Tools Appl. **80**(5), 8091–8126 (2020). https://doi.org/10.1007/s11042-020-10139-6
18. Koenig, S., et al.: The power of sequential single-item auctions for agent coordination. In: Proceedings of AAAI, vol. 2 (2006)
19. Konak, A., Coit, D.W., Smith, A.E.: Multi-objective optimization using genetic algorithms: a tutorial. Reliabil. Eng. Syst. Safety **91**(9), 992–1007 (2006). https://doi.org/10.1016/j.ress.2005.11.018
20. Kootstra, G., Wang, X., Blok, P.M., Hemming, J., Van Henten, E.: Selective harvesting robotics: current research, trends, and future directions. Curr. Rob. Rep. **2**, 95–104 (2021)
21. Kurtser, P., Edan, Y.: Planning the sequence of tasks for harvesting robots. Rob. Auton. Syst. **131** (2020)
22. Larranaga, P., Kuijpers, C.M.H., Murga, R.H., Inza, I., Dizdarevic, S.: Genetic algorithms for the travelling salesman problem: a review of representations and operators. Artif. Intell. Rev. **13**(2), 129–170 (1999). https://doi.org/10.1023/A:1006529012972
23. McIntire, M., Nunes, E., Gini, M.: Iterated multi-robot auctions for precedence-constrained task scheduling. In: AAMAS 2016 - Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems, pp. 1078–1086. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS) (2016)
24. Nunes, E., Gini, M.: Multi-robot auctions for allocation of tasks with temporal constraints. In: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI 2015, pp. 2110–2216. AAAI Press (2015). https://dl.acm.org/doi/10.5555/2886521.2886614
25. Nunes, E., McIntire, M., Gini, M.: Decentralized allocation of tasks with temporal and precedence constraints to a team of robots. In: 2016 IEEE International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAR), pp. 197–202 (2016)

26. Pallin, M., Rashid, J., Ögren, P.: A decentralized asynchronous collaborative genetic algorithm for heterogeneous multi-agent search and rescue problems. In: 2021 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR), pp. 1–8 (2021). https://doi.org/10.1109/SSRR53300.2021.9597856
27. Peerlinck, A., Sheppard, J., Pastorino, J., Maxwell, B.: Optimal design of experiments for precision agriculture using a genetic algorithm. In: 2019 IEEE Congress on Evolutionary Computation (CEC), pp. 1838–1845 (2019). https://doi.org/10.1109/CEC.2019.8790267
28. Pelham, J.: The Impact of Brexit on the UK Soft Fruit Industry. British Summer Fruits, London (2017)
29. Potvin, J.Y.: Genetic algorithms for the traveling salesman problem. Ann. Oper. Res. **63**(3), 337–370 (1996). https://doi.org/10.1007/BF02125403
30. Raphael, J., Sklar, E.I.: Towards dynamic coalition formation for intelligent traffic management. In: Belardinelli, F., Argente, E. (eds.) EUMAS/AT -2017. LNCS (LNAI), vol. 10767, pp. 400–414. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01713-2_28
31. Schneider, E., et al.: The Application of Market-based Multi-Robot Task Allocation to Ambulance Dispatch. arxiv.org/abs/2003.05550 (2020)
32. Schneider, E., Sklar, E.I., Parsons, S.: Evaluating multi-robot teamwork in parameterised environments. In: Proceedings of the 17th Towards Autonomous Robotic Systems (TAROS) Conference, pp. 301–313 (2016)
33. Schneider, E., Sklar, E.I., Parsons, S., Özgelen, A.T.: Auction-based task allocation for multi-robot teams in dynamic environments. In: Dixon, C., Tuyls, K. (eds.) TAROS 2015. LNCS (LNAI), vol. 9287, pp. 246–257. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-22416-9_29
34. Seyyedhasani, H., Peng, C., Jang, W.J., Vougioukas, S.G.: Collaboration of human pickers and crop-transporting robots during harvesting - part i: model and simulator development. Comput. Electron. Agric. **172**, 105324 (2020)
35. Shamshiri, R.R., Hameed, I.A., Karkee, M., Weltzien, C.: Robotic harvesting of fruiting vegetables: a simulation approach in V-REP, ROS and MATLAB. In: Proceedings in Automation in Agriculture-Securing Food Supplies for Future Generations (2018)
36. Shapiro, S.S., Wilk, M.B.: An analysis of variance test for normality (complete samples). Biometrika **52**(3/4) (1965)
37. Sklar, E.I.: CEL: a framework for enabling an internet learning community. Ph.D. thesis, Brandeis University (2000)
38. Sullivan, N., Grainger, S., Cazzolato, B.: Sequential single-item auction improvements for heterogeneous multi-robot routing. Rob. Auton. Syst. **115**, 130–142 (2019). https://doi.org/10.1016/j.robot.2019.02.016
39. Wilhelmstötter, F.: Jenetics (2022). http://jenetics.io/. Accessed 21 Mar 2022