

Multiobjective Routing in Sustainable Mobility-On-Demand

Mengya Liu, Vahid Yazdanpanah, Sebastian Stein and Enrico Gerding

Agents, Interaction and Complexity Research Group, University of Southampton, SO17 1BJ, Southampton, UK

Abstract

It is estimated that smart on-demand mobility services can significantly reduce emissions caused by urban transportation, especially when combined with the use of low emission vehicles and ridesharing. While current research on sustainable routing typically focuses on *economic* sustainability (as cost minimisation), this paper also considers the other pillars of sustainability, i.e., the *environmental* and *social* aspects of what we call sustainable and equitable Mobility-On-Demand (MOD). To that end, we apply multiobjective genetic algorithms and generate routing options that balance all three pillars of sustainability. We envisage that a diverse set of routing solutions allows participation of end-users in determining an equitable route (e.g., through voting processes) and strongly supports widespread adoption of sustainable MOD and ridesharing services. This work follows principles of human-centred intelligent systems and provides a foundation for building participatory, dynamic, and explainable MOD systems.

Keywords

Ridesharing, Multiobjective Algorithm, Mobility-on-demand, Sustainable Transportation, Evolutionary Computation

1. Introduction

Mobility-On-Demand (MOD) services traditionally aim to reduce the economic and environmental cost of transportation [1, 2]. Roughly speaking, MOD promises to utilise the mobility capacity in a more efficient way (leading to reductions in the collective and individual economic costs), while also reducing emissions caused (e.g., by avoiding single-driver journeys). Therefore, MOD systems and methods to support their implementation directly contribute to Sustainable Development Goal (SDG) 11 and 13 on *sustainable cities and communities* and *climate action*, respectively. If MOD is widely adopted, cities will enjoy its environmental and economic benefits and take a step towards mitigating climate change. However, if such services merely focus on what is optimal from the service providers' perspective and ignore users' requirements, it will be difficult to encourage users to move towards this service and ignore the comfort of using personal vehicles. Such a sustainability-oriented transition necessitates looking not only at operators' (economic) criteria, but also evaluating how a particular routing choice may affect

ATT'22: Workshop Agents in Traffic and Transportation, July 25, 2022, Vienna, Austria


✉ Mengya.Liu@soton.ac.uk (M. Liu); v.yazdanpanah@soton.ac.uk (V. Yazdanpanah); ss2@ecs.soton.ac.uk (S. Stein); eg@ecs.soton.ac.uk (E. Gerding)

🌐 <https://www.ecs.soton.ac.uk/people/vy1y20> (V. Yazdanpanah); <https://www.ecs.soton.ac.uk/people/ss1y07> (S. Stein); <https://www.ecs.soton.ac.uk/people/eg> (E. Gerding)

🆔 0000-0002-4468-6193 (V. Yazdanpanah); 0000-0003-2858-8857 (S. Stein); 0000-0001-7200-552X (E. Gerding)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

individuals, e.g., via their total travel or waiting times. In this work, we argue that *sustainable MOD* needs to capture all the three pillars of economic, environmental, and social sustainability [3] and aim for routing solutions that are balanced with respect to all three aspects. The economic and environmental aspects call for minimising respective costs, both at a collective level, but also (to ensure fairness) for all the riders. Finally, the social aspect requires considering fairness in distributing tasks among drivers such that riders receive a balanced workload. Thus, addressing the routing problem in sustainable MOD requires a multiobjective approach that captures potential trade-offs among different aspects of sustainability for generating sustainable routing options.¹

As discussed in related work, e.g., [5, 2, 6], the multidimensionality and complexity of the MOD routing problem, and, in our case, in view of the three pillars of sustainability, result in inapplicability of exact multiobjective optimisation techniques and justifies using genetic algorithms (GA). While [7] explored adopting reinforcement learning for multiobjective optimisation, this work considers GA to provide a diverse range of solutions (for a diverse set of users). Multiobjective GA allows capturing various objectives with fewer compromises regarding scalability [8]. In particular, we use a form of the Non-dominated Sorting Genetic Algorithm (NSGA) that is proven to be effective in various mobility settings [9].

Against this background, for the first time in this work, we capture the *social* and *environmental* aspects of sustainable routing in MOD and use genetic algorithms for generating routing options that consider all three pillars of sustainability. This is the first approach that integrates these two pillars of sustainability into traditional models of (purely) *economic* sustainability and generates routing solutions under six sustainable ridesharing objectives: travelling time, waiting time, overall/excess distance, travel cost, total emission, and working time balance. The list of sustainable routing options can be used in a “*user participation*” phase (e.g., in ridesharing services), where riders can select their desired route from a list of options. Using our approach, service providers can generate routing options that balance all the three pillars of economic, environmental, and social sustainability. In addition, they can provide routing options that reflect riders’ preferences (e.g., by focusing on the environmental dimension). This is a step towards integrating equitability and user participation [10] into sustainable MOD practices.

2. Sustainable MOD

The MOD routing problem is to allocate a given set of riders to vehicles with respect to different objectives. Each rider requires a ride from its starting point to its destination, along with a specified earliest departure time for taking a ride. Each vehicle can take a limited number of riders aboard at the same time, excluding the driver, which we refer to as its capacity. And the driving costs and capacity of a vehicle are associated with its type. A solution to the MOD routing problem is an arrangement that sends vehicles to pick up riders at their starting point, and then drops them off at their destination. The objectives evaluate the efficiency of a solution. In the following, we present the mathematical notations used for modelling the MOD problem

¹In view of human-centred AI techniques and the need for developing trustworthy human-AI partnerships [4], we see sustainable MOD as an inherently sociotechnical problem and argue that its acceptance by society depends on the ability to capture all the three aspects of economic, environmental, and social sustainability.

and developing our approach.

2.1. Ride Requests

In the MOD routing problem, all the riders post their ride requests at the beginning. Let R be the set of posted requests and r represent a single request. To locate riders, we adopt a graph structure to model the real-world map, i.e., a map graph is $G = (N, E)$, where nodes in N represent the intersections on a real-world map and edges in E that link nodes together represent the roads between intersections. In general, to represent an intersection, a node is in the form of a tuple marking the latitude and longitude of the intersection. Thereby, let $r(s, t, u)$ denote a rider's request for a ride from node s to node t along with an earliest time for the rider to leave, u .

2.2. Features of Vehicles

To model the sustainability of the MOD routing problem from economic, environment and social aspects, this work considers 3 features of a vehicle as well as its location. Let $v(s, t, p, e, c)$ denote a vehicle that starts working at node s , returns to node t at the end of its service with a physical capacity of p , emission level of e and a travelling cost per kilometre, c , including the driver wage, vehicle maintenance and fuel costs. Regardless of the difference of vehicle brands, there is a positive correlation between c and p . Hence, we assume that $c = \alpha \times p$ for a vehicle and α varies according to the type of vehicle. For a pessimistic estimation, we use $\alpha = 1$ in our experiments. Besides, the emission level of a vehicle e is a vector, and the i th element in the vector, e_i , denotes the emission rate of a vehicle when there are i riders aboard, since different numbers of riders aboard cause different emission rate. Specifically, in this work, the emission of a vehicle generally includes greenhouse gas (GHG) and air pollution. With respect to the reports from the UK's Department for Transport [11] and National Atmospheric Emission Inventory (NAEI) [12], and the EU standard vehicle emissions calculator, COPERT [13], the GHG and air pollution emission of a vehicle per kilometre are mainly related to the fuel and type of a vehicle, but the emission of GHG also depends on the number of passengers. Therefore, we model the emission level of a vehicle as a vector, where each element represents a emission rate associated with a number of riders on board. Notice that we also assume that all vehicles will drive at the same speed to simplify the problem. This can be simply extended to simulate a dynamic speed by varying the speed in a range of minimum to maximum urban/legal speed.

Table 1
Estimated Features for Different Types of Vehicle

	Vehicle Type	Fuel	Capacity	Vector of Emission Level	Cost
1	Car	Petrol	1	$\langle 1, 2 \times 0.9 \rangle$	1
2	Medium Car	Electric	3	$\langle 1/3.4, 2/3.4 \times 90\%, 3/3.4 \times 90\%^2, 4/3.4 \times 90\%^3 \rangle$	3
3	Large Car	Petrol	10	$\langle 1, 2 \times 90\%, \dots, 11 \times 90\%^{11} \rangle$	10

The features of a vehicle, such as capacity, emission level and travelling cost, depend on its

type. We consider 3 types of general passenger vehicles according to the vehicle categories specified on the UK Driving Licence Categories [14]: Small cars, medium-sized vehicles, large vehicles². Table 1 lists their features. To simulate the emission level of different types of vehicles, we use the *car* type with a petrol engine and one passenger on board as the standard, and assume it emits 1 unit of greenhouse gas and air pollution per kilometre (e.g., 1 unit could be 100g of CO₂) and costs 1 price unit per kilometer (e.g., \$1). According to the Transport and Environment Statistics [11], “an average petrol car emits around 4 times more per passenger than the equivalent journey by coach, or 3.4 times more per passenger emitted by the average electric car”, and “maximising the number of people per vehicle can reduce emissions per person”. Hence, we assume that the average emission per passenger reduces by 10% when the number of passengers aboard increases and the cost of vehicles are related with its capacity. The emission column in Table 1 lists the emission vector for each vehicle type where elements in the vector are emissions of a vehicle in the order of 0 passengers to its full capacity.

Our focus in this work is to demonstrate the impact of emissions of vehicles, and we are aware of the existence of other types of vehicles, such as motorcycles, and different types of emission calculators [16]. The types of vehicles and the emission estimation considered in this work are standard types and presented for the purpose of showing the performance of the approach.

2.3. Solution Design

Our intelligent routing approach is based on genetic algorithms [17]. First, we model an arrangement that a vehicle picks up and drops off a rider as a genetic chromosome: (vehicle, weight of picking up priority, weight of dropping off priority), and a solution to the MOD routing problem of arranging vehicles for m riders as:

$$\begin{bmatrix} r_1 : & (v^1, w^{1[s]}, w^{1[t]}) \\ r_2 : & (v^2, w^{2[s]}, w^{2[t]}) \\ \dots & \dots \\ r_m : & (v^m, w^{m[s]}, w^{m[t]}) \end{bmatrix} \quad (1)$$

For rider r_i , v^i denotes the vehicle that serves r_i a ride, $w^{i[s]}$ is a positive real number that represents the priority weight of picking up r_i , and $w^{i[t]}$ is a real number that indicates the priority weight of r_i to get off v^i at its destination. A higher priority weight implies a greater sense of urgency to start or finish a ride. Thus, for a vehicle that offers a ride to multiple riders, it will stop at nodes to pick up and drop off the riders with respect to their priority weights. In reality, the priority weight can be the time of a ride request, arriving time or even promotion tips.

2.4. Solution to Route

To calculate the routes for vehicles there are two factors to satisfy: feasibility and uniqueness. Feasibility requires that when a vehicle follows a route to pick up and drop off riders, the number

²we exclude minibuses and buses [15]. Since they have stable routes and we do not consider asking riders to change vehicles during their rider, they left no space for picking riders at their starting points.

of riders at any given time must not exceed the capacity of the vehicle. Uniqueness requires that, given a solution, one should be able to derive one and only one way to route the vehicles from it. The uniqueness criterion is necessary because it is the solutions that the genetic algorithm evaluates and optimises while objectives in the evaluation and optimisation are based on the routes for vehicles. Thus, to be able to evaluate a solution, we require a 1-1 correspondence with routes that the solution entails. We will explain our method to map a solution to a feasible and unique routing in the following, and introduce the objectives afterwards.

First, to capture the meaning of the priority weights in a solution, we use the following two rules when comparing the priority weights of different riders:

1. No consideration of dropping off priority for riders who are waiting for pick-up.
2. For riders with equal priority weights, the rider with a lower index number is prioritised, considering that the rider posted its ride request earlier.

Regarding a solution, let $Passenger(v)$ denote the set of riders to whom vehicle v offers a ride, and $Passenger(v) = \{r_i | v^i = v\}$. The route of v is a sequence of nodes, $Path(v)$, that are either starting points or destinations of riders in $Passenger(v)$, and the path that a vehicle travels from a node to another is the shortest path calculated by Dijkstra's algorithm [18]. While a vehicle travels, let $Board(v, n)$ denote the riders that are on the vehicle v when it visits node n , and $Hold(v, n)$ be the riders that are still waiting for the vehicle for a pick-up.

Figure 1 displays our routing algorithm that maps a solution to the routes for a vehicle.³ The very first node in the route of a vehicle is its starting point, and at that node, the boarding passengers is null and all the riders assigned to it are on hold. Then, the next node in the route of the vehicle depends on the priority weights of the aboard and waiting riders. First, we compare the waiting riders' weights of picking up priority and get the top one waiting rider, and then compare the aboard riders' weights of dropping off priority and get the top one aboard rider. If the aboard rider's weight of dropping off priority is greater than the waiting rider's weight of picking up priority, the next node in the route of the vehicle is the destination of the aboard rider. Otherwise, we check whether picking up the waiting rider violates the vehicle's capacity constraint. If not, the vehicle will travel to the starting point of the waiting rider and pick it up. If yes, the vehicle still needs to drop off the aboard rider first. Until there is no rider aboard or waiting, the vehicle will travel to its destination and terminates its route.

This routing algorithm guarantees the feasibility of the travel paths of vehicles generated from a solution and the uniqueness of the generation dynamically. Note that, regardless of the uniqueness, it is still possible that different solutions generate the same routes for one or even more vehicles. This is because the different weights of either picking up priority or dropping off priority can result in the same ranking of the riders in the algorithm. Note that the potential redundancies are left to be resolved in the genetic algorithm.

2.5. Objectives

Regarding the travel paths of vehicles and their loaded riders, we introduce and minimise 6 objectives from 3 aspects of sustainability: economic, environmental and social. The economic

³For a complete implementation of our routing algorithm, please refer to <https://github.com/Miya-Liu/equitable-ridesharing>.

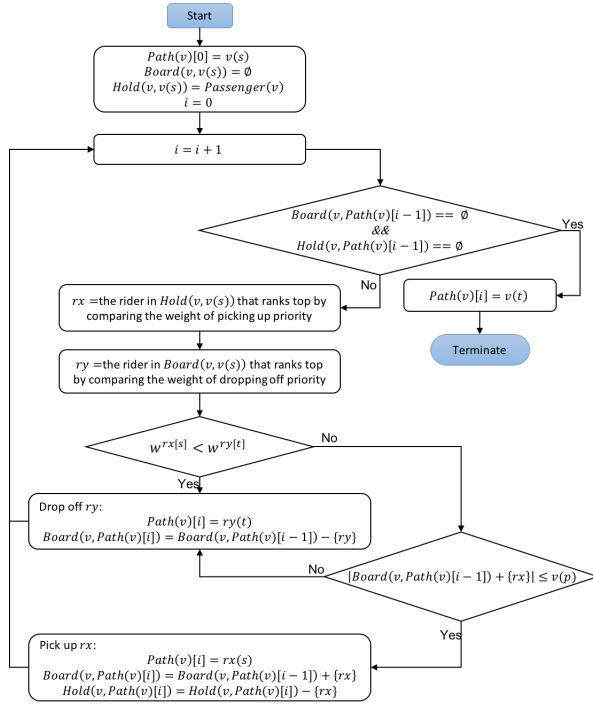


Figure 1: Algorithm for Mapping Solution to Routes.

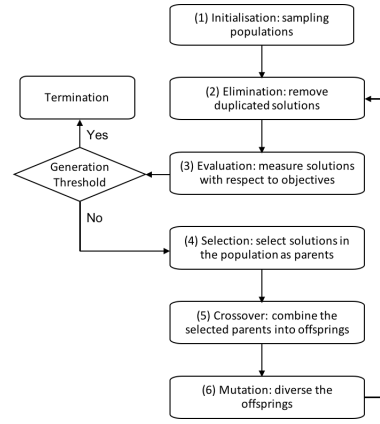


Figure 2: Modified NSGA3 for Sustainable MOD.

aspect evaluates the efficiency of the routes generated from a solution with respect to travelling time, waiting time, travel cost and excess distance. Then, the environmental aspect of sustainability considers the impact of vehicles' emission and tries to minimise the total emission of rides. Finally, the social aspect concerns the working time of drivers and aims at reducing the differences among the working time of all drivers.

Economic - Travelling Time (et): This objective measures the total travelling time of individual riders. The measurement of the travelling time for one rider is the time that it takes from the moment the rider gets on a vehicle until the vehicle drops off the rider at her destination. This includes the time that the vehicle travels and waits to pick up and drop off other riders while the rider is on board. The waiting time of a vehicle includes time periods when the vehicle stops at a starting point of a rider to pick her up. Such a waiting takes place when a vehicle arrive (too) early, i.e., when the arriving time of the vehicle is earlier than the earliest leaving time of the rider. Less travelling time means that the riders entails a more efficient trip.

Let $wait(v, n)$ denote the time that a vehicle v waits for picking riders up at node n along its route. Let $d(n_i, n_j)$ be the shortest distance between node n_i and n_j , and $arrive(v, n)$ represent the time that the vehicle arrives at node n along its route. Hence, $wait(v, n_0) = 0$, $arrive(v, n_0) = 0$, and

$$arrive(v, n_i) = arrive(v, n_{i-1}) + wait(v, n_{i-1}) + \frac{d(n_{i-1}, n_i)}{speed},$$

where $n_i = Path(v)[i]$, and $speed$ is a given average speed. Assume that v will pick up k riders at n_i , thus, $wait(v, n_i) = \max\{0, r_x(u) - arrive(v, n_i)\}$, where $r_x(u)$ is the greatest earliest leaving time among the k riders. Therefore,

$$ET = \sum_r (arrive(v^r, r(t)) - arrive(v^r, r(s))). \quad (2)$$

Economic - Waiting Time (EW): This objective is to evaluate the waiting time for all riders before vehicles pick them up with respect to their earliest leaving time.

$$EW = \sum_r \max\{0, arrive(v^r, r(s)) - r(u)\} \quad (3)$$

Economic - Excess Distance (ED): This objective measures the extra distance that a vehicle travels when it needs to pick up and drop off riders compared to the distance of directly driving from its starting point to the destination. Let n_i be the i th node in a route,

$$ED = \sum_v \left(\sum_{i=0}^{|Path(v)|-1} d(n_i, n_{i+1}) - d(v(s), v(t)) \right) \quad (4)$$

Economic - Travel Cost (EC): This objective measures the cost of all rides in total.

$$EC = \sum_v (v(c) \times \sum_{i=0}^{|Path(v)|-1} d(n_i, n_{i+1})) \quad (5)$$

Environmental - Emission (SE): This objective is designed to measure the emission of all the vehicles. By minimising this objective, sharing a ride can reduce pollution. Recall that the emission rate of a vehicle is related to the number of passengers on the vehicle. Therefore, the emission of all the vehicles regarding one solution is

$$SE = \sum_v \sum_{i=0}^{|Path(v)|-1} v(e)[l(v, n_i)] \times d(n_i, n_j). \quad (6)$$

where $l(v, n) = |Board(v, n)|$.

Social - Working Time (SW): The working time is calculated from the moment a vehicle leaves its starting point until it arrives at its destination, which is $arrive(v, v(t))$. This objective demonstrates the workload of a vehicle. Regarding the social sustainability, this work tries to balance the workload among all drivers and ensure a sustainable MOD service that is fairness-aware. Hence, this objective is defined as the Gini coefficient [19] of all vehicles' working time, $W = \{arrive(v_1, v_1(t)), arrive(v_2, v_2(t)), \dots, arrive(v_m, v_m(t))\}$ as follow.

$$SW = Gini(W). \quad (7)$$

With the above-defined multiobjectives, we will later explain our algorithm that generates multiple routing options that balance the six objectives of all three pillars of sustainability.

3. NSGA3 for Sustainable MOD

This work adopts an existing genetic algorithm called Non-dominated Sorting Genetic Algorithm 3 (NSGA3) [20] for dynamic routing in the sustainable MOD problem. Figure 2 shows the workflow of the NSGA3 with modifications for the MOD setting.

NSGA3 requires no configurations of the importance or weights of multiple objectives in the optimisation, but balance them automatically. The optimisation procedure includes:

1. **Sampling:** It generates an initial sample population. In this step, for each solution in the population, we randomly assign vehicles to serve riders, and assign random values as the weights of the picking up and dropping off priority of riders.
2. **Selection:** It selects some solutions as parents for generating offspring in next generation. NSGA3 uses a reference points [21] based selection operator. As we applying this genetic algorithm for multiobjective optimisation, this selection is ideal as it is guided by specifying a set of well-maintain diversity in the population regarding different objectives.
3. **Crossover:** It combines the selected parents to generate offspring. We define the crossover as two parent solutions generating one offspring. The pattern to generate an offspring for each pair of parents is to use the first half chromosomes from a parent and the second half chromosomes from the other parent to generate an offspring solution. Note that our implementation supports splitting both parents into any number of slices and then selecting the same number of slices to generate an offspring. However, the efficiency evaluation of crossover patterns is out of the scope of this work.
4. **Mutation:** It mutates offspring to increase the diversity of the current population. The modified mutation is: for each offspring, we select half riders and change the value of its corresponding chromosomes by (1) changing the vehicle assigned to a rider; (2) increasing its weight of picking up priority by a positive number; (3) increasing its weight of dropping off priority by a random positive number.
5. **Elimination:** It deletes duplicate solutions. And if the size of the current population after elimination is smaller than the initial population, the crossover process is repeated until the desired number of offspring is fulfilled.

We set the threshold of the number of generations as the condition to terminate the optimisation.

This approach will automatically generate multiple routing solutions when we set the size of population greater than 1. In addition, the routing solutions are feasible and balance the economic, environmental and social sustainability, while the algorithm optimise the six objectives. Note that we did not consider the objectives when calculating the shortest path among all nodes of a map graph. This is because the objectives are defined and calculated based on the paths of all the vehicles that they will drive, pick up, and drop off riders. For instance, we cannot get the override distance just for the edge between two nodes.

4. Experimental Evaluation

The main goal of this work is to demonstrate the impact of sustainable objectives in MOD routing and the efficiency of our GA-based routing approach. We present 4 groups of instances with various numbers of riders, vehicles, objectives and generations to illustrate the performance of our approach. This section evaluates the performance with respect to the standard metrics in the field of vehicle routing [22, 23] and includes social and environmental metrics such as the waiting time of a rider to start a ride.

4.1. Data Sources

We use the Cargo benchmark dataset [24], which takes data from the MOD ridesharing company Didi. The instances have maximum 65,500 riders and 50,000 vehicles over a long time horizon and a scale of 876km² area. Since we focus on one-shot routing, we take slices from the dataset for our evaluation. Note that in practice, new routes can be calculated as more requests come in.

- Road Map: We use the road map of Manhattan from Cargo [24]. It has 12,320 nodes, 15,722 edges in an area of 59km².
- Instances: We design 4 groups of instances as detailed in Table 2. The highlighted cell in each row is the parameter that we vary in each group. Regarding the extracted instances with the same number of vehicles and riders, we vary the capacity and type of vehicles by setting all vehicles to be one same type from Table 1.
- Objectives: With the 6 above-mentioned objectives from 3 aspects of sustainability, we optimise routes for setting up experiments either based on one aspect and then considering all of the 3 aspects, as the instances in group 4 represent.

Table 2

Information about Group Instances. The blue color marks the parameter that we vary in each group and the range of the variation.

Group Instances	Parameters				
	#Vehicles	Capacity	# Riders	# Objectives	# Generation
1	10	3	[10, 30]	6	20
2	100	1	150	6	[0, 90]
3	10	1, 3, 10	30	6	20
4	10	3	30	[1, 6]	20


4.2. Experiment Setting

The implementation of the GA algorithm in the Python programming language (using pymoo⁴) allows configuring the population size, number of generations, the offspring rate, and muting/enabling different objectives. To present a first evaluation of our approach, we set the population size and offspring size to 10 and 5, respectively.


⁴<https://pymoo.org/>

4.3. Results

For each solution (that assigns riders to be picked up by a vehicle), we calculate the waiting time per rider. Figure 3 presents the waiting time per rider according to the generated solutions for 2 instances in group 1. When the number of riders increase from 10 to 30, the median waiting times per rider only increase from [160, 220] to [190, 290]. The increase ratio is less than the ratio of riders, which indicates the capability of our routing approach for handling larger populations of riders. Besides, Figure 3b has fewer outliers than Figure 3a, but a wide range of waiting time per rider. This is because when the number of riders increases (without any increase in the number of vehicles), more riders need to wait longer. And when there are fewer riders, the location of vehicles has a greater effect on riders' waiting time. Note that having more riders implies that vehicles travel around more frequently, hence, it is more probable to get close to the starting point of riders.



(a) Instance 1 with 10 riders



(b) Instance 2 with 30 riders

Figure 3: Waiting Time for Each Rider.

To understand the balancedness and diversity of GA-based routing solutions with respect to the economic, environmental and social aspects of sustainability, we focused on instance group 2 (in Table 2) and plotted the results in 4 sampling generations. Figure 4 presents the evolution of the populations at generations 0, 10, 40 and 80. The solutions in each population have diverse effectiveness against the three aspects of sustainability. In general, as the optimisation proceeds, the social inequality decreases. The initial generation (in blue) has greater social inequality than other generations while the 80th generation has the lowest inequality. However, from the perspective of economics and environment, the costs and emissions do not improve for all solutions. However, it is observable that we have a more diverse set of solutions. For instance, solution *A* (in the 80th generation) has a low economic cost and social inequality which compensates for its high environmental emission level. The other notable solution is labelled with a boxed *B* which performs well against all the three dimensions. We argue that such a diverse set of solutions provides an ideal voting pool for what we call participatory route selection (see Section 5).

We further evaluate the effectiveness and diversity of our routing approach in optimising individual objectives with group 4 instances. Figure 5 presents the radar plots of the effectiveness of the 10 generated solutions with respect to the 6 objectives. The smaller number of an objective




Figure 4: Evolution of GA Solutions.




Figure 5: Routes for Optimising Against All Objectives.

implies a better performance of a solution on that objective. Among these 10 solutions, the effectiveness regarding social sustainability varies greater than the others. This is because the changes in allocating riders from a vehicle to another directly affects the working time of the vehicles. The population offers diverse solutions that improve social inequality between 0.1 to 0.3. We expect to take advantage of populations' diversity in the next phase of our work on sustainable mobility and allow riders to vote on routes with respect to their concerns, such as economic costs, environment emissions, or social inequality.

5. Conclusions

In this work, we presented a multiobjective evolutionary approach based on GA algorithm for generating routing options in sustainable MOD. Although there are well-studied multiobjective optimisation methods [7], a GA algorithm generates a diverse set of solutions naturally for the mobility-on-demand problem. Our method is not only sustainability-aware but also establishes

a foundation for explainable, participatory, and dynamic MOD services.

Explainability for Riders, Drivers, and Operators: In comparison to data-driven techniques with black-box optimisation components, in our approach, stakeholders can be provided with visualisations to see how different objectives (e.g., minimising emissions) affect routing solutions. For instance, they can be presented using graphs as in Figure 5 and with explanations on how waiting a bit more (in comparison to using private rides) can benefit the environment or the fairness of the service for drivers.

Participatory Route Selection: Building on this approach, MOD operators can present routing options to riders (or autonomous agents that represent riders) and allow voting among them. This way, users can directly participate in the route selection process and opt for the most collectively equitable route. Our diverse set of GA solutions are not ranked. Thus, a set of riders may prefer one over another and to allow that, we aim to extend our work by adding a preference/vote-based route ranking module in the future.

Dynamic Fine-Tuning: Our approach allows dynamic fine-tuning over time. Users and service operators can inspect routing solutions, evaluate if they are realistic and feasible, and participate in fine-tuning the route generation algorithm and the objective weights to set trade-offs. One can use focus groups for such a tuning over time—e.g., as a city and its citizens change—to enable dynamic fine-tuning of sustainable MOD services.

We aim to extend our work by integrating a participatory route selection process and allowing users to vote over a diverse set of routing solutions with all the objectives and then also muting one or two objectives to provide solutions that match diversity in users' preferences. With a better understanding of users' preferences, we aim to explore other methods for multiobjective optimisation in the context of MOD service. For example, we can define the assignment of a rider to a vehicle as a move and evaluate the move with respect to the multiple objectives, and then adopt reinforcement learning for this problem. Moreover, we plan to test the efficacy of our approach in larger datasets and investigate simulation-based methods to analyse how different map structure and spatio-temporal properties of requests affect the optimality and equitability of solutions.

Data access statement. This study was a re-analysis of data that are publicly available from the the Cargo benchmark dataset [24]. Implementations and data derived through the re-analysis undertaken in this study are available from the public GitHub repository at <https://github.com/Miya-Liu/equitable-ridesharing>.

Acknowledgements. We thank the anonymous reviewers for their incisive comments that were most useful in revising this paper. This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) through a Turing AI Fellowship (EP/V022067/1) on Citizen-Centric AI Systems (<https://ccaais.ac.uk/>) and the platform grant entitled “AutoTrust: Designing a Human-Centred Trusted, Secure, Intelligent and Usable Internet of Vehicles” (EP/R029563/1). For the purpose of open access, the author has applied a creative commons attribution (CC BY) licence to any author accepted manuscript version arising.

References

- [1] M. Furuhata, M. Dessouky, F. Ordóñez, M.-E. Brunet, X. Wang, S. Koenig, Ridesharing: The state-of-the-art and future directions, *Transportation Research Part B: Methodological* 57 (2013) 28–46.
- [2] B. Atasoy, T. Ikeda, X. Song, M. E. Ben-Akiva, The concept and impact analysis of a flexible mobility on demand system, *Transportation Research Part C: Emerging Technologies* 56 (2015) 373–392.
- [3] B. Purvis, Y. Mao, D. Robinson, Three pillars of sustainability: in search of conceptual origins, *Sustainability science* 14 (2019) 681–695.
- [4] S. D. Ramchurn, S. Stein, N. R. Jennings, Trustworthy human-ai partnerships, *iScience* 24 (2021) 102891.
- [5] G. Simonin, B. O’Sullivan, Optimisation for the ride-sharing problem: a complexity-based approach, in: *ECAI 2014*, IOS Press, 2014, pp. 831–836.
- [6] Y. Yu, Y. Wu, J. Wang, Bi-objective green ride-sharing problem: Model and exact method, *International Journal of Production Economics* 208 (2019) 472–482.
- [7] C. F. Hayes, R. Rădulescu, E. Bargiacchi, J. Källström, M. Macfarlane, M. Reymond, T. Verstraeten, L. M. Zintgraf, R. Dazeley, F. Heintz, et al., A practical guide to multi-objective reinforcement learning and planning, *Autonomous Agents and Multi-Agent Systems* 36 (2022) 1–59.
- [8] C. M. Fonseca, P. J. Fleming, Multiobjective genetic algorithms, in: *IEE colloquium on genetic algorithms for control systems engineering*, Iet, 1993, pp. 6–1.
- [9] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: Nsga-ii, *IEEE transactions on evolutionary computation* 6 (2002) 182–197.
- [10] E. Bardaka, L. Hajibabai, M. P. Singh, Reimagining ride sharing: Efficient, equitable, sustainable public microtransit, *IEEE Internet Computing* 24 (2020) 38–44.
- [11] S. Millen, E. Page, *Transport and environment statistics 2021 annual report*, 2021. URL: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/984685/transport-and-environment-statistics-2021.pdf.
- [12] UK NAEI - National Atmospheric Emissions Inventory, Primary NO₂ emission factors for road vehicles, 2022. URL: https://naei.beis.gov.uk/resources/Primary_NO2_Emission_Factors_for_Road_Vehicles_NAEI_Base_2021_v3.pdf.
- [13] EMISIA SA Corporate, Copert the industry standard emissions calculator, 2018. URL: <https://www.emisia.com/utilities/copert/>.
- [14] UK Government, Driving licence categories, 2022. URL: <https://www.gov.uk/driving-licence-categories>.
- [15] Colorado Department of Transportation, Overview of transit vehicles, 2022. URL: https://www.codot.gov/programs/innovativemobility/assets/commuterchoices/documents/trandir_transit.pdf.
- [16] B. A. Weigel, F. Southworth, M. D. Meyer, Calculators to estimate greenhouse gas emissions from public transit vehicles, *Transportation research record* 2143 (2010) 125–133.
- [17] T. Murata, H. Ishibuchi, et al., Moga: multi-objective genetic algorithms, in: *IEEE international conference on evolutionary computation*, volume 1, IEEE Piscataway, NJ, USA, 1995, pp. 289–294.

- [18] T. H. Cormen, Section 24.3: Dijkstra's algorithm, *Introduction to algorithms* (2001) 595–601.
- [19] M. T. Catalano, T. L. Leise, T. J. Pfaff, Measuring resource inequality: The gini coefficient. *numeracy 2* (2): Article 4, 2009.
- [20] H. Jain, K. Deb, An evolutionary many-objective optimization algorithm using reference-point based nondominated sorting approach, part ii: Handling constraints and extending to an adaptive approach, *IEEE Transactions on evolutionary computation* 18 (2013) 602–622.
- [21] J. Blank, K. Deb, P. C. Roy, Investigating the normalization procedure of nsga-iii, in: *EMO*, 2019, pp. 229–240.
- [22] G. Laporte, Fifty years of vehicle routing, *Transportation Science* 43 (2009) 408–416. doi:10.1287/trsc.1090.0301.
- [23] S. Ben Cheikh-Graiet, M. Dotoli, S. Hammadi, A tabu search based metaheuristic for dynamic carpooling optimization, *Computers & Industrial Engineering* 140 (2020) 106217. URL: <https://www.sciencedirect.com/science/article/pii/S0360835219306862>.
- [24] J. J. Pan, G. Li, J. Hu, Ridesharing: simulator, benchmark, and evaluation, *Proceedings of the VLDB Endowment* 12 (2019) 1085–1098.