Exploring representations for optimising connected autonomous vehicle routes in multi-modal transport networks using evolutionary algorithms.

HAN, K., CHRISTIE, L.A., ZAVOIANU, A-C. and MCCALL, J.A.W.

2024

Exploring Representations for Optimizing Connected Autonomous Vehicle Routes in Multi-Modal Transport Networks Using Evolutionary Algorithms

Kate Han, Lee A. Christie, Alexandru-Ciprian Zăvoianu, and John A.W. McCall, *Member, IEEE*

***Abstract*— The past five years have seen rapid development of plans and test pilots aimed at introducing connected and autonomous vehicles (CAVs) in public transport systems around the world. While self-driving technology is still being perfected, public transport authorities are increasingly interested in the ability to model and optimize the benefits of adding CAVs to existing multi-modal transport systems. Using a real-world scenario from the Leeds Metropolitan Area as a case study, we demonstrate an effective way of combining macro-level mobility simulations based on open data with global optimisation techniques to discover realistic optimal deployment strategies for CAVs. The macro-level mobility simulations are used to assess the quality of a potential multi-route CAV service by quantifying geographic accessibility improvements using an extended version of Dijkstra's algorithm on an abstract multi-modal transport network. The optimisations were carried out using several popular population-based optimisation algorithms that were combined with several routing strategies aimed at constructing the best routes by ordering stops in a realistic sequence.**

***Index Terms*— Multi-modal public transport, macroscopic simulations, reachability isochrones, evolutionary algorithms.**

I. Introduction

IN LIGHT of the urgent need to balance environmental and economic development goals, the establishment of sustainable low-carbon mobility systems has been identified as a key development goal by numerous local and regional transport authorities across the globe [1], [2]. In most cases, the envisioned backbone of such environmentally friendly mobility policies is an effective multi-modal *public transport* (PT) system that can promote a shift away from private car use. However, the costs associated with introducing, expanding and operating PT cannot be understated and are often the key constraint when (re-)designing a PT system [3], [4], [5].

Manuscript received 28 February 2023; revised 3 October 2023 and 22 January 2024; accepted 24 February 2024. This work was supported in part by the Automated Road Transport (ART)-Forum, an Interreg Project through the North Sea Programme of the European Regional Development Fund of the European Union. The Associate Editor for this article was K. Gao. *(Corresponding author: Kate Han.)*

Kate Han is with Salford Business School, University of Salford, M5 4WT Manchester, U.K. (e-mail: k.han3@salford.ac.uk).

Lee A. Christie, Alexandru-Ciprian Zăvoianu, and John A.W. McCall are with the National Subsea Centre, Robert Gordon University, Aberdeen, AB21 0BH Scotland, U.K.

Concurrently, in the field of road transportation, the past decade has been marked by two technological developments that have the potential to change mobility paradigms:

- the development and rapid adoption of commercially-ready electric [6], [7] and hydrogen-fueled [8], [9], [10] zero-emission-vehicles [11], [12];
- the emergence of *connected and autonomous vehicles* (CAVs) [13], [14].

Unsurprisingly, the possibility to improve existing and/or design new PT systems using zero-emission CAVs has attracted the interest of transport authorities and early test pilots indicating a relatively high level of public acceptance [15], [16]. This is because, apart from the positive environmental impact, fleets of autonomous buses and shuttles are expected to have lower deployment costs (when compared with light rail alternatives) and lower operational costs (when compared with classical buses). In turn, this enables a niche deployment on routes with expected passenger volumes that are too low to be economically viable otherwise, but can nevertheless bring important social benefits to local communities [17].

The main motivation for the present research stems from a call to provide public transport authorities with an effective means of quickly identifying the most promising options for deploying CAVs to improve existing PT systems. Specifically, Section VI describes our initial real-life application scenario to support the *West Yorkshire Combined Authority* (WYCA) with the automatic discovery (scoping) of optimal routes suitable for a CAV-centered pilot project that aims to improve urban mobility in the Leeds Metropolitan Area.

For this type of exploratory analysis, the key optimality criterion is the ability of a potential CAV-serviced route to improve accessibility (generally reduce commuting times) in the serviced area. In Section VI-A we describe how, based on feedback from WYCA domain experts, we formalized this objective for the considered case study.

II. Background

Transportation system research plays a crucial role in addressing the complex challenges of modern mobility systems as effective transportation is vital for economic growth, social connectivity, and environmental sustainability. Historically, research in this field has driven an overall improvement in the quality of life contributing to developing efficient transportation networks, increasing safety, reducing congestion, minimizing environmental impact, enhancing public transport systems and advancing new automotive propulsion technologies.

In the last decade, multi-modal transport has drawn an increasing amount of research interest. Litman [18] provides a comprehensive exploration of the principles, methodologies, and challenges involved in integrating diverse transportation modes into an efficient and sustainable framework, addressing the needs of both passengers and freight. Benefiting from recent advances in computing, multi-modal transport models are more likely to be analysed using micro- and macro-scopic simulations that can inform different key performance indicators (KPIs). For instance, in [19], Kelle et al. developed an evaluation model which integrates both simulation approaches to assess the tradeoffs between operational efficiency and environmental sustainability in the planning of multi-modal freight transportation. With regard to PT systems, Zheng et al. [20] investigates how to promote a modal shift via a time-dependent area-based pricing scheme optimisation in a congested urban multi-modal transportation network. The authors demonstrate the smart pricing scheme and its ability to reduce congestion using a macroscopic fundamental diagram (MFD) as well as an agent-based simulation approach that shows smooth behavioral stabilisation across two user groups.

Emerging research on connected and autonomous vehicles (CAVs) focuses on the immense inherent potential to improve overall traffic efficiency via optimal autonomous decision making. For example, Yao et al. [21] propose a two-level model to optimize scheduling and trajectories for CAVs in a conflict zone (e.g. ramp, intersection, work-zone) using mixed integer linear programming (MILP) and nonlinear programming (NLP) that can reduce delays by up to 54% and fuel consumption by up to 34%. The cooperative autonomous traffic organization method for CAVs in multi-modal road networks proposed by Wang et al. [22] integrates an autonomous crossing strategy at intersections, improved trajectory optimisation in road segments, and a composite strategy for route planning in road networks to reduce overall delays and ensure fairer travel times.

However, as the Wang et al. [22] highlight, there are general obstacles that must be surpassed in order to achieve many of the estimated CAV-based efficiencies related to optimal autonomous driving (e.g. trajectories, scheduling, dynamic re-routing). The two most important ones are related to efficiency within hybrid traffic flows (on roads shared by both human-driven and autonomous vehicles) and the need for high-quality vehicle-to-vehicle and vehicle-to-infrastructure communications. Furthermore, Yap et al. [23] indicate that, based on a stated preference experiment, only certain types of travelers prefer automated vehicles for the last-mile segment of multi-modal train journeys.

The novelty of the present research is that it focuses on investigating the potential of CAVs to augment accessibility and efficiency within urban public transportation irrelevant of CAV penetration levels.

The only major assumption is that CAVs are able to safely operate in hybrid traffic. Furthermore, in order to mitigate travel preference risks, we focus on improving community accessibility goals by incorporating guidance from West Yorkshire Combined Authority (WYCA) council experts into our simulation framework and experiment design. Thus, our work aims to address the gap between what CAV technology can deliver in the near future and what PT systems need by illustrating how optimally-designed CAV-serviced peripheral routes can significantly improve mobility.

The approach we propose is grounded in efficiently coupling two key components:

1) The first is described in Section III and leverages our bespoke macro-level mobility simulation of multi-modal public transport systems to provide an accurate estimation of commuting times over a given geographical area (i.e. reachability isochrones).
2) The second applies proven nature-inspired global optimisation techniques to efficiently explore the decision space and discover optimal CAV-serviced routing options that can improve mobility in the targeted geographical area(see Sections IV and V).

The current work builds on a initial approach by Han et al. [24] aimed at discovering optimal CAV deployment options for an area in the North of Leeds. While constrained to be circular in nature and always include a pre-determined stop, the type of optimised single-route services analysed in [24] could ultimately deliver area-wide average mobility improvements of up to 13.0% over the baseline. Encouraged by this result, the present approach aims to integrate more CAV route and service design flexibility in the hopes that the extended design space also contains better solutions. Given the increased complexity of the search space that must be efficiently explored, the current work focuses on analysing the comparative performance of different options for representing (encoding) the extended design space when using metaheuristic solvers. Our numerical results reported in Section VIII indicate that the extra deployment flexibility can deliver improvements of up to 16.1% over the baseline.

III. CAV-Based Multi-Route PT Services

This section describes how multi-route CAV services can be evaluated using our macro-level simulation approach.

A. Abstract Multi-Modal Public Transport Network

Graph-based and data-driven techniques for simulating the accessibility/reachability (at a macro level) provided by multi-modal transport systems have been proposed by several researchers over the years [25], [26], [27]. As our stated aim is to use PT accessibility assessments to inform fitness computation within evolutionary algorithms, there was a need to develop a bespoke lightweight simulation that can capitalize on the particularities of the application domain. Our implementation achieves efficiency gains by incorporating three key concepts:

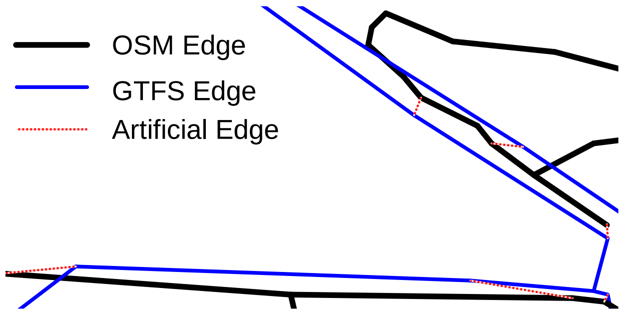

Fig. 1. Example of multi-modal PT network. Black edges indicate roads (OSM data) and walking pathways, blue edges mark PT connections (GTFS data), and red edges denote artificial connections that ensure the connectivity of the black and blue sub-graphs.

1) A Spatial-Temporal Graph: The data sets required by our simulation approach come from Open Street Maps (OSM) in the case of road data and from General Transit Feed Specification (GTFS) files for PT timetable data. As we proceed with constructing an abstract PT network (see Fig.1), artificial edges are added to the network to represent moving from the OSM-based vertices to GTFS-based vertices or vice versa, since the GPS coordinates do not line up exactly between the two data sets. For each GTFS-based vertex, the nearest OSM-based vertex was identified and the Haversine distance was used to calculate the cost associated with the newly introduced edge.

The implementation of the abstract PT network data structure stores the OSM road network as an undirected graph (with edges storing road-length distances) alongside GTFS timetable information (departure and arrival times and locations). In this (largely static) spatial-temporal graph representation, a vertex represents a time and a location. The cost of travelling on the road was based on the length of the contour of the road and a given travel speed. Travel cost on GTFS data is calculated based on timetable information plus any required stop waiting time.

All-in-all, spatial-temporal modeling allows for:

- modeling reachability as a graph-search problem.
- separation of the static (existing network) and transient (candidate CAV routes) parts of the network. and
- a reversible graph structure for computation of both inbound and outbound accessibility.

2) An Efficient Graph Traversal: Dijkstra's algorithm [28] is a well-known standard algorithm for computing the shortest path between two points on a directed graph with non-negative edge costs. We use an implementation of the algorithm without a defined target vertex to produce a shortest-path tree (outbound accessibility isochrones), and export a list of all vertices and their shortest-path distance from a given origin point. As the spatial-temporal graph structure is reversible, without any modification to the algorithm, inbound isochrones can also be easily computed.

3) A Transient CAV Layer: Since the complete spatial-temporal graph is not stored in memory and only exposed to the traversal algorithm via a decorator pattern [29], when the decorator is applied any additional layer of timetabling information (i.e. timetables associated with new services) can be easily added for evaluation. This way, the large underlying static network (based on OSM and GTFS data) can be safely shared as an immutable resource during parallel evaluation with only the CAV layers differing between different decorator views of the network, with each decorator layer local to an individual thread.

B. CAV Service Simulation Parameters

In our approach, in order to simulate the effect of adding a new CAV service to an existing (static) PT network, one must first fix a set of parameters that inform how the new service is to be integrated with the spatial-temporal graph that underpins the overall macro-scopic simulation.

A first decision concerns the *maximum number of routes* contained within the CAV service plan. Whilst the approach in [24] fixes this to one route per service, the stated aim of this work is to investigate if deploying available CAVs along multiple routes can drive increased mobility improvements. The second decision relates to the *selected stops* for each route. These *selected stops* are usually chosen from a larger set of *candidate (bus) stops* in the study area that are suitable for CAVs. Once the selected stops have been identified, it is important to also decide the *order* in which they are to be visited. As the combination of stop selection and ordering directly determines a CAV deployment option, these two topics are discussed at length in Section IV.

The type of routes we focus on in the current work are non-circular in nature. This means that that each valid route has two terminal stops T_1 and T_2 (with $T_1 \neq T_2$). Each route is serviced by two CAVs, initially positioned at T_1 and T_2, with the vehicles travelling back and forth between the terminals and stopping in all intermediate stops along the way. For a route to be valid, the maximum journey time between its terminal stops must be less than one hour apart. Our expectation is that a set containing several such lightweight direct routes offers more flexibility in deploying CAVs in a demand responsive manner than the classical circular routes (i.e. where $T_1 \neq T_2$) previously investigated in [24].

CAV routes can be further constrained by setting three more parameters: *stop waiting time*, *service time range*, and *minimum frequency*. These values are usually set based on feedback from transport planners. In our case, based on feedback from WYCA and given the stated desire to compare with previous literature results concerning the study area, we opted for a stop waiting time of 10 seconds, a service that operates between 06:00 AM and 07:00 PM and a minimum frequency of one bus every 10 minutes. Furthermore, the CAV speed is matched to the road speed limit up to a maximum of 32km/h. As we also maintain the [24] limit of using a maximum of 8 CAVs in total, each CAV service plan will contain up to 4 different routes. Frequency, speed, stop waiting time and the spatial distribution of candidate stops act as constraints that limit the number of viable route designs.

A CAV service plan is fully defined by its associated minimal GTFS format that contains stop geolocation and timetabling information. For example the data in Table I and II

TABLE I
CAV STOP GEOLOCATION INFORMATION

route ID	stopID	Latitude	Longitude
0	111	53.848854	-1.580331
0	72	53.856162	-1.607061
0	43	53.858352	-1.611814
1	8	53.897615	-1.583419
1	64	53.897749	-1.591057
1	102	53.898096	-1.603259
1	21	53.882644	-1.61095

TABLE II
CAV TIMETABLE IN GTFS FORMAT

fromLat	**fromLng**	**deptime**	**toLat**	**toLng**	**arrTime**
53.84885	-1.58033	06:00:30	53.85616	-1.60706	06:04:16
53.85616	-1.60706	06:04:46	53.85835	-1.61181	06:05:22
53.85835	-1.61181	06:05:52	53.85621	-1.61357	06:06:28
53.89762	-1.58342	06:00:30	53.89775	-1.59106	06:01:20
53.89775	-1.59106	06:01:50	53.89810	-1.60326	06:03:14
53.89810	-1.60326	06:03:44	53.88264	-1.61095	06:07:02
53.88264	-1.61095	06:07:32	53.88145	-1.60833	06:08:00

is sufficient to a create a transient layer for a CAV service plan that has 2 routes and a total of 7 stops.

IV. SOLUTION REPRESENTATION

To model service planning as an optimisation problem, a suitable choice of representation, fitness function, and algorithms is essential. This section covers representations used for modelling CAV service plans containing multiple routes.

A representation is a way of mapping a service planning solution to a point in a mathematical space that can be searched by an optimisation algorithm. The algorithm explores the search space to find optimal points which in turn correspond to the best CAV allocations (i.e. planning solutions).

For the problem considered in this study, multiple routes can be complicated to represent. Each representation has advantages and disadvantages, but any good representation should depict the problem thoroughly while allowing an efficient search. In this work, we investigate three common representations which are well-established in the optimisation community: bit-string, real number vector (continuous) and integer vector. We compare these representations in terms of the diversity of services they can represent and the ease by which they can be searched to discover high-quality solutions.

A. Binary Representation

In the binary representation of the service planning, a single bit-string is used to represent the selected stops for multiple routes. Each bit in the string $x_j \in \{0, 1\}$ denotes if a stop is selected (i.e. $x_j = 1$) or not (i.e. $x_j = 0$). The dimension of the bit string is equal to $n \times m$, where n denotes the total number of candidate stops and m indicates the number of routes that we wish to represent.

Fig. 2 contains an example of a binary representation and of how it can be easily translated into a corresponding route plan. The maximum number of routes that can be encoded in the example is 4 (marked from route no. 0 to route no. 3). The set $\{A_1, A_2, A_3, \ldots, A_n\}$ to contains all the n candidate

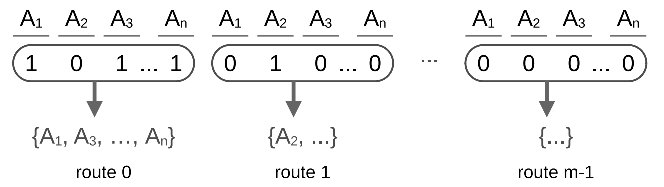

Fig. 2. Example of translation steps for a binary bit-string representation of size $n \times 4$ that can encode up to 4 sets of selected stops (one for each route we aim to represent).

stops. Within the first n bits of the representation, we find 1s in positions corresponding to candidate stops A_1, A_3, and A_n denoting that this stops will be part of route no. 0. In the next set of n bits, the value of 1 associate with candidate stop A_2 indicates that A_2 will be part of route no. 1.

Once the stops are selected, the order in which they are to be visited for each route will be decided by an (external) ordering procedure. The main aim of the ordering is to generate routes that are both realistic and efficient. Fortunately, these two goals are well aligned, as shorter routes also require linkages between stops that are geographically close to each other. Thus, whilst it is not necessarily the case that an optimally short route around a given set of stops is also optimal for the CAV planning problem, it is reasonable to generally prefer shorter routes as they closely resemble existing (human-generated) route designs. In light of these considerations, we have chosen to treat the stop ordering as a classic travelling salesman problem (TSP) and we experimented with three classic TSP algorithms from literature: a dynamic programming (DP) algorithm [30], the Lin-Kernighan heuristic (LKH) [31] and a Greedy ordering strategy [24]. As the solution to a TSP problem is the shortest cycle that visits all the nodes (i.e. selected stops) and we are deliberately searching for non-circular routes, we simply delete the longest edge from each TSP solution to obtain the final routing.

Finally, it important to note that within our wider approach, the routing TSPs are sub-problems that must be solved very fast in order to allow for an evaluation of any CAV service plan. This is particularly problematic for the DP approach as its execution time increases exponentially with the number of stops in the TSP instance. Fig. 3 illustrates that if the number of stops is greater than 20, the DP execution time becomes impractical given the need to evaluate tens of thousands of TSP instances during a single optimisation run. Based on these findings, throughout our experiments, the DP ordering is only applied on routes that have a maximum of 15 selected stops – i.e. we only use DP(15) sorting.

B. Continuous Representation

The previously described binary representation has the disadvantage of increasing the natural dimension of the problem (i.e. n the number of candidate stops), making it more difficult for evolutionary solvers that rely on standard genetic operators. As an alternative to this, we proposed a new indirect continuous representation that is more compact.

The continuous representation requires a vector of size n as a single real value r in the interval $[0, 1]$ is used to denote

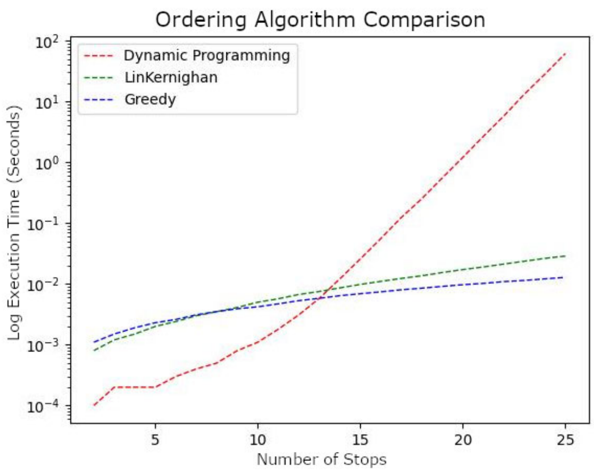

Fig. 3. Comparison of TSP instance size vs execution run time (log plot) across the three ordering algorithms.

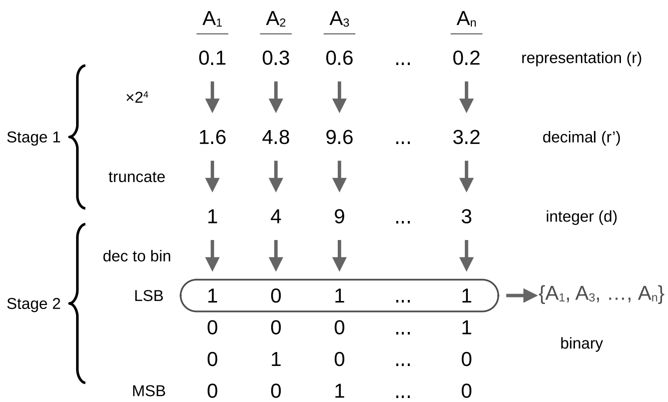

Fig. 4. Example of decoding step for the continuous representation when wishing to generate a maximum of $m = 4$ sets of stops.

a candidate stop across all routes. The translation of this indirect representation to a set of routes is a two-stage process. In the first stage each real value r_i in the representation vector is scaled up and then truncated to an integer value using Equation 1. The notation m is the maximum number of routes to be used.

$$d\left(r_i\right) = \begin{cases} 2^m - 1 & \text{if } r_i = 1 \\ \left\lfloor r_i \times 2^m \right\rfloor & \text{otherwise} \end{cases} \tag{1}$$

In the second stage of the translation process, the integer obtained in the first stage is converted to its m-bit binary representation. The least significant bit (LSB) bit value of the binary representation will denote if candidate stop i is selected on route 0. Accordingly, the second least significant bit will denote if candidate stop i is selected on route 1 and the most significant bit (MSB) will denote selection on route $m - 1$.

Fig. 4 illustrates the two-stage decoding process of the proposed continuous representation when considering a maximum number of $m = 4$ routes and a total of n candidate stops: $A_1, A_2, A_3, \ldots, A_n$. For stop A_1, $r_1 = 0.1$ is the original representation value. 0.1 is first multiplied by 2^4 giving 1.6 and then truncated to $d(r_1) = 1$. After that, 1 is converted to its 4-bit binary value 0001. Based on the LSB $\rightarrow$ MSB

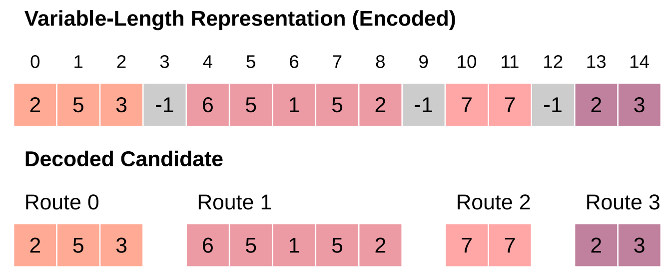

Fig. 5. Example decoding of variable-length representation. It is important to note that route no.2 is infeasible and would not be included in the CAV service simulation.

interpretation of this binary value, stop A_1 is only used by route no. 0. Similarly, stops A_3 and A_n are also selected for inclusion on route no. 0 in the given example.

As with the binary representation, a TSP ordering algorithm must be finally applied to determine the order in which stops should be visited along each route.

C. Variable-Length Integer Representation

The binary representation is high dimensional, and the continuous one has a complex decoding process. We therefore also propose a low-dimensional variable-length (VL) representation with simple decoding. In this case, CAV service plan solutions are represented by a list of integer symbols. These uniquely identify PT stops (i.e., they are candidate stop IDs), with the symbol -1 delimiting different routes. The decoding process does not need a separate ordering algorithm. The order of the stops is dictated by placement within the representation and is thus exposed directly to the optimisation algorithms.

In the example in Fig. 5 the number of potential stops is 10 and the number of routes is 4. At indices 3, 9, and 12, the delimiters separates the CAV service plan into 4 routes. This representation allows for more route design diversity, since the ordering is unconstrained and one stop can be selected multiple times in a route, enabling the creation of routes that can have both circular and non-circular parts. However, the VL representation can also allow the creation of unfeasible routes as for example route no. 2 contains only contains two identical candidate stops (A_7 and A_7). The size of a variable-length representation can increase drastically during the optimisation process. Appropriate operators need to be used to avoid the generation of ill-formed solutions. These is detailed in section V-C.

V. Optimization Algorithms

This section discusses the optimisation algorithms we applied for the three proposed representations. We sought to couple each representation with suitable well-known solvers in order to ascertain a fair view of relative representation performance on our multi-route CAV service planning problem.

A. Binary

For the binary representation described in section IV-A, we applied a standard genetic algorithm (GA) [32] and the population-based incremental learning (PBIL) [33] approach.

The standard genetic algorithm (GA) was selected as it is a widely-used and robust optimiser capable of exploring the rugged and irregular fitness landscapes typically encountered when there is a complex translation from the representation to the realised solution. The GA workflow is well-known and we refer the reader to the pseudocode for Algorithm 1. An initial population of candidate solutions (in this case bitstrings representing the availability of stops in a multi-route CAV service) is generated by an INITIALISATION operator. The GA then iteratively applies a sequence of genetic operators, SELECTION, CROSSOVER, MUTATION, and REPLACEMENT, evolving the population at each iteration, until a stopping criterion is reached and the final population is returned. A range of choice is available to configure and parameterise the genetic operators and we explored several configuration choices, as described in Section VII, in order to improve optimisation performance.

Algorithm 1 Genetic Algorithm

```
function GA(pop_size, stopCriterion)
    pop ← INITIALISATION(pop_size)
    while ¬stopCriterion do
        parents ← SELECTION(pop)
        children ← CROSSOVER(parents)
        children ← MUTATION(children)
        pop ← REPLACEMENT(pop, children)
    end while
    return pop
end function
```

PBIL is a simple estimation of distribution algorithm (EDA). EDAs are evolutionary algorithms that learn probabilistic models of the relationship between candidate representation values and fitness. Their behavior differs from a GA by generating new solutions by sampling a probabilistic model, rather than applying genetic operators to specific candidates. This makes them an attractive choice alongside a GA for a real-world investigation because they explore the space in a different way.

Across both GA and PBIL, the infeasible candidate solutions were handled during evaluation and routes with fewer than 2 stops were not evaluated.

B. Continuous

As mentioned in section IV-B, the continuous representation of the multi-route CAV service planning problem is an indirect representation and we might lose some intuitive understanding of how the algorithm operators affect the generated routes. In particular, what appears to be a small numerical change may turn out to have quite radical effects. We rely therefore on experiments with two very well-known algorithms to determine the efficacy of the proposed continuous representation.

Differential evolution (DE) is a stochastic, population-based global optimisation heuristic widely applied on continuous optimisation problems arising from diverse domains of science and engineering in light of its good performance and relative simplicity with regard to implementation and parameter tuning. DE performs local exploration through the recombination of usually 3 or more different real-value parameter vectors (i.e. candidate solutions). DE has multiple variants in the literature which differ in the mutation, recombination, and selection strategies used. In [34], [35], and [36], multiple DE variants are introduced and compared on different optimisation problems. No variant performed best across all problems. As our multi-route planning problem is derived from a complex real-world application and the continuous representation proposed for this problem is an indirect compact representation of a combinatorial problem, we have chosen to apply the DE/rand/1/bin variant in lights of its robust performance across multiple highly challenging continuous problems [37].

Particle swarm optimisation (PSO), initially proposed by Kennedy and Eberhart [38], has emerged as a very competitive alternative to evolutionary algorithms for complex real-value optimisation problems as evidenced by several recent reviews and surveys of metaheuristic solvers [39], [40], [41]. In PSO, a set of particles iteratively explores the search. Each particle is defined by a position vector in the solution space and a velocity vector. The position vector represents a candidate solution while the velocity vector controls the search speed and direction of its corresponding particle. Each particle explores the search space and saves knowledge as a personal best position vector. At the same time, all particles will share information through the global best position vector that stores the overall best-found solution in the swarm (across all iterations). During the search process, the velocity of each particle is first updated based on its current velocity, personal best position, and swarm global best position. Afterwards, the position of each particle is updated based on the updated velocity and its current position. This two-step update process ensures that each particle incorporates the swarm search result before performing its own next search step. The personal best position of each particle will be replaced if the updated position has better fitness. At the end of each iteration, the global best positions of the swarm will be updated whenever the personal best position of any particle has better fitness.

C. Variable-Length Integer

A variable-length GA (VLGA) has been implemented to solve the multi-route CAV service planning problem using the variable-length integer representation described in Section IV-C. VLGA variants have been proposed since the 90s and the approach has been used across multiple complex problems ranging from classical routing [42], [43] to ensemble selection for deep learning [44]. An important characteristic of VLGA is that it requires specially-designed operators that are largely application specific as they aim to always generate meaningful new candidate solutions.

The VLGA contains an INITIALIZATION step, after which the following steps are iterated until the stopping criteria are met: SELECTION , CROSSOVER , one of three MUTATION operators, and REPAIR . We continue with a brief description of these steps / operators and provide in Fig. 6 a schematic view of their mechanics.

- INITIALIZATION The initial population consists of candidate solutions of randomly initialized lengths up to a

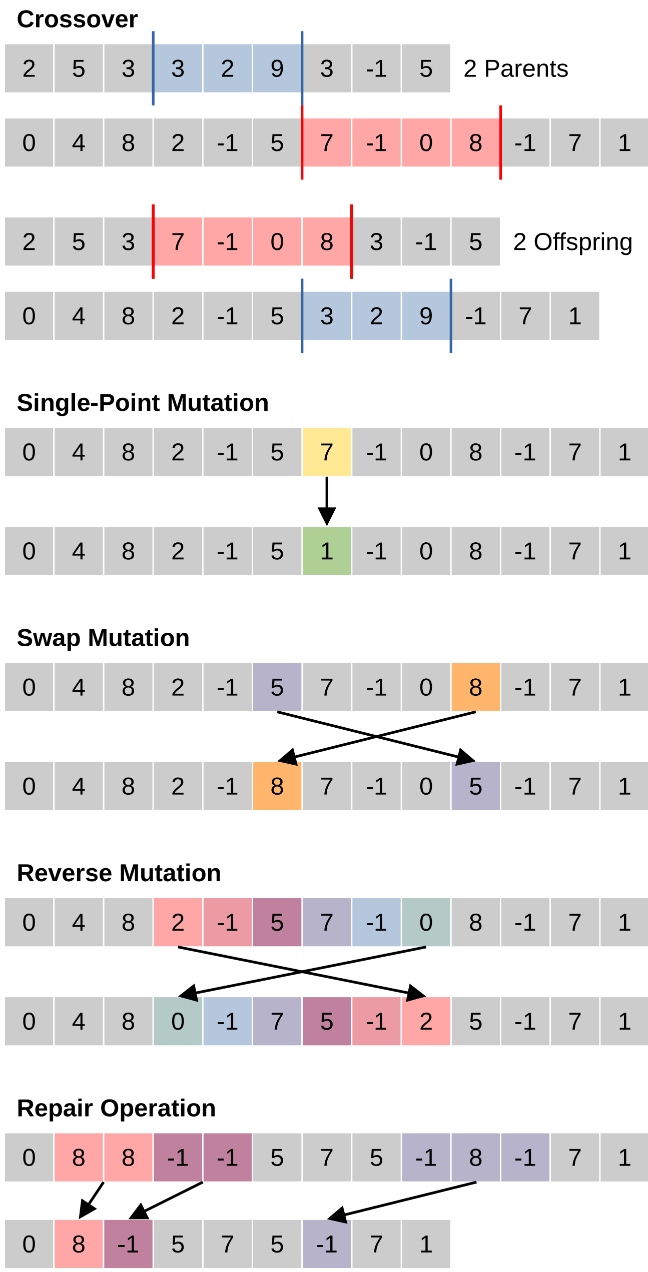

Fig. 6. VLGA operators.

maximum length. For each candidate, vector values are randomly initialized from the set of candidate stop IDs plus the route delimiter (i.e. “-1”).

- **Selection**: We apply tournament selection where a candidate with lower constraint violation is preferred. The constraint value is the excess number of routes above the predefined maximum route number m. Ties are broken in favour of solution candidates with a better fitness value.
- **Crossover**: For each parent, two random cut points are selected, these may be different for each parent and of different lengths. The cut part of the representation is swapped to produce offspring candidates. which may be longer, shorter, or the same length as the parents.
- **Mutation**: Mutation is performed with a preset probability and operates on a single candidate solution. The following three mutation operators were trialed: **Single-Point Mutation** replaces a single value with a new random value. **Swap mutation** randomly selects two values in the same candidate solution and swaps them. **Reverse Mutation** randomly chooses a start point and end point and reverses the order of values between these two points.
- **Repair**: If a candidate solution after crossover and mutation contains two identical adjacent values, they are merged. If there is a route with only one stop, it is deleted.

Length change is a distinguishing feature of VLGA as variable length representations can be prone to bloating under evolutionary pressure [45]. To control this phenomenon, we impose a penalty for having more than $maxStops$ selected stops across all viable routes in a candidate solution. The magnitude of the penalty is proportional to the difference between the actual number of selected stops and $maxStops$.

VI. Case Study

The wider geographical region proposed for investigation by the WYCA is situated in the North-West of the Leeds Metropolitan Area and covers 10 districts (electoral wards as defined by the UK Office of National Statistics). It is serviced by 1,015 public transport (PT) stops for local and regional buses and includes three multi-modal PT hubs centered on key railway stations. The right-hand plot from Fig.7 shows the position of these stops and their spread indicates that they provide very-good walking accessibility to the PT system for all populated areas across the 10 districts.

It is however important to notice that not all 1,015 PT stops are classified as high-frequency by the WYCA. This is illustrated in the left-hand plot from Fig. 7. This discrepancy between regular and high-frequency PT stops is especially visible in the north of the study area and is largely explained by underlining population density levels that drive passenger volumes. Nevertheless, the presence of infrastructure (physical bus stops) is a strong argument for focusing the case study on the northern district (Adel and Wharfedale). The goal suggested by WYCA transport officers is to investigate the ability of a potential multi-route CAV service to improve reachability to Leeds Central Train Station by 10:00 AM on a work day (i.e. improve average morning commute times).

A. Grid Map Evaluation of Area-Wide Reachability

Since the goal is to improve average commuting time via a candidate CAV service x, to compute solution fitness $f(x)$ we rely on a grid G of sample points that cover the entire target area. The spacing of this grid is 0.005 degrees on both latitude and longitude. For any grid point $g \in G$, the shortest multi-modal (baseline PT system + walking + multi-route CAV service encoded in the candidate) travel time $t(g)$ to Leeds Central Station is obtained by computing an inbound isochrone centred on the destination using our bespoke macro-level simulation strategy outlined in Section III. It is noteworthy to underscore that the computation of $t(g)$ depends on the spatial-temporal graph constructed from OSM and GTFS data as the graph traversal algorithm that underpins the macro-simulations can only discover locations that exist in

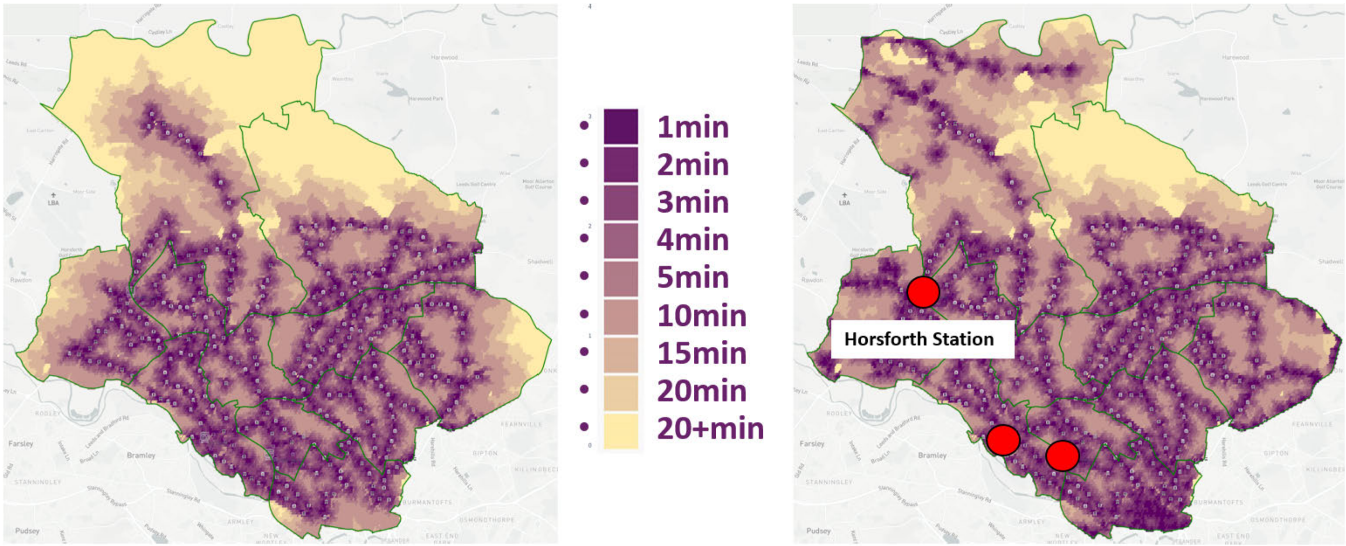

Fig. 7. Positioning and walking accessibility to 759 high-frequency PT stops (left) and to all PT stops (right).

the graph. The implications of this are two-fold. Firstly, grid points in areas associated with disconnected components in the spatial-temporal graph cannot be evaluated (illustrated in blue in Fig. 8). Secondly, in order to calculate the travel between a given grid point A and a general destination point B, we need to first identify vertices in the spatial-temporal graph that are closest to A and to B. Assuming these graph vertices are marked with A' and B', Fig.9 illustrates the calculation process. Finally, fitness is computed as the average commuting time (in seconds) across the entire geographic study area using Eq. 2.

$$f_1(x) = \frac{1}{|G|} \sum_{g \in G} t_g(x) \tag{2}$$

We use a penalty scheme on fitness calculation to direct the algorithms to search for solution with practical CAVs route size. For a CAV route(s) plan solution, we use a $maxTotalStop$ constraint. If after translation, the sum number of stops $sumStops$ of all individual routes CR_i in a route plan CR is bigger than $maxStops$, a penalty value is calculated and added to the fitness value. $f(e)$ is updated using according to (3).

$$f'(e) = \begin{cases} f(e) & \text{if } sumStops \leq maxStops \\ f(e) \times \dfrac{sumStops}{maxStops} & \text{otherwise} \end{cases} \tag{3}$$

For our Adel and Wharfedale case study, the baseline fitness value corresponding to the isochrone in Fig.8 is 3379 seconds (56 minutes and 19 seconds).

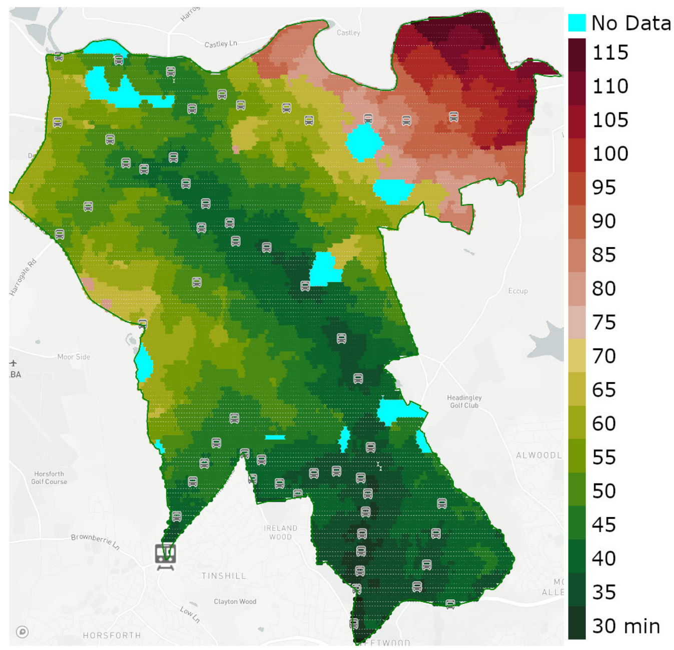

Fig. 8. Accessibility via current PT to Leeds (Central) station by 10:00 AM on a workday. This inbound isochrone was computed using the technique described in section VI-A.

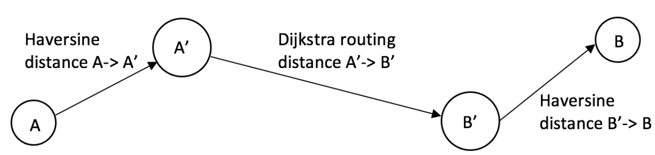

Fig. 9. Grid point travel time calculation. When using the Haversine distance, we assume a walking speed of 5 km/h.

VII. Experimental Setup

In order to obtain a fair comparison of representation and solver performance, we first need to tune parameters that are known to influence the search behaviour of the optimisation algorithms described in Section V.

A. Representation Setup and General Parameter Tuning

As detailed at the end of Section III-A, the maximum number of routes m used across all representations and algorithms is 4. It is noteworthy that m is only a soft constraint within the variable-length representation while strictly enforced (by design) in the binary and continuous representations.

All the global solvers applied in this study are (nature-inspired) population-based algorithms. As such, the size of the population and the number of iterations are critical parameters. For each numerical experiment, we have performed 500 iterations using a population size of 200, resulting in a total computational budget of 10,000 fitness evaluations for each optimisaiton run. Using these settings, all tested algorithms appear to have converged by the end of their runs.

In general, the purpose of population initialisation is to create a highly-exploratory initial set of solutions that (i) maximises opportunities for the optimisation algorithm to sample solutions of better than average fitness and (ii) enables the search to benefit from high population diversity. In our binary and continuous representations, (translated) solution bits denote the presence of particular stops in CAV routes. The more stops on the route, the longer the transit times and total stop times, decreasing route quality. While this effect is captured in the fitness evaluation process described in Section VI-A, we remark that in general the number of desirable stops is much lower than the number of candidate stops. If opting for a roughly equal number of "0"s and "1"s in randomly initialised populations, there is a risk that the algorithms spend too much time searching routes with too many stops to ever produce improvements. We therefore explored the extent to which a specialised initialisation process can focus the algorithms on routes involving realistic numbers of stops while not biasing or impeding the search process. Preliminary results indicated that for GA, PBIL, DE and PSO better optimisation outcomes can be achieved when randomly generated initial candidate solutions have a 90% to 10% ratio of "0"s to "1"s. As such, this skewed initialisation was used for all numerical experiments carried out with these solvers.

B. Individual Solver Parameterisation

Given the long execution time of individual optimisation runs, across all five solvers we could only run limited best parameter grid searches to complement literature recommended settings. Nevertheless, for each tested algorithm, we were able to identify parameterisations capable of delivering competitive results on our complex multi-route optimisation problem.

Thus, based on previous results reported in [24], for the GA we opted for a mutation rate of 0.01 for the single-point bit flip mutation operator and a crossover rate of 0.9 for the single point crossover operator. We conducted limited tests with both roulette and tournament selection, and the latter genetic operator delivered better outcomes.

In the case of PBIL, based on the experiments in [24], we set the mutation probability to 0.01 and conducted comparative tests with three learn rate values: 0.001, 0.01, and 0.02. The lowest learn rate produced the best results.

For DE, we fixed the crossover probability to 0.2 and real constant scale factor to 0.5 based on recommendations in [46], [47] and limited testing with higher crossover probability values.

The PSO solver was parameterised based on insights from [48] with an inertia weight value of 0.729 and an equal particle and swarm acceleration coefficient values of 1.49445. We ran tests with three different velocity ranges, $[-0.2, 0.2]$, $[-0.3, 0.3]$, and$[-0.4, 0.4]$ and the latter setting produced the best results.

The VLGA solver can be seen as more tightly coupled to our multi-route CAV service planning problem on account of the more specialised representation and genetic operators it uses. Furthermore, this solver doesn't benefit from the 90% to 10% skewed initialisation strategy described in the last paragraph of Section VII-A. To alleviate this downside, we experimented with limiting the route sizes of randomly initialised candidate solutions to a maximum of 10, 20, and 30 selected stops. We also ran a more extensive best parameter grid search that varied among the three mutation operators described in Section V-C, three mutation rates (0.001, 0.05, and 0.1) and two selection operators: roulette and tournament. Finally, to prevent the emergence of bloated solutions we also experimented with three different values for the $maxStops$ parameter: 40, 80, and 120. The best results were obtained by the version that initialises a maximum of 10 stops per route, uses single-point mutation with a mutation rate of 0.01, applies tournament selection and limits $maxStops$ to 80.

C. Comparison Baseline

While the accessibility plot in Fig.8 and its associated baseline fitness value of 3379 seconds is what we aim to improve on (i.e. minimise) by discovering optimal multi-route CAV services, as we are dealing with a real-life problem, the best possible solution for our problem is unknown. As such, the quality of an optimisation result can only be assessed based on its improvement with respect to the baseline. However, it is reasonable to expect that any deployment of new transport resources (i.e. CAVs across new routes) in the study area is highly likely to bring some fitness improvement. Therefore, in order to better estimate the actual impact of employing nature-inspired global optimisation algorithms to solve the problem, it would make sense to compare their performance to that of randomly exploring the search space associated with each representation.

In light of these considerations we have set three random search routines: BRS (binary random search, for binary representation), CRS (continuous) and VRS (variable-length). We coupled the binary and continuous versions with the three stop ordering procedures described in Section IV-A. During each search, we randomly generated 10,000 candidate solutions and subsequently returned as outcome the candidate solution with the best fitness value.

VIII. Result Analysis

We summarise in Table III the best fitness values achieved by each combination of solver + stop ordering heuristic. As stop ordering is not required for VLGA, for this solver we present the best average commuting times achieved for different settings of the $maxStops$ parameter.

TABLE III
BEST OPTIMISATION RESULTS AFTER 5 INDEPENDENT RUNS FOR EACH SOLVER AND ORDERING HEURISTIC / *maxStop* COMBINATION

Representation	Solver	Ordering heuristic		
		LKH	Greedy	DP(15)
Continuous	DE	2978	2926	**2908**
	PSO	2934	2882	**2878**
	CRS	3021	**2993**	3380
Binary	GA	2860	2845	**2835**
	PBIL	3037	**2990**	3003
	BRS	**3159**	3259	3380
		$maxStops =$		
		40	80	120
Variable-length	VLGA	2907	**2894**	2902
	VRS	3212	**3101**	3152

The best results table indicates that all tested nature-inspired solvers are able to explore the problem design space and discover high-quality solutions that are better than both the baseline (i.e. 3379) and the best solution discovered by random search (i.e. CRS and greedy with 2993). The two best performing solvers are GA and PSO (both with DP(15) ordering) when considering average result quality over the 5 independent runs: 2860.41 for GA and 2892.70 for PSO. These average results are approx. 5.5% and 4.5% better than the best average results for optimised single-route services reported in [24]. We've extended the number of independent runs to 30 each for these two top multi-route solvers and applied a one-sided Mann-Whitney-Wilcoxon Test [49] with a present significance level of 0.05 to test the statistical significance of the observed differences between average results. The test confirmed (with a p-value > 0.9999) the better performance of the GA on our multi-route CAV service problem.

With regard to stop ordering heuristics, the best results for all nature-inspired solvers using the continuous and binary representations were obtained by DP(15) dynamic programming and Greedy.

In terms of solution representation for multi-route services, Table III shows that all the three options proposed in this work can be used to obtain high-quality solutions with a value lower than 2940 – the best result achieved by optimised single-route services in [24]. However, it is important to note the importance of finding the right solver for each representation. For example the best two results across all experiments were obtained using GA and the binary representation with DP(15) ordering (2835) and Greedy ordering (2845). At the same time, when focusing on the LHK and Greedy orderings, the best results obtained by PBIL using the same binary representation are only marginally better the the best results obtained by CRS (i.e. random search on a continuous representation).

In Fig.10 and Fig.11 we illustrate the 4-route layouts that correspond with the best and second best overall optimisation results. It is important to notice that, for both options, the inclusion of an 8 vehicle CAV service along these routes has an important impact and improves overall accessibility (when compared with the baseline in Fig. 8). In order to achieve this, both high-quality solutions seek to deploy CAVs to provide better connection to the Northern and Mid-Western parts of the study area. The

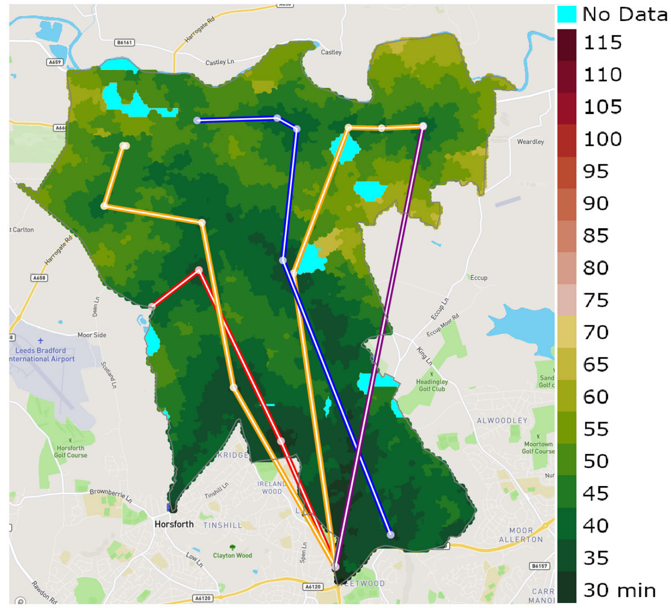

Fig. 10. Best optimisation result: GA with DP(15) ordering.

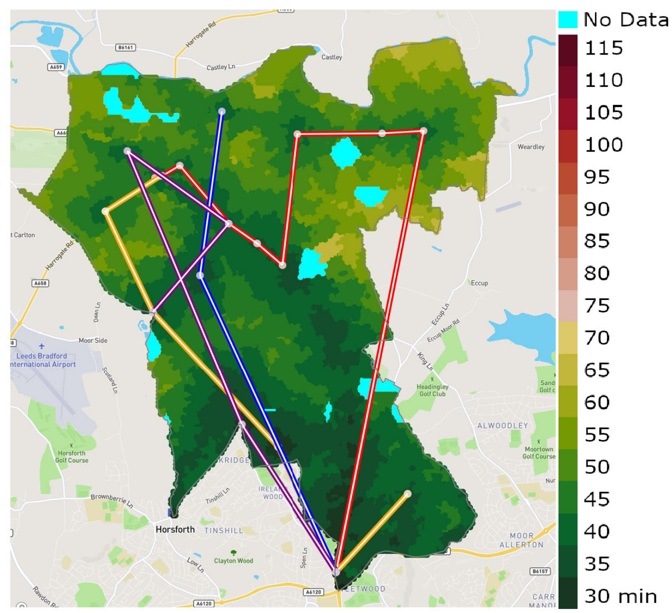

Fig. 11. Second best optimisation result: GA with Greedy ordering.

main differences come in the type of routes used to achieve this:

1) The best solution uses very direct routes to link small but well-delimited clusters of stops to the existing high-frequency services/stops in the South. The poorly served North-East corner is serviced by two routes.
2) The second best solution has routes that tend to overlap more (i.e. use the same stops), especially in the Western part and addresses the North-East with a route that links to the existing high-frequency service both in the South and in the center of the study area.

Whilst the best solution proposes a route layout that is more intuitive, it is important to notice that the difference in quality is virtually insignificant (i.e., 2835 vs 2845) and when considering factors outwith our simulation parameters (e.g. impact on traffic congestion, proximity to service/charging depots), a human decision maker could opt for the second best overall result. This underlines one of the main strengths of our approach based on stochastic solvers: the ability to present decision makers with several distinct high-quality multi-route service solutions suitable for further analysis.

IX. Conclusion and Future Work

In this paper, we have demonstrated an application for evolutionary algorithms in support of introducing Connected Autonomous Vehicles (CAV) to a public transport system. Our approach uses three different representations to explore the planning decision space coupled with a macro-level simulation of the effect of planning decisions on accessibility isochrones for the affected area. Our results show that evolutionary algorithms can be deployed effectively in this space with minimal customization. Standard approaches including proven routing heuristics and a fitness penalty for bloat were effective in ensuring that practicable routes with realistic applicability are produced. We conclude that data-driven simulation-optimisation approaches, using nature-inspired algorithms in general (and evolutionary algorithms in particular) are a promising approach to the adaptation challenges facing public transport systems.

A key element of our approach with potential for future research concerns the data assumptions on which the macro-level simulations rest. Thus, further information on passenger volumes and destinations can be added to better inform outcomes. Furthermore, objectives involving the reachability of more than one destination or facility can also be considered. Finally, where suggested new CAV routes interact with other traffic, as may increasingly be the case in the future, the effects of new route addition on other traffic behavior need to be assessed. This can be addressed through several micro-level traffic simulation technologies already available, very naturally combining macro-level search with detailed evaluation of the good solutions it identifies.

The technology proposed here is only a tool, but one which we hope in time can be integrated with planning processes, automating the production of realistic options to face the increasingly complex challenges of public transport planning.

Acknowledgment

The authors would like to thank the support and constructive feedback provided by West Yorkshire Combined Authority transport policy officers.

References

[1] K. Nakamura and Y. Hayashi, "Strategies and instruments for low-carbon urban transport: An international review on trends and effects," *Transp. Policy*, vol. 29, pp. 264–274, Sep. 2013.

[2] C. Brand, J. Anable, and M. Tran, "Accelerating the transformation to a low carbon passenger transport system: The role of car purchase taxes, feebates, road taxes and scrappage incentives in the U.K." *Transp. Res. A, Policy Pract.*, vol. 49, pp. 132–148, Mar. 2013.

[3] A. Nurdden, R. A. O. K. Rahma, and A. Ismail, "Effect of transportation policies on modal shift from private car to public transport in Malaysia," *J. Appl. Sci.*, vol. 7, no. 7, pp. 1013–1018, Mar. 2007.

[4] L. Redman, M. Friman, T. Gärling, and T. Hartig, "Quality attributes of public transport that attract car users: A research review," *Transp. Policy*, vol. 25, pp. 119–127, Jan. 2013.

[5] D. van Lierop, M. G. Badami, and A. M. El-Geneidy, "What influences satisfaction and loyalty in public transport? A review of the literature," *Transp. Rev.*, vol. 38, no. 1, pp. 52–72, Jan. 2018.

[6] J. Miles and S. Potter, "Developing a viable electric bus service: The Milton Keynes demonstration project," *Res. Transp. Econ.*, vol. 48, pp. 357–363, Dec. 2014.

[7] M. Rogge, E. van der Hurk, A. Larsen, and D. U. Sauer, "Electric bus fleet size and mix problem with optimization of charging infrastructure," *Appl. Energy*, vol. 211, pp. 282–295, Feb. 2018.

[8] H. C. Frey, N. M. Rouphail, H. Zhai, T. L. Farias, and G. A. Gonçalves, "Comparing real-world fuel consumption for diesel- and hydrogen-fueled transit buses and implication for emissions," *Transp. Res. D, Transp. Environ.*, vol. 12, no. 4, pp. 281–291, Jun. 2007.

[9] J. Ally and T. Pryor, "Life-cycle assessment of diesel, natural gas and hydrogen fuel cell bus transportation systems," *J. Power Sources*, vol. 170, no. 2, pp. 401–411, Jul. 2007.

[10] T. O'Garra et al., "Is the public willing to pay for hydrogen buses? A comparative study of preferences in four cities," *Energy Policy*, vol. 35, no. 7, pp. 3630–3642, Jul. 2007.

[11] S. Potter, "Transport energy and emissions: Urban public transport," in *Handbook of Transport and the Environment*. Bradford, U.K.: Emerald Group, 2003.

[12] G. Collantes and D. Sperling, "The origin of California's zero emission vehicle mandate," *Transp. Res. A, Policy Pract.*, vol. 42, no. 10, pp. 1302–1313, Dec. 2008.

[13] L. Ye and T. Yamamoto, "Modeling connected and autonomous vehicles in heterogeneous traffic flow," *Phys. A, Stat. Mech. Appl.*, vol. 490, pp. 269–277, Jan. 2018.

[14] A. Talebpour and H. S. Mahmassani, "Influence of connected and autonomous vehicles on traffic flow stability and throughput," *Transp. Res. C, Emerg. Technol.*, vol. 71, pp. 143–163, Oct. 2016.

[15] K. Mouratidis and V. C. Serrano, "Autonomous buses: Intentions to use, passenger experiences, and suggestions for improvement," *Transp. Res. F, Traffic Psychol. Behaviour*, vol. 76, pp. 321–335, Jan. 2021.

[16] C. Bernhard, D. Oberfeld, C. Hoffmann, D. Weismüller, and H. Hecht, "User acceptance of automated public transport: Valence of an autonomous minibus experience," *Transp. Res. F, Traffic Psychol. Behaviour*, vol. 70, pp. 109–123, Apr. 2020.

[17] S. R. Borg and D. B. Lanng, "Automated busses to mobilize a suburb," in *Proc. Roy. Geographical Soc. Annu. Int. Conf.* Cardiff, U.K.: Cardiff Univ., 2018, pp. 75–105.

[18] T. Litman, *Introduction To Multi-modal Transportation Planning*. Victoria, BC, Canada: Victoria Transport Policy Institute, 2017.

[19] P. Kelle, J. Song, M. Jin, H. Schneider, and C. Claypool, "Evaluation of operational and environmental sustainability tradeoffs in multimodal freight transportation planning," *Int. J. Prod. Econ.*, vol. 209, pp. 411–420, Mar. 2019.

[20] N. Zheng, G. Rérat, and N. Geroliminis, "Time-dependent area-based pricing for multimodal systems with heterogeneous users in an agent-based environment," *Transp. Res. C, Emerg. Technol.*, vol. 62, pp. 133–148, Jan. 2016.

[21] Z. Yao, H. Jiang, Y. Cheng, Y. Jiang, and B. Ran, "Integrated schedule and trajectory optimization for connected automated vehicles in a conflict zone," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 3, pp. 1841–1851, Mar. 2022.

[22] Y. Wang, P. Cai, and G. Lu, "Cooperative autonomous traffic organization method for connected automated vehicles in multi-intersection road networks," *Transp. Res. C, Emerg. Technol.*, vol. 111, pp. 458–476, Feb. 2020.

[23] M. D. Yap, G. Correia, and B. Van Arem, "Preferences of travellers for using automated vehicles as last mile public transport of multimodal train trips," *Transp. Res. A, Policy Pract.*, vol. 94, pp. 1–16, Dec. 2016.

[24] K. Han, L. A. Christie, A.-C. Zăvoianu, and J. McCall, "Optimising the introduction of connected and autonomous vehicles in a public transport system using macro-level mobility simulations and evolutionary algorithms," in *Proc. Genetic Evol. Comput. Conf. Companion*, Jul. 2021, pp. 315–316.

[25] D. Kirchler, "Efficient routing on multi-modal transportation networks," Ph.D. dissertation, Dept. Comput. Sci., Lab. École Polytechn., Laboratoire d'Informatique de l'Ecole Polytechn., Palaiseau, France, 2013.

[26] M. Innerebner, M. Böhlen, and J. Gamper, "ISOGA: A system for geographical reachability analysis," in *Proc. Int. Symp. Web Wireless Geographical Inf. Syst.* Berlin, Germany: Springer, 2013, pp. 180–189.

[27] T. Chondrogiannis, M. A. Nascimento, and P. Bouros, "Relative reachability analysis as a tool for urban mobility planning: Position paper," in *Proc. 12th ACM SIGSPATIAL Int. Workshop Comput. Transp. Sci.*, Nov. 2019, pp. 1–4.

[28] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Math.*, vol. 1, no. 1, pp. 269–271, Dec. 1959.

[29] E. Gamma, R. Helm, R. Johnson, and J. Vlissides, "Decorator pattern," in *Design Patterns: Elements of Reusable Object-oriented Software*. Reading, MA, USA: Addison-Wesley, 1995, p. 175.

[30] R. Bellman, "Dynamic programming," *Science*, vol. 153, nos. 37–31, pp. 34–37, 1966.

[31] S. Lin and B. Kernighan, "An effective heuristic algorithm for the traveling-salesman problem," *Oper. Res.*, vol. 21, no. 2, pp. 498–516, Oct. 1973.

[32] J. H. Holland, "Genetic algorithms," *Sci. Amer.*, vol. 267, no. 1, pp. 66–73, 1992.

[33] S. Baluja, "Population-based incremental learning. A method for integrating genetic search based function optimization and competitive learning," Dept. Comput. Sci., Carnegie-Mellon Univ., Pittsburgh, PA, USA, Tech. Rep. AD-A282 654, CMU-CS-94-163, 1994.

[34] K. Price, R. M. Storn, and J. A. Lampinen, *Differential Evolution: A Practical Approach to Global Optimization*. Berlin, Germany: Springer, 2006.

[35] K. Fleetwood, "An introduction to differential evolution," in *Proc. Math. Statist. Complex Syst. (MASCOS) Symp.*, Brisbane, QLD, Australia, Nov. 2004, pp. 785–791.

[36] R. Storn and K. Price, "Differential evolution—A simple and efficient heuristic for global optimization over continuous spaces," *J. Global Optim.*, vol. 11, no. 4, pp. 341–359, 1997.

[37] K. V. Price, "Differential evolution," in *Handbook of Optimization*. Berlin, Germany: Springer, 2013, pp. 187–214.

[38] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proc. Int. Conf. Neural Netw. (ICNN)*, vol. 4, Nov. 1995, pp. 1942–1948.

[39] P. W. Shaikh, M. El-Abd, M. Khanafer, and K. Gao, "A review on swarm intelligence and evolutionary algorithms for solving the traffic signal control problem," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 1, pp. 48–63, Jan. 2022.

[40] K. Gao, Z. Cao, L. Zhang, Z. Chen, Y. Han, and Q. Pan, "A review on swarm intelligence and evolutionary algorithms for solving flexible job shop scheduling problems," *IEEE/CAA J. Autom. Sinica*, vol. 6, no. 4, pp. 904–916, Jul. 2019.

[41] K. Z. Gao, Z. M. He, Y. Huang, P. Y. Duan, and P. N. Suganthan, "A survey on meta-heuristics for solving disassembly line balancing, planning and scheduling problems in remanufacturing," *Swarm Evol. Comput.*, vol. 57, Sep. 2020, Art. no. 100719.

[42] C. W. Ahn and R. S. Ramakrishna, "A genetic algorithm for shortest path routing problem and the sizing of populations," *IEEE Trans. Evol. Comput.*, vol. 6, no. 6, pp. 566–579, Dec. 2002.

[43] Z. Qiongbing and D. Lixin, "A new crossover mechanism for genetic algorithms with variable-length chromosomes for path optimization problems," *Exp. Syst. Appl.*, vol. 60, pp. 183–189, Oct. 2016.

[44] K. Han, T. Pham, T. H. Vu, T. Dang, J. McCall, and T. T. Nguyen, "VEGAS: A variable length-based genetic algorithm for ensemble selection in deep ensemble learning," in *Proc. 13th Asian Conf. Intell. Inf. Database Syst. (ACIIDS)*, Phuket, Thailand. Berlin, Germany: Springer, 2021, pp. 168–180.

[45] W. B. Langdon and R. Poli, "Fitness causes bloat," in *Soft Computing in Engineering Design and Manufacturing*. London, U.K.: Springer, 1998.

[46] R. Gämperle, S. D. Müller, and P. Koumoutsakos, "A parameter study for differential evolution," *Adv. Intell. Syst., Fuzzy Syst., Evol. Comput.*, vol. 10, no. 10, pp. 293–298, 2002.

[47] D. Zaharie, "Critical values for the control parameters of differential evolution algorithms," in *Proc. 8th Int. Conf. Soft Comput. (MENDEL)*, 2002, pp. 62–67.

[48] Y. Shi, "Particle swarm optimization," *IEEE Connections*, vol. 2, no. 1, pp. 8–13, Feb. 2004.

[49] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *Ann. Math. Statist.*, vol. 18, no. 1, pp. 50–60, Mar. 1947.

Kate Han received the Ph.D. degree in software engineering from Queen's University, Belfast, U.K., in 2018. She holds the position of Lecturer at Salford Business School, where she engages in research that spans scheduling, optimization, ensemble learning, and transportation simulation and optimization. Her work is primarily dedicated to using metaheuristics and evolutionary algorithms for solving real-world optimization problems. Her Ph.D. research focused on the application of hyper heuristics and reinforcement learning for solving timetabling issues. Her research contributions have significantly improved job scheduling and timetabling processes and have advanced the modeling and optimization of multi-modal transportation networks. Her academic achievements include contributions to prestigious AI conferences and journals such as GECCO, Eurocast, ACIIDS, and Knowledge-based Systems, showcasing her impact in the field.

Lee A. Christie has worked a Research Fellow on themes of optimisation and AI at Robert Gordon University and University of Stirling. He has been working in this area of research since completing his Ph.D. degree in 2016, on the topic of pseudo-Boolean optimisation. Recent work with industrial partners includes various projects building efficient transport models and optimising transport networks. He has interests in decentralised computing, machine learning, functional programming, and API design. He also teaches Python for business analytics and is part of the steering group for the Aberdeen Python User Group.

Alexandru-Ciprian Zăvoianu received the Ph.D. degree in computer science from the Johannes Kepler University Linz, Austria, in 2015. He was a Research Assistant at the department of Knowledge-Based Mathematical System, Johannes Kepler University Linz, between 2011 and 2019. Since 2019, he has been a Senior Research Fellow within the School of Computing at Robert Gordon University Aberdeen, U.K. He has authored over 40 academic publications and is a reviewer for several conferences and journals in the field of AI. His research interests include stochastic optimization, evolutionary computation, machine learning, data mining, and distributed computing.

John A.W. McCall (Member, IEEE) is a Professor of Computing and the Director of the National Subsea Centre at Robert Gordon University. He has researched in machine learning, search and optimisation for over 30 years, making novel contributions to a range of nature-inspired optimisation algorithms and predictive machine learning methods, including EDA, PSO, ACO and GA. He has 170+ peer-reviewed publications in books, international journals and conferences. These have received 3200+ citations with an h-index of 25. He specialises in industrially-applied optimization and decision support (Industry 4.0), working with major international companies in energy and telecommunications as well as a diverse range of SMEs. Major application areas for this research are: vehicle logistics, fleet planning and transport systems modelling; predictive modelling and maintenance in energy systems; and decision support in industrial operations management. He and his team have atracted £Ms in direct industrial funding and grants from U.K. and EU research councils and technology centres. He is the Founding Director of Celerum, which specialises in freight logistics and of PlanSea Solutions, which focuses on marine logistics planning.

Dr. John is a Member of ACM. He has served as a member of the IEEE Evolutionary Computing Technical Commitee, an Associate Editor of IEEE COMPUTATIONAL INTELLIGENCE MAGAZINE and the IEEE SYSTEMS, MAN AND CYBERNETICS JOURNAL, and is currently an Editorial Board member for the *Journal Complex and Intelligent Systems*. He frequently organises workshops and special sessions at leading international conferences, including several ACM GECCO workshops in recent years. Most recently, he has organized workshops in Evolutionary Computation and Explainability (ECXAI) at GECCO and is co-editing a Special Issue on Explainability for ACM Transactions in Evolutionary Learning and Optimisation (TELO).