Vehicle-Pedestrian Optimization Framework for Exposure-Aware Routing

Gurban Aliyev[1,2*] and Mirco Nanni[2]

[1*]Department of Computer Science, University of Pisa, Pisa, 56127, Italy.
[2]KDD Lab, ISTI-CNR, Pisa, 56127, Italy.

*Corresponding author(s). E-mail(s): g.aliyev@studenti.unipi.it;
Contributing authors: mirco.nanni@isti.cnr.it;

Abstract

Vehicular traffic is a major source of air pollution in urban areas, exposing pedestrians and residents to harmful emissions. Recent works have proposed exposure-aware pedestrian routing strategies based on static emission maps. In this study, we extend this approach to a dynamic, multi-agent simulation framework involving both cars and pedestrians. Starting from the initial fastest-path routing, we simulate the co-evolution of vehicular emissions and pedestrian exposure over multiple steps, where pedestrian flows dynamically influence car emissions, and vice versa. Two routing strategies are explored: global weighting, where a shared trade-off between travel time and exposure is selected, and local weighting, where each trip independently chooses its optimal trade-off. Experiments on real-world urban data of a medium-sized city in Italy show that both strategies achieve significant reductions in pedestrian exposure, but differ in their impact on vehicle emissions and travel times. Global weighting yields more coordinated adaptation but at a higher systemic cost, while local weighting achieves more balanced outcomes with lower disruption. These results provide insights into designing urban routing policies that jointly optimize mobility efficiency and environmental sustainability.

Keywords: vehicular emissions, pedestrian routing, emission exposure, vehicular routing

1 Introduction

Air pollution is a major issue for human health, especially in urban settings, where large communities of residents and working people meet strong concentrations of emission sources, in particular circulating cars and other transportation means [1]. Recent studies [2–4] have explored the spatial and temporal trends of vehicular emissions by analyzing vehicle trajectories and providing insights on the relations between mobility features and how their emissions affect the urban area. These efforts support the evaluation of current green mobility initiatives [4], the development of next-generation vehicle routing strategies [3], and the potential impact of electrifying portions of circulating traffic [2].

The objective of this work is to study the interplay between vehicle and pedestrian mobility in an urban setting from the viewpoint of human exposure, designing an efficient optimization strategy to assign routes to both populations, in such a way that the resulting trips are still efficient and yet reduce exposure. Leveraging our previous work in [5] we observe that the effects on human health of existing emissions strongly depend on where exactly they are generated and dispersed, and thus a careful planning of pedestrian routes might allow to avoid areas with higher concentrations of emissions with small or moderate impact on the trip efficiency. In this work, we move a step forward, and consider the more challenging problem of optimizing both pedestrian and vehicle routes (the latter were considered fixed in [5]). This new setting introduces an element of complexity due to the interaction of the two sets of routes: on one hand, the routes of vehicles determine how polluted are roads in the city, and thus which ones should be preferred or avoided by pedestrians; on the other hand, the routes of pedestrians determine which areas of the city have a denser walking population, and thus vehicles aiming to reduce exposure should try to avoid them and favor alternative paths. In addition, as in the simpler pedestrian-only problem, we need to solve the issue of deciding for each trip what is the route with

the best trade-off between travel time and consequent exposure.

We address this difficult routing optimization problem through an efficient algorithm inspired by many two-step iterative approaches, like Expectation-Maximization [6], k-means [7] and several gradient descent schemata. The idea is to break down the problem into two smaller ones, in our case the pedestrian routing and vehicle routing problems; each problem fixes some optimization variables (the routes of the other subproblem), searching optimal values for the remaining ones; then, the optimization of the two halves is alternated and repeated along several iterations. The balance between time and exposure, then, is treated through a simple weighting schema in line with [5].

Finally, we carefully build a use case based on a medium-sized Italian city, defining the mobility demand of vehicles and pedestrians, taking into consideration the different categories of pedestrians (workers, students, residents, locals vs. commuters, etc.), their share and their potential origin and destination areas. A simulation experiment is then performed on this use case, analyzing the impact of our approach in terms of exposure, travel times, and overall emissions, as well as providing insights about the functioning of the method.

In summary, this paper provides the following novel contributions:

- we introduce and formally define a joint pedestrian-vehicular routes optimization problem based on multiple objectives (pedestrian exposure and travel times for both pedestrians and vehicles);
- we carefully analyze the mathematical definition of the problem to highlight its components from two viewpoints – pedestrian- and vehicle-oriented – exploited to break down the problem into two (still mutually dependent) parts;
- we design an iterative, two-step heuristic algorithm to find approximate solutions to the problem, based on the insights mentioned above and on two simple vectorization approaches (adapted from our previous work [5]) to combine time and exposure costs;
- we design a use case by adopting a methodology easy to replicate in other cities, which simulates a mobility demand considering different categories of pedestrians, and is based on open data (though additional sources can be deployed, as we do for vehicle mobility demand);
- we perform experiments on the use case, evaluating the trade-offs between time and exposure obtained, analyzing in detail the spatial impact of exposure-aware routing as well as the different vectorization approaches adopted.

In the rest of the paper, we discuss the background and relevant literature (Section 2), formally introduce our exposure-aware pedestrian-vehicle routing problem and analyze it (Section 3), define an iterative, two-step heuristics to address the optimization problem (Section 4), and show experimental results assessing the impact of our strategy in terms of exposure-travel time trade-off (Section 5). Finally, we provide some conclusive remarks (Section 6).

2 Related Work

In this section, we discuss existing approaches for estimating emissions at the road level from mobility data, and routing of vehicles and pedestrians, taking into consideration the amount of exposure to emissions.

2.1 Emissions Inference

Several works in the literature estimate vehicular emissions by exploiting the spatial precision and relative abundance of mobility data coming from vehicle fleets [4], typically applying microscopic emission models to GPS traces [1, 2] based on emission factors (EF) relating pollutants emitted vs kilometer traveled or liter of fuel consumed [8, 9].

Initial studies [10] focus only on the assessment of exposure in static locations (workplace, school, and home) using population census data. Later studies [11, 12] use more accurate mobile data to quantify average mobility exposure. More recent studies investigate the impact of travel behavior on traffic-induced emissions and subsequent exposure to traffic emissions [13]. However, these studies have a key limitation in not considering the dispersion of vehicular emissions that can thus affect areas not too far from where the emissions were generated. There are various dispersion models to be considered, and the main approaches used in these models include Computational Fluid Dynamics (CFD), Gaussian, Lagrangian, and Box models [14, 15]. Depending on the dispersion levels, exposure to car emissions on roads and around roads can change. The concentration of on-road vehicular emissions in certain places can increase or decrease depending on different factors, such as weather conditions (temperature, precipitation, etc.) and building profile data.

2.2 Emission-aware routing

Eco-routing has been explored in multiple domains, with prior studies focusing on vehicle routing for emissions mitigation and pedestrian routing for pollution exposure reduction.

Vehicle Routing. Several studies have proposed eco-routing algorithms to reduce the total emissions from private and commercial vehicles. Demir et al. [16] introduced the bi-objective pollution-routing problem (PRP), which optimizes vehicle routes based on fuel consumption and CO_2 emissions, balancing cost and environmental impact. However, this approach focuses on reducing total emissions rather than minimizing

exposure. Luo et al. [17] extended this concept by developing pollutant exposure-aware vehicle routing, where road segments with high human exposure were penalized in route selection.

Pedestrian Routing for Exposure Reduction. A more relevant line of research investigates pedestrian routing to reduce personal exposure to pollutants. Davies and Whyatt [18] introduced a least-cost path approach to minimize pedestrian exposure by integrating GIS-based routing with pollutant dispersion modeling. Yoon et al. [19] further refined this by using a network centrality-based approach to identify low-exposure pedestrian routes in Seoul, whereas Nadal et al. [20] start from sensor information to infer emissions. Sinnott and Zhong [21] apply machine learning to predict emissions before computing pedestrian routes, and in our previous work [22] a similar procedure is applied by combining a physical emission model and graph imputation methods. Willberg et al. [23] applied multi-exposure optimization, integrating air pollution, noise, and greenery exposure to optimize cyclist routes. Finally, in our work [5] we applied optimization strategies to identify the best trade-off between travel time and exposure in pedestrian trips, whereas previous solutions assume a user-provided fixed importance weight.

All the solutions mentioned above address the routing problem from a single angle, typically unilaterally re-routing pedestrians under the assumption that environmental conditions cannot be changed. In this work we propose a more integrated view, where both the pedestrian and vehicular mobility are part of the optimization problem, aiming to route them to minimize exposure of pedestrians and yet without affecting too much the efficiency of trips.

3 Preliminaries

Generally speaking, our objective consists in finding the routes that a given pool of vehicles and pedestrians need to follow to minimize the exposure of pedestrians to cars' CO_2 emissions without affecting too much the travel times of both groups. In this section, we provide a detailed formulation of the problem that will help later to design a heuristic algorithm to efficiently find an approximate solution.

3.1 Mobility demand and routes

The basic inputs of our problem are the mobility demand P_{OD} of a population of pedestrians and that (V_{OD}) of a pool of vehicles, each defined as a set of origin-destination pairs of locations. Namely: $P_{OD} = \{(o_1^p, d_1^p), \ldots, (o_n^p, d_n^p)\} \subset L \times L$, and similarly: $V_{OD} = \{(o_1^v, d_1^v), \ldots, (o_m^v, d_m^v)\} \subset L \times L$, where L denotes the set of locations considered (in our case, all the nodes in the detailed road graph provided by OpenStreetMap), $n = |P_{OD}|$ and $m = |V_{OD}|$.

Following standard formulations, then, we define a routing for mobility demands P_{OD} and V_{OD} as a pair (P_{route}, V_{route}) where $P_{route} = \{pr_1, \ldots, pr_n\}$ is a set of routes, one for each origin-destination pair, each route pr_i is a sequence of consecutive road segments $pr_i =< r_1, \ldots, r_k >$ that starts from o_i^p (namely $start(r_1) = o_i^p$) and ends in d_i^p (namely $end(r_k) = d_i^p$); and V_{route} is defined similarly.

3.2 Time cost

The time cost of routes for pedestrians and vehicles is defined at the granularity of the single road segment r, based on its length $r.len$ and the average travel speed of pedestrians/vehicles. In this work, we assume that vehicles always travel at the speed limit associated with the road (denoted by $r.max\_speed$) and that pedestrians walk at a fixed average speed avg_p_speed that, following common choices in literature [18], we fixed to 1.4 m/s or 5 km/h. Thus, the time cost for a pedestrian route pr is computed as

$$T_P(pr) = \sum_{r \in pr} \frac{r.len}{avg\_p\_speed} \tag{1}$$

and that for a vehicle route vr as

$$T_V(vr) = \sum_{r \in vr} \frac{r.len}{r.max\_speed}. \tag{2}$$

3.3 Exposure cost

Basic models for estimating the exposure of pedestrians traversing a road segment typically express exposure as proportional to emission density, breath rate, and exposure time. In our study, we assume that breath rate can be approximated to a constant, and since walking speed is also considered uniform, the exposure time will be proportional to the length of the road traversed. Thus we can express the exposure of a pedestrian traversing road segment r simply as: $e(r) = c \cdot ed(r) \cdot r.len$, where $ed(r)$ is the emission density on r and c is a constant factor[1].

3.3.1 Emissions density and dispersion

Following the work in [5], we adopt a standard emission model that computes the instantaneous CO_2 emissions generated by a vehicle at speed v and acceleration a as: $f_E(v, a) = f_1 + f_2 v + f_3 v^2 + f_4 a + f_5 a^2 + f_6 va$, where the emission factors f_i are determined by the type of pollutant and the engine type (petrol, diesel, etc.). Since the acceleration information is not available, we assume it to be zero, thus simplifying the formula to the first three terms, also considering that we assume vehicles always move at the speed limit: $f_E(r) = f_E(r.max\_speed, 0) = f_1 + f_2 \cdot r.max\_speed + f_3 \cdot r.max\_speed^2$. This allows us

[1] Since all our evaluations will be based on percentage variations, the actual value of c has no impact, thus we will fix $c = 1$.

to compute a first approximation of $ed(r)$ that considers only emissions generated directly on the road r, and thus we will denote that $ed_{loc}(r)$. Notice that $ed_{loc}(r)$ represents the sum of the emissions generated by all the vehicles passing on r. Since we assume that all vehicles traverse a road at the speed limit, we can also write: $ed_{loc}(r) = f_E(r) \cdot n\_cars(r)$, where $n\_cars(r) = |\{vr \in V_{route} | r \in vr\}|$.

As emissions created on one road can propagate in the surrounding areas, we consider a dispersion mechanism based on [5], which was a simplified version of the standard Gaussian plume model [24]. The model computes a drop of emissions starting from the generation point as a function of the distance, following a Gaussian trend. Since directly applying the model to estimate dispersed emissions for each point in space would be computationally very heavy – which is not sustainable in our context, where intensive simulations will be involved – we introduce the following simplified process to propagate the emissions of every road r to the surrounding ones:

1. for each road segment r:
   (a) build a buffer R around r (default radius: 100 m)
   (b) select all roads $r_1, \ldots, r_n$ that intersect R, including r itself, referenced as r_1, and denote $r'_i = r_i \cap R$ for each r_i (in general, the intersection can be partial)
   (c) compute the midpoint c_i for each r'_i
   (d) for each r_i, define $\delta(r, r_i) := e^{-d(c'_1, c'_i)^2/\sigma^2}$ (default σ value: 50 m)
2. where not defined, set $\delta(.,.) := 0$
3. return $\delta()$

The result is a function $\delta(r, r')$ that computes the fraction of emission density that propagates from r to r'. Thus, the final emission density of road r can be obtained by summing up all the contributions of local emissions of roads in the surrounding (including r itself): $ed(r) = \sum_{r'} ed_{loc}(r') \cdot \delta(r', r)$.

3.3.2 Pedestrians' exposure

The procedure defined above allows us to estimate the emissions for each road segment r, and thus to compute its corresponding exposure $e(r)$ for the pedestrians passing through it. From here it is easy to infer the overall exposure $E(pr)$ of a whole pedestrian route pr as a sum of exposures of all roads traversed; and the global exposure $E(P_{route}, V_{route})$ of pedestrians, by summing up all exposures for all the routes. Formally, $E(pr) = \sum_{r \in pr} e(r)$ and thus:

$$E(P_{route}, V_{route}) = \sum_{pr \in P_{route}} \sum_{r \in pr} e(r) \tag{3}$$

Notice that $e(.)$ depends only on the emissions generated by vehicles, thus the objective of exposure minimization from the viewpoint of pedestrians consists in finding a set of walking routes with the best tradeoff between exposure cost in (3) and time cost in (1).

3.3.3 Vehicles impact on exposure

While the formulations provided above already introduce all the quantities and functions needed to define the overall problem, it is useful to rewrite the exposure cost in (3) to better emphasize the role of vehicle routes:

$$\begin{aligned} &E(P_{route}, V_{route}) = \textstyle\sum_{pr \in P_{route}} \sum_{r \in pr} e(r) \\ &= \textstyle\sum_{pr} \sum_{r \in pr} c \cdot r.len \cdot ed(r) \\ &= \textstyle\sum_{pr} \sum_{r \in pr} c \cdot r.len \cdot \sum_{r'} ed_{loc}(r') \cdot \delta(r', r) \\ &= \textstyle\sum_{r'} ed_{loc}(r') \cdot c \cdot \sum_{pr} \sum_{r \in pr} r.len \cdot \delta(r', r) \\ &= \textstyle\sum_{r'} n\_cars(r') \cdot \underbrace{f_E(r') \cdot c \cdot \sum_{r} n\_ped(r) \cdot r.len \cdot \delta(r', r)}_{e\_impact(r')} \end{aligned} \tag{4}$$

where, in the last passage, we introduced the term $n\_ped(r)$ to denote the number of pedestrian routes that pass through r, i.e. $n\_ped(r) = |\{pr \in P_{route} | r \in pr\}|$. The expression over the bracket represents the potential impact that a unit of emission on road r' can have on the population that is moving in the surrounding roads. This depends only on pedestrian routes and road features, and not on the vehicle movement, thus from the viewpoint of vehicle optimization that can be represented as a constant value, here denoted $e\_impact(r')$. This allows us to rewrite $E(P_{route}, V_{route})$ simply as:

$$E(P_{route}, V_{route}) = \sum_{vr \in V_{route}} \sum_{r \in vr} e\_impact(r) \tag{5}$$

This formulation is very similar to the one obtained in (3), now expressed in terms of vehicle routes. Thus, we see that the objective of exposure minimization from the viewpoint of vehicles consists in finding a set of routes with the best tradeoff between overall exposure impact in (5) and time cost in (2).

3.4 Problem definition

Exploiting the definitions of the costs involved in the problem, we can formally state the optimization task we aim to address as a multi-objective minimization, as follows:

Input: mobility demands P_{OD} and V_{OD}
Objective:

$$\begin{aligned} \min_{P_{route}, V_{route}} \quad & E(P_{route}, V_{route}), \\ & T_P(P_{route}), \\ & T_V(V_{route}) \end{aligned} \tag{6}$$

Constraints: (P_{route}, V_{route}) is a routing for mobility demands P_{OD} and V_{OD}.

The optimization of $T_P(P_{route})$ and $T_V(V_{route})$ alone would be a rather simple task, since here we are not considering interferences between vehicles on the same road, nor crowds of pedestrians that can slow down walks, thus each trip (either pedestrian or vehicle) can be optimized isolated from the others. The main complexity of the problem comes from the $E()$ term, which conflicts with the time costs (indeed, shortest paths might be more exposure-prone than some longer ones) and also creates an interaction between vehicle and pedestrian routes, due to the causal connection between vehicles' emissions and pedestrians' exposure. In addition, multi-objective problems by design admit multiple optimal solutions, thus some mechanism to disentangle the different criteria needs to be defined. In the next section, we will sketch an algorithmic heuristic to search for approximate solutions to the overall optimization problem efficiently.

4 Methodology

4.1 Problem decomposition

From the mathematical formulations of $E(P_{route}, V_{route})$, $T_P(P_{route})$ and $T_V(V_{route})$ it is clear that the routes of vehicles have an impact on the optimal choice of pedestrian routes, and vice versa, while each vehicle route is independent of the other vehicles, and similarly for pedestrian routes. This suggests a partitioning of the problem into two steps, namely:

$$\min_{P_{route}} E(P_{route}, V^*_{route}), T_P(P_{route}) \tag{7}$$

and

$$\min_{V_{route}} E(P^*_{route}, V_{route}), T_V(V_{route}) \tag{8}$$

where P^*_{route} and V^*_{route} represent fixed values for the corresponding routes. This rewriting lets us read the optimization as an iterative two-step process in the style of Expectation-Maximization [6], where one subset of the output variables is handled at each step (e.g. P_{route} in (7)), and the optimal values yielded become fixed parameters for the second step (e.g. P^*_{route} in (8)).

In the rest of this section, we will first describe how to solve the two minimization sub-problems above, each involving an exposure cost and a time cost, then we will provide a procedure to tackle both problems through an iterative approach.

4.2 Optimizing Time and Exposure

Following [5], we approach the dual-objective minimization problems through a vectorization technique that projects the two dimensions of the problem into a single one, and a selection strategy for the best projection based on Pareto fronts and hyper-rectangle volume maximization. In the rest of this subsection, we summarize the key details.

Vectorization. To simplify notation, given a road segment r we denote the normalized values (min-max rescaling to the $[0, 1]$ interval) of its exposure cost as $C_{exp}(r)$ and its associated normalized travel time as $C_{time}(r)$. Then, given a parameter $w \in \mathcal{R}, w \in [0, 1]$, we compute the trade-off cost $C_{tradeoff}(r, w)$ as:

$$C_{tradeoff}(r, w) = w \cdot C_{time}(r) + (1 - w) \cdot C_{exp}(r)$$

By definition, $C_{tradeoff}(r, 0) = C_{exp}(r)$ and $C_{tradeoff}(r, 1) = C_{time}(r)$, while values $w \in (0, 1)$ represent different blendings of the two components into a single cost value. Higher values of w will give more importance to the time cost, while lower values will weigh exposure costs more. Once a w value is fixed, the overall optimal routes for a set of origins and destinations can be found by applying a standard Dijkstra's algorithm using $C_{tradeoff}(r, w)$ as the costs of road segments.

Global trade-off selection. Given a set of OD pairs, we simulate their optimal routes based on $C_{tradeoff}(r, w)$ for some values of w and then measure the time and exposure costs for each route. These costs are compared against the reference results obtained by computing a simple fastest path, to express them as percentage deviation for ease of aggregation. The sum of such values will produce a pair $(time\_deviation_w, exp\_deviation_w)$ associated with w, and repeating the process for several values of w we will obtain several deviation pairs that form a Pareto front of locally optimal choices. To select one solution as the best trade-off between the two dimensions, we apply the so-called *hyper-rectangle volume* method [16], which computes for each point on the Pareto front the area $HV(w)$ of the (hyper-)rectangle having opposing vertices at the point itself and an ideal point representing the theoretical worst possible solution. Then, the HV-based optimum value is chosen as $w^* = \arg\max HV(w)$.

Local (trip level) trade-off selection. The selection mechanism described above seeks a single value of w that overall fits best all the trips considered, which might lead to single trips adopting an unbalanced choice, namely a path that is either much longer with yet only slightly less exposure cost, or (on the opposite) just slightly shorter yet with much more exposure. An alternative local selection strategy is thus introduced, that applies the same selection approach separately to each trip, building its corresponding Pareto front and selecting the w that maximizes $HV(w)$. In general, that leads to having different optimal w^* values for different origin-destination pairs, better adapting to the single trip. In the experiments, we will evaluate both the global and the local selection strategies.

4.3 Optimization algorithm

We are now ready to introduce a more detailed implementation of the general schema summarized by equations (7) and (8). The *PV optimization* procedure is illustrated in Algorithm 1.

Algorithm 1 PV optimization

Require:

1: P_{OD} pedestrian mobility demand,
2: V_{OD} vehicles mobility demand

Ensure:

3: P_{route} routes for pedestrians,
4: V_{route} routes for vehicles

5: $\forall r.e^{new}(r) = 0$
6: $\forall r.e\_impact^{new}(r) = 0$
7: **repeat**
8: $\quad e() = e^{new}()$
9: $\quad e\_impact() = e\_impact^{new}()$
10: $\quad P_{route} = SimulPedestrians(P_{OD}, e)$
11: $\quad e\_impact^{new} = InferEImpact(P_{route})$
12: $\quad$ **if** first iteration **then**
13: $\quad\quad V_{route} = SimulVehicles(V_{OD}, e\_impact)$
14: $\quad$ **else**
15: $\quad\quad V_{route} = SimulVehicles(V_{OD}, e\_impact^{new})$
16: $\quad$ **end if**
17: $\quad e^{new} = InferE(V_{route})$
18: **until** $e^{new}() \simeq e()$ & $e\_impact^{new}() \simeq e\_impact()$

The two core phases of the algorithm are function *SimulPedestrians* on line 10, which optimizes the pedestrian routes keeping *frozen* the vehicle routes found at the previous step, and *SimulVehicles* on lines 13 and 15, which in turn optimize the vehicle routes based on the previously fixed pedestrian routes. Both phases apply Dijkstra's algorithm to compute the least cost path for their origin-destination pairs, using their specific road cost function.

The algorithm is initialized in lines 5-6 by setting the exposure cost functions e (ref. Eq. (3)) and e_impact (ref. Eq. (5)) to zero for each road segment, which is used for the first iteration of the loop (the if-clause in line 12 ensures that) and which corresponds to not consider exposure at all. After the optimal routes for pedestrians are computed (line 10), the exposure impact of roads for vehicles is re-computed for each road segment, considering its length, the number of people visiting it, and its proximity to other roads (Eq. (4)). These costs will be an input for the vehicle-oriented optimization step in line 15 (except the first iteration, which uses the null-exposure values in line 13), whose output (V_{route}) is in turn used to re-compute the exposure costs e in line 17. Finally, the test of the repeat-until cycle stops the computation when the output of the last iteration does not change significantly from the previous one.

5 Experiments and Results

In this section, we will describe the experimental process adopted to study the impact of routing on exposure and travel times over a use case centered in a medium-sized city in Italy.

We evaluate the performance of the proposed EM-style iterative optimization procedure (Algorithm 1), assessing how exposure and time metrics evolve over simulation steps. We compare along several dimensions the solutions obtained with different strategies: the standard fastest-path routing, the global trade-off strategy, and the local (trip-specific) trade-off strategy shown in Section 4.2.

5.1 Experimental setup

Use case. We conduct our experiments in the city of Pisa, a mid-sized Italian city with a dense urban core. The historical center includes a few pedestrian-only streets and restricted-access zones while surrounding areas feature major vehicular corridors with varied traffic intensity. This combination makes Pisa a suitable testbed for studying pedestrian-vehicle interaction under exposure-aware routing. In addition to being a university hub with heavy commuter flows between campuses and transport nodes, Pisa is a major tourist destination, attracting large numbers of visitors without fixed mobility routines. Our simulations reflect conditions on a typical weekday in May, focusing on the late morning peak hours between 08:00 and 10:00.

Generation of OD Pairs. To analyze the impact of our approach on travel times and pollution exposure, we generated a total of 4,000 pedestrian and 1,000 vehicular origin-destination (OD) pairs across the city of Pisa. To make the analysis easy to reproduce and to apply in any other city, we devised the following process, which considers different strata of the population and mainly leverages open data:

- **Pedestrian OD pairs** were created using a group-specific gravity model based on likely origin and destination zones for three population segments: 37.5% residents (background pedestrian traffic), 37.5% workers and students, and 25% tourists. This partitioning was based on general statistics about the city, in this case assigning a very significant weight to tourism and students, since Pisa is a university city and also a very popular tourist hub.
  Origins and destinations. For background pedestrian trips, potential origins and destinations were sampled uniformly from residential and general-purpose urban locations. For workers and students, residential buildings and transport hubs (used to capture daily

commuters) were used as origins, while schools, universities, and workplaces were used as destinations. For tourists, common origins included stations and parking areas (for one-day visits) as well as hotels (for longer-term visits), while destinations were sampled from popular landmarks and cultural attractions such as the Leaning Tower and Piazza dei Miracoli.
Implementation details. The locations are retrieved from OpenStreetMap, which provides both points of interest and buildings. The size of buildings (area and number of floors) is used to compute the location weight. An exception is made for touristic attractions, for which the weight is computed as the square root of the number of reviews they received on the Google platform, which we retrieved through their standard Places APIs. Finally, the gravity model that we adopted to sample origin-destination pairs assigns selection probabilities proportional to the product of source and target weights, modulated by a distance decay function (more exactly: $f(d) = d^{-2}$). To ensure realistic walking distances, pedestrian OD pairs were constrained to have a Haversine distance between 200 and 1,500 meters.

- **Vehicular OD pairs** were derived from a partitioning of the city into a regular grid of cells having size 200m × 200m.
Flows estimation. In our experiments, the probability of movement between each pair of cells was inferred through aggregation of real GPS trajectories used in our previous work [22], covering private car movements in Pisa. Where real travel data are not available, mobility generation models like gravity or radiation can be adopted, based on population density or similar information easy to access from open sources – e.g. NASA's Socioeconomic Data and Applications Center (SEDAC).
Trips generation. The sampling process of specific origin-destination pairs starts from randomly selecting a pair of origin and destination cells, following the probability described above, and then uniformly randomly selecting one node of the OpenStreetMap road network inside each cell. No distance constraints were applied, allowing a broader distribution of travel lengths representative of urban driving behavior.

Preprocessing time and exposure costs. Exposure values at each simulation step are computed using the emission map derived from vehicle routes of the same step (as defined in Section 3.3). This map accounts for dispersion from neighboring roads (within 100 m, weighted by Gaussian decay) and reflects the total emissions each pedestrian is exposed to while walking. For vehicles, exposure-aware routing uses a population-weighted term: emissions multiplied by the number of pedestrians in each segment (Eq. (4)).

Combining the emission map with the road network (from OpenStreetMap), we derive mode-specific attributes for each road segment (see Section 3). For pedestrians, walking time is computed as segment length over a fixed speed (5 km/h), and personal exposure as the product of walking time and per-meter CO_2 emission. For vehicles, driving time is based on segment length and speed limit, while pedestrian exposure impact is computed as vehicular emissions multiplied by the number of pedestrians in each segment.

Computing Trade-Off Cost Routes. For each OD pair, we apply Dijkstra's algorithm based on cost values assigned to each road segment. In the general case, such costs are a trade-off between travel time and exposure costs, based on a weight w. For **pedestrians**, the weighted cost function $C_{tradeoff}(w, r)$ combines normalized walking time and per-meter exposure, allowing routes to deviate from the fastest path to reduce pollution intake. For **vehicles**, the cost function combines driving time (Eq. (2)) with the exposure impact metric (Eq. (5)). To identify the **best trade-off** for each OD pair, we selected 10 representative values of the weight parameter w using an exponential scaling scheme of the form 0.8^{2k}, with $k \in \mathcal{N}$ ranging from 0 to 9. This provides a diverse set of time–exposure trade-offs while ensuring smooth coverage of the Pareto front. Then, as explained in Section 4.2, the best value for w is chosen through hyper-volume computation on the Pareto front, either globally (one single w is adopted for all trips) or locally (each trip chooses its own w).

5.2 Insights on pedestrian optimization

Spatial distribution of walking time and potential exposure levels. As the first insight of the use case, Fig. 1 provides two views of the city map. In one, it visualizes the log-scaled values of walking time (left) and in the other pedestrian exposure (right) across the road network in central Pisa. In the left panel, darker segments correspond to longer walking durations, typically found on indirect or lengthy connections, while lighter roads indicate quicker paths. The right panel highlights segment-level pedestrian exposure based on emissions from car flows – namely, what exposure a pedestrian would suffer by traversing the single road segment. Since such emissions depend on vehicular traffic, here we adopted a basic fastest-path routing for all the trips simulated. High-exposure segments appear in red and tend to align with high-traffic areas, whereas blue areas represent lower exposure levels, often coinciding with pedestrian-priority streets or less trafficked zones. The figure is the result of our earlier work [5], which was based on static vehicular traffic, whereas the simulations in this study recompute exposure-aware vehicle paths dynamically, then considering the emissions they generated and their spatial dispersion to condition pedestrian routes.

Time and exposure distribution. Standard routing strategies typically aim to minimize travel

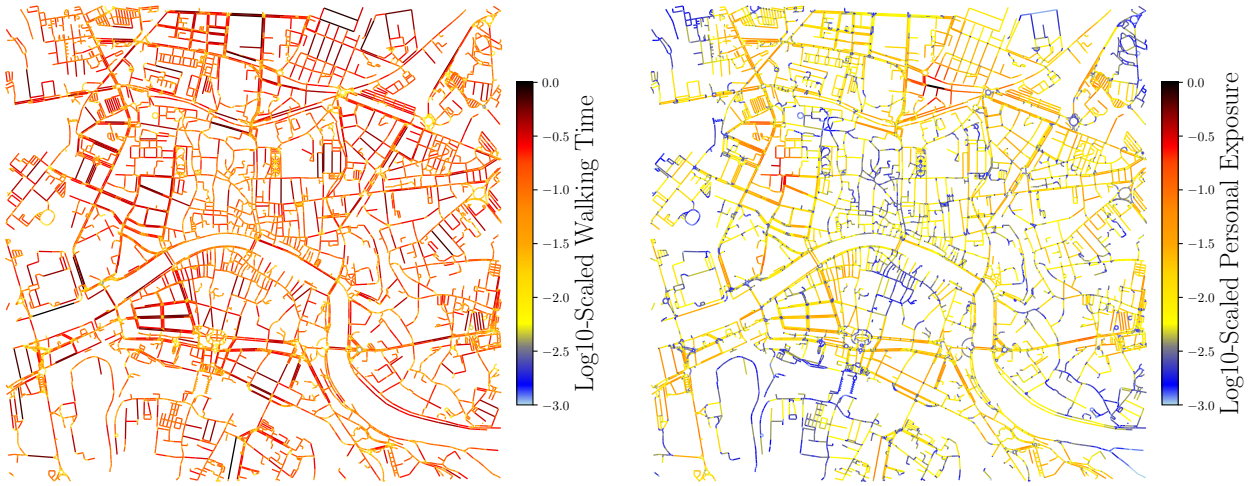

***Fig. 1**: (Log-scaled) walking time (left) and personal exposure (right) over the road network. Red roads indicate longer travel times (left) or higher exposure levels (right). Figure from [5].*

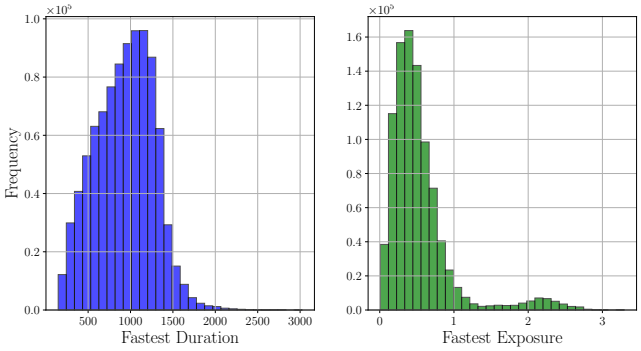

***Fig. 2**: Distribution of pedestrian trips durations (left) and exposure (right) when fastest-time routes are applied to all OD pairs. Figure from [5].*

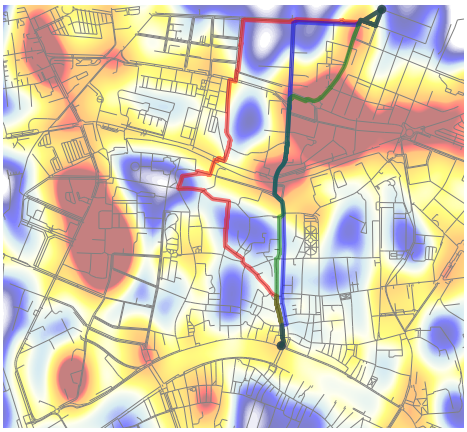

***Fig. 3**: Three routes for the same origin-destination pair, overlaid on a complete CO_2 concentration map. Each route corresponds to a different weight (red=0.01, green=0.50, blue=0.99). Figure from [5].*

time, without considering environmental or health-related impacts. As a baseline for comparison, Fig. 2 presents the distributions of pedestrian travel time (left) and exposure (right) under fastest-path routing across all origin-destination pairs. Travel durations are left-skewed, with most trips concentrated between 500 and 1500 seconds, while longer trips are less common. Exposure values follow a similar pattern, though a secondary peak around 2.0 mg reveals a subset of trips with substantially higher exposure. In many cases, faster routes incidentally yield lower exposure due to shorter traversal times. However, in areas with spatially uneven emissions, these routes may pass through pollution hotspots, resulting in avoidable exposure. This underscores a key limitation of exposure-blind routing, particularly in dense or mixed-use urban environments.

Road usage pattern. To provide a spatial perspective, the step 0 panels of Figures 5 and 8 show road usage patterns under fastest-path routing. Both car and pedestrian flows converge on primary corridors, particularly in central areas and key connectors such as the bridges across the river. This concentration leads to overlapping movement patterns and an increased risk of co-located exposure as anticipated in the joint optimization framework of Section 4.3.

Effects of cost weights. As discussed before, routing a pedestrian or a vehicle trip requires a trade-off between travel time and associated exposure (suffered or caused), thus results can change depending on the importance we give to the two terms. Fig. 3 shows an example of how different values of the trade-off parameter w can change the route assigned to a pedestrian: the red route ($w = 0.01$) prioritizes exposure avoidance by bypassing polluted areas; the blue route ($w = 0.50$) balances exposure and time by slightly deviating from the shortest path; and the green route ($w = 0.99$) minimizes travel time with little regard for emissions.

The Pareto front in Fig. 4 depicts the trade-off between average time difference and average exposure reduction across different route optimization weight choices (w) – also in this case, vehicle routes are fixed

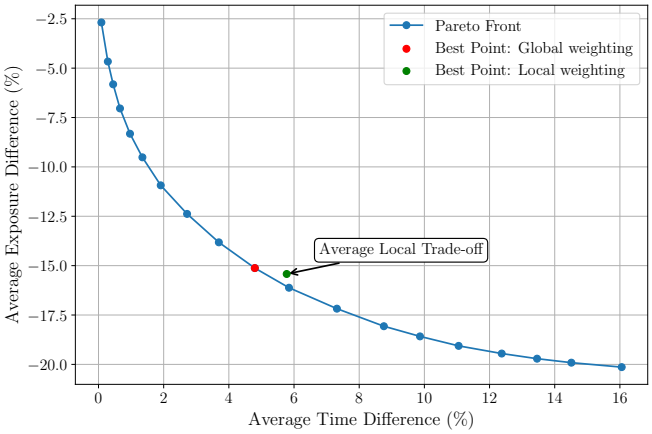

***Fig. 4**: Sample Pareto front of average exposure vs. average travel differences obtained with different trade-off weights. Figure from [5].*

(fastest routes) and optimization modifies pedestrian routes only. As expected, lower exposure values are achieved at the cost of increased travel time, forming a convex frontier. The highlighted best point (in red) represents the optimal trade-off w^* found by maximizing the hypervolume metric $HV(w)$, balancing a 15% reduction in exposure with a moderate increase in travel time (approximately 5%) at $w = 0.107$. The green point shows the average results obtained with local weight choices (different origin-destination pairs can be assigned to different optimal values of w), which in this case slightly improves exposure costs and increases travel times.

5.3 Impact of PV optimization

In this section we analyze the results obtained with the iterative, two-step procedure summarized by Algorithm 1 applied to the use case described in Section 5.1, where both pedestrian and vehicle routes are optimized, and where both the weight optimization strategies introduced (global and local) are tested. The step-by-step summary results are reported in Table 1, and more detailed views are provided in Figures from 5 to 10.

5.3.1 Impact of Global Trade-Off

We first analyze the results obtained by adopting a simpler global weight optimization criterion for the time vs. exposure trade-off. Statistical results are shown in the upper part of Table 1, labeled *Global*.

Already after one iteration (Step 1), the exposure-aware routing yields a very large reduction in total pedestrian exposure (-97.15%) at the cost of slightly larger walking times (+5.35%) and rather significantly larger vehicle travel times (+36.29%). Also, the overall production of vehicle emissions (thus, also considering those that do not affect the pedestrians simulated) is increased significantly (+36.63%). Successive iterations improve walking times, which at the last iteration reach a small +1.61%, yet slightly increasing pedestrian exposure (-95.94%) and further increasing total emissions and driving times (both around 55%). These results indicate that global routing successfully reduces pedestrian exposure with minimal disruption to walking routes, though at a notable cost to car traffic efficiency and overall pollution. The system converges after 6 iterations to $w = 0.262$ for pedestrians and $w = 0.018$ for cars. These values reflect the fact that pedestrians tend to maintain near-fastest paths, while cars shift routes more aggressively to reduce their impact on pedestrians.

To understand how this shift affects spatial mobility behavior, we examine how road usage evolves across simulation steps. Fig. 5 (top line) shows that from step 1 onward, vehicular traffic becomes increasingly concentrated on the main arteries surrounding the historical center of Pisa. In contrast, traffic within the city walls decreases, particularly around pedestrian-only streets and high-tourism zones. Pedestrian routing (bottom line), on the other hand, remains largely stable over time. The main visible changes occur along pedestrian-priority streets (*Borgo Stretto*, *Corso Italia*, etc.) which see increased usage, and along paths bending toward the core touristic location around the Leaning Tower (*Piazza dei Miracoli*, in the middle of the North-West quadrant), which shift slightly to reduce exposure.

To further assess the effect of exposure-aware routing, Fig. 6 presents scatter plots of pedestrian count versus vehicle emissions per road segment at steps 0, 1, and 6. In the Step 0 plot we can see that, while most points cluster in the low-to-mid range, a visible cloud in the top-right corner reveals segments with both high pedestrian traffic and high emissions – zones that present elevated health risks. This "danger zone" highlights the unintended consequences of uncoordinated routing, which exposure-aware strategies are designed to address in later simulation steps. The high-risk "danger zone" is significantly sparser at Step 1, and then further reduced by Step 6. This confirms that global routing successfully separates pedestrian flows from emission hotspots, largely reducing exposure.

Fig. 7 shows the contribution of each road segment to the total exposure and illustrates how the global trade-off strategy reshapes pedestrian and vehicular routing to reduce overall exposure. As observed in Fig. 5(left), roads surrounding the city walls – particularly those just inside or along the perimeter – carry significant car traffic and emissions. The global optimization (center plot) reroutes pedestrians away from these high-exposure corridors and concentrates them on central, pedestrian-priority streets within the city walls. The effectiveness of this shift is also visible in the scatter plot (Fig. 7(right)), where most points fall below the diagonal, indicating reduced exposure on many segments by Step 6. Despite the overall improvement, a notable area of persistent exposure remains in the southwest corner of the city. This is largely due to *Via Aurelia*, a regional

Fig. 5*: Road usage evolution for cars (top) and pedestrians (bottom) at simulation steps 0, 1, and 6 using* ***global*** *weighting. Line width and color intensity reflect the square root–transformed traffic volume: thicker, redder lines indicate heavier usage, while thinner, greener lines indicate lighter traffic.*

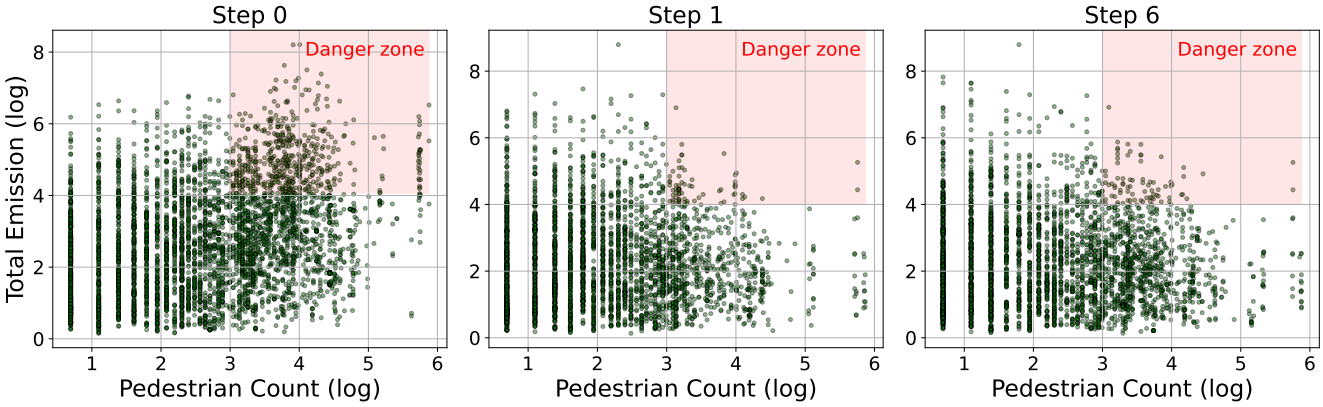

Fig. 6*: Scatter plots showing the relationship between pedestrian density and road-level vehicular emissions at three simulation steps (0, 1, and 6) under the* ***global*** *trade-off strategy. Each point represents a road segment; the x-axis shows the number of pedestrians on the road, while the y-axis shows the total emissions on the same road, normalized using a power transformation.*

highway that connects Rome to northern Italy, as well as the presence of industrial facilities and schools in the neighborhood. Both the importance of this corridor for northbound vehicular traffic and the limited alternative routes to reach key destinations in this area constrain the system's ability to fully divert flows.

5.3.2 Impact of Local Trade-Off

Under local weighting, each OD pair independently selects its optimal trade-off between exposure and travel time. As shown in Table 1, this strategy converges much faster, and achieves a 97.14% reduction in total pedestrian exposure, mirroring the effectiveness of global routing. However, it does so with minimal disruption:

***Fig. 7**: Evolution of total pedestrian exposure across simulation steps under the **global** trade-off strategy. The maps (left and center) show exposure per road segment at steps 0 and 3, respectively, where both color and line width indicate the magnitude of exposure (log-transformed). The scatter plot (right) compares log-transformed exposure values between step 0 and step 3, with each point representing a road segment.*

car routes remain largely unchanged, with both emissions and driving time increasing by less than 0.1%. The trade-off is instead absorbed by pedestrians, who exhibit a moderate increase in walking time (5.21%).

The road usage patterns in Fig. 8 reflect this behavior. Car traffic remains relatively stable across steps, with road usage distributed across both the inner and outer city. In contrast, some shifts are visible in pedestrian flows. In particular, streets along the western city wall – such as *Via Porta a Mare*, *Ponte della Cittadella*, and *Via Bonanno Pisano* – see reduced usage. Meanwhile, pedestrian density increases along the two main routes toward a major touristic hub (*Piazza dei Miracoli*): one via *Corso Italia* and *Borgo Stretto* (for the most part closed to car traffic), and another slightly westward via *Via Crispi* and *Via Roma*. The turning path from *Borgo Stretto* to the square also shifts similarly to what was observed under global weighting.

The scatter plots in Fig. 9 further illustrate the effect of local weighting. As in the global strategy, the high-emission/high-density "danger zone" observed at step 0 is largely eliminated by Step 3. The remaining points are more tightly concentrated in areas of low to moderate emissions, indicating that local routing also succeeds in steering pedestrians away from polluted segments. While the effect is slightly less pronounced than in the global case, the trend confirms that even decentralized, trip-specific rerouting is effective in reducing exposure overlap.

Fig. 10 shows the evolution of pedestrian exposure under the local trade-off strategy. Overall, the results are qualitatively similar to those obtained with the global strategy, with a general reduction in exposure across the network and a visible concentration of pedestrian flows on lower-emission streets. However, two notable differences emerge. First, unlike the global case, some exposure persists in the central pedestrianized areas of the city, suggesting that cars are less rerouted under the local weighting – a finding consistent with Fig. 8. Second, the exposure reduction in the southwest corner of the city, particularly around *Via Aurelia*, is more pronounced compared to the global case.

Overall, the contrast between global and local strategies highlights two distinct behaviors. Global weighting promotes coordinated rerouting, yielding sharp reductions in exposure at the cost of higher emissions and driving times. Local weighting, by contrast, provides a more conservative alternative: it distributes trade-offs more flexibly, avoids systemic disruption, and preserves vehicle efficiency, while still achieving near-identical exposure reduction.

6 Conclusion

In this paper, we investigated the impact of vehicular emissions on pedestrian exposure and proposed a dynamic, exposure-aware routing framework involving both pedestrians and cars. We integrated emission-aware cost functions into a multi-objective optimization process involving travel times of vehicles and pedestrians, and exposure of pedestrians. We proposed an iterative heuristic algorithm to efficiently find solutions to the problem above, equipped with two weighting strategies to identify a trade-off between time and exposure costs. Finally, we carefully built a simulation experiment on a medium-sized city capturing the main mobility flows of pedestrians and vehicles, over which we studied the effects of our framework.

Our results show that significant reductions in pedestrian exposure (up to 97%) can be achieved with minimal impact on walking efficiency, but that the cost to vehicle efficiency varies depending on the coordination strategy. Global weighting, which applies a single trade-off per mode to all trips, yields more coordinated rerouting and a cleaner separation between pedestrian flows and vehicle emission hotspots, but at the expense

Fig. 8*: Road usage evolution for cars (top) and pedestrians (bottom) at simulation steps 0, 1, and 6 using* ***local*** *weighting. Line width and color intensity reflect the square root–transformed traffic volume: thicker, redder lines indicate heavier usage, while thinner, greener lines indicate lighter traffic.*

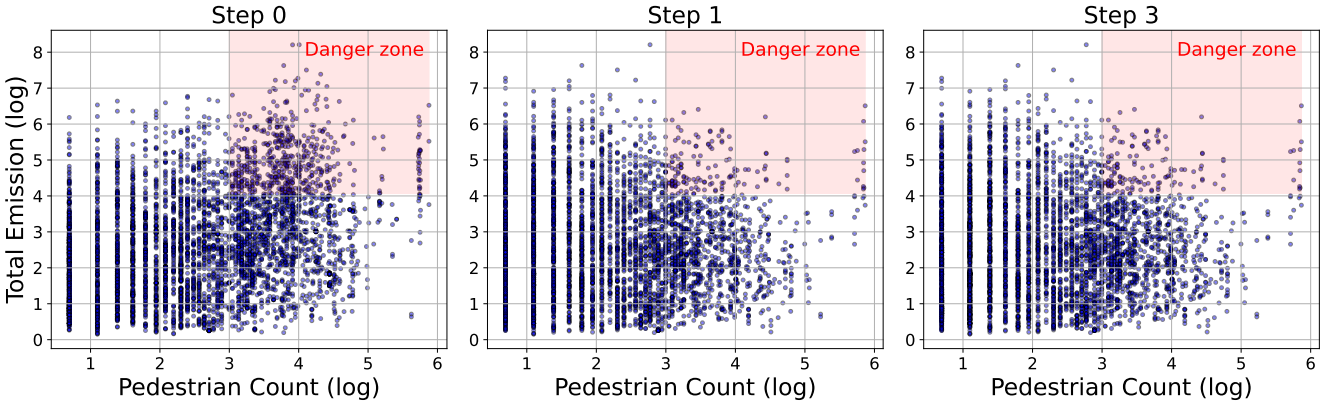

Fig. 9*: Scatter plots showing the relationship between pedestrian density and road-level vehicular emissions at three simulation steps (0, 1, and 3) under the* ***local*** *trade-off strategy. Each point represents a road segment; the x-axis shows the number of pedestrians on the road, while the y-axis shows the total emissions on the same road, normalized using a power transformation.*

of increased car emissions and travel times. In contrast, local weighting preserves vehicle efficiency while achieving similar exposure reduction, thus largely improving the direct impact on health (local exposure of pedestrians) without affecting environment sustainability at large (total emissions) and the practical usability of the routing proposed both for vehicles and pedestrians.

The spatial analysis highlights that the highest-risk areas – where pedestrian flows and vehicle emissions overlap – can be effectively mitigated through coordinated routing, especially near historic and high-foot-traffic areas. These findings emphasize the importance of jointly optimizing pedestrian and vehicular mobility for environmental exposure mitigation.

Future work will try to better explore how changes in vehicle speed affect emissions and pedestrian exposure, extending our approach to consider more realistic speed profiles of car routing, as well as making speed

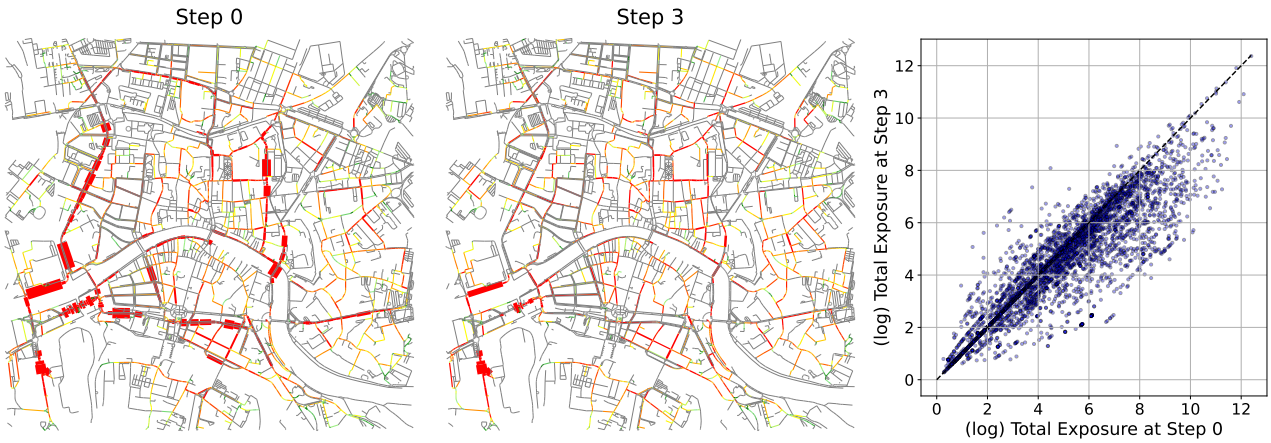

***Fig. 10**: Evolution of total pedestrian exposure across simulation steps under the **local** trade-off strategy. The maps (left and center) show exposure per road segment at steps 0 and 3, respectively, where both color and line width indicate the magnitude of exposure (log-transformed). The scatter plot (right) compares log-transformed exposure values between step 0 and step 3, with each point representing a road segment.*

Mode	Step	Exposure Δ (%)	Emission Δ (%)	Walk Time Δ (%)	Drive Time Δ (%)
Baseline	0	0.00	0.00	0.00	0.00
Global	1	-97.15	36.63	5.35	36.29
	2	-96.09	47.80	2.10	48.30
	3	-96.07	52.65	1.77	53.10
	4	-96.07	54.40	1.66	54.94
	5	-95.94	55.34	1.61	55.90
	6	-95.94	55.34	1.61	55.93
Local	1	-97.15	0.00	5.19	0.04
	2	-97.14	0.00	5.21	0.04
	3	-97.14	0.00	5.21	0.04

***Table 1**: Percentage changes in exposure, emissions, walk time, and drive time across simulation steps, relative to the baseline (Step 0 – using fastest-path routing). The 'global' strategy applies a single optimized weight to all pedestrian trips at each step, while the 'local' strategy selects the best weight for each OD pair.*

at the level of road segment another optimization variable. Another important direction is validating the pedestrian simulation with real-world mobility data, such as counts or trajectories, to assess how well the routing model reflects actual movement. These extensions would help increase the realism of the framework and support potential applications in urban mobility planning.

Supplementary information. The code to reproduce our experiments will be available at https://github.com/baygaliyev/exawaro upon acceptance.

Declarations

- Funding: this work was partially supported by the NextGeneration EU programme, grant n. PE00000013.
- Conflict of interest/Competing interests: the authors have no competing interests as defined by Springer, or other interests that might be perceived to influence the results and/or discussion reported in this paper.
- Ethics approval and consent to participate: Not applicable.
- Consent for publication: not applicable.
- Data availability: part of the data used is open, part is proprietary yet accessible for free, part is proprietary and not publicly accessible.
- Materials availability: not applicable.
- Code availability: the code will be made open source.
- Author contribution: All authors contributed to the study conception and design. Material preparation, data collection and analysis were performed by G.A. Both authors contributed to the discussion of results and writing of the manuscript.

References

[1] Nyhan, M., al.: Predicting vehicular emissions in high spatial resolution using pervasively measured transportation data and microscopic emissions model. Atmospheric Environment **140**, 352–363 (2016) https://doi.org/10.1016/j.atmosenv.2016.06.018

[2] Böhm, M., Nanni, M., Pappalardo, L.: Gross polluters and vehicle emissions reduction. Nature Sustainability **5**(8), 699–707 (2022) https://doi.org/10.1038/s41893-022-00903-x

[3] Cornacchia, G., Böhm, M., Mauro, G., Nanni, M., Pedreschi, D., Pappalardo, L.: How routing strategies impact urban emissions. In: Proceedings of the 30th International Conference on Advances in Geographic Information Systems. SIGSPATIAL '22. Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3557915.3560977

[4] Rahman, M.N., Idris, A.O.: TRIBUTE: Trip-based urban transportation emissions model for municipalities. International Journal of Sustainable Transportation **11**(7), 540–552 (2017) https://doi.org/10.1080/15568318.2016.1278061

[5] Aliyev, G., Nanni, M.: Exploiting vehicular data for exposure-aware pedestrian routing. In: 26th IEEE International Conference on Mobile Data Management (MDM) (2025)

[6] Meng, X.-L., Van Dyk, D.: The em algorithm—an old folk-song sung to a fast new tune. Journal of the Royal Statistical Society: Series B (Methodological) **59**(3), 511–567 (2002) https://doi.org/10.1111/1467-9868.00082 https://academic.oup.com/jrsssb/article-pdf/59/3/511/49588939/jrsssb_59_3_511.pdf

[7] Lloyd, S.: Least squares quantization in pcm. IEEE Transactions on Information Theory **28**(2), 129–137 (1982) https://doi.org/10.1109/TIT.1982.1056489

[8] Huang, Y., al.: Remote sensing of on-road vehicle emissions: Mechanism, applications and a case study from Hong Kong. Atmospheric Environment **182**, 58–74 (2018) https://doi.org/10.1016/j.atmosenv.2018.03.035

[9] Gately, C.K., Hutyra, L.R., Peterson, S., Wing, I.S.: Urban emissions hotspots: Quantifying vehicle congestion and air pollution using mobile phone GPS data. Environmental Pollution **229**, 496–504 (2017) https://doi.org/10.1016/j.envpol.2017.05.091

[10] Souza, P., al.: Quantifying disparities in air pollution exposures across the united states using home and work addresses. Environmental science & technology **58**(1), 280–290 (2024) https://doi.org/10.1021/acs.est.3c07926

[11] Nyhan, M., al.: Exposure track - the impact of mobile-device-based mobility patterns on quantifying population exposure to air pollution. Environmental science & technology **50**, 9671–81 (2016) https://doi.org/10.1021/acs.est.6b02385

[12] Dewuf, B., al.: Dynamic assesment of exposure to air pollution using mobile data. International Journal of Health Geographics **15** (2016) https://doi.org/10.1186/s12942-016-0042-z

[13] Li, Q., Liang, S., Xu, Y., Liu, L., Zhou, S.: Assessing personal travel exposure to on-road PM2.5 using cellphone positioning data and mobile sensors. Health & Place **75**, 102803 (2022) https://doi.org/10.1016/j.healthplace.2022.102803

[14] Leelőssy, A., al.: Dispersion modeling of air pollutants in the atmosphere: a review. Open Geosciences **6**(3), 257–278 (2014) https://doi.org/10.2478/s13533-012-0188-6

[15] Liang, M., Chao, Y., Tu, Y., Xu, T.: Vehicle pollutant dispersion in the urban atmospheric environment: A review of mechanism, modeling, and application. Atmosphere **14**(2) (2023) https://doi.org/10.3390/atmos14020279

[16] Demir, E., Bektaş, T., Laporte, G.: The bi-objective pollution-routing problem. European Journal of Operational Research **232**, 464–478 (2014) https://doi.org/10.1016/j.ejor.2013.08.002

[17] Luo, J., Barth, M.J., Boriboonsomsin, K.: Vehicle routing to mitigate human exposure to traffic-related air pollutants. In: 2018 21st International Conference on Intelligent Transportation Systems (ITSC), pp. 2765–2770. IEEE, NY (2018). https://doi.org/10.1109/ITSC.2018.8569501

[18] Davies, G., Whyatt, D.: A least-cost approach to personal exposure reduction. Transactions in GIS **13**, 229–246 (2009) https://doi.org/10.1111/j.1467-9671.2009.01150.x

[19] Yoon, S., Moon, Y., Jeong, J., Park, C.-R., Kang, W.: A Network-Based Approach for Reducing Pedestrian Exposure to PM2.5 Induced by Road Traffic in Seoul. Land **10**, 1045 (2021) https://doi.org/10.3390/land10101045

[20] Nadal, S., Jordán, J., Sanchez-Anguix, V., Alberola, J.M., Julián, V., Botti, V.: Optimizing pedestrian paths to minimize exposure to urban pollution through traffic data analysis. In: Julian, V., Camacho, D., Yin, H., Alberola, J.M., Nogueira, V.B., Novais, P., Tallón-Ballesteros, A. (eds.) Intelligent Data Engineering and Automated Learning – IDEAL 2024, pp. 196–207. Springer, Cham (2025)

[21] Sinnott, R., Zhong, S.: Real-time route planning to reduce pedestrian pollution exposure in urban settings. In: Proceedings of the IEEE/ACM 10th International Conference on Big Data Computing, Applications and Technologies. BDCAT '23. Association for Computing Machinery, New York, NY, USA (2024). https://doi.org/10.1145/3632366.3632381 . https://doi.org/10.1145/3632366.3632381

[22] Aliyev, G., Nanni, M.: From gps traces to individual emission exposure: A data-driven four-step process. In: Kocian, A., Milazzo, P., Henriques Martins, A.L., Nanni, M., Pappalardo, L. (eds.) Intelligent Transport Systems, pp. 64–82.

Springer, Cham (2025). https://doi.org/10.1007/978-3-031-86370-7_5

[23] Willberg, E., Poom, A., Helle, J., Toivonen, T.: Cyclists' exposure to air pollution, noise, and greenery: a population-level spatial analysis approach. International Journal of Health Geographics **22**, 5 (2023) https://doi.org/10.1186/s12942-023-00326-7

[24] Briggs, G.A.: Diffusion estimation for small emissions. preliminary report (1973) https://doi.org/10.2172/5118833