Does a 15-Minute City Model Limit Freedom of Choice? Amenity Distribution in Urban Areas Across the UK

ECMM428 Individual Research Project

Thomas Newbold
Student ID: 710001264

Abstract

Pedestrian accessibility can be measured by inspecting isochrones - polygons containing the area reachable from a point within an x-minute walk. Carlos Moreno popularised the *15-minute* standard for benchmarking walkability. This concept has been explored at length globally - but when comparing the most walkable areas against each other, it falls short, offering little meaningful insight. Previous studies consider the closest instance of any one amenity category in their analysis - a measure which fails to fully capture accessibility in developed nations, such as the United Kingdom. This work instead counts Points of Interest within a set of time ranges, using the result to develop a metric for quantifying not just proximity but the breadth of available choice, coined Average Amenities Per Capita (AAPC). Freely available data is chosen for the analysis: OpenStreetMap provides the street data and amenity information, and geospatial population distributions are sourced from the Global Human Settlement Layer. Calculated isochrones and resulting metrics are displayed in an interactive web application. Comparisons are made between cities, with reference to existing sustainable development schemes, and between amenity categories (such as `sustenance`, `education` and `transportation`). Principal Component Analysis is used to discern city clusters, and Pearson correlation coefficients for contrasting against other variables like land area.

☑ **I have not used any GenAI tools in preparing this assessment.**

☑ **I certify that all material in this dissertation which is not my own work has been identified.**

Signature:

T Newbold

1 Introduction

Sustainable development comes in many forms; the United Nations defines 17 Sustainable Development Goals (SDG) [28]. Most pertinent here is goal 11 - sustainable cities and settlements - though this work also has relevance to goals 8 (economic growth) and 13 (climate action). Devised by Carlos Moreno [13], and presented at the same 2015 UN Climate Conference [27] where the SDG were finalised, the 15-minute city model defines an urban area to be "accessible" if essential services can all be reached with a 15-minute walk. With car-centric infrastructure leading to higher average commute times [22], the advantage of prioritising active mobility [23] is obvious. This model was encouraged due to its environmental, economic and social benefits [9].

An assortment of studies have been carried out (and interactive tools produced as a result [2, 16]) analysing the minimal distances to any category of service, particularly in the resurgence of the concept post-pandemic [9]. While this is a good baseline, it struggles to identify room for improvement in already accessible environments. According to Nicoletti [16], on average 96% of the urban population in the UK lives within a 15 minute walk of essential services. With a number of cities even being 100% accessible by this criterion, it needs revisiting.

What this model fails to address is the idea of "choice" with respect to amenities, as only the closest in any given category is considered by prior works. Accounting for this shortcoming may help to combat the "hostile reception" [4] received by sustainable urban planning concepts. This project tackles amenity "choice" by counting the number of Points of Interest (POIs) in the time range - building on the work of Scalas [24] and their isochrone crossover counting method - with the aim of determining which cities offer the most options to residents.

This research aims to determine where populations spatially correlate with different categories of amenities. A metric quantifying "choice" is defined (for a generalised x-minute city [11]), and calculated using freely-available data: amenity locations from OpenStreetMap [20] and population data from the Global Human Settlement Layer (GHSL) [25]. 25 UK cities - classified as 100% accessible by Nicoletti et al. [16] - are chosen as case studies, and compared with this new measure. An interactive visualisation is created to present the analysis.

Section 2 briefly covers existing studies into urban accessibility, along with an introduction to urban emissions policies, then Section 3 summarises and outlines project aims. Section 4 highlights the technologies employed by this project, while Section 5 describes specifics of the system, commenting on decisions made during the implementation process. Section 6 provides results, figures and analysis, discussing their validity, drawing connections to another scheme. Finally, the success of the project is evaluated critically in Section 7, and a conclusion is reached in Section 8.

2 Background

An isochrone [24] is defined as the region accessible from a point[1], within a certain time or distance constraint. In the context of this work, and others similar, the distance is calculated given a time interval, and an assumed average pedestrian walking speed (≈ 5 km/h)[2]. The following section summarises previous works in the urban data science field, related to accessibility.

2.1 Existing Studies

Isochrone Crossover Scalas [24] uses the intersection of isochrones to identify "basins of pedestrian accessibility". As a case study, they survey the region of Novara, Italy. While they define 7 categories (across 4 clusters: commercial, healthcare, education and entertainment), these are small[3], resulting in only 463 POIs being assessed across all categories[4]. For each time interval, overlaps are calculated between the merged isochrones for each category, identifying regions where all 7 are reachable. Like all other methods presented here, it effectively solely considers the closest POI in each category.

[1] It can also be considered as set of points from which the centre can be reached, without loss of generality.

[2] See giscience.github.io/openrouteservice/technical-details/travel-speeds/.

[3] Mostly consisting of a single tag, with only "*Food and Drink*" and "*Sport*" having more.

[4] Though this may simply be due to the sparsity of the region chosen.

Accessibility Score Nicoletti [16] investigated the accessibility inequality across different socio-economic backgrounds, deriving a score from average walking distance. Their results are publicly available, presented on an interactive map (available at www.cityaccessmap.com) showing an average travel time heatmap. Categories chosen for their study were: "Mobility, Active Living, Entertainment, Food Choices, Community Space, Education and Health and Well-being" [16, p. 835]. They concluded that there was a clear gap for "vulnerable" profiles, being situated in the least accessible areas.

Their study provided the inspiration for both the interactive visualisation, and the use of Global Human Settlement Layer (GHSL) dataset for population information. In addition, the case studies chosen for this project are those that Nicoletti's work finds to be 100% accessible.

Proximity Time Bruno *et al.* [2] introduces Proximity Time. The work provides an interactive map for exploring the results (see whatif.sonycsl.it/15mincity) per city, and within each on a 200m hexagonal grid [2, p. 8]. Focussing on optimising accessibility by relocating POIs, they prove that car-centric urban layouts (such as those in North American cities) could significantly benefit reorganising the majority of amenities. In contrast, European cities are "well-optimised" [2, p. 3]. As an example, improved POI locations are proposed to improve overall service accessibility in Rome.

Next Proximity Index Olivari *et al.* [19] apply a new accessibility measure, NEXI, to Ferrara and Bologna. This metric was primarily developed to focus intervention efforts in reducing service inaccessibility. Variations on the index are developed, gleaning insights into "discomfort" - which is high for low "proximity" areas with large population [19, p. 5] - as well as weighting categories according to "perceived relative importance" (similar to WalkScore$^{\mathrm{TM}}$; see below). Discussed in the paper, among other, are two interactive maps (mentioned above) of walkability evaluations. Research is conducted on the publicly available street data archive OpenStreetMap [20].

WalkScore$^{\mathrm{TM}}$ walkscore.com was created to help people looking for an area to live, by measuring its suitability for pedestrians, cyclists and users of public transport. On their website, US states are ranked according to 3 scores for those transport categories, with a search function allowing scores to be calculated for any area in the world. In an impartial study, Carr [5] found a correlation between WalkScore$^{\mathrm{TM}}$ and various indicators of a "activity-friendly environment". In spite of this, that study did still caution its use, as crime rate was discovered to correlate positively with WalkScore$^{\mathrm{TM}}$.

***x*-minute** Logan *et al.* [11] presents another analysis - like those discussed above - for cities in both USA and New Zealand. A generalised x-minute city is leveraged, allowing areas to be compared against self-imposed walkability targets. Similarly to [2] and [16], an interactive tool is provided[5].

Walkability Score Kim *et al.* [10] adapts WalkScore$^{\mathrm{TM}}$, awarding a score between 0 and 100 according to distance from closest amenity in nine categories: "grocery stores, restaurants, shopping centers, coffee shops, banks, parks, schools, books, and entertainment". The goal of their research was to investigate the correlation between service categories (and uptake of active transportation [23]).

2.2 Acceptance Of Sustainable Development Schemes

Walkability of an area may impact the success of similar urban development policies. Clean Air and Low Emission Zones encourage the use of active or public transport [23] by requiring a charge to be paid to drive in the area if the vehicle does not meet certain emissions standards, aiming to reduce air pollution (related to SDG-13). They have nonetheless faced friction from the public. When ULEZ[6] was introduced in the Greater London area, it saw its CCTV cameras - required for determining entry and exit times - vandalised, and in some cases destroyed, in protest of the scheme [4]. Khavarian-Garmsir *et al.* [9] mentions barriers to implementations of 15-minute cities, and Caprotti *et al.* [4] discusses opposition to conceptually-adjacent schemes.

[5] research.uintel.co.nz/x-minute-city

[6] tfl.gov.uk/modes/driving/ultra-low-emission-zone

3 Project Specification

3.1 Research Goals

To summarise Section 2, the *15-minute city* model has been the foundation of a collection of works which seek to measure walkability (or, interchangeably, accessibility). They all use the time to closest amenity as basis for their scores, yet this fails to consider how many instances are within the 15-minute range. This amenity count is the focus of this study, in an attempt to quantify the "choice" of services available to residents, by correlating the spatial distributions of amenities to the population. This study also attempts to determine if existing sustainable urban transport policies have been implemented effectively, by inspecting an example - specifically, restricted emission zones like ULEZ - in relation to the produced scores.

3.2 Success Criteria

Functional Requirements

The following functionality should be present in the data processing pipeline:

- Geocoding (city name to bounds)
- Fetching POI locations within region
- Generating isochrones from a set of points
- Aggregating population statistics from polygons

For the visualisation, the functionality should include:

- Displaying isochrones
- Presenting calculated metrics graphically
- Different regions (case study sites) being selectable

Non-Functional Requirements

Four desirable qualities of the visualisation are:

1. Interactive
   (a) Toggleable layers (using both mouse and keyboard input)
   (b) Ability to zoom/pan map
2. Responsive
   (a) Map manipulation should feel continuous, without obvious loading delay
   (b) Multiple sequential inputs should not impede each other
3. Robust
   (a) Sequential interactions should be handled properly, preventing unexpected behaviour
4. Informative
   (a) All presented information should remain relevant to aims
   (b) Legends and graphs should be clear and easy to understand

The data processing step is expected to run in a *reasonable* amount of time - suggested here to be 60 seconds, per city per category.

4 Design

4.1 Chosen Case Studies

The United Kingdom acts as this project's scope. 25 cities have been chosen where all residents are less than a 15-minute walk from essential services, according to [16]. 8 of these, italicised in the list, have implemented a Clean Air or Low/Zero Emission Zone [15].

- *London*
- *Oxford*
- *Bristol*
- *Southampton*
- *Newcastle*
- *Bradford*[7]
- *Edinburgh*
- *Aberdeen*
- Exeter
- Bournemouth
- Brighton
- Margate
- Maidstone
- Crawley
- Reading
- Milton Keynes
- Cambridge
- Colchester
- Ipswich
- Peterborough
- Coventry
- Nottingham
- Manchester
- York
- Preston

4.2 Data Sources

Two types of data are required for this project: map data - including POIs and streets for routing calculations - and population data - for metric calculation. With the desire to encourage reproducibility, building off of free-to-use data was preferred.

4.2.1 Street Data

OpenStreetMap (OSM) [20] is an open-source[8] street registry. Being a crowdsourced dataset, sometimes referred to as Volunteered Geographical Information (VGI) [7], its accuracy can be (and has extensively been) debated. But, while it cannot challenge a map like Ordnance Survey [21] in quality, for instance, its coverage is substantial - being updated regularly - and sufficient for most research purposes [3, 7]. Additionally, the many labelled POIs [3] contained in the OSM database are a crucial resource, providing the centre points for isochrone calculation. As OSM is used by multiple other studies [2, 3, 11, 16, 19, 24], it has been selected as the primary data source for this project.

A benefit to OSM being so popular is that route-finding tools integrating with the data already exist. In particular, openrouteservice.org - abbreviated to ORS - is leveraged by this project, as it contains an endpoint capable of generating the isochrones required later. Given this API will be run locally (to prevent rate limits impacting development), an archive of UK street data is downloaded from download.geofabrik.de/europe.html. Running ORS locally also minimises the risk of results becoming irreproducible, though repeating with up-to-date date would require the OSM/ORS services (and population data source discussed below) remain accessible.

4.2.2 Population Data

As mentioned in Section 2.1, Nicoletti [16] uses GHSL[9] [25] population data. For consistency and comparability with existing research, this project uses the same source. GHSL data is free to download and use, following the open-source methodology shared by this work and many others in the field. For this specific application, the GHS-POP 2025, at a resolution of 3 arc seconds[10], and coordinate system WGS84, is downloaded. This resolution is the best available GHS-POP resolution - for any year - so resampling to a smaller resolution is required; covered in more detail in Section 5.3.1. GHSL is derived from census data, but remains completely anonymised, meaning there are no ethical considerations required for its use.

[7] [16] aggregates the area of Bradford under Leeds - which has 100% access - so is included here.
[8] OSM functions under an *Open Data Commons Open Database License*.
[9] human-settlement.emergency.copernicus.eu/download.php?ds=pop
[10] Approximately 93m at the equator, and approximately 58m at 51.5°N (London).

4.3 Dependencies

Five sequential stages make up the data processing stage, to be implemented in Python. These are outlined in Figure 1 below. The first three require, respectively, the three APIs outlined in section 4.3.1 below. The fourth and fifth steps are covered in Section 4.5 and 4.6 respectively. By combining these tools, a fully automated pipeline can be created which integrates POI location collection as well as isochrone processing.

Figure 1: The five main stages of the data processing pipeline.

4.3.1 APIs

Nominatim Geocoding is facilitated by the Nominatim API, taking city names and returning the bounding box of the area - required for later queries. While not essential, as areas can be selected and input manually, this step produces a more intuitive user experience.

Overpass Overpy - a Python wrapper of Overpass Turbo - can search the OSM database for POIs matching a set of feature tags. Groups of feature tags - or categories - are defined in Section 4.4. The latitude/longitude of returned nodes are saved. If any *ways* are returned, their centre points can be taken. Only named[11] data points should be taken - improving quality, as well as reducing the overall POI count which would reduce processing times.

Openrouteservice Isochrone calculation is done with ORS (see Section 4.2.1 above), using a sample[12] of the centre points obtained from the prior stage. In addition to using OSM data, this API can be downloaded and run locally to allow more requests than the limited public version - maximum 2,500 daily requests to Isochrones V2 on the *Collaborative* plan[13] for "academic" applications. Both geocoding and POI query services are also offered as by ORS, but the chosen APIs mentioned above are more purpose-built for this, where ORS is designed primarily for route-finding.

4.3.2 Libraries

Data Processing GeoPandas (geopandas.org) is chosen to integrate all these services together into a coherent processing pipeline - outlined in Figure 1 - integrating well with packages like pandas [6]. This library can manage the `geojson` files, with functions like `dissolve` for merging polygons by their properties. As a Python package, the skill barrier to entry associated with a desktop tool - such as QGIS [8] - is avoided, and automating the data processing will be much easier.

Visualisation A JavaScript library, D3 [18], is selected to produce geospatial visualisations. Being a web library, the produced visualisations could be easily embedded into a web page, promoting widespread sharing of results. Zooming and panning functionality is built-in to `Leaflet.js`[14] (the library D3 uses to render maps), encouraging interaction. `Leaflet.js` is also strongly customisable, with many configurable UI elements, along with support for GeoJson rendering. While this library can work in real-time[15], this is best on a small scale - so here, the processing pipeline is kept separate, with D3/Leaflet handling only the final presentation.

[11] Containing a `name` tag.
[12] See Section 5.3.2.
[13] account.heigit.org/info/plans
[14] leafletjs.com/reference.html
[15] See observablehq.com/@targomo/{exploring-the-15-minute-city, deep-dive-on-location-scoring}

4.4 Categories

Overpass uses OSM's feature tags, listed at wiki.openstreetmap.org/wiki/Map_features. Amenity categorisation varies between studies in the literature, so seven desire categories are manually defined here, given in Table 1 below. These will be stored in `.json` format and used to construct the composite Overpass queries (see Section 4.3.1). Though the number of tags in each is different, an attempt has been made not to bloat any one specific category too significantly, as this could impact the comparability of metric values.

Category	OSM Features
`sustenance`	amenity={bar, biergarten, cafe, fast_food, food_court, ice_cream, pub, restaurant} **or** shop={general, supermarket, convenience}
`education`	amenity={college, dancing_school, driving_school, first_aid_school, kindergarten, language_school, library, surf_school, research_institute, training, music_school, school, university, public_bookcase, planetarium}
`transportation`	amenity={bicycle_parking, bicycle_rental, bicycle_wash, boat_rental, bus_station, car_rental, car_sharing, car_wash, vehicle_inspection, ferry_terminal, fuel, motorcycle_parking, parking, taxi} **or** aeroway=terminal **or** public_transport={platform, station}
`healthcare`	amenity={baby_hatch, clinic, dentist, doctors, hospital, nursing_home, pharmacy, social_facility, veterinary}
`culture_and_religion`	amenity=place_of_worship **or** tourism={attraction, gallery, museum}
`leisure`	amenity={arts_centre, cinema, community_centre, conference_centre, events_venue, exhibition_centre, music_venue, nightclub, social_centre, stage, theatre} **or** leisure={dog_park, fitness_centre, park, pitch, playground, stadium, swimming_pool} **or** shop={department_store, mall, charity, beauty, cosmetics, hairdresser, art, music, video}
`public_service`	amenity={courthouse, fire_station, police, post_box, post_depot, post_office, ranger_station, townhall, telephone, toilets} **or** emergency={defibrillator, fire_extinguisher, life_ring}

Table 1: Features tags contained in each category.

Mainly, amenity tags were targeted, as they categorise the majority of services well, though other features with relevance to the group were included: shop, aeroway, public_transport, tourism, leisure and emergency. Only "essential" amenities were selected - in an attempt to moderate the number of results - but this will always be subjective to a degree. Correctly categorising all tags into comprehensive categories of what everyone might find important is impossible - e.g. someone who doesn't own a bike would not be affected by any cycling related amenities, but they are included for the populous who do.

Duplicate Tags Care had to be taken when categorising certain overlapping tags. For example, `amenity:parking_space` identifies individual spaces, skewing the distribution; `amenity:parking` is preferred in this scenario, identifying a car park as a single instance, rather than many. The linked wiki page provides helpful sub-categories, which acted as guidance for those defined here.

Relevance To Choice Though *work* is a category in Moreno's 15-minute model [14], it is omitted here, given that "choice" is not really applicable. By this reasoning, `public_service` does not make sense as a category where "choice" is relevant - however, in this context, AAPC is intended to approximate access to a range of tags within the category (assuming two instances of the same tag are unlikely to be situated close together). `healthcare` is another interesting category, as typically the closest instance is most desirable, so again "choice" represents different sub-tags.

4.5 Isochrone Crossover

Overlap polygons form the basis for the metric calculation. Producing these polygons is conceptually quite easy - like constructing a Venn diagram, only with non-uniform polygons rather than circles. Algorithm 1 shows the steps required to do so. Unfortunately, this algorithm is likely to be quite slow to execute - being at least $O(|P| \cdot |I|)$ from the loop on lines 5-7. Additionally, both the union on line 3 and splitting on line 4 may become slow for more complex isochrones (proportional to the total number of edges). The intersection counts produced on line 7 are for constructing the weighted sum defined in Section 4.6 below.

Algorithm 1 Isochrone Crossover

Input: List of Isochrone Polygons, I

Output: List of Crossover Polygons, P, with "overlap count" value

1: **procedure** COUNT_OVERLAPPING_FEATURES(I)
2: $B \leftarrow$ List of all edges from all isochrones
3: $P \leftarrow$ Union of all isochrone geometries
4: $P \leftarrow$ Split up P using B
5: **for** each polygon $p \in P$ **do**
6: Count number of elements of I intersecting with p
7: Store the intersection count within the polygon object p
8: **return** P

4.6 Metrics

Choice is quantified in this work by counting the number of POIs which can be accessed within a given range - rather than simply the distance to closest POI instance, as in existing studies. By first calculating the isochrone overlap polygons (using the algorithm above...), then analysing the population distribution across these regions, nuanced insight can be gained into how services are distributed compared to residents - and the possible impacts that might have. Scores are calculated using these distributions, per city per category, with the following formulas. Let $\{w_n\}$ define a set of weights, where w_n acts as a multiplier on the population within the n-overlap polygon. For service category $\lambda \in \Lambda$ and time $t \in \Gamma = \{5, 10, 15\}$, the Population curve $P_\lambda(t)$ is given by Equation 1:

$$P_\lambda(t) := \sum_{n=0}^{n_{\max}(\lambda)} w_n p_\lambda(t, n) \tag{1}$$

where $p_\lambda(t, n)$ is the population with access to precisely n instances of service λ within an t-minute walk. $n_{\max}(\lambda)$ is the largest n for which $p_\lambda(t, n) \neq 0\ , \forall t \in \Gamma$. Plots containing $P_\lambda(t)$ for different categories will act as the main focus of evaluation in this work.

To turn these population distributions into scalar values, for more intuitive comparison, the following novel metric is defined. Taking $w_n = n$ as weights, and normalising the sums by the city's total population P_{total} gives the Average Amenities per Capita (AAPC) score in Equation 2:

$$\bar{A}_\lambda(t) := \frac{P_\lambda(t)}{P_{total}} \qquad \text{where} \qquad P_{total} := \sum_{n=0}^{n_{\max}(\lambda)} p_\lambda(t, n) \tag{2}$$

By measuring per person, cities with vastly different populations can be more easily compared.

Time Intervals The number of times for which this analysis is carried out is a trade-off. Using $t \in \{x : x \in \mathbb{N} \wedge x \leq 15\}$ is computationally infeasible, and unnecessarily fine-grained. To create a repeatable study, $t \in \{5, 10, 15\}$ is chosen - as in [24] - to include the base 15-minute range along with a couple smaller times for comparing the most accessible cites.

4.7 Visualisation

After the isochrone overlap polygons have been generated for each time interval, they can be displayed on a map. Seven independent categories (`sustenance`, `education`, `transportation`, `healthcare`, `culture_and_religion`, `leisure` and `public_service`) need to be displayed, each having integer values associated. Therefore, seven distinct colours (Figure 2) are required - one for each category. The value of the polygon (i.e. number of accessible services per category) will be indicated with its transparency, mimicking lightness [17, p. 266], with darker tones symbolising higher values; this helps to draw attention to these amenity-dense areas.

Figure 2: Colour Palette

Information density is another consideration. To control the visual complexity of the map, each category should be toggleable by the user, so will need to be drawn on separate layers. This allows specific categories of interest to be isolated, removing visual interference from others, which would be prominent in the dense urban areas being studied. To prevent any confusion associated with this, the legend should only display the colours associated with visible layers.

Interactivity is key, so there should be multiple ways to interact with the application, from mouse controls for map movement, to keyboard inputs for layer control.

Metric Graphs Ideally, the metric graph should display both the population distribution and the calculated AAPC value for every category. This provides the desired information alongside the context used to calculate it, aiding comprehension. As with the layers, each category will have its lines drawn in the associated colour. AAPC will be drawn as a vertical line, to stand out against the population curves.

To improve intentionality when utilizing the application, static plots are favoured over dynamic ones - as a constantly changing graph might be confusion to the user, and would act as an obstacle to understanding the information being presented.

5 Implementation

5.1 Flowcharts

Data Processing Figure 3 contains the intensive calculations, so can take a while to run - particularly if many cities are selected (or more categories are added). It is split into three stages (scripts). The first obtains the required bounds of specified regions. The second queries Overpass to find POI locations, saving the result to a `json` file to be sampled at the next step. Finally, the third script runs ORS to generate the isochrones, returned as geojsons to be visualised, and used for metric calculation.

Visualisation Both the metric calculation and visualisation flow are outlined in Figure 4. `bucketed_geojsons.py` (added for performance reasons; see Section 5.4) needs to be executed after the data processing pipeline outlined in Figure 3, before the visualisation will run. Furthermore, note that for metric graphs to be shown in the web visualisation, `metrics.py` has to have been executed (though isochrones can still be explored without this step).

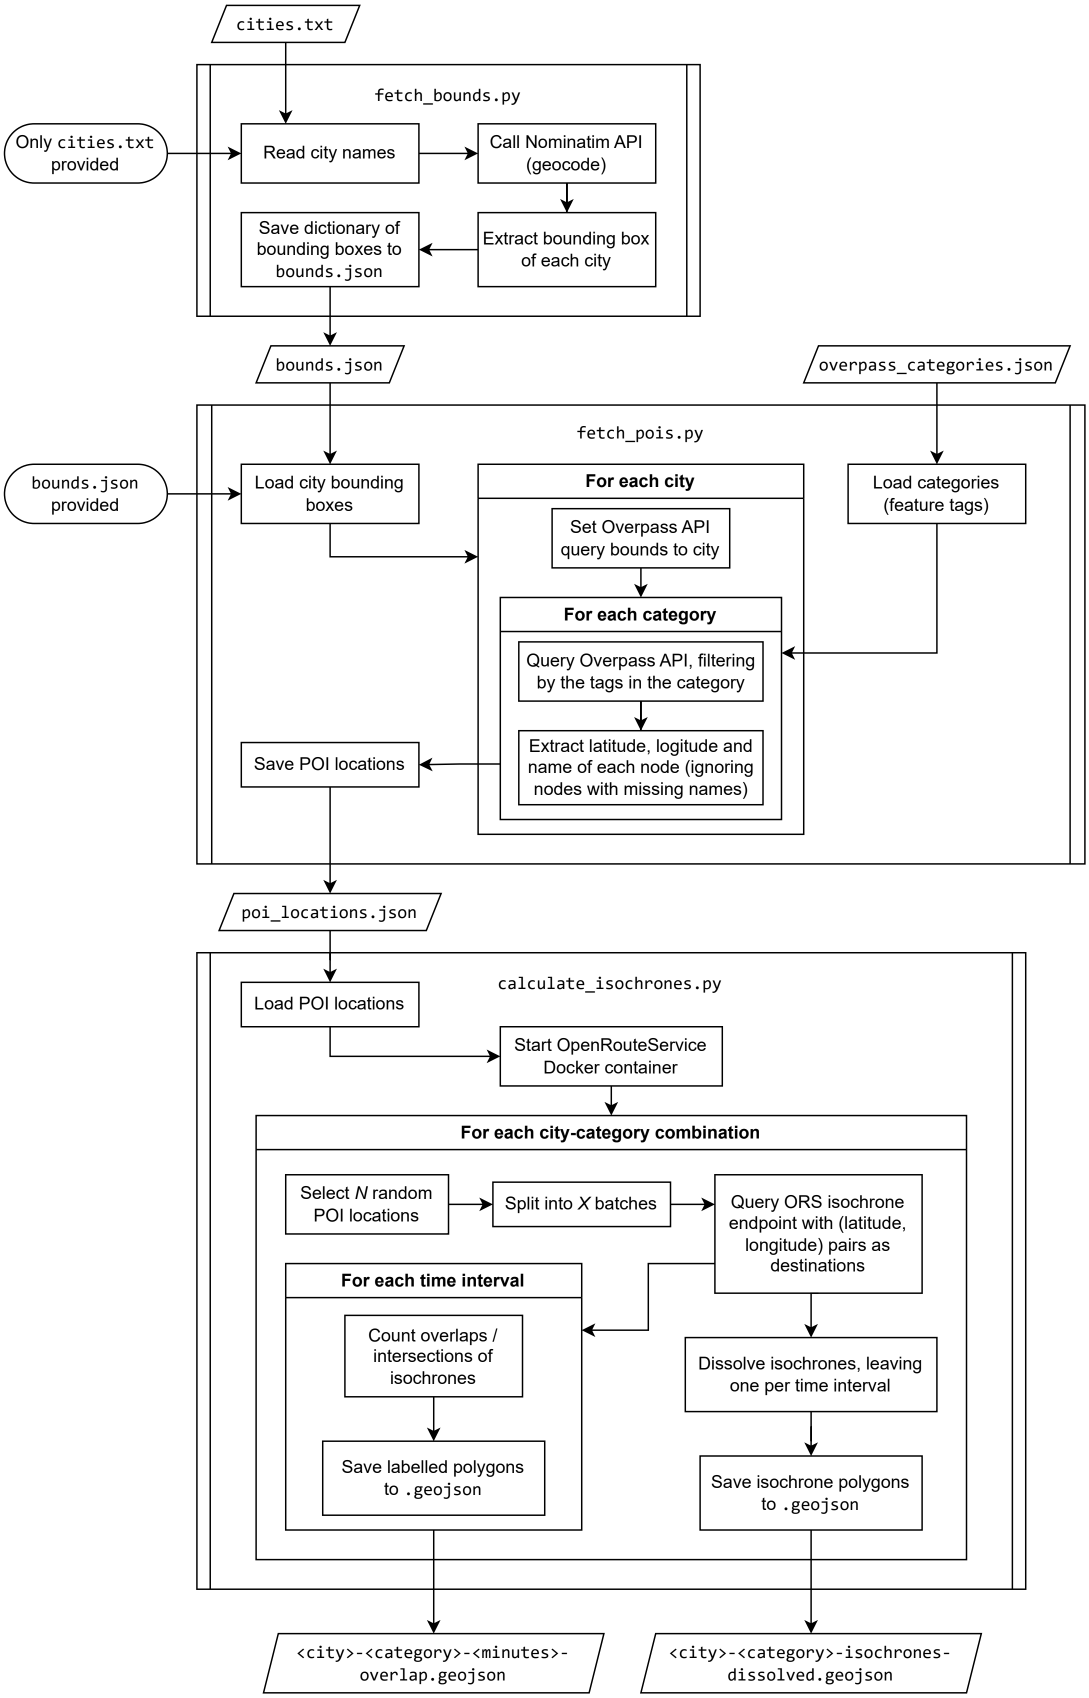

Figure 3: Pre-visualisation flowchart. N is number of samples and X is number of batches.

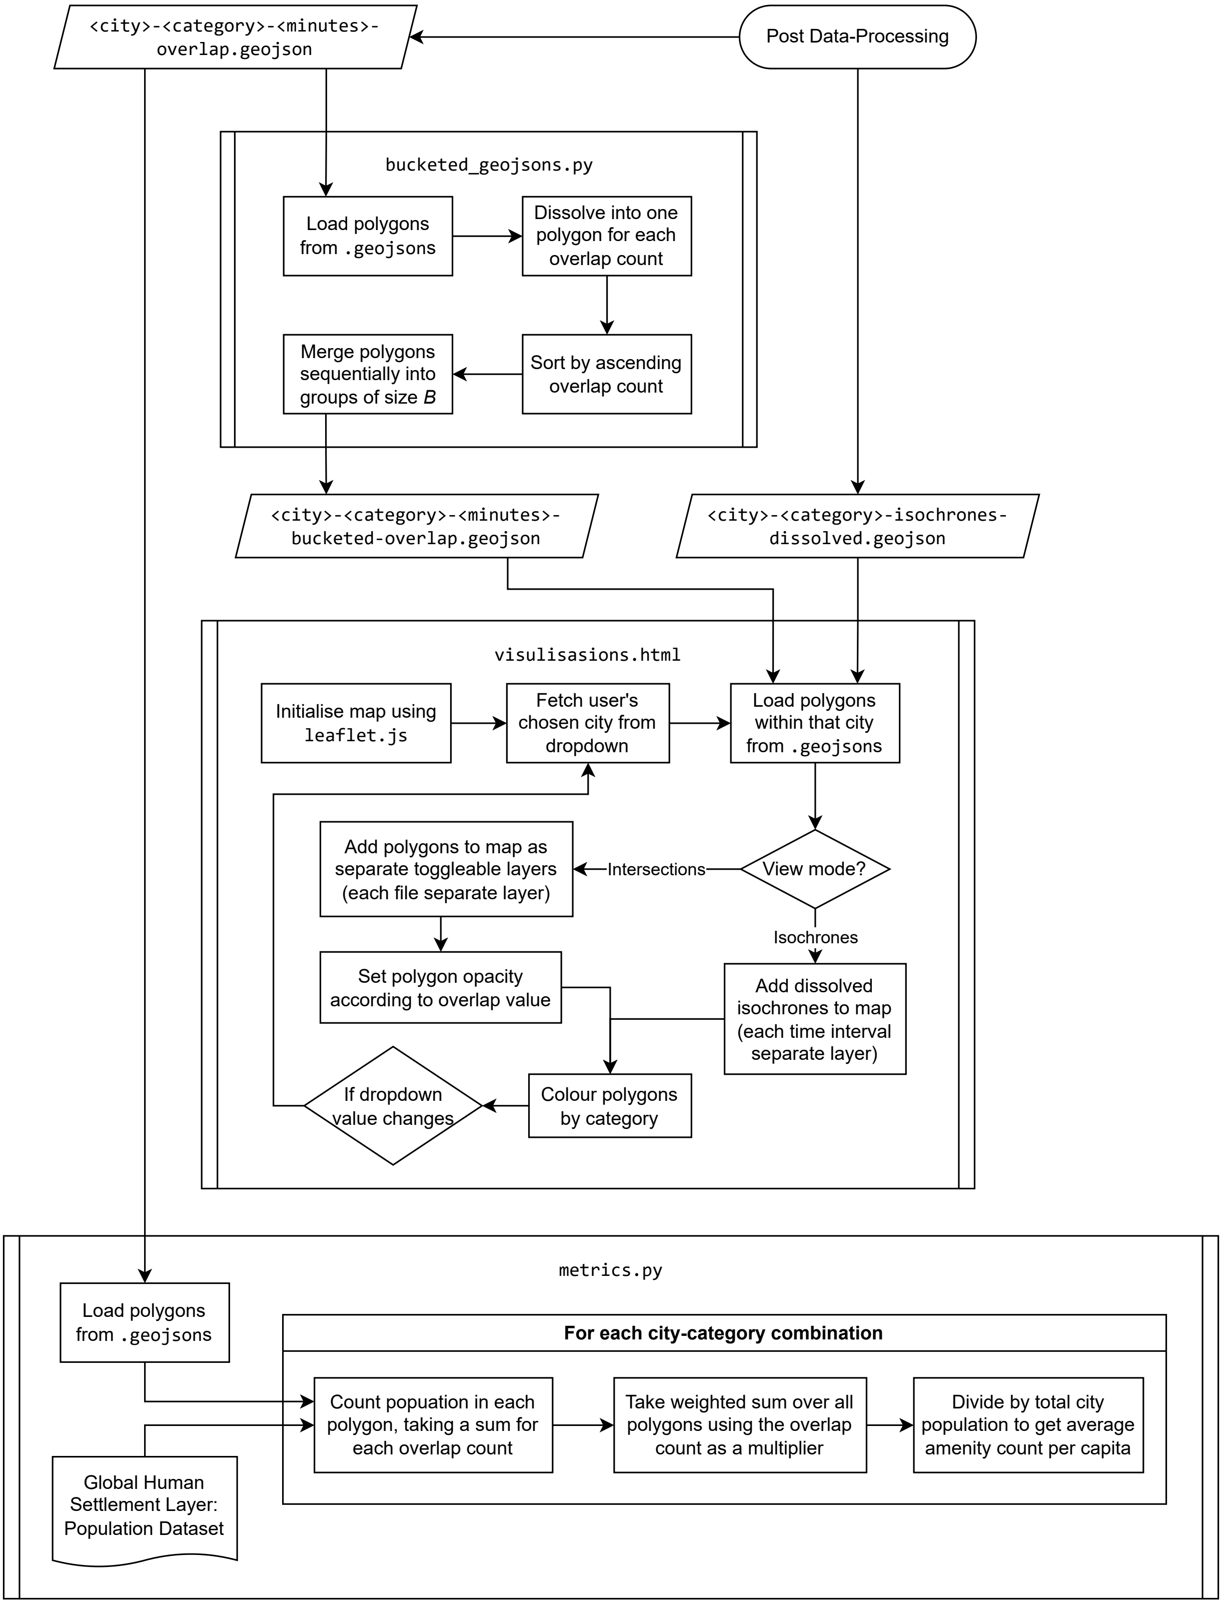

Figure 4: Post-processing flowchart. B is the size of each bucket.

5.2 API Considerations

Detailed here are the endpoints and query formats capable of fetching the necessary information, their associated rate limit policy, use of Docker, etc.

5.2.1 Nominatim

A script was used to fetch bounding boxes of all cites being studied; given the limited number, using the existing API (nominatim.openstreetmap.org) was desirable, so a 1 second pause between queries had to be added to account for rate limiting. Specifically, the returned bounds are a list of 4 floats, denoting South and North Latitude, followed by West and East Longitude.

The query is sent to `https://nominatim.openstreetmap.org/search?q=<city_name>,+<search_region>&format=json`, replacing `<city_name>` with the city being evaluated. `<search_region>` is used to refined the search - first it is set to *UK*, then - if no result of `addresstype` *city* or *town* is returned - set to *Ireland* and the query resent. Python's `requests` library handles the API interactions. This API's fair use policy includes requiring a `User-Agent` with the query, to identify it - this script specifies `"fetching bounds for 15 minute cities"`.

5.2.2 Overpass

Bounds returned by Nominatim acted as the search area for *overpy* queries. This API allows filtering of features by any of their tags/properties. Filtering by tags required formatting a composite query. For demonstrative purposes, an example query is given in Figure 5. Returned `(latitude,longitude)` pairs were used as "destination" locations in calls to ORS (next step). Those without identifying name tags were removed from the response. Rate limits were not a problem with this API - any returned `OverpassTooManyRequests` exceptions could be caught and resolved by waiting a few seconds to retry the request.

```
[timeout:60][bbox:51.3972838,-2.7183704,51.5444317,-2.5104192];
(
    nwr["amenity"="place_of_worship"][name];
    nwr["tourism"="attraction"][name];
    nwr["tourism"="gallery"][name];
    nwr["tourism"="museum"][name];
);
out center;
```

Figure 5: An example Overpass API query, for the `culture_and_religion` category, within the area of Bristol. Adding `[name]` returns only features with identifying name labels.

5.2.3 Openrouteservice

The default docker compose file had to be modified to increase memory allocation given the size of the street data store; as well as the `config` file, to increase the location and interval allowance (3 intervals and up to 25 locations) for each call to the endpoint. As the sample size (discussed in Section 5.3.2 below) needed to be greater than 25 to suitably represent the area, the list of node locations was split into batches and isochrones evaluated using multiple queries.

The endpoint sits at `ors/v2/isochrones/foot-walking`. Figure 6 presents the structure of the query, specifying the target (POI) locations and time intervals. These isochrone can be then passed through the overlap algorithm in Section 4.5 to produce the polygons required for the remaining steps - namely metric calculation and visualisation.

```
{
    "locations": [... <(lat, long) pairs> ...],
    "range": [900,600,300],
    "range_type": "time",
    "location_type": "destination"
}
```

Figure 6: `json` passed into the ORS API to calculate isochrones. `locations` takes the list of centre points as (latitude, longitude) pairs.

5.3 Metric Calculation

In order to obtain AAPC scores, two prerequisites must be met. For the calculation to even be feasible, a representative sample of POI locations has to be taken. Alongside this, for accuracy, a high resolution population dataset is required - relative to the precision of the overlap isochrones. Both of these considerations are addressed below.

5.3.1 Data Resolution

Spatial accuracy of the population dataset was a problem, as mentioned in Section 4.2.2. Some of the polygons formed at the isochrone intersections were small enough as to not cover a full pixel of the GHSL dataset, so were considered empty; others were packed dense enough to cause double or triple counting. These artefacts would lead to either runtime errors[16], or the total population being too large[17]. In an attempt to fix this, the population data was resampled from 3 arc seconds to 0.3 arc seconds. Shown in Figure 7, values were retained (divided by 100 to correct for the new square area).

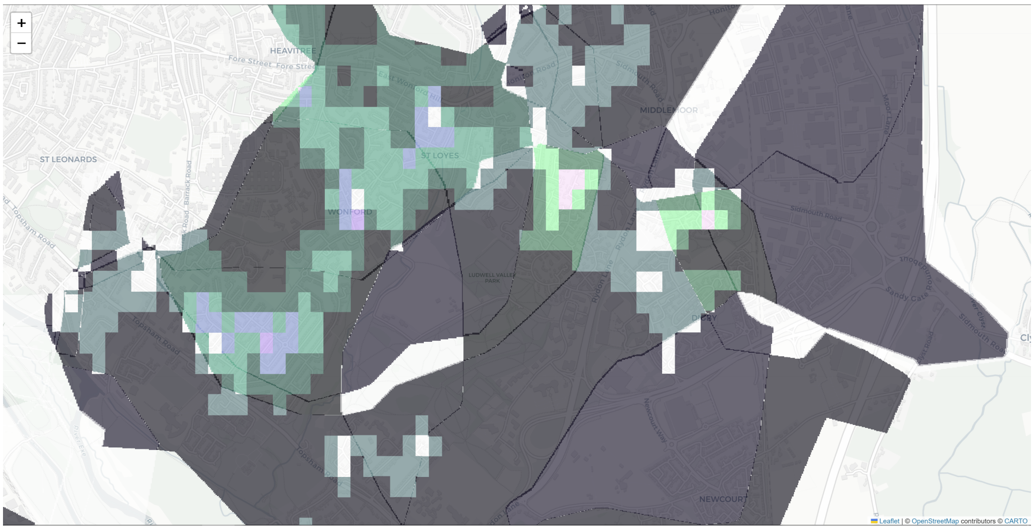

Figure 7: Resampling at 10x resolution.

The only assumption here which reduces the approximation accuracy slightly is that the population was uniformly distributed within each square. Another, slower method - like bilinear filtering - could instead have been used to smooth out the boundaries, but this equally makes assumptions about the distribution, so is just as approximate (without extra population information being available). Doing this drastically reduced both the number of null (empty) polygons, along with the overestimation due to double counting, improving the precision of the final calculations.

[16] `NoDataInBounds` error thrown by `rioxarray`, the library being used to load the GHSL tiles.

[17] At times exceeding the recorded population of the city; in extreme cases, this could result in a negative AAPC.

5.3.2 Sampling

Using all fetched POIs to calculate AAPC would have been infeasible - as obtaining the intersections is a computationally intense process which scales poorly with the number of isochrones - so a sample of the POIs in each city had to be taken. Repeatability will have been reduced slightly by this choice, due to the introduction of non-seeded stochasticity into the calculation, though using an equivilent sample size should yield similar values[18]. In this study, the number of samples was determined as a function of population: the nearest integer to 3 times the cube root[19].

$$\left\lfloor 3 \times \sqrt[3]{\text{population}} + \frac{1}{2} \right\rfloor \tag{3}$$

The cube root was chosen to be a balance between population representation and computation feasibility, though this is a hard balance to strike. Replacing with either a square root, or even linear, function (scaling down accordingly) might improve representation. However, depending on the chosen scaling, doing this would either: exclude areas with larger populations from the study - due to computational requirements - or reduce the sample size of smaller areas to a point where the metric is no longer representative.

This sampling process acts as one of the main areas for improving this project in the future. Different sample size functions, or more nuanced sampling methods (e.g. spatial stratification or clustering), could be introduced to more fairly represent all cities.

5.4 Visualisation Performance Optimisations

Performance also became an issue during testing. Even for a relatively small number of samples (e.g. 100 POIs), a large amount of overlap was found for the more densely-populated cities. For example, the `leisure` in Exeter, with a 15 minute interval, had up to 65 overlapping isochrones one point - the result of which was many small and precise polygons. The inherent metric accuracy issue was solved by resampling the population dataset - discussed in Section 5.3.1 - but an issue was also encountered when attempting to visualise the polygons. Drawing these polygons in *Leaflet.js* suffered from obvious aliasing (see Figure 8) - and slow performance for even 1 category, let alone drawing all 7 simultaneously. As responsiveness was a key goal for the visualisation, this issue had to be resolved.

Polygon Simplification An attempt was made to tackle both of these problems using polygon simplification. *topojson*[20], a "topologically-aware" simplification algorithm, was chosen as many of the target polygons were touching. Where a regular simplification algorithm would likely leave gaps, this algorithm retained the structure of adjacent polygons. However, a suitable threshold could not be found. A value between 0.25 and 0.5 was the best, but this significantly misrepresented some of the protruding geometry, while only minorly improving performance; The visualisation still noticeably struggled when rendering, so a different approach was taken.

Isochrone Bucketing It was clear that rendering all polygons was not feasible, so instead, a script was created to group the polygons into "buckets" (i.e. 0-4 overlaps, 5-9 overlaps, etc.). Using the `dissolve` functionality provided by the `geopandas` library, this was easy to implement. Polygons were first dissolved, then sorted, by overlap count. Indices were then assigned in sequential groups of 5, allowing the polygons to be further dissolved by their index. This naturally simplifies the polygons, reducing the number of vertices without changing the surrounding shape. Figure 9 shows these merged polygons on a map.

[18] Unless the number of available POIs, or sample size, is very small.

[19] Adding 0.5 then applying the floor operation is equivalent to rounding to the nearest integer.

[20] github.com/topojson/topojson

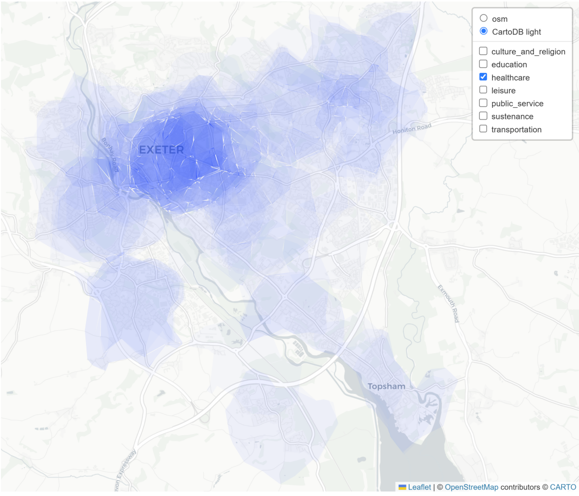

Figure 8: Visible aliasing artifacts on an early prototype of the data visualisation.

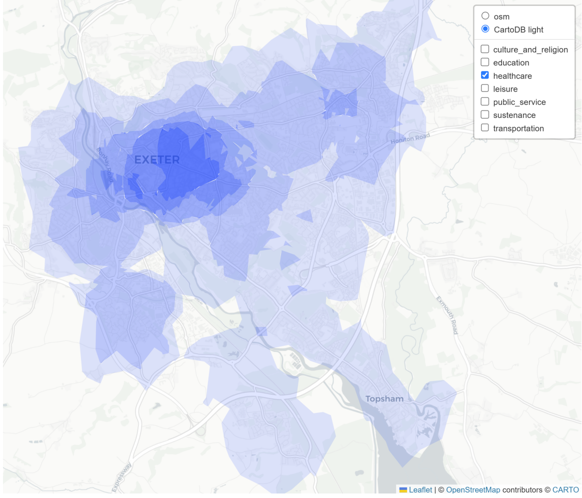

Figure 9: Polygons grouped by intervals of 5 overlaps.

5.5 Final Visualisation

The produced visualisation tool can be seen in Figures 10 and 11. The former displays many isochrones with union applied by time interval (as was used in [24]). The latter shows the isochrone overlaps, with a legend in the bottom right, and graph of the metrics for that time interval in the top left. By default, this is how the visualisation application appears to the user. The two figures also show the two sets of background tiles available (`osm` and `CartoDB`). Important conclusions drawn from the data, with the help of this tool, will be covered in the following section.

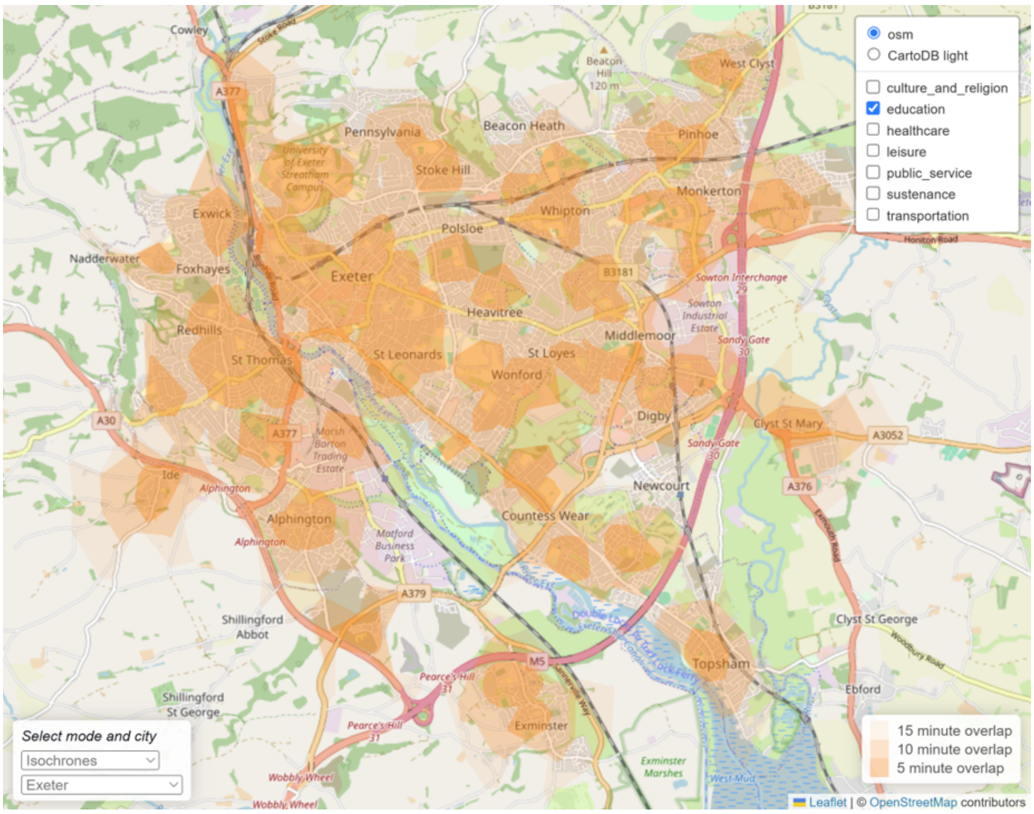

Figure 10: Isochrones for the `education` category in Exeter overlaid onto a map.

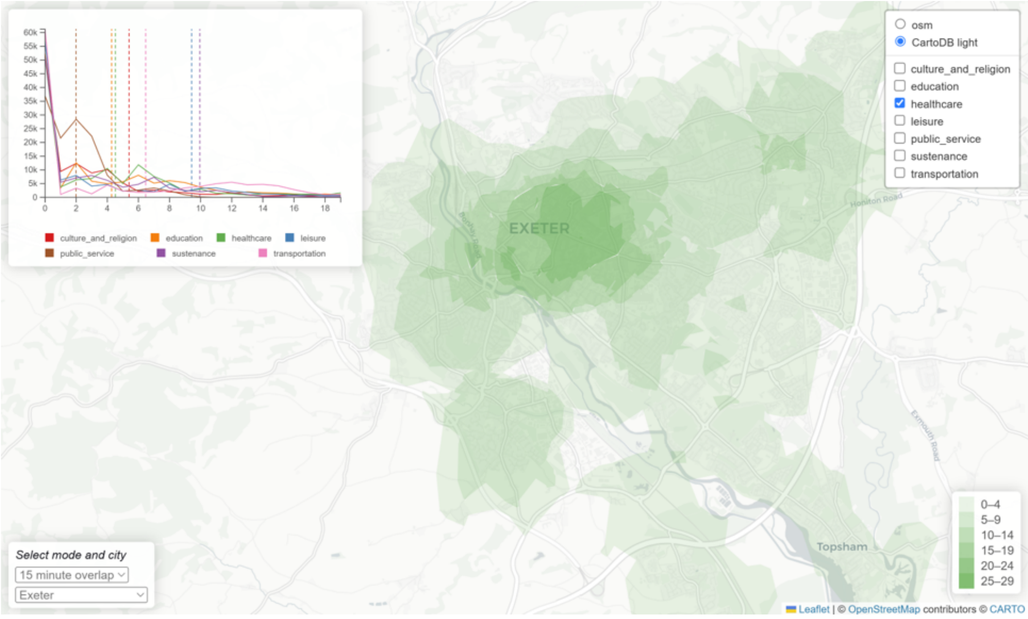

Figure 11: 15-minute isochrone intersections for `healthcare` category in Exeter.

6 Results

6.1 Rank Tables

Provided in Tables 2, 3 and 4 are the city rankings in each category - with lower number indicating better ranks. These are totalled up and ranked again to provide an overall order, shown in the final column. Exact metric values are provided in Tables 8, 9 and 10 which can be found in Appendix A.

City	sus.	edu.	tra.	hea.	cul.	lei.	pub.	Totals	Overall
London	25	25	25	25	25	25	22	172	25
Oxford	4	1	3	1	1	6	10	26	4
Bristol	17	6	15	14	11	15	11	89	14
Southampton	9	11	9	6	10	7	12	64	8
Newcastle upon Tyne	19	14	18	13	14	19	20	117	18
Bradford	24	22	24	22	22	24	19	157	23
Edinburgh	16	8	17	11	12	13	2	79	11
Aberdeen	10	18	13	8	9	9	16	83	12
Exeter	3	5	5	12	5	3	7	40	5
Bournemouth	12	16	10	7	8	8	9	70	9
Brighton	7	4	11	5	13	11	4	55	6
Margate	1	9	1	2	3	1	5	22	1
Maidstone	23	24	23	24	24	23	21	162	24
Crawley	8	21	7	21	20	18	1	96	15
Reading	2	3	2	4	4	2	6	23	2
Milton Keynes	21	15	19	20	23	21	25	144	21
Cambridge	5	2	4	3	2	5	3	24	3
Colchester	18	19	21	17	18	14	24	131	19
Ipswich	6	13	6	15	6	4	13	63	7
Peterborough	20	23	20	23	21	20	23	150	22
Coventry	15	12	12	18	17	12	17	103	16
Nottingham	13	7	14	10	7	16	8	75	10
Manchester	22	20	22	19	19	22	14	138	20
York	14	17	16	9	15	17	18	106	17
Preston	11	10	8	16	16	10	15	86	13

Table 2: AAPC city ranks: 15 minute

City	sus.	edu.	tra.	hea.	cul.	lei.	pub.	Totals	Overall
London	25	25	25	25	25	25	23	173	25
Oxford	2	1	3	1	1	4	7	19	1
Bristol	17	8	15	14	9	14	11	88	13
Southampton	10	9	8	7	10	8	12	64	7
Newcastle upon Tyne	18	14	18	10	14	17	20	111	18
Bradford	24	23	24	22	21	24	15	153	23
Edinburgh	16	11	19	11	13	13	2	85	12
Aberdeen	9	18	13	8	11	10	19	88	13
Exeter	5	5	5	13	5	3	8	44	5
Bournemouth	12	16	9	5	6	7	9	64	7
Brighton	7	4	11	6	12	11	4	55	6
Margate	1	6	1	2	3	1	5	19	1
Maidstone	23	24	23	24	24	22	21	161	24
Crawley	8	21	7	19	20	19	1	95	15
Reading	3	3	2	4	4	2	6	24	4
Milton Keynes	21	12	17	20	23	20	25	138	20
Cambridge	4	2	4	3	2	5	3	23	3
Colchester	19	19	20	17	19	15	24	133	19
Ipswich	6	13	6	16	8	6	13	68	9
Peterborough	20	22	21	23	22	21	22	151	22
Coventry	14	15	12	18	17	12	17	105	16
Nottingham	13	10	14	12	7	16	10	82	11
Manchester	22	20	22	21	18	23	16	142	21
York	15	17	16	9	16	18	18	109	17
Preston	11	7	10	15	15	9	14	81	10

Table 3: AAPC city ranks: 10 minute

City	sus.	edu.	tra.	hea.	cul.	lei.	pub.	Totals	Overall
London	25	25	25	25	25	25	24	174	25
Oxford	2	1	2	2	1	4	6	18	1
Bristol	16	6	16	13	8	14	10	83	12
Southampton	11	11	9	8	9	8	13	69	8
Newcastle upon Tyne	18	14	18	12	13	15	22	112	18
Bradford	24	21	24	22	22	24	12	149	23
Edinburgh	15	7	17	10	12	12	2	75	9
Aberdeen	12	19	11	9	14	10	23	98	14
Exeter	5	9	4	11	5	5	9	48	5
Bournemouth	10	16	8	6	6	7	7	60	7
Brighton	7	4	10	5	11	11	4	52	6
Margate	1	13	1	1	3	1	5	25	3
Maidstone	22	24	22	24	24	22	19	157	24
Crawley	9	23	6	19	21	20	1	99	15
Reading	3	3	3	4	4	2	8	27	4
Milton Keynes	20	10	19	20	23	19	25	136	20
Cambridge	4	2	5	3	2	3	3	22	2
Colchester	21	18	20	17	19	16	20	131	19
Ipswich	6	12	7	16	10	6	18	75	9
Peterborough	19	22	21	23	20	21	21	147	22
Coventry	13	15	12	18	15	13	16	102	16
Nottingham	14	8	14	14	7	18	11	86	13
Manchester	23	20	23	21	18	23	17	145	21
York	17	17	15	7	16	17	15	104	17
Preston	8	5	13	15	17	9	14	81	11

Table 4: AAPC city ranks: 5 minute

6.2 Analysis

Metric value distributions across cities were very similar for each time interval (note the colour similarity between the three heat maps in Figure 12) - with ranks differing by at most three across the three tested times, for any one city. Implying that AAPC is (almost) independent of time, any conclusions will generalise well from the results of any time interval. The heatmaps also accentuate the extremes of the data, namely Margate and London, as well as the `public_service` category. Key observations are split per city, per category, then in context of other existing policy. A statistical analysis is also presented to solidify the significance of results.

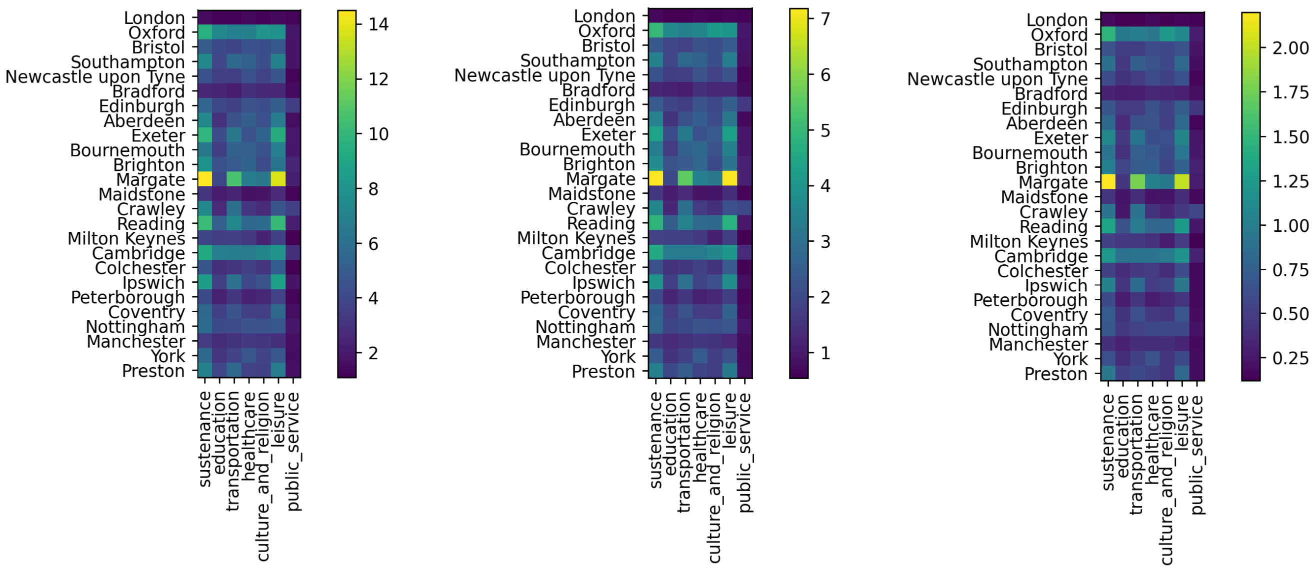

(a) 15-minute

(b) 10-minute

(c) 5-minute

Figure 12: Heatmaps of metric values. Uses values from tables in Appendix A.

6.2.1 Cities

Margate exhibits the greatest AAPC values for any city or category by a large margin - 14.5 for `sustenance` and 13.8 for `leisure`. Overall, it ranks the highest, though coupled with the regions small population and land area this could imply the result to be an outlier (revisited below). At the 5-minute range it falls to 3rd (as its `education` ranking falls), and at 10-minutes it is joint first with Oxford.

Oxford ranked repeatedly among the highest over all categories (being ranked 1st in `education` and `culture_and_religion` at all time intervals). It is 1st overall for 5 and 10 minutes, but drops to 4th (behind Margate, Reading and Cambridge) at the 15 minute threshold. Apart from the `public_service` category, AAPC scores are 8 ± 2, making it the most consistent scoring city.

Greater London ranked consistently lowest across all categories and time intervals, underscoring the challenge of amenity distribution for the largest urban areas. This shows that the spatial distribution of residents is poorly correlated with service distribution - i.e. many services are concentrated in the centre, whereas the bulk of the population is spread further around the outskirts. The metric estimates that, on average, those living in London do not have access to more than one instance of any amenity category, meaning their "choice" is hindered.

Metric graphs for the three contrasting cities discussed above can be seen in Figure 13, showing how isochrone distribution relates to the average number of accessible amenities. London covers a larger area, but with fewer overlaps, and the high concentration of amenities discussed above can be clearly seen. Two lines, intersecting at the centre, are formed by the overlap polygons for Oxford. Margate, with AAPC scores around 14 for two categories (specifically `sustenance` and `leisure`), shows a single main hub where services are concentrated. With London and Margate being the metric, population and area extremes of the study, the correlation to AAPC is investigated in Sections 6.3 and 6.4.

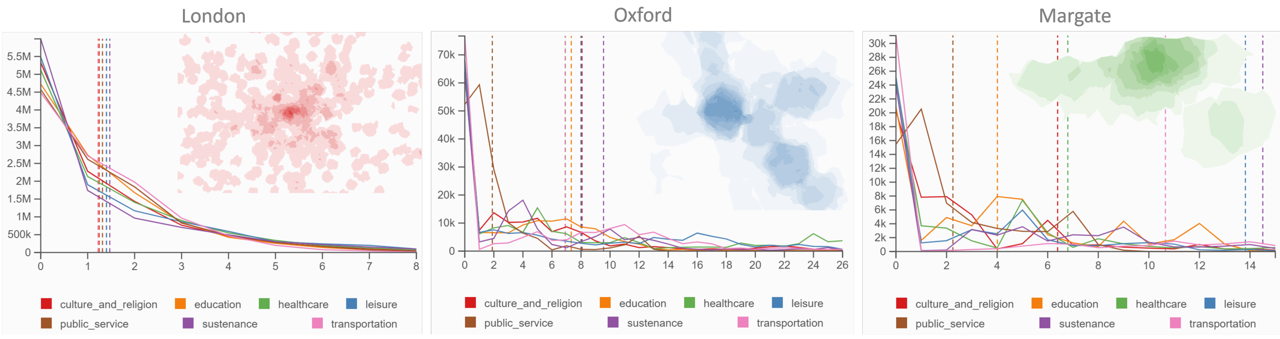

Figure 13: Various amenity choice population distributions for a time interval of 15 minutes. The intersection polygons for one category in each city are overlaid in the top right of each graph.

Principal Component Analysis To further interpret the results, Principal Component Analysis (PCA) is applied to the rankings, reducing the dimensionality down from 7 to 2. On these new points, K-means clustering can be applied to find distinct clusters. Figure 14 shows this approach on the 15 minute rankings in Table 2. Roughly speaking, the first component could represent overall service accessibility, with lower values relating to higher total ranks; the second can be thought to capture rank variance, with low magnitude values indicating better consistency.

The blue cluster contains the best ranking cities, high across all categories - including cities like Oxford, Exeter and Margate. Green contains other cities with consistent but lower scores. To contextualise, these cities are located near the coast, with some southern (Southampton/Brighton) and some northern (Aberdeen/Edinburgh). Cities with mixed performance - such as Bristol and Newcastle - make up the yellow cluster. These have some strong and some weak categories, so may be in the process of decentralising their services. Last of all, the red cluster highlights underperforming regions - including the aforementioned London, and others like Manchester. Having poorer rankings across most categories, it is clear that larger, sprawling urban areas struggle - and would benefit most from localising services, evident in the abundance of congestion and inequality. Clusters are similar for 10 and 5 minutes - omitted here, but included in Appendix B - as expected (see Section 6.2).

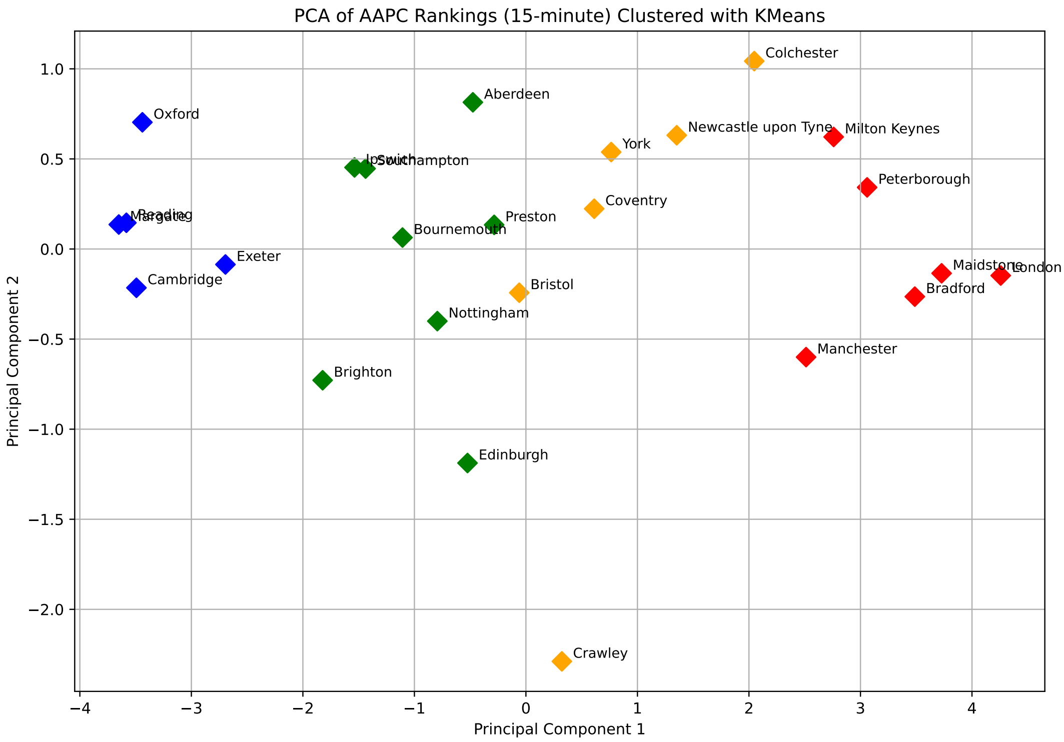

Figure 14: PCA on 15 minute ranks, split into 4 clusters.

Overall AAPC Total 15-minute AAPC values, coloured by category, are presented in Figure 15. This solidifies the observations above and facilitates easy comparison between cities. It also provides another alternative ranking to that in Table 2. On average, Oxford/Cambridge residents can access almost 50 different amenities, across the studied categories; while in somewhere like Manchester or Peterborough, this number is only 20 - accentuating the discrepancy between these areas that prior studies have failed to acknowledge.

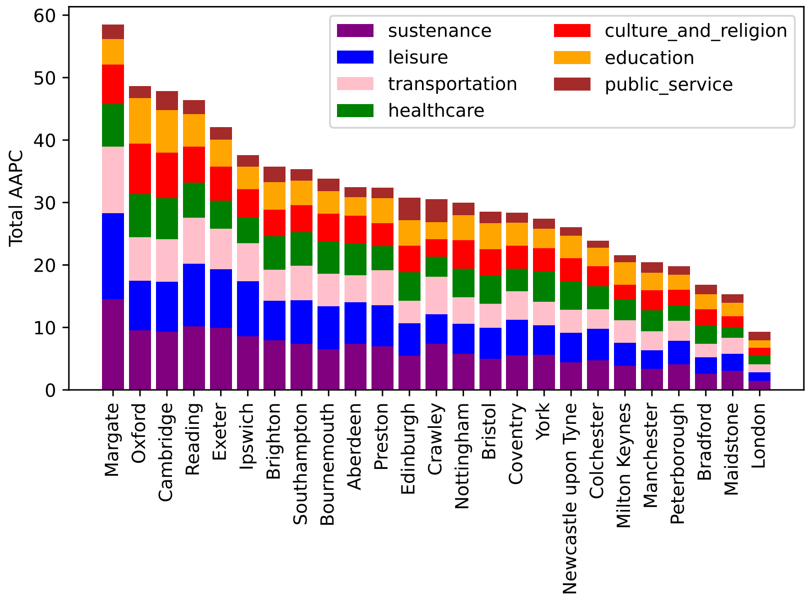

Figure 15: Stacked bars of average metric values per city, for 15 minute interval.

6.2.2 Categories

The `sustenance` and `leisure` categories offered the most choice, which is to be expected as these amenities are in the business of competition. Decentralising across residential zones to meet demand boosts their business, so is a desirable venture.

Public services were considerably less accessible than all other categories - with an average AAPC score approximately 30% of the maximum, as seen in Figure 16. This is despite containing over 3 times as many distinct tags as the smallest category, `culture_and_religion`, showing that AAPC is not exclusively determined by category size. Amenities in this category require lengthy planning processes, making it difficult to keep up with urban sprawl, and fundamentally leading them to fall behind other services.

It should be highlighted that certain amenities holding personal significance (e.g. places of worship) may have a higher likelihood of being thoroughly tagged in a crowdsourced dataset, like OSM. Unfortunately, this is nearly impossible to verify without manual ground-truthing of study sites.

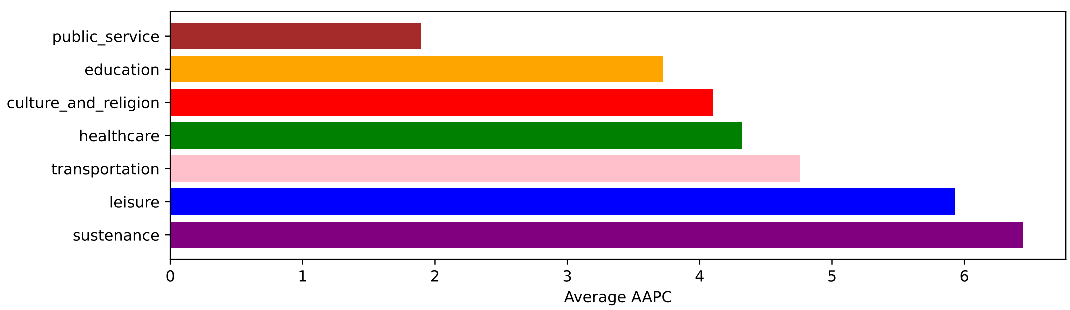

Figure 16: Average metric values per category, sorted, for 15 minute interval.

Category Profiles Individual distributions for some categories are given in Figure 17 below. The `healthcare` category is given a similar weight by many cities, having the least pronounced peak, exhibiting less overall city-to-city variance. Apart from Margate, the `sustenance` category is similarly distributed, though the absolute value is higher. `education` peaks in Oxford and Cambridge, which is to be expected given their globally-recognised academic prestige. The least overall variance, as well as smallest values, are seen for the `public_service` category - accentuating the issue discussed above. Omitted from Figure 17 are `transportation` and `leisure` (distributed most similarly to `sustenance`), and `culture_and_religion` (closely following the `education` distribution); these plots are included in Appendix B.

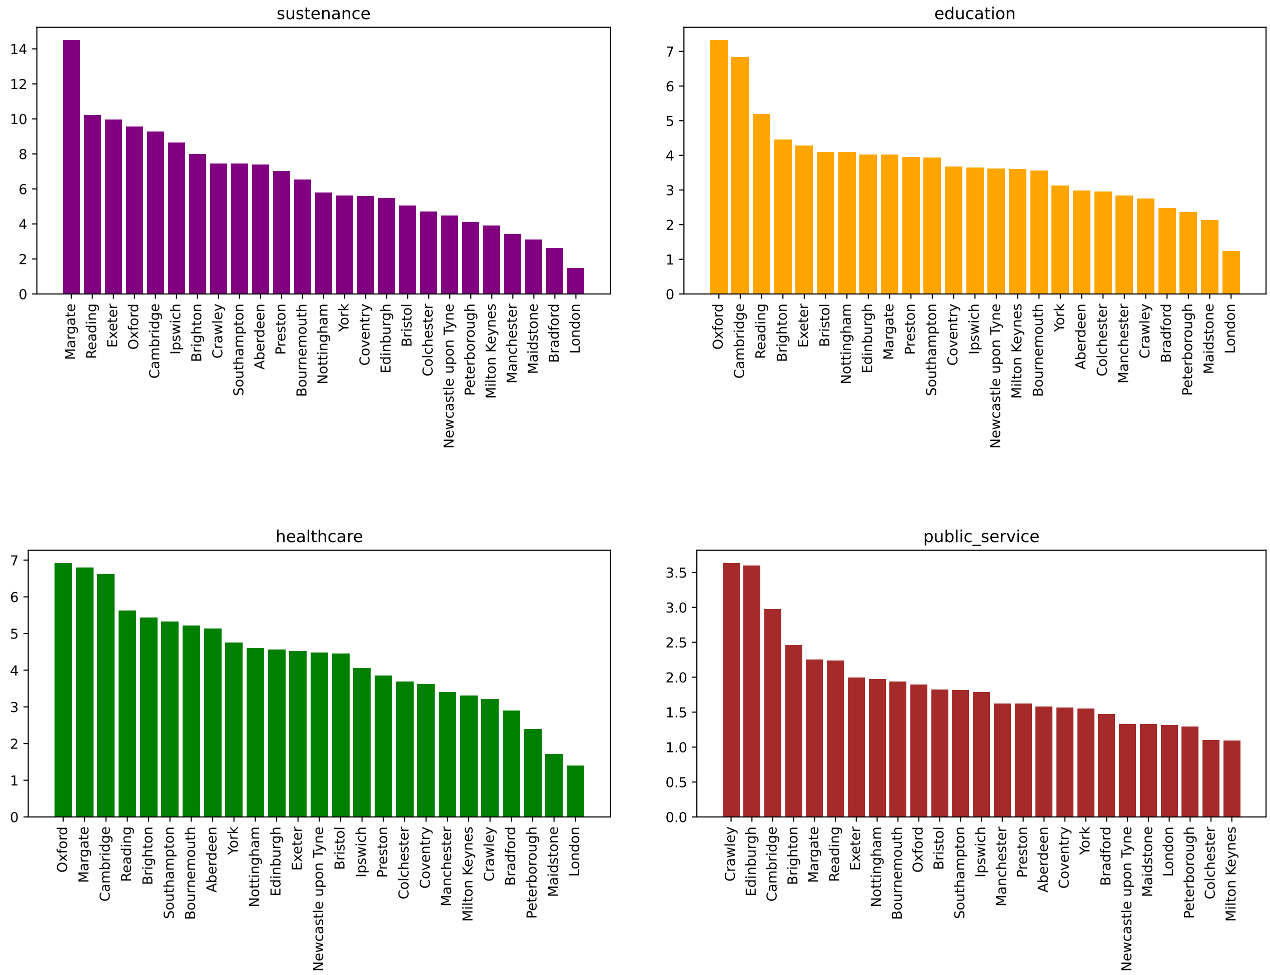

Figure 17: Average metric values per city, sorted, for 15 minute interval. Similar distributions omitted.

Clustering From these distributions, a few informal clusters can be suggested. These groupings indicate three distinct types of amenity, which could be used to guide investment: the first requires little intervention; the second could benefit from development, targeted at the lowest scoring cities; the last requires intentional re-evaluation of resource allocation, even amenity relocation in some cases.

- Commercial: `sustenance`, `transportation` and `leisure` - the top performing categories, all of which are distributed similarly across cities.
- Institutional: `education`, `healthcare`, and `culture_and_religion` - all exhibiting values of similar magnitude.
- Public: `public_service` falls into a group by itself, being managed locally by councils.

6.3 Population-to-Metric Correlation Study

Section 5.3.2 discusses the requirement to sample POIs for metric calculation, with the number of samples being related to a city's population. Most effected by this was likely the Greater London area, as the discrepancy between fetched and sampled POIs would have been greatest.

To test if this sampling methodology had induced any bias in the results, a correlation study was undertaken. Pearson correlation coefficients between city population and average AAPC were computed for each time interval, and are shown in Table 5. The `public_service` category has a p-value much greater than 0.05 (a typical significance threshold), making the correlation non-robust. Across the remaining statistically significant results, a correlation of between -0.38 and -0.45 is found, suggesting a slight negative correlation. As this is only a weak to moderate correlation, it is clear this is not the primary factor influencing AAPC scores, and that other factors are also at play. Consequently, the conclusions presented above can still be considered valid.

City	15 minute	10 minute	5 minute
`sustenance`	-0.419 (0.0373)	-0.414 (0.0395)	-0.384 (0.058)
`education`	-0.422 (0.0358)	-0.433 (0.0306)	-0.4 (0.0476)
`transportation`	-0.429 (0.0323)	-0.428 (0.0327)	-0.405 (0.0445)
`healthcare`	-0.454 (0.0227)	-0.458 (0.0215)	-0.434 (0.0303)
`culture_and_religion`	-0.407 (0.0435)	-0.404 (0.0455)	-0.376 (0.0637)
`leisure`	-0.406 (0.044)	-0.397 (0.0496)	-0.38 (0.061)
`public_service`	-0.187 (0.371)	-0.196 (0.349)	-0.196 (0.348)

Table 5: Pearson correlation between AAPC and population. Confidence (p-value) in brackets, rounded to 3 significant figures.

Figure 18 plots AAPC values against the logarithm of the population (such that London, furthest right, can be included) for the 15 minute interval, visualising this slight negative correlation. Plots for 10 and 5 minutes are visually very similar, so can be found in Appendix B. It is important that any future work into modifying the sampling method or size function continues this analysis to avoid the pitfall of sampling bias.

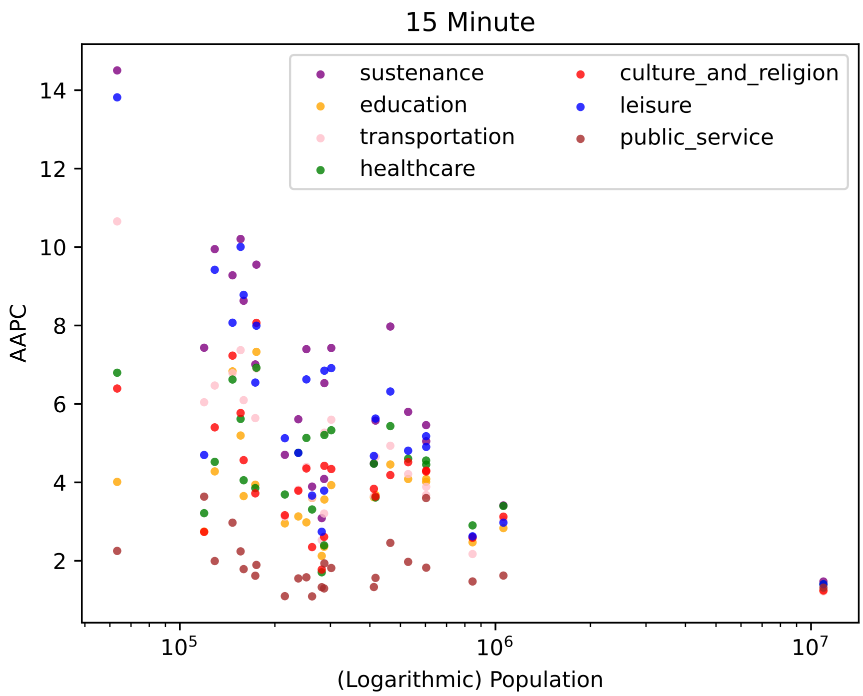

Figure 18: AAPC plotted against the logarithm of the population.

This negative correlation is logical, as cities with larger populations would require denser residential areas in order to match the AAPC values of a less populated city. It is apparent that larger cities choose to centralise more services, attempting to satisfy as many residents as possible by minimising the average distance to all residents. Yet, this has the opposite effect for cities covering a large geographical area, as, when this minimum distance is large, the resulting service distribution is not walkable for anyone. Most strongly correlated is the `healthcare` category, while `leisure` is the least[21]. However, the variance between these two extremes is small, so no conclusion is drawn from this observation.

6.4 Area-to-Metric Correlation Study

Land area varies drastically between study sites - from London, covering 2640km^2, to Margate, only 27km^2 (calculated according to API query bounding boxes). A similar correlation analysis was carried out to examine the relation between it and AAPC, with coefficients given in Table 6. The correlation exhibited is greater, averaging -0.56 across all intervals, with much stronger significance (though `public_service` is still over the threshold so not considered valid[22]).

City	15 minute	10 minute	5 minute
`sustenance`	-0.567 (0.00313)	-0.566 (0.00316)	-0.533 (0.00611)
`education`	-0.564 (0.0033)	-0.558 (0.00373)	-0.505 (0.00996)
`transportation`	-0.62 (0.00095)	-0.61 (0.00119)	-0.565 (0.00328)
`healthcare`	-0.577 (0.00254)	-0.574 (0.00269)	-0.547 (0.00468)
`culture_and_religion`	-0.573 (0.00275)	-0.57 (0.00295)	-0.54 (0.00534)
`leisure`	-0.565 (0.00322)	-0.553 (0.00415)	-0.533 (0.00607)
`public_service`	-0.278 (0.179)	-0.275 (0.184)	-0.247 (0.233)

Table 6: Pearson correlation between AAPC and city area. Confidence (p-value) in brackets, rounded to 3 significant figures.

The visual plot is provided in Figure 19, once more using the logarithm of the independent variable to compress the maxima of the scale horizontally. 10 and 5 minute plots placed in Appendix B.

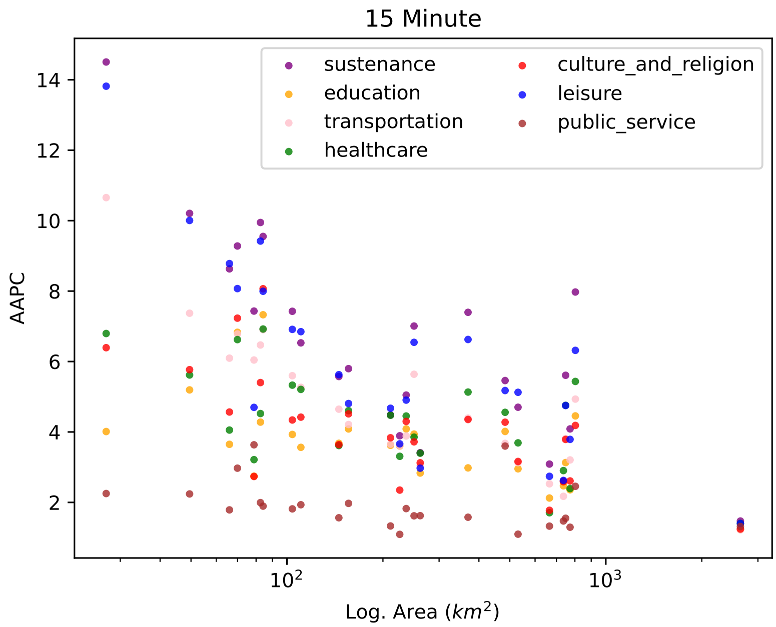

Figure 19: AAPC plotted against the logarithm of the area.

[21]Excluding `public_service` as its correlation coefficients are not statistically significant.
[22]Being a reoccurring observation, this could indicate that this category is poorly recorded within the OSM dataset.

This analysis shows that AAPC is often lower for larger cities. Again, this correlation is logical, as producing a walkable environment is harder when the area to cover is greater; centralising amenities no longer works once the residential sprawls beyond a certain point. Against area, `transportation` has the strongest correlation - with larger cities requiring longer-distance transport, so stops are less densely arranged - and `education` has the weakest. Interestingly, this is different to the order of categories discussed above for population, though again the difference between these extremes is less only ≈ 0.05 so is not considered notable.

6.5 Comparison With Restricted Emission Zones

Being similarly climate-driven, the 15-minute city concept is often faced with the same negative reception as policies like "congestion charging, carbon pricing, or low-emission zones" [12]. Specifically, the correlation of this new metric to existing restricted emission zones is inspected.

In the UK, there are a few different types of emission policy: Clear Air (CAZ), Low Emission (LEZ), Ultra Low Emission (ULEZ) and Zero Emission (ZEZ) Zone. The most strict is the ZEZ, exclusive to Oxford, applying a charge even for hybrid cars - though the physical area it encompasses is small, covering only a handful of streets in the city's centre[23]. London contains an ULEZ nested within a LEZ - both now covering the majority of the interior of the M25 - with differing charges. There are multiple classes of CAZs, from A (most lenient) to D (most restrictive) - where the Scottish LEZ is similar to a class D zone [15]. Studied zones - and associated cites - are summarised in Table 7; the only class A CAZ is situated in York, which is not covered in this work.

Zone	Cities
CAZ-B	Southampton[24]
CAZ-C	Bradford, Newcastle
CAZ-D	Aberdeen, Bristol, Edinburgh, London
ULEZ	London
ZEZ	Oxford

Table 7: Studied emission zones, by type [15].

AAPC appears not to have as strong a relation as expected to the cities where restricted emission zones are implemented: London and Bradford rank 25th and 23rd respectively; Oxford ranks very highly (see Section 6.2.1); other such cities rank around 8th-14th. Quite a few cities without these urban restrictions rank higher - e.g. Cambridge (2nd/3rd) and Reading (2nd/4th).

This may indicate that emissions policies are unpopular due to them being implemented without due regard to the current walkability of the area. Regions are chosen where the scheme would most significantly reduce air pollution (i.e. those with high private vehicle use), though not necessarily where they are well supported by the infrastructure. If an area's amenities are poorly accessible on foot - as is the case for London, Newcastle and Bradford according to this analysis - then introducing restrictions on private transport can feel limiting, and may compound existing inequalities. Residents may feel forced to use public transport, rather than it being an option, contributing to the overall lacklustre acceptance of CAZ/LEZs. Disproportionately affected are lower-income families [1], for which the financial aspect is non-trivial; those unable to afford a newer, low-emission vehicle would find the financial commitment to public transport unviable, hence limiting them to a restricted set of services.

However, a counter-argument could be made with Oxford. From the above analysis, it is the most successful implementation of a policy restricting emissions. Its Zero Emission Zone (ZEZ), coupled with an urban environment designed and optimised for pedestrian accessibility, serves as an example which might inspire more cities to develop sustainably. AAPC analysis could see applications as a prerequisite study to the implementation of these policies in the future, ensuring walkability and equitable access is prioritised, and laying the groundwork for strong public support.

[23] www.oxfordshire.gov.uk/transport-and-travel/oxford-zero-emission-zone-zez/view-map-zez

[24] Non-charging for private vehicles [15].

7 Evaluation

7.1 Results Summary

Of the case study cities, Margate scored highest overall, and London lowest. The difference in their total available “choice” is vast (with almost 6 times as many services accessible in Margate than London), despite prior studies characterising these regions all with the same “100% walkable” label. This is an striking new insight, as it demonstrates that despite the extensive and wide-spread infrastructure present in urban hubs - like London - this development translates into little tangible benefit for residents. Over-centralised amenities, all too common in sprawling urban areas, is a strong barrier to sustainable development in the UK.

Other sustainable policies (i.e. restricted emission zones) have had trouble gaining traction, and these results reveal a justification for this. London’s ULEZ appears to have been implemented before evaluating city walkability (see Section 6.5), which explains the disapproval associated with the scheme. On the contrary, Oxford ranks very highly, with an effectively implemented Zero Emission Zone (ZEZ). Cambridge and Reading rank highly as well - implying they are strongly walkable - despite having no CAZ, LEZ or equivalent implemented.

In terms of categories, sustenance and leisure both offered the greatest choice of walkable services for residents. However, public services - such as emergency services and facilities - do not correlate as well spatially with population, though this may be due to a general lack of such amenities rather than an inefficient layout.

These results accentuate a research vacuum with great potential for future analysis. Introducing more nuanced measurements, such as this AAPC metric for quantifying “choice”, can give better guidance for making improvements in urban developments. Not only this, but such metrics could be used to complement other sustainable development schemes - for example, emissions regulations, which are not currently being deployed as effectively as they could be.

7.2 Requirements

Functionally, the developed application meets all the goals of the project - a full pipeline from city names to metrics (including geocoding, which was initially an optional feature). However, the success on non-functional requirements was more mixed.

Visualisation On the successful side, the visualisation web app exhibited all the desired qualities: interactive, responsive, robust and informative. Interactivity was facilitated by integrating both mouse and keyboard inputs, streamlining navigation. Map responsiveness was primarily improved using isochrone bucketing (see Section 5.4), ensuring polygons would load fast enough not to drop inputs; though the pre-processing pipeline was crucial. In regard to robustness, inputs were blocked while the map was loading to prevent inconsistent behaviour, and thanks to the quick map loading, this only had an impact in practice when multiple keys were held down. Last of all, including dropdowns for selecting between a range of display options, implementing dynamic legends, and showing full metric graphs all contributed towards the overall visualisation, communicating key information to the user.

Data Processing Pipeline The processing pipeline did not meet all its goals. Accuracy was high, thanks to the resampling procedure (Section 5.3.1); by using basic nearest neighbour sampling, the GHSL dataset remained numerically unmodified, meaning no extra uncertainty from was introduced. A reasonable number of POIs were sampled for this, at the expense of execution speed, with most cities taking a few minutes per category, and the largest areas - i.e. Greater London - taking up to 15-20 minutes (on the tested hardware). Repeatability is impacted somewhat, as a slow pipeline is inconvenient to rerun after small modifications. Moreover, having a stochastic element (mentioned in Section 5.3.2) adds a slight variance to repeat runs.

7.3 Limitations and Improvements

Computation speed was the main limitation of this work. A parallel processing setup could improve upon the isochrone and metric calculations, as they are entirely independent between cities and categories. Although making the isochrone crossover algorithm (see Section 4.5) more efficient is difficult, the metric calculation is optimisable - aggregating the population within polygons could be threaded, resulting in considerable speed-up on the right hardware.

The second limitation is the sampling process. Over 14 thousands POIs were found for some categories in London, being reduced to less than half that amount by the implemented sampling method - evidently, a fair sampling method is paramount for dependable results. Some areas might benefit from a spatially-stratified approach, where the chosen sample attempts to match the physical distribution of all POIs as closely as possible. For instance, each POI could be assigned a probability according to its proximity to others (i.e. selection probability inversely proportional to distance to closest neighbour). Another approach - easily integrable in conjunction with the above - is stratifying the sample by feature tag. This preserves the internal distribution of categories, preventing less common tags from being over-represented in the sample. Such improvements would increase the objectivity of results, and give more trustworthy conclusions.

Walking speed is assumed constant in ORS isochrone calculations. Instead, a custom routing engine could be purpose-built to harness terrain data (elevation and gradient), with a variable speed according to if sections of the route are uphill or downhill. This would be a more resource-intensive calculation, extending the runtime of the pipeline even further, but would model terrain differences - most noticeable for pedestrians - accurately. OSM allows users to upload GPS traces[25], which could be queried, finding the closest routes in order to approximate route segments, and produce a speed map. These traces can be either public (unordered) or identifiable (ordered points, with time stamps); due to isochrones being calculated on times rather than distance, only the latter is suitable. Additionally, it would require the tracks to be tagged by method of transport[26] in order estimations to be valid.

7.4 Practical Applications

The practical applications of this project are shared with many of the works discussed in Section 2.1. Integrating into property evaluation websites - like walkscore.com - by providing per-category scores is one such application. Some categories might be more important to a user than others, so areas could be sorted by relevant scores.

Another example use-case is urban planning - like NEXI or Proximity Time - where the analysis is invaluable for prioritising infrastructure investment, both city and category-wise, and driving sustainable policy integration towards SDG-11. Hypothetical scenarios can be proposed and tested - be that projected population distributions, introducing or rearranging amenities (like in [2]), etc. Because AAPC is normalised across population, inter-city comparison is also facilitated well - critical for planning at the national level.

Being a lightweight web-app, this analysis can be easily shared by hosting as a stand-alone website, like www.cityaccessmap.com [16], offering visual and interactive access to the data for other researches alongside the public. In addition, with a global archive of OSM street data, this analysis could be extended to any desired research area - Europe being a good example.

Interdisciplinary research carries promise too. Integrating these results with other socio-economic data, like active transportation[27] uptake [23] or social demographics [16], could help to further highlight lifestyle or equality issues raised by other works. Alternatively, modifications to this methodology could be made (e.g. changing weights or introducing thresholds) to better suit the nature of the research being undertaken.

[25] wiki.openstreetmap.org/wiki/GPS_tracks

[26] Alternatively, identifiable traces could be categorised based on the maximum, minimum, and average journey speed, requiring these values to fall within an acceptable range for a pedestrian.

[27] Only 8 of the 25 studied cities are contained in github.com/rafaelprietocuriel/ModalShare/blob/main/ModalShare.csv [23], so further data collection is required.

7.5 Critical Reflection

Overall, this study was successful in quantifying the concept of "choice", and introduces a new point of comparison for well-developed cities - beyond existing metrics - and advancing the urban data science field. It serves as another example of how open-source data can be used to derive reproducible insight into urban environments, guiding future policies and encouraging continued progress towards the UN Sustainable Development Goals [28].

Nevertheless, there were some limitations. A sample of POI locations had to be taken due to the complexity of metric calculation, and this will have influenced the results - from the function determining sample size to the sampling method itself, both of which could be improved. Metric values are only claimed to be correct up to the given resolution of the available population data - 3 arc seconds - and even though the data set was resampled, this required assumptions that will have impacted accuracy of results. This work also relies on the consistency in feature tagging (with regard to both qualitative tags and location precision) of OSM, a dataset of Volunteered Geographic Information. Finally, this project would be strengthened by a bespoke routing engine, taking terrain into account when calculating isochrones.

8 Conclusion

In essence, Moreno's model [13, 14] was intended to design settlements to benefit everyone - both socially, promoting health, happiness and equitable access, and environmentally, by decarbonising transport and cutting air pollution. So far, creating practical implementations of this vision has been a challenge, with many urban areas lacking pedestrian infrastructure, and demonstrating over-reliance on private transport.

The goal of this research was to assess walkability across UK cities in more detail than the prior studies listed in Section 2.1, to aid sector-targeted development. A pipeline for calculating the "choice" available, within 3 time intervals, for different categories of amenity is presented, following an open-source methodology (i.e. using freely-available data and APIs). 25 UK cities are analysed through the lens of this novel metric, with London being most notably lacking in number of available amenities. Margate scored highly in specific categories, but this may have been an outlier due to its small population and area. Therefore, this work concludes that Oxford (closely followed by Cambridge and Reading) offers the best range of walkable amenities to its residents. While this work limits its scope to the UK, any region within the bounds of the two core datasets (OSM and GHSL) can theoretically have its AAPC score computed using the application.

Using Principal Component Analysis, 4 city clusters were identified, based on patterns emerging when comparing amenity and population distributions. Certain common characteristics - including location and level of urban sprawl - were grouped together, translating to either high or low rankings, and consistency or inconsistency across categories.

Emissions regulations are shown to be uncorrelated with AAPC, given both the highest and lowest ranking cities implement some variation of this policy - ULEZ in London and ZEZ in Oxford. This seems to indicate their introduction was not well informed by any pre-assessment of walkability - rationalising the scepticism associated with them.

It was also shown that sustenance and leisure services were best distributed across the board, requiring little to no improvement, while public services seemed to be significantly less accessible on average. What this shows is that commercial services - built around competition - naturally decentralise; on the contrary, public services require more planning and intention.

Other applications for this methodology include commercial and public-sector uses - such as property evaluation, or guiding sustainable urban planning. Academic investigations across various disciplines could also incorporate the presented results. By successfully quantifying choice in relation to pedestrian walkability, this study has paved the way for quality-of-life improvements in urban areas throughout the developed world.

Future Work Improvements could be made by increasing the number of samples taken in the data processing step, or by implementing different sampling methods (e.g. stratifying spatially or by feature tag). Another avenue for future research would be exploring weight functions in metric calculation. For example, taking the square root of intersection count - as each new amenity instance in range is of diminishing significance - or applying negative weights to values below some preset “acceptable” threshold.

Beyond accuracy, subjective research could be integrated, such as resident sentiment or perception of accessibly (like in [26], through a Walkability app), as whether or not they correlate with AAPC would be interesting to explore. Given worldwide OSM and GHSL data remains publicly available in the future, this work could be easily extended to a global scale.

References

[1] Andriolli, A. C., and Silva, L. T. Are Low Emission Zones Truly Embraced by the Public? *Environments 11*, 6 (2024), 106.

[2] Bruno, M., Monteiro Melo, H. P., Campanelli, B., and Loreto, V. A universal framework for inclusive 15-minute cities. *Nature Cities* (2024), 1–9.

[3] Camargo, C. Q., Bright, J., and Hale, S. A. Diagnosing the performance of human mobility models at small spatial scales using volunteered geographic information. *arXiv preprint arXiv:1905.07964* (2019).

[4] Caprotti, F., Duarte, C., and Joss, S. The 15-minute city as paranoid urbanism: Ten critical reflections. *Cities 155* (2024), 105497.

[5] Carr, L. J., Dunsiger, S. I., and Marcus, B. H. Walk score™ as a global estimate of neighborhood walkability. *American journal of preventive medicine 39*, 5 (2010), 460–463.

[6] Graser, A. Movingpandas: efficient structures for movement data in python. *GIForum 1* (2019), 54–68.

[7] Haklay, M. How good is volunteered geographical information? A comparative study of OpenStreetMap and Ordnance Survey datasets. *Environment and planning B: Planning and design 37*, 4 (2010), 682–703.

[8] Jovanović, D. From Qgis To Python: Comparison of Free and Open Tools for Statistical Analysis of Cultural Heritage and Data Representation. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences 48* (2022), 229–236.

[9] Khavarian-Garmsir, A. R., Sharifi, A., and Sadeghi, A. The 15-minute city: Urban planning and design efforts toward creating sustainable neighborhoods. *Cities 132* (2023), 104101.

[10] Kim, E. J., Kim, J., and Kim, H. Neighborhood walkability and active transportation: a correlation study in leisure and shopping purposes. *International journal of environmental research and public health 17*, 7 (2020), 2178.

[11] Logan, T., Hobbs, M., Conrow, L., Reid, N., Young, R., and Anderson, M. The x-minute city: Measuring the 10, 15, 20-minute city and an evaluation of its use for sustainable urban design. *Cities 131* (2022), 103924.

[12] Marquet, O., Mojica, L., Fernández-Núñez, M.-B., and Maciejewska, M. Pathways to 15-minute city adoption: Can our understanding of climate policies' acceptability explain the backlash towards x-minute city programs? *Cities 148* (2024), 104878.

[13] Moreno, C. La ville du quart d'heure: pour un nouveau chrono-urbanisme!

[14] Moreno, C. The 15-minute city, 2020. `https://www.ted.com/talks/carlos_moreno_the_15_minute_city` [Accessed: 17/10/2024].

[15] Motorway Online Ltd. `https://motorway.co.uk/sell-my-car/guides/uk-clean-air-zones` [Accessed: 30/10/2024].

[16] Nicoletti, L., Sirenko, M., and Verma, T. Disadvantaged communities have lower access to urban infrastructure. *Environment and Planning B: Urban Analytics and City Science 50*, 3 (2023), 831–849.

[17] Nöllenburg, M. *Human-Centered Visualization Environments: GI-Dagstuhl Research Seminar: Revised Lectures, Dagstuhl Castle, Germany.* Springer, 2006, ch. Geographic Visualization, pp. 257–294.

[18] Observable. D3, 2024. https://d3js.org [Accessed: 16/10/2024].

[19] Olivari, B., Cipriano, P., Napolitano, M., and Giovannini, L. Are Italian cities already 15-minute? Presenting the Next Proximity Index: A novel and scalable way to measure it, based on open data. *Journal of Urban Mobility 4* (2023), 100057.

[20] OpenStreetMap contributors. Planet dump retrieved from https://planet.osm.org , 2017. https://www.openstreetmap.org [Accessed: 02/10/2024].

[21] Ordnance Survey Limited, 2024. https://www.ordnancesurvey.co.uk/about/what-we-do [Accessed: 23/11/2024].

[22] Prieto-Curiel, R., González Ramírez, H., and Bishop, S. A ubiquitous collective tragedy in transport. *Frontiers in Physics 10* (2022), 882371.

[23] Prieto-Curiel, R., and Ospina, J. P. The ABC of mobility. *Environment International 185* (2024), 108541.

[24] Scalas, M. Pedestrian Isochrones Facilities Overlapping with Openrouteservice. An Easy, Fast and Opensource Indicator in Novara, Italy. In *International Conference on Computational Science and Its Applications* (2023), Springer, pp. 293–307.

[25] Schiavina, M., Freire, S., Carioli, A., and MacManus, K. GHS-POP R2023A - GHS population grid multitemporal (1975-2030). European Commission, Joint Research Centre (JRC), 2023. https://data.jrc.ec.europa.eu/dataset/2ff68a52-5b5b-4a22-8f40-c41da8332cfe [Accessed: 22/01/2025].

[26] transform transport. UX Mobility: Multi-User Walkability Route Planner , 2025. ; https://transformtransport.org/research/inclusive-mobility/ux-mobility-multi-user-walkability-route-planner/ [Accessed: 22/04/2025].

[27] United Nations. COP 21, 2015. https://unfccc.int/event/cop-21 [Accessed: 17/10/2024].

[28] United Nations. Transforming our World: The 2030 Agenda for Sustainable Development, 2015. https://sdgs.un.org/2030agenda [Accessed: 27/03/2025].

A Exact Metric Tables

All values rounded to 2 decimal places. Totals are calculated city-wise and category-wise.

City	sustenance	education	transportation	healthcare	culture_and_religion	leisure	public_service	Totals
London	1.471	1.232	1.221	1.401	1.244	1.402	1.315	9.29
Oxford	9.551	7.328	6.904	6.923	8.067	7.995	1.894	48.66
Bristol	5.047	4.087	3.886	4.453	4.299	4.905	1.825	28.5
Southampton	7.426	3.929	5.597	5.33	4.34	6.913	1.816	35.35
Newcastle upon Tyne	4.481	3.621	3.654	4.475	3.835	4.673	1.331	26.07
Bradford	2.605	2.471	2.173	2.903	2.587	2.624	1.473	16.84
Edinburgh	5.46	4.014	3.677	4.558	4.277	5.177	3.6	30.76
Aberdeen	7.397	2.98	4.393	5.133	4.354	6.626	1.579	32.46
Exeter	9.948	4.279	6.469	4.523	5.402	9.421	1.994	42.04
Bournemouth	6.53	3.563	5.268	5.208	4.421	6.848	1.933	33.77
Brighton	7.974	4.456	4.933	5.435	4.186	6.317	2.457	35.76
Margate	14.505	4.012	10.654	6.795	6.393	13.818	2.252	58.43
Maidstone	3.088	2.124	2.527	1.707	1.776	2.741	1.326	15.29
Crawley	7.433	2.745	6.043	3.214	2.736	4.698	3.635	30.5
Reading	10.206	5.194	7.371	5.618	5.769	10.005	2.24	46.4
Milton Keynes	3.892	3.602	3.61	3.309	2.349	3.663	1.092	21.52
Cambridge	9.28	6.83	6.783	6.622	7.232	8.071	2.972	47.79
Colchester	4.702	2.954	3.127	3.691	3.162	5.126	1.098	23.86
Ipswich	8.631	3.649	6.096	4.054	4.566	8.781	1.787	37.56
Peterborough	4.087	2.357	3.206	2.393	2.609	3.789	1.295	19.74
Coventry	5.575	3.679	4.647	3.615	3.64	5.627	1.564	28.35
Nottingham	5.797	4.086	4.211	4.604	4.514	4.808	1.973	29.99
Manchester	3.412	2.834	3.047	3.397	3.125	2.972	1.623	20.41
York	5.609	3.131	3.812	4.755	3.788	4.751	1.55	27.4
Preston	7.008	3.94	5.639	3.853	3.72	6.545	1.619	32.32
Totals	161.11	93.1	118.95	107.97	102.39	148.29	47.24	779.05

Table 8: AAPC: 15 minute

City	sustenance	education	transportation	healthcare	culture_and_religion	leisure	public_service	Totals
London	0.69	0.556	0.558	0.656	0.572	0.659	0.614	4.3
Oxford	4.956	3.706	3.444	3.563	4.17	4.049	0.979	24.87
Bristol	2.398	1.955	1.874	2.195	2.138	2.463	0.918	13.94
Southampton	3.461	1.948	2.749	2.613	2.126	3.353	0.831	17.08
Newcastle upon Tyne	2.202	1.762	1.789	2.287	1.913	2.291	0.667	12.91
Bradford	1.251	1.176	1.069	1.435	1.302	1.286	0.765	8.28
Edinburgh	2.458	1.92	1.784	2.274	1.979	2.552	1.668	14.63
Aberdeen	3.475	1.52	2.179	2.518	2.048	3.183	0.704	15.63
Exeter	4.379	2.093	3.227	2.222	2.451	4.306	0.975	19.65
Bournemouth	3.362	1.722	2.667	2.766	2.235	3.463	0.956	17.17
Brighton	3.78	2.235	2.458	2.721	2.032	3.082	1.213	17.52
Margate	7.182	2.011	5.524	3.404	3.213	7.138	1.163	29.64
Maidstone	1.589	1.08	1.328	0.892	0.911	1.414	0.664	7.88
Crawley	3.588	1.244	2.986	1.708	1.424	2.153	1.972	15.07
Reading	4.878	2.552	3.498	2.849	2.761	4.787	1.004	22.33
Milton Keynes	1.896	1.817	1.817	1.658	1.129	1.783	0.54	10.64
Cambridge	4.546	3.292	3.277	3.31	3.562	4.047	1.335	23.37
Colchester	2.196	1.503	1.629	1.825	1.502	2.44	0.584	11.68
Ipswich	3.984	1.79	3.019	1.926	2.167	3.829	0.779	17.49
Peterborough	2.027	1.244	1.59	1.148	1.294	1.781	0.637	9.72
Coventry	2.673	1.746	2.244	1.754	1.722	2.635	0.731	13.5
Nottingham	2.723	1.939	2.042	2.259	2.173	2.34	0.93	14.41
Manchester	1.593	1.35	1.478	1.625	1.507	1.361	0.752	9.67
York	2.486	1.521	1.849	2.466	1.773	2.192	0.718	13.0
Preston	3.423	1.982	2.51	1.928	1.783	3.238	0.768	15.63
Totals	77.2	45.66	58.59	54.0	49.88	71.82	22.87	380.03

Table 9: AAPC: 10 minute

City	sustenance	education	transportation	healthcare	culture_and_religion	leisure	public_service	Totals
London	0.171	0.12	0.125	0.162	0.131	0.155	0.144	1.01
Oxford	1.473	0.959	0.992	0.946	1.234	1.116	0.279	7.0
Bristol	0.648	0.504	0.502	0.598	0.577	0.649	0.233	3.71
Southampton	0.89	0.462	0.753	0.658	0.57	0.83	0.192	4.35
Newcastle upon Tyne	0.609	0.425	0.48	0.6	0.523	0.64	0.156	3.43
Bradford	0.326	0.282	0.294	0.35	0.337	0.311	0.2	2.1
Edinburgh	0.673	0.487	0.493	0.636	0.524	0.735	0.46	4.01
Aberdeen	0.816	0.372	0.639	0.657	0.511	0.828	0.153	3.98
Exeter	1.084	0.466	0.928	0.623	0.63	1.067	0.246	5.04
Bournemouth	0.911	0.417	0.759	0.698	0.59	0.935	0.255	4.57
Brighton	1.018	0.563	0.733	0.741	0.553	0.78	0.314	4.7
Margate	2.199	0.43	1.779	1.022	0.963	2.012	0.286	8.69
Maidstone	0.429	0.24	0.388	0.221	0.235	0.381	0.168	2.06
Crawley	0.923	0.252	0.91	0.426	0.344	0.477	0.563	3.89
Reading	1.323	0.636	0.972	0.807	0.8	1.253	0.253	6.04
Milton Keynes	0.524	0.464	0.476	0.415	0.29	0.48	0.129	2.78
Cambridge	1.229	0.898	0.918	0.92	0.975	1.187	0.323	6.45
Colchester	0.523	0.385	0.474	0.485	0.383	0.602	0.165	3.02
Ipswich	1.036	0.436	0.848	0.5	0.556	0.952	0.174	4.5
Peterborough	0.584	0.28	0.442	0.287	0.348	0.433	0.163	2.54
Coventry	0.726	0.423	0.624	0.464	0.443	0.682	0.18	3.54
Nottingham	0.689	0.47	0.547	0.569	0.58	0.581	0.225	3.66
Manchester	0.404	0.313	0.385	0.395	0.391	0.337	0.175	2.4
York	0.628	0.39	0.526	0.658	0.441	0.586	0.187	3.42
Preston	0.944	0.533	0.59	0.524	0.44	0.829	0.191	4.05
Totals	20.78	11.21	16.58	14.36	13.37	18.84	5.81	100.94

Table 10: AAPC: 5 minute

B Omitted Figures

Extra figures omitted from the main body are included here for completeness. Mainly consists of 10 and 5 minute plots, which were similar to the 15 minute plot.

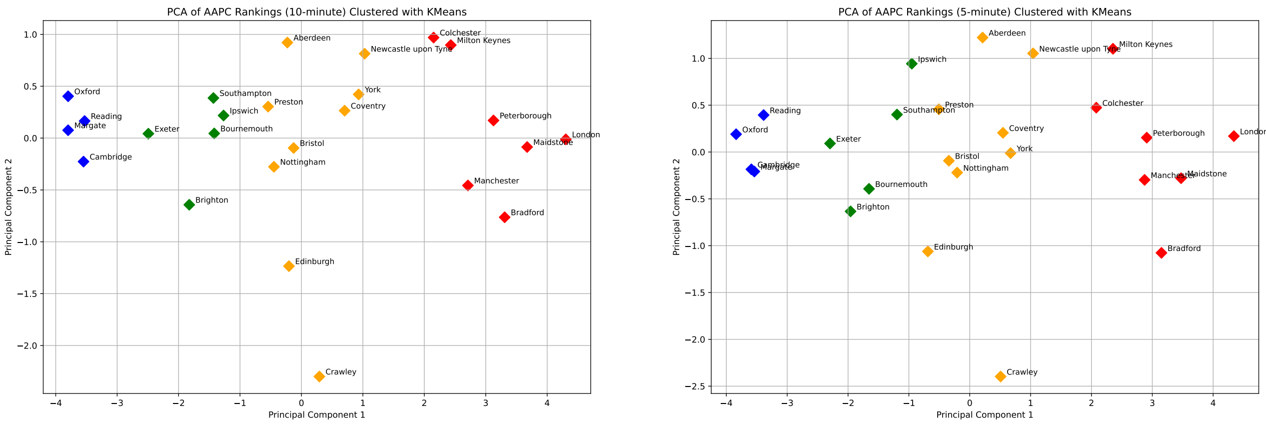

Figure 20: PCA on 10 and 5 minute ranks.

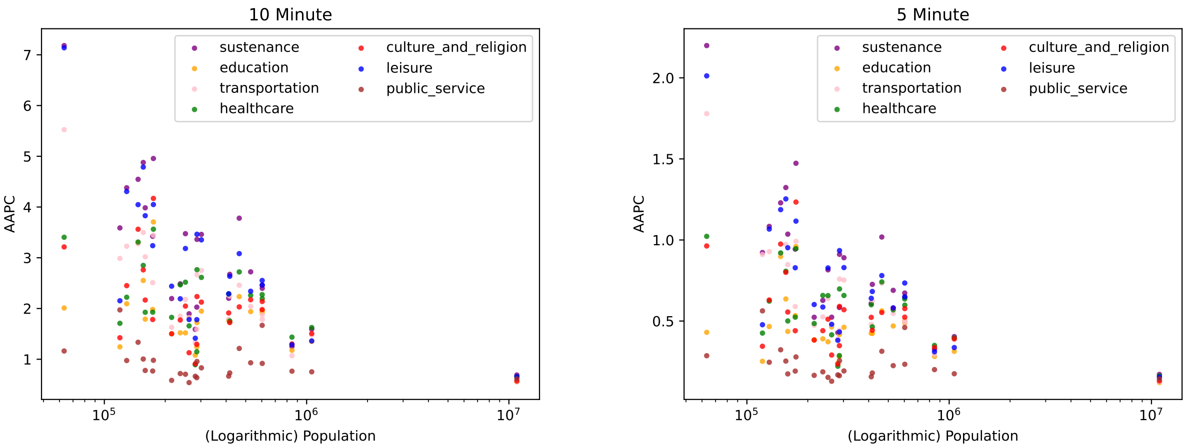

Figure 21: AAPC plotted against the logarithm of the population.

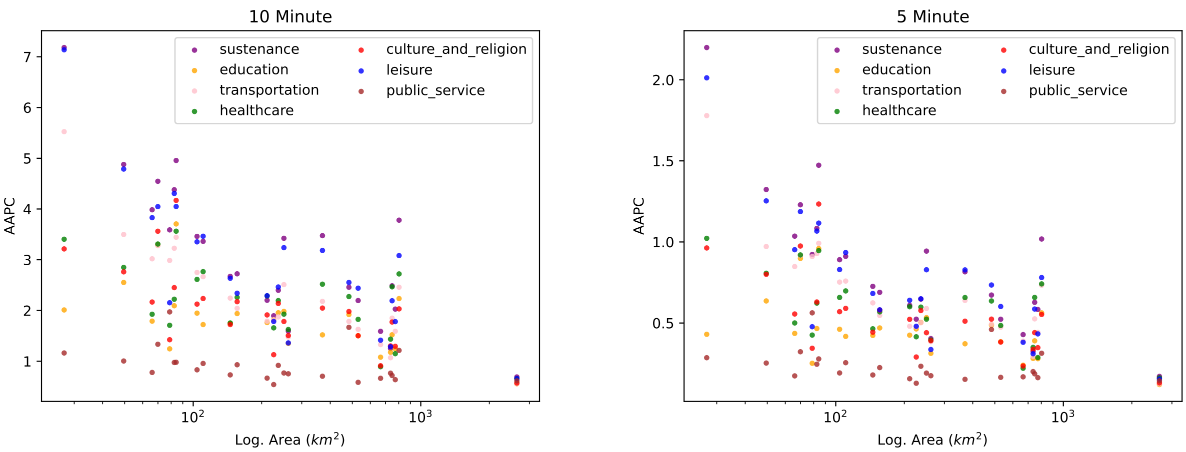

Figure 22: AAPC plotted against the logarithm of the area.

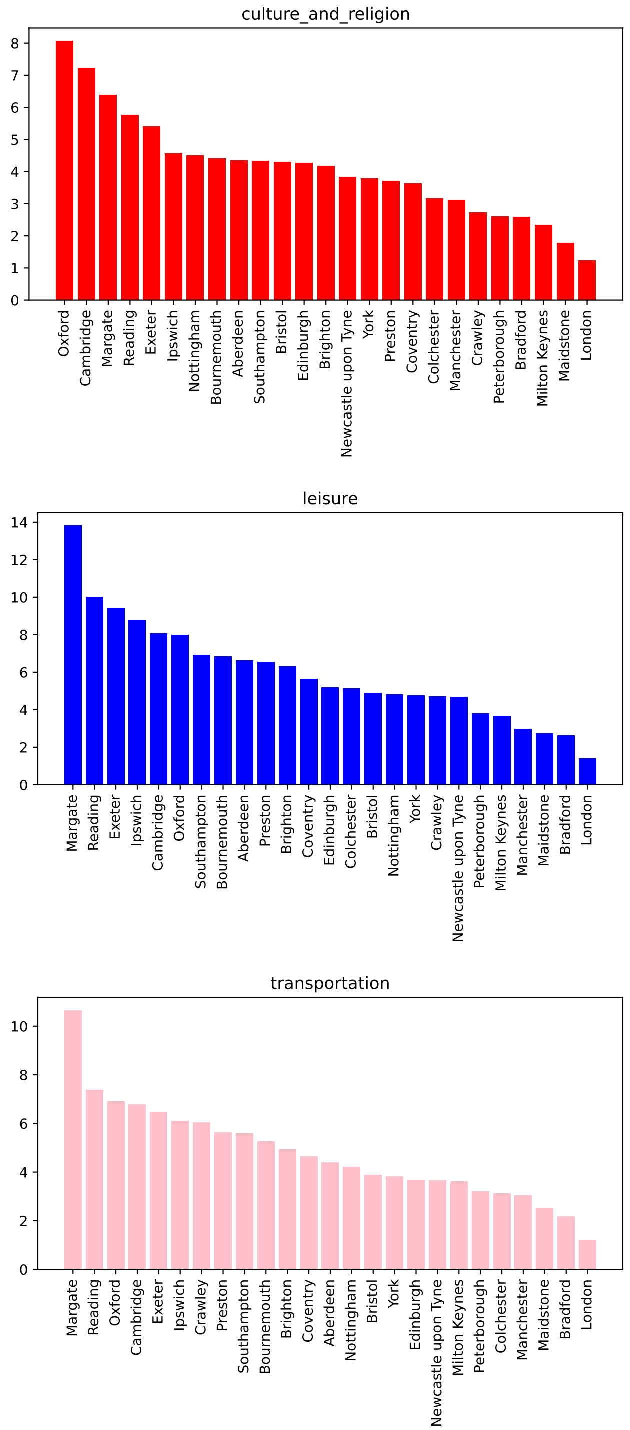

Figure 23: Average metric values per city, sorted, for 15 minute interval.