Knowledge-based generation of plausible air quality maps in the absence of sensor data

Duarte Vital, Instituto Universitário de Lisboa (ISCTE-IUL) (duarte_vital@iscte-iul.pt)

Pedro Mariano, Information Sciences and Technologies and Architecture Research Center (ISTAR-IUL), Instituto Universitário de Lisboa (ISCTE-IUL) (plsmo@iscte-iul.pt)

Susana Marta Almeida, Centro de Ciências e Tecnologias Nucleares, Instituto Superior Técnico (smarta@ctn.tecnico.ulisboa.pt)

Pedro Santana, Information Sciences, Technologies and Architecture Research Center (ISTAR-IUL), Instituto Universitário de Lisboa (ISCTE-IUL) (pedro.santana@iscte-iul.pt)

Abstract

Industrialization increased air pollution sources, which is a cause of major health problems. As such, air pollution became a growing concern and there is a need to monitor and easily visualize air pollution data. There are thousands of air quality monitoring stations throughout the world that are used to measure air quality. Moreover, there are plenty of applications that have been developed to visualize air pollution that use information gathered by these air quality monitoring stations as well as other sources of information, such as traffic intensity or weather forecasts. This paper introduces a novel graphical tool that taps on a new source of information: expert knowledge of air pollution sources. This tool allows experts to represent air pollution sources and their dynamics, and to assign them to different map elements (e.g., buildings, roads). We have performed tool's usability and viability tests with 30 participants of which 6 are from environmental sciences. The obtained results and the provided feedback show that the proposed approach is a promising complement to sensor-based mapping approaches.

1 Introduction

Air pollution is a serious health issue and is of great concern for Europeans as reported by the European Commission (2017). This has led to various efforts to monitor air quality. Historically this has been done with static stations equipped with complex and costly equipment. Additionally, these stations are usually located in cities where most population live and work. More recently, the advent of low-cost sensors (LCS) has given rise to networks of these sensors massively increasing the coverage of monitored areas, as demonstrated by the work of Becnel et al. (2019), Hasenfratz et al. (2015), and Lin et al. (2020). Despite lower accuracy and resolution, Morawska et al. (2018) have shown that low-cost sensors complement the more complex monitoring stations. Hasenfratz et al. (2015) have also shown that the data collected by these low-cost sensor networks can be used to construct highly detailed pollution maps. However, despite the cheaper deployment cost, there are still major population centers that are not monitored.

There are several tools and programming libraries, such as Breezometer (2020), IQAir (2021), PlumeLabs (2020), which can be used for air pollution estimation and visualization. In general, these libraries use data collected by air quality monitoring stations in their estimation models. In areas not covered by these stations, it is difficult to obtain an accurate air pollution map, given that estimation models that are not based on sensory information are rather incomplete. Lakhani et al. (2010), reviews different monitoring methods, all of them requiring *in situ* measurements of the target pollutant.

To improve the representation of air pollution maps, we propose the representation of knowledge-based pollution maps, where an environmental expert can represent pollution sources. To realize this, we developed a novel graphical tool that taps on exploiting explicit expert knowledge for solving the task at hand. Through this tool, experts can graphically and, thus, intuitively represent dynamic air pollution profiles. A profile can depend on weekday, hour, or any other parameter that the expert finds suitable (e.g., wind orientation). The expert can then associate a pollution profile to different map elements (e.g., building, road), thus specifying which map elements are the actual pollution sources of the associated pollution profile. Once the expert assigns a set of air pollution values to the pollution source profile's parameters, users can visualize the corresponding predicted pollution diffusion pattern through a heat map.

The design of our tool was influenced by existing map-based applications. For instance, in Google Maps users can select two types of maps: one that displays polygons and polylines representing different features such as buildings or roads among other; a second that displays satellite imagery. On top of these two information layers, users can view additional layers (e.g., traffic, terrain elevation). Another example, Open Street Map, is a programming library that provides the user with different pictorial representations. The selection of a particular library to develop our graphical tool was based on the associated learning curve and map data availability.

This article is an extended version of a previous conference paper (Vital et al., 2021). The focus of this article is to introduce the concept of knowledge-based generation of pollution maps, to present a graphical tool and to assess its usability to represent expert knowledge. To achieve this goal, we performed a set of usability tests with 30 participants, of which 3 were environmental experts and 3 were environmental sciences master students. Participants performed a set of tasks and filled a usability questionnaire.

The work presented in this article is part of the ExpoLIS project. The goal of this project is to develop a network of mobile LCS to monitor air quality and to implement a set of graphical tools to present the gathered sensory data to users (Santana et al., 2021). The LCS will be installed on the top of city buses and will supply real-time air quality data. As mentioned, user-centered graphical tools will use these data to provide up-to-date air quality information for environmental experts and to the public, thus increasing the population's awareness to air pollution. While the tool presented in this article is an example of the first type, Santana et al (2021) have implemented a mobile app to display air quality to the public and Teles et al. (2020) have developed a 3D game-like virtual environment to visualize air pollution data in an engaging way, targeting youngsters.

This article is organized as follows. We start with a survey of related work in Section 2. Then in section 3, we present the developed tool. The testing phase is detailed, and the obtained results are presented and analyzed subsequently in Section 4. We finish we the discussion, conclusions and future work directions in section 5.

2 Related Work

There is not a universal air pollution measuring index as different countries or applications consider different sets of pollutants, time periods, and concentrations. One of the most used in scientific studies, such as in Bumm (2019), Karuppasamy et al. (2020), Kumar and Goyal (2011), or Selvam et al. (2020), is the American version of the Air Quality Index (AQI). This index uses five major pollutants: ground-level ozone, particulate matter, carbon monoxide, sulfur dioxide, and nitrogen dioxide. It consists of six levels, as presented in Table 1.

health concern level	value range	colour
Good	0-50	
Moderate	51-100	
Unhealthy for sensitive groups	101-150	
Unhealthy	151-200	
Very unhealthy	201-300	
Hazardous	301-500	

Table 1: Meaning, range of values and color of Air Quality Index levels.

There are two main types of air pollution maps. The first type only displays data from air quality monitoring stations. The World Air Quality Index is such an example. It collects data from more than 30000 monitoring stations around 2000 major cities. It combines the data to compute an AQI. Since this map depends on high-cost stations, it does not provide information with a city block resolution and there are many regions in Earth that are note covered.

The second type of air pollution maps uses data from varying sources, such as monitoring stations, weather, and traffic information. These data are processed using proprietary algorithms, machine learning techniques, big data analytic, and air pollution dispersion modeling to produce heat maps. An example of this kind of map is Breezometer (2020), which claims a city block level resolution air pollution accuracy of 90%. However, the crux of this modeling is the availability of data form high-cost stations. In the absence of these, Breezometer is unable to provide forecasts, as is the example of Cuba, despite presenting weather and traffic information.

Recently, we have witnessed an increase of using LCSs to measure air pollution of risen in popularity (Morawska et al., 2018). PlumeLabs is a company that sells a LCS called Flow. It is considerably cheap when compared to a typical monitoring station. It collects data on particulate matter with diameter less than 1µm (PM1), less than 2.5µm (PM2.5), or less than 10µm (PM10), nitrogen dioxide (NO_2), and volatile organic compounds every 60 seconds. These data are sent to PlumeLabs's databases, which are then processed to create AQI maps with street level resolution. Although Flow is deployed in hundreds of major cities, it still lacks global coverage.

One of the drawbacks of current air pollution maps produced with data obtained from modeling techniques is the lack of individual pollution emission sources. If we use the Breezometer map to examine the Bełchatów power station in Poland, it looks that it has good air quality, despite being one of the most polluting factories in Europe, as reported by the European Environmental Agency (2017). Without sensors installed *in situ*, current models cannot accurately infer the existence of pollution in every situation. Either we increase the number of deployed sensors, or mathematical models have to take into account every single pollution source. The first solution has physical costs, in terms of sensor hardware, while the second has increased spatial and time computational complexity costs. This paper addresses these limitations as it allows an environmental expert to express one's knowledge on pollution sources. This knowledge can then be used to generate air quality maps in locations where there are no installed sensors.

There are several programming libraries that can be used to develop an air pollution mapping application. As selection criteria we chose: access to multiple data providers; intuitive user interface; fast development; cross-platform deployment; and customization ability. We have reviewed MapBox GL JS, OpenLayers, and Leaflet. All of them rely on Open Street Map (OSM) as data provider. OSM is an open-source geographical dataset, to which anyone can make contributions (Haklay and Weber, 2008). Maps can be rendered using raster or vector tiles. Every time a user selects a new map resolution or style, the application must download a raster tile set which impacts interaction. Moreover, anytime a change is done in the geographical dataset, all raster tile sets must be recomputed. Vector tile sets do not suffer these problems: the same geographical dataset can be used for any resolution and style combination. Leaflet requires a plug-in to work with vector tiles, but it still needs a larger network bandwidth. MapBox GL has better documentation compared to OpenLayers. In terms of performance, MapBox GL is slightly better according to Netek et al. (2020).

An important feature of the graphical tool that we developed is the ability to assign pollution sources to map entities. This is a process similar to image annotation. This depends on selecting and identifying key areas in an image. Image annotation is used in supervised machine learning (Weston et al., 2011). LabelMe is an example of a web-based application that annotates images (Russell et al., 2008). The availability of vast quantities of image data has led to research on the topic of semi-automatic segmentation. The result of this effort can be seen in popular applications features, such as Microsoft's Paint3D or Adobe Photoshop's MagicWand. A basis for semi-automatic segmentation is the graph-cut optimization algorithm by Boykov and Jolly (2001). This technique has been extended by Rother et al. (2004), by minimizing the amount of interaction that the user must perform. Despite advances in automatic image segmentation, there are still edge cases that are not recognized. Consequently, there are a couple of examples in OSM data where there is no map entity for specific satellite images.

3 Tool Description

Given the state of art reviewed in the previous section, we have developed a graphical tool with the following features:

1. It enables environmental experts to represent their knowledge of air pollution sources;
2. It generates dynamic air pollution maps that can depend on any expert defined parameter;
3. It allows the user to create map entities whenever the corresponding OSM data is absent.

We have developed the tool in node.js. Figure 1 shows the class diagram focused on the above features. The source code of the graphical tool is freely available at http://github.com/ExpoLIS-project/pollution-maps.

3.1 Map Entities

The tool assumes that maps are composed of objects that can represent roads, buildings, or green areas. These maps are automatically created by the image segmentation feature that OSM uses when processing satellite images. Then, users can add or edit the automatically created map entities in the OSM databases.

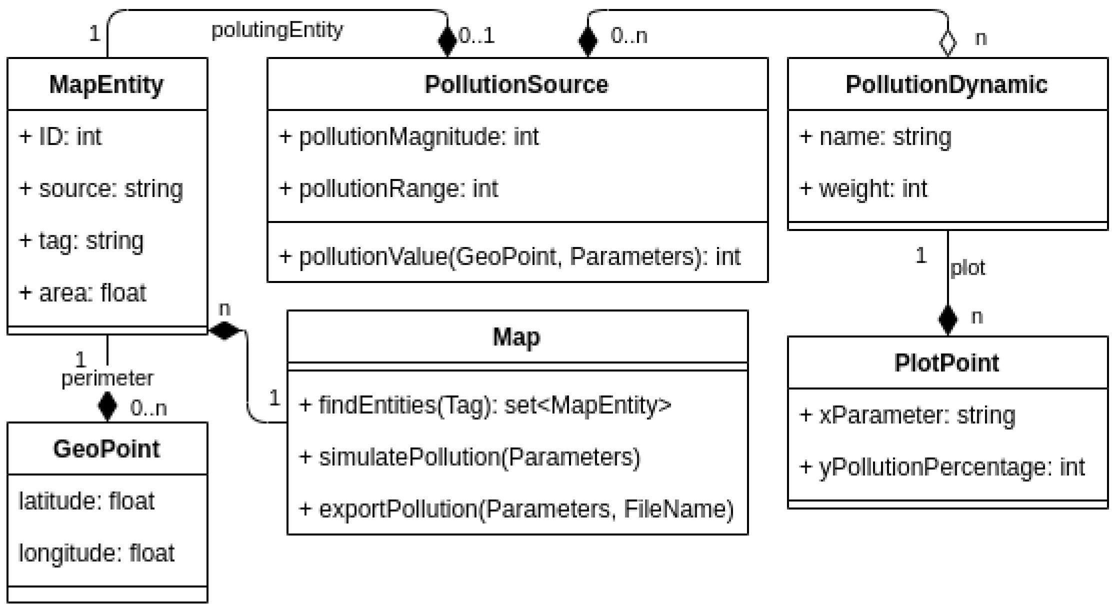

Figure 1: Class diagram illustrating the concepts and features of the graphical tool.

For each OSM object that our tool imports, a map entity is created and represented by the class `MapEntity`, as depicted in Figure 1. OSM tags, such as *building*, and *natural*, as well as OSM relations, such as *line*, *point*, *polygons*, and *roads*, are used by our tool to classify each entity as a building, road, or vegetation. The tool samples the perimeter (outline) of every OSM object to obtain the locations where virtual pollution sources are placed. These virtual pollution sources are then used to compute the pollution map over all created map entities. Map entities can be searched by type or by tag. This can be used by the user to associate a pollution source profile to multiple map entities.

Figure 2 shows a screenshot of the graphical tool. On the left side of the image, a selected map entity is highlighted, and on the right side of the screenshot there is a panel with properties that correspond to the entity's information imported from OSM. Map entities are colored in blue, magenta or green if their type is building, road or vegetation, respectively.

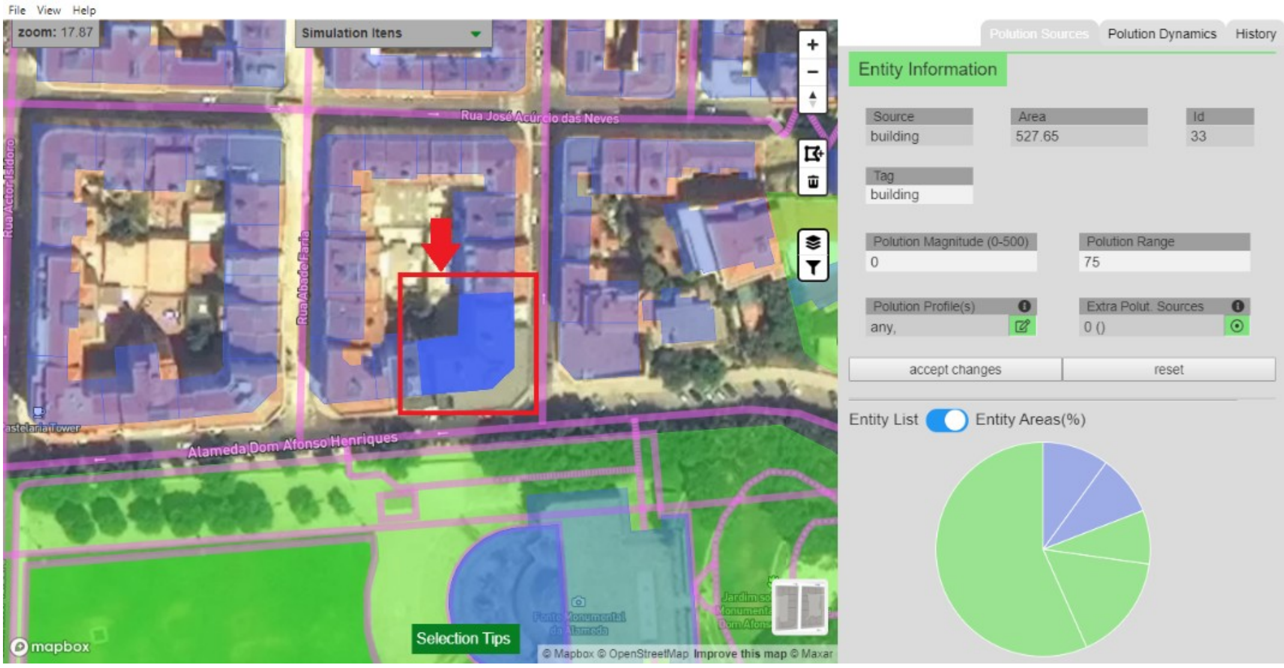

Figure 2: Screenshot of the graphical tool showing a selected map entity (red square) with its properties displayed on the right side.

3.2 Pollution Heat Map

Pollution sources are represented by a function that decays exponentially and radially, outwards the location of the pollution source. A heat map is used to represent pollution levels. We use the AQI scale and colors as shown on Table 1. An example of a point source pollution is displayed in Figure 3.

Figure 3: Single point pollution source with 250 AQI units.

Objects in OSM are represented by either polygons or polylines. These are represented in our graphical tool as the perimeter of map entities. Class `GeoPoint` and relation perimeter in Figure 1 represent this concept. To create a pollution source, we sample several points along this perimeter. Each point will be the center of a Gaussian function:

$$p_p(\boldsymbol{g}, \boldsymbol{a}, r, p) = p \cdot \exp\left(\frac{\|\boldsymbol{g}-\boldsymbol{a}\|^2}{r^2}\right), \qquad (1)$$

where $\boldsymbol{g}$ is the geographical location where we want to know the pollution, $\boldsymbol{a}$ is a point belonging to the perimeter of the associated OSM object, r and p are the pollution range and pollution level of the pollution source, respectively. The user can control the pollution range and level of the pollution source via the interface, which in turn will affect the result of the Gaussian functions assigned to each point in the perimeter.

3.3 Pollution Profile

As can be seen in class diagram in Figure 1, every map entity can be associated with a single pollution source. This, in turn, only specifies a base pollution level and range, which are represented in class `PollutionSource` in Figure 1. However, pollution can vary depending on many other different factors, such as hour of the day, day of the week, weather conditions, or traffic intensity. The expert can express this knowledge as a set of k functions $f_i : \mathbb{R} \to \mathbb{R}$, with $1 \leqslant i \leqslant k$, one per factor. Hence $f_i(x_i)$ maps dynamic factor value x_i (e.g., a day of the week, a wind direction) to a pollution influence, which will then be used to modulate the base (static) pollution level. Functions f_i are defined by the user graphically, as depicted by the plot in Figure 4. In addition to these functions, the expert can also specify the weight of each factor to the overall pollution influence level, $\boldsymbol{w} \stackrel{\text{def}}{=} (w_1, w_2, \ldots w_k)$, with w_i representing the weight of factor i.

Jointly, functions f_i and the vectors $\boldsymbol{x}$ and $\boldsymbol{w}$ constitute a set of pollution profiles. Figure 4 illustrates how the expert sets up a pollution dynamic profile. It also shows where the expert can specify the relevance or weight of a pollution source in the map entities to which it is associated. On the top of the screenshot, the expert can specify the value of all factors that influence pollution. Pollution is represented using the color code of the AQI index has shown in Table 1.

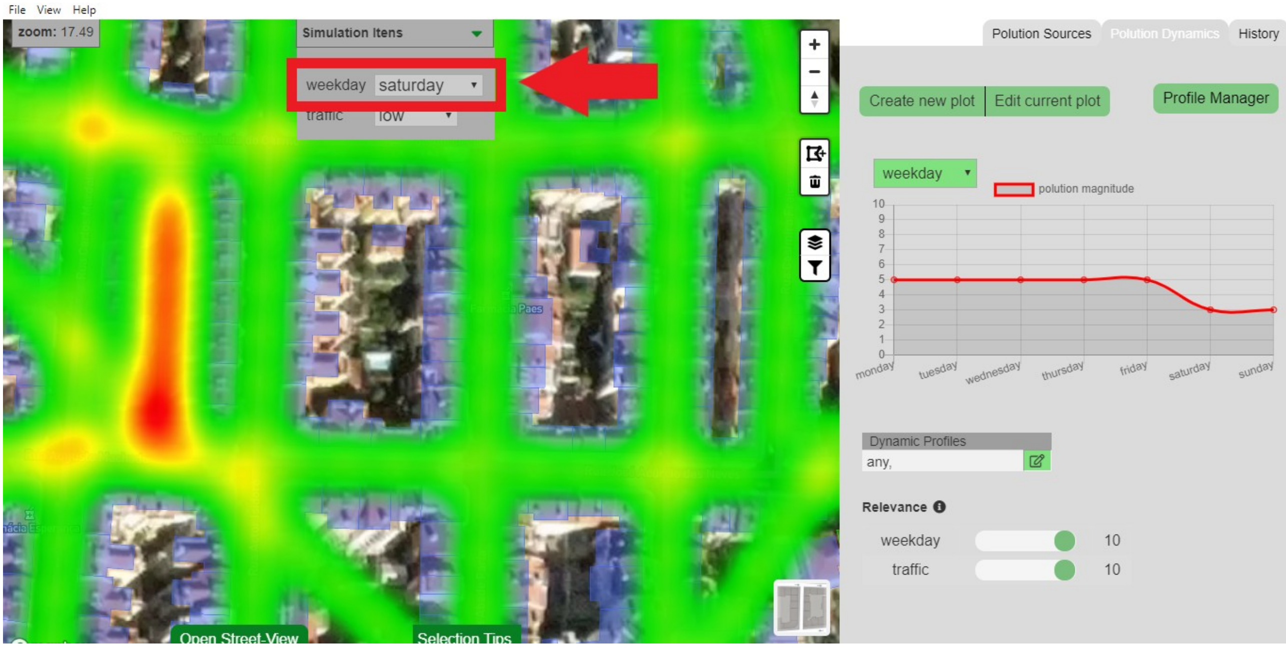

Figure 4: Pollution profile that depends on day of week. Pollution is lower on weekends compared to weekdays. The user can also specify the relevance of weight of existing pollution sources. On the left panel, the highlighted box shows where the expert can select a set of pollution factor values.

Given the set of functions f_i, the vector of factor values in a given moment n, $\boldsymbol{x_n}$, the vector of factor weights $\boldsymbol{w}$, and assuming the Gaussian function in equation (1), the pollution contribution of a given point $\boldsymbol{a}$, belonging to the perimeter of a given OSM object, to a position $\boldsymbol{g}$ in the map is given by:

$$p_m(\boldsymbol{g}, \boldsymbol{a}, r, p, \boldsymbol{x_n}, \boldsymbol{w}) = p_p(\boldsymbol{g}, \boldsymbol{a}, r, p) |\boldsymbol{x_n}|_1^{-1} (\boldsymbol{f}(\boldsymbol{x_n}) \cdot \boldsymbol{w}) \quad , \qquad (2)$$

where r and p are the pollution range and pollution level of the pollution source, respectively, $\boldsymbol{f}(\boldsymbol{x_n}) \stackrel{\text{def}}{=} (f_1(x_1), f_2(x_2), \ldots, f_k(x_k))$ is the vector containing the pollution influences for the input vector $\boldsymbol{x_n}$ factor functions results, and $|\boldsymbol{x_n}|_1^{-1}$ is the L-1 norm of vector $\boldsymbol{x_n}$.

3.4 Pollution Analysis

Pollution scenarios can be created when the expert creates pollution sources, constructs pollution dynamic profiles, associates them to map entities, assigns values to the pollution factors, and attributes weights of factors. The result is similar to the image shown on the left side of the screenshot in Figure 4.

The user can export the produced pollution heat map to an image file. To achieve this purpose, the tool considers a rectangular grid that covers a predefined area. Since the ExpoLIS project is targeted to the city of Lisbon, the area is located between the geographical coordinates (38° 40' 54.3216"N, 9° 18' 1.8216"W) and (38° 48' 46.9254"N, 9° 4' 41.0268"W). We then sample points inside this rectangular grid with a step of 0° 0' 2.0052" latitude degrees and 0° 0' 2.2716" longitude degrees. The pollution $p(\boldsymbol{g},\boldsymbol{x_n},\boldsymbol{w})$ computed for each grid point $\boldsymbol{g}$, given a vector of factor values $\boldsymbol{x_n}$ and the vector of factor weights $\boldsymbol{w}$, is:

$$p_f(\boldsymbol{g},\boldsymbol{x_n},\boldsymbol{w})=\sum_{\Phi\in E}\sum_{\boldsymbol{a}\in\Phi_A} p_m(\boldsymbol{g},\boldsymbol{a},\Phi_r,\Phi_p,\boldsymbol{x_n},\boldsymbol{w}) \quad , \qquad (3)$$

where E is the set of map entities with pollution sources, Φ_A is the set with the points sampled from the perimeter of map entity $\Phi\in E$, Φ_r is the user-defined pollution range of map entity Φ, and Φ_p is the user-defined pollution base of map entity Φ. Each perimeter point $\boldsymbol{a}$ only contributes to the pollution in grid point $\boldsymbol{g}$ if the distance between these points is less than the pollution range Φ_r. A heat-map resulting from the sum of all these functions, one per pollution source, is exemplified in Figure 5.

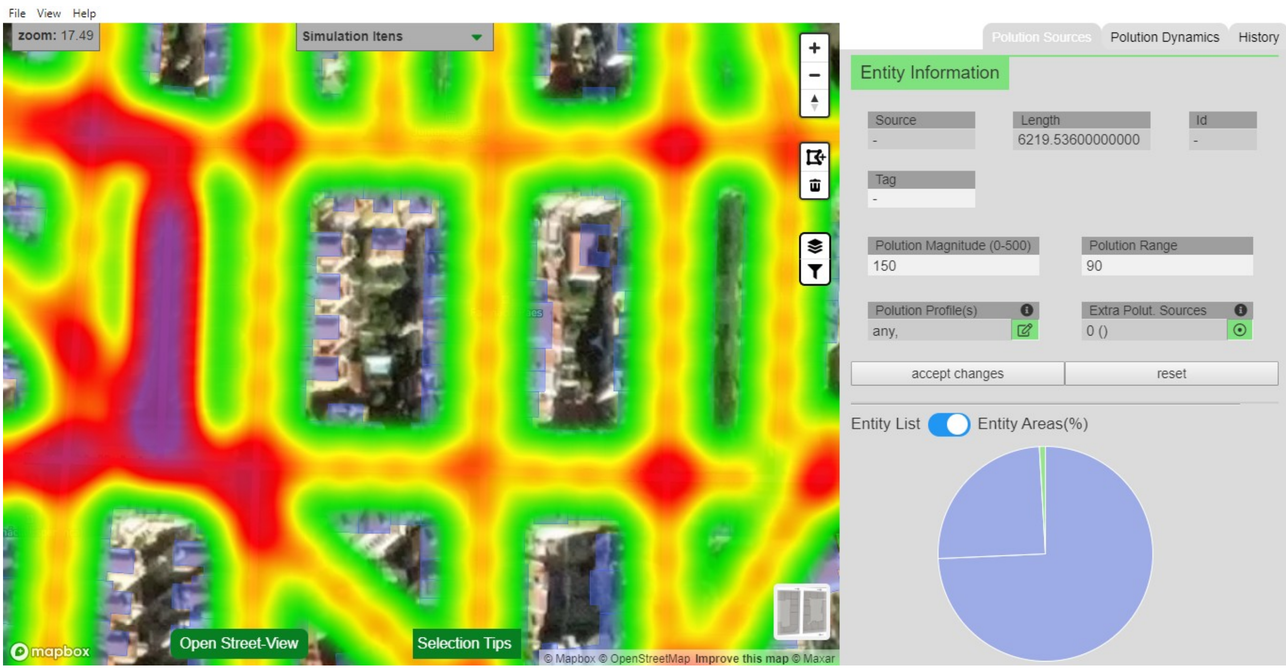

Figure 5: Resulting pollution heat map after assigning pollution to all roads.

Equation 3 represents a different pollution model from the one presented in our earlier work (Vital et al., 2021). Pollution in our previous work was implemented using the heat map functionality of the MapBoxGL programming library. The previous model was limited to a fixed zoom level. In the current version of our graphical tool, pollution is represented by a grid lattice, which operates independently of zoom level.

3.5 Urban Topology

When characterizing a pollution source, experts need to know the topology of the surrounding area, which can be partially determined based on the height of surrounding buildings. The OSM database holds this information for some buildings, which is exploited by MapBox library to display a 3D representation of a surrounding area. However, due to the incompleteness and inaccuracy of the resulting representation, we opted for an alternative approach, that is, to embed the Google Street View application into our tool. This way the user is provided with high quality street-level imagery data to determine the topology of the environment. Figure 6 shows a screenshot of this functionality. There is a button on the bottom of the screen, that whenever the user presses, a new window is opened with a web page of Google Street View.

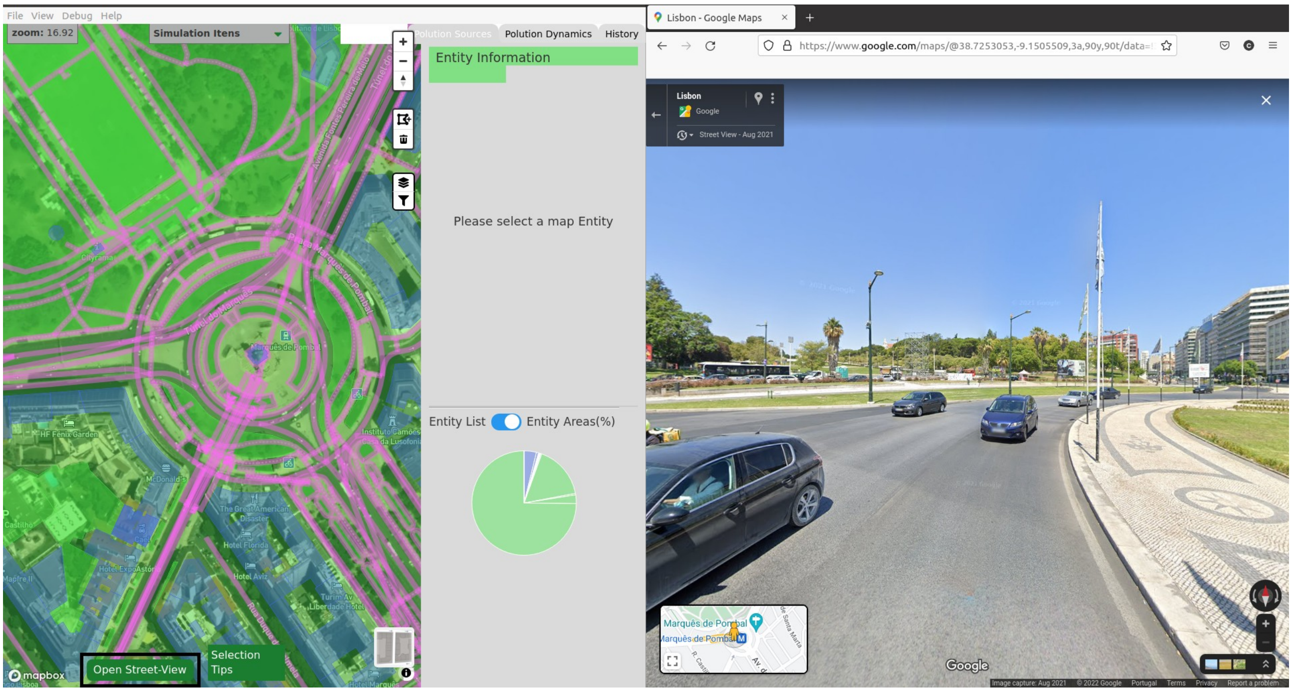

Figure 6: Example of using Open Street View to examine the urban topology of the area in the center of the map.

4 Evaluation

The first phase of evaluating an application consists in discussing and assessing the main functionalities with the target users using paper mock-ups, before engaging in its actual implementation. In this sense, we conducted a session with environmental experts where interface sketches were discussed to assess the potential of the application. These sketches focused on selecting map elements, creating a variable pollution profile, and associating map elements with pollution profiles to create pollution sources.

In a second evaluation phase, we employed the Cognitive Work Analysis framework (Vicente, 1999) to assess and organize all the interface elements. This framework specifies a set of constraints that restrict what the user knows and can do in a particular interface. Heuristic evaluation, as proposed by Nielsen (1994), was also carried out to detect potential usability issues in the early development phases. We also performed formative evaluations with small groups of testers (six people with an average age of 20.7) to identify earlier bugs or misconceptions in the interface. After concluding the tool's development, an extensive summative evaluation session was carried out, whose results are discussed in the following sections.

4.1 Summative Evaluation

A total of 30 participants took part of the summative evaluation. Of these, 6 are from environmental science area. They are either researchers, Ph.D., or M.Sc. students. This group of 6 people formed the *environmental knowledgeable* (EK) group and only their results and feedback were taken into consideration when assessing the potential of our graphical tool. We considered environmental science MSc students as belonging to the knowledgeable group, because they have more knowledge about the environment compared to other students. The remaining 24 participants are MSc or BSc students, 3 of them from economics area and 21 of them from computing science area. They formed the non-EK group. To evaluate the usability of our graphical tool, we considered the results and feedback from all 30 participants.

The evaluation session that each participant went through had the following steps:

1 in the beginning the experimenter gave a brief overview of the graphical tool and its relevance to the ExpoLIS project;

2 the experimenter asked the participant to express his/her reasoning while performing the tasks, in order for the experimenter to understand which interface element the participant was trying to use;

3 the third step consisted in presenting a tutorial to the participant with seven slides describing the functionalities of the graphical tool. This step complemented the brief overview given in the first step;

4 the fourth step marked the evaluation of the graphical tool, where the participant had to perform seven tasks related to the main functionalities. Table 2 shows a description of these tasks;

5 the fifth step consisted in requesting the participant to perform an additional final task, which was more complex than the previous ones, as it required the participant to create a pollution map from scratch in an unfamiliar location. The purpose was to avoid any bias by the participant regarding any pollution knowledge he may have;

6 the sixth step consisted in filling the user interface evaluation questionnaire, or System Usability Scale (SUS), described by Sauro (2011) and Lewis and Sauro (2018). The questionnaire had three additional questions for assessing the ability of participants from the EK group to create a pollution map using the tool. Tables 3 and 4 show the description of these two sets of questions;

7 the seventh and last step consisted in a final debriefing, where participants could comment on the user experience and representation capabilities of the tool.

#	**task description**
1	Associate a pollution magnitude of 100 AQI units to a specific road
2	Select and associate a pollution magnitude of 150 AQI units to every building alongside a specific road
3	Represent week-long variation in the viewport considering that at the weekend pollution is reduced to half
4	Select and associate a pollution magnitude value of 100 AQI units to all roads currently in the viewport
5	Associate a pollution magnitude of 500 AQI units to a specific non-segmented building
6	Represent day-long variation of the previous task's building, considering that this variation has half the impact of the week-long variation
7	Create a pollution map in an unfamiliar location

Table 2: Description of the tasks each participant had to perform.

From the fourth and fifth step we obtained task completion rates. For the sixth step, the SUS questionnaire uses a five-point Likert scale that ranges from strongly disagree to strongly agree. We also used this scale in the last three questions. While these results represent quantitative data, we also obtained qualitative data. The qualitative data was obtained thanks to the fact that most participants thought aloud throughout the evaluation session as they were trying to complete all the tasks. In the final debriefing, participants also provided qualitative data, as they expressed positive and negative aspects of their experience with the graphical tool.

#	SUS question description
1	I think that i would like to use this system frequently.
2	I found the system unnecessarily complex.
3	I thought the system was easy to use.
4	I think that I would need the support of a technical person to be able to use this system.
5	I found the various functions in this system were well integrated.
6	I thought there was too much inconsistency in this system.
7	I would imagine that most people would learn to use this system very quickly.
8	I found the system very cumbersome to use.
9	I felt very confident using the system.
10	I needed to learn a lot of things before I could get going with this system.

Table 3: Description of the ten questions of the SUS questionnaire.

#	SUS question description
11	In the last task, the tool allowed me to create a pollution map given a zone and its properties
12	In the last task, the tool allowed me to express my knowledge on pollution emission given a zone and its properties
13	The pollution map I created given a zone, and its properties, is representative of what I expect is correct

Table 4: Description of the three last questions that are related to the usability of the graphical tool as being able to represent pollution maps.

4.2 Results

4.2.1 Tasks 1-6: Assessing the Main Functionalities

The main results of the first six tasks performed by the 30 participants in the fourth step of the evaluation proceeding (see Table 2), which aimed to assess the main functionalities of the graphical tool, consisted in completion rates. In order to evaluate and compare these completion rates, we created three categories: *yes* meaning the participant completed the tasks within five minutes; *middling* meaning the participant required more than five minutes; and *no* meaning the participant was unable to complete the task. Figure 7 shows the results of the completion rates of the first six tasks. We have divided the results between the one obtained with the normal group and with the EK group. As it can be seen, in 80 % of the occasions, participants were able to complete the first five tasks within five minutes. If we do not impose a time limit, then, in 94 % of occasions these tasks were completed. If we focus on the participants from the EK group, which contains the targeted audience of this graphical tool, the percentage raises to 100 %. In task six, participants had more difficulties completing it, with more people requiring more than five minutes to finish.

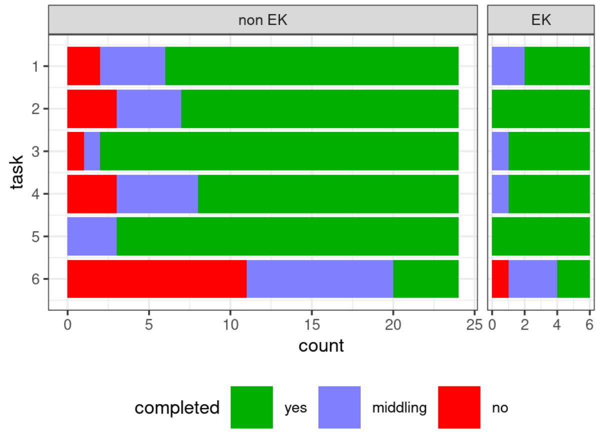

Figure 7: Completion rates of tasks.

The difficulties that participants encountered in the sixth task can be explained by the fact that it requires the completion of three sub-tasks. These sub-tasks were not explicitly described to the participant. They consisted in creating a new pollution dynamic, adjusting the weight of the pollution, and associating the pollution dynamic to the pollution profile that was created in the fifth task. Table 5 shows the percentage of participants that successively performed each of these three sub-tasks. In task six a completion rate of *yes* required that all three sub-tasks were successful. A completion rate of *middling* was attributed if only two sub-tasks were successful. Although most participants were able to create a pollution dynamic, editing the pollution profile weight and the pollution factor values was more difficult. This then impacted the association of the new pollution dynamic to the pollution profile. In the debriefing step, participants expressed doubts on how to manage pollution profiles. The tutorial that was shown on step three was not clear on what was the purpose of pollution profiles and what could be done with them. Moreover, when a participant selects the pollution profile of a map entity, there is only one pollution profile visible and it is not clear that it can select other pollution profiles. Perhaps the inclusion of predefined pollution profiles could clarify this functionality.

sub-task	success
create a new pollution variation chart	76 %
adjust the chart's relevance to half	43 %
associate the chart to profile of previous task	33 %

Table 5: Success rate of task six sub-tasks.

To assess the usability of the graphical tool by any person from the target population, we computed 95 % confidence intervals for the completeness rate. We computed one confidence interval that only used the *yes* completion rate, and a second confidence interval that used both *yes* and *middling* completion rates. Figure 8 shows that if users are given enough time, the confidence intervals for the first five tasks are above the 75 % completion rate mark. If we are strict on the time users are allowed to complete a task, then only tasks 3 and 5 are above 75 % mark. As for task 6, the confidence intervals are low, which is consistent with what we have previously discussed.

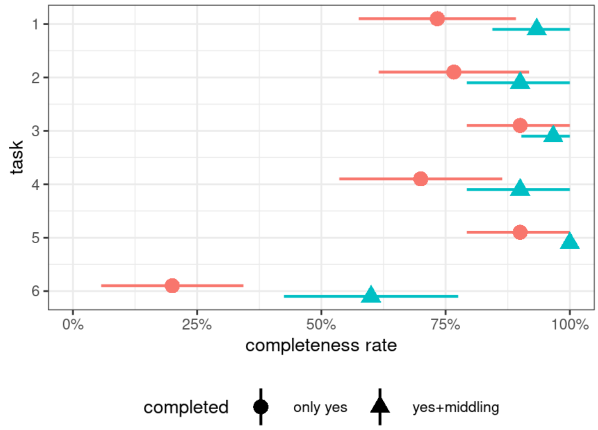

Figure 8: Confidence intervals for completion rates of tasks, for two sets of completion rates: one considering only the yes category (red circle), a second with the yes and middling categories (cyan triangle).

4.2.2 Task 7: Assessing Creation of a Pollution Map

The evaluation of task 7 which involved the creation of a pollution map in step five of the evaluation proceeding, is more complex, as participants had to build on the knowledge they obtained in the previous six tasks. Task 7 required the creation of a pollution map in a location unfamiliar to the participant. In terms of completion rate, all participants did create a pollution map, but the kind of map created differed from participant to participant. Most participants (53 %) selected multiple entities and just assigned the same pollution magnitude to all of them. Only 13 % of the participants assigned different pollution magnitudes to the several map entities.

Note that the experimenter did not tell participants to assign different pollution values. The experimenter just asked participants to create a pollution map that reflected what the participant thought what pollution levels could occur in an unfamiliar location. Participants from the EK group reported that, when trying to assign a particular pollution level, felt they needed more data to make an informed decision regarding pollution values.

Given that the target audience of the graphical tool is environmental experts, the feedback they provided is especially relevant. They mentioned that, rather than assigning an AQI value to pollution entities, they would prefer assigning individual pollutants, such as nitrogen oxides (NOx), carbon monoxide (CO), or PM. They also lacked other sources of information that are relevant for profiling air quality, such as traffic data, which is a considerable source of air pollution. Other found issues include the inability of the tool to allow the experts to express the relation between air pollution dispersion and buildings height, for instance, that street canyons do not disperse pollutants as easily as open areas. Overall, the experts' provided feedback was useful to set a road map for future improvements of the graphical tool.

4.2.3 Usability Tests

After performing the several tasks discussed in the previous sections with the graphical tool, participants filled a two-part questionnaire that focused on the usability issue. The first part contains generic usability items from SUS, whereas the second part contains non-SUS items focused on assessing the viability of the graphical tool to be used by environmental experts to express their knowledge. The top and bottom panels of Figure 9 show the results for the SUS and viability questionnaire parts, respectively. Results are also divided by non-EK group and EK group participants (left and right panels, respectively).

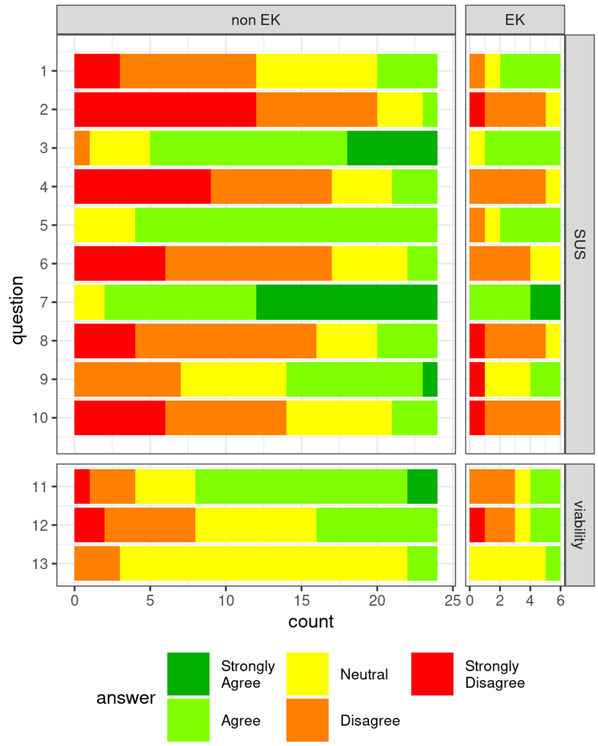

Figure 9: Results from the questionnaire.

To evaluate the results of a SUS questionnaire it is customary to compute a score following the method of Lewis and Sauro (2018). This yields a value between 0 and 100. Sauro (2011) has analyzed several usability questionnaires looking for a threshold value above which a tool can be considered with an acceptable usability level, reporting it to be 68. The graphical tool that we developed achieved a value of 68.83 for the SUS part of the questionnaire, which sits in the edge of acceptability. This value can explain the fact that participants had some difficulties using pollution profiles and thus this feature can be improved to increase its usability.

Lewis and Sauro (2018) pointed out that the SUS questionnaire has two main dimensions: usability and learning. If we compute the value of these two dimensions separately for the results we obtained for the SUS part of the questionnaire, we get the values 68.13 and 71.67 for usability and learning, respectively. Again, usability is just above the acceptance threshold, whereas the learning component fared a little better but it is still close to the 68 threshold.

As for the questionnaire's last three non-SUS items, which are focused on assessing the viability of our graphical tool, most participants either give a neutral answer or disagreed on the ability of being able to represent a pollution map with our tool. On the other hand, the feedback and comments participants gave, are useful to fix and improve the viability of our graphical tool.

Since the learning dimension is above the 68 threshold, this means the functionalities of our tool can be learned. However, as we have already mentioned, the pollution profile feature needs further adjustments to improve its usability. The questions related to the viability of our tool, show that additional features are needed to increase the usability of our tool. The ability to see the urban topology was implemented after conducting the summative evaluation. As we have said previously, it was implemented using Google Street View. Other functionalities such as different pollutants, require changing the data model.

5 Discussion, Conclusions and Future Work

In this paper we have presented the concept of knowledge-based generation of air pollution maps. Existing map solutions are either based on sensory data from static ground stations or on mathematical models. In this paper we propose tap expert knowledge and use this knowledge to create air pollution maps. The aim of this approach is to complement static ground stations, which are the major source of information on air pollution. Despite their wide deployment, there are still major parts of the world that are not monitored. In order to represent expert knowledge, we have developed a graphical tool aimed at air pollution experts. With this tool we aim at predicting air pollution in non-monitored areas.

This article has two major parts. In the first part, we described the engineering approach we took, in particular how we model pollution. In our case we have used multiple single point sources, where each point source is described by a Gaussian function. We described the functionalities that we have implemented. One of them is the ability to create pollution profiles that can depend on user defined parameters (e.g. day of week). This allows an environmental expert to express how pollution varies with day of week, hour of day, traffic intensity, wind direction, and so forth.

In the second part of this article, we focused on assessing the usability and viability aspects of the graphical tool. We have performed a series of experiments in which participants were asked to perform a set of tasks and fill a questionnaire afterwards. Most participants were able to complete the tasks of simple or average complexity. Some of these tasks involve user-interface interactions that are quite common in an array of commonly used applications. For instance, in most applications users must select an object and modify it (e.g., image editing tools). Task one is an example of this action, where the user was asked to select a map entity and assign a pollution value, and in our case, almost all participants were able to complete it. The task in which participants had to create a pollution map in an unknown location revealed a couple of issues with the graphical tool. Nevertheless, the score of SUS part of the questionnaire was 68.83, which according to previous empirical evidence indicates a usable interface.

Participants provided useful comments on desired functionalities. Examples include expanding the tool to provide the user with better map entity search capabilities, as well as with the ability to express expert knowledge as a function of additional sources of information, such as traffic, weather, and urban 3D topology. Pollution profiles allow a user to express pollution as a function of user defined parameters. However, this feature was not easily perceived during usability tests. As such improvements in this functionality are required. In particular, the ability to access such data (traffic, weather, topology) for a given geographical location was a functionality that users asked for in the final debriefing.

Concerning the access to urban topology, we have implemented a feature that allows a user to access Google Street View. This allows the user to see surrounding buildings and get a sense of street canyons. MapBox has a functionality that allows a 3D representation of surrounding buildings. As for traffic data, MapBox also has a functionality that presents data in major regions such as North America, Western Europe, Middle East, Australia, New Zealand, Japan, and Korea.

Experts also commented on the need to represent different pollutant sources. A specific geographical location may be a high emitter of particulate matter but not of carbon monoxide. This feature requires changing the data model, as currently, a map entity can have at most a single pollution entity. This also entails the modification of existing user interfaces. The most obvious one is the ability to select the pollutant when creating a new pollution source in a map entity.

Besides implementing the additional features requested by the participants, we also plan to study how the pollution maps generated by the graphical tool can be fruitfully fused with air pollution sensor data. We also plan on using machine learning to elicit knowledge from environmental experts to predict air pollution on unclassified areas. Other future work includes the use of the graphical tool as a gamified educational asset for children to learn how city entities (e.g., roads, buildings) correlate with air pollution sources.

Acknowledgments

This work was developed within the ExpoLIS Project (LISBOA-01-0145-FEDER-032088), funded by FEDER and national funds, through FCT-Fundação para a Ciência e Tecnologia.

References

Becnel, T., Tingey, K., Whitaker, J., Sayahi, T., Lê, K., Goffin, P., Butterfield, A., Kelly, K., & Gaillardon, P.-E. (2019). A distributed low-cost pollution monitoring platform. IEEE Internet of Things Journal, 6 (6), 10738–10748. https://doi.org/10.1109/JIOT.2019.2941374

Boykov, Y., & Jolly, M.-P. (2001). Interactive graph cuts for optimal boundary region segmentation of objects in n-d images, In Proceedings eighth IEEE international conference on computer vision. ICCV 2001. https://doi.org/10.1109/ICCV.2001.937505

Breezometer. (2020). Global air quality map. Retrieved June 14, 2021, from https://breezometer.com/air-quality-map

Bumm, S. (2019). Using satellite images to determine AQI values in California [Course project]. Course project. Retrieved June 14, 2021, from http://cs229.stanford.edu/proj2019spr/report/22.pdf

European Commission. (2017). Special eurobarometer 468: Attitudes of European citizens towards the environment. http://data.europa.eu/euodp/en/data/%20dataset/S2156 88 1 468 ENG

European Environmental Agency. (2017). Coal-fired power plants remain top industrial polluters in Europe. Retrieved March 8, 2022, from https://www.eea.europa.eu/highlights/coal-fired-power-plants-remain

Haklay, M., & Weber, P. (2008). Openstreetmap: User-generated street maps. IEEE Pervasive Computing, 7 (4), 12–18. https://doi.org/10.1109/MPRV.2008.80

Hasenfratz, D., Saukh, O., Walser, C., Hueglin, C., Fierz, M., Arn, T., Beutel, J., & Thiele, L. (2015). Deriving high-resolution urban air pollution maps using mobile sensor nodes [Selected Papers from the Twelfth Annual IEEE International Conference on Pervasive Computing and Communications (PerCom 2014)]. Pervasive and Mobile Computing, 16, 268–285. https://doi.org/10.1016/j.pmcj.2014.11.008

IQAir. (2021). Live animated air quality map (AQI, PM2.5...) Retrieved July 14, 2021, from https://iqair.com/air-quality-map

Karuppasamy, M. B., Seshachalam, S., Natesan, U., Ayyamperumal, R., Karuppannan, S., Gopalakrishnan, G., &Nazir, N. (2020). Air pollution improvement and mortality rate during covid-19 pandemic in India: Global intersectional study. Air Quality, Atmosphere & Health, 13 (11), 1375–1384. https://doi.org/10.1007/s11869-020-00892-w

Kumar, A., & Goyal, P. (2011). Forecasting of daily air quality index in Delhi. Science of The Total Environment, 409 (24), 5517–5523. https://doi.org/10.1016/j.scitotenv.2011.08.069

Lakhani, A., Balasubramanian, R., & Gurjar, B. R. (2010). Air pollution monitoring and source characterization. In B. R. Gurjar, L. T. Molina, & C. S. Ojha (Eds.), Air pollution, health and environmental impacts (pp. 493–518). CRC Press. https://doi.org/10.1201/EBK1439809624

Lewis, J. R., & Sauro, J. (2018). Item benchmarks for the system usability scale. Journal of Usability Studies, 13 (3), 158–167.

Lin, Y.-C., Chi, W.-J., & Lin, Y.-Q. (2020). The improvement of spatial-temporal resolution of pm2.5 estimation based on micro-air quality sensors by using data fusion technique. Environment International, 134, 105305. https://doi.org/10.1016/j.envint.2019.105305

Morawska, L., Thai, P. K., Liu, X., Asumadu-Sakyi, A., Ayoko, G., Bartonova, A., Bedini, A., Chai, F., Christensen, B., Dunbabin, M., Gao, J., Hagler, G. S., Jayaratne, R., Kumar, P., Lau, A. K., Louie, P. K., Mazaheri, M., Ning, Z., Motta, N., … Williams, R. (2018). Applications of low-cost sensing technologies for air quality monitoring and exposure assessment: How far have they gone? Environment International, 116, 286–299. https://doi.org/10.1016/j.envint.2018.04.018

Netek, R., Masopust, J., Pavlicek, F., & Pechanec, V. (2020). Performance testing on vector vs. raster map tiles—comparative study on load metrics. ISPRS International Journal of Geo-Information, 9 (2). https://doi.org/10.3390/ijgi9020101

Nielsen, J. (1994). Enhancing the explanatory power of usability heuristics, In CHI'94: Conference companion on human factors in computing systems. https://doi.org/10.1145/259963.260333

Rother, C., Kolmogorov, V., & Blake, A. (2004). "GrabCut": Interactive foreground extraction using iterated graph cuts, In ACM SIGGraph 2004 papers, Los Angeles, California, Association for Computing Machinery. https://doi.org/10.1145/1186562.1015720

Russell, B. C., Torralba, A., Murphy, K. P., & Freeman, W. T. (2008). LabelMe: A database and web-based tool for image annotation. International Journal of Computer Vision, 77 (1), 157–173. https://doi.org/10.1007/s11263-007-0090-8

Santana, P., Almeida, A., Mariano, P., Correia, C., Martins, V., & Almeida, S. M. (2021). Air quality mapping and visualisation: An affordable solution based on a vehicle-mounted sensor network. Journal of Cleaner Production, 315, 128194. https://doi.org/10.1016/j.jclepro.2021.128194

Selvam, S., Muthukumar, P., Venkatramanan, S., Roy, P., Manikanda Bharath, K., & Jesuraja, K. (2020). Sars-cov-2 pandemic lockdown: Effects on air quality in the industrialized Gujarat state of India. Science of The Total Environment, 737, 140391. https://doi.org/10.1016/j.scitotenv.2020.140391

Teles, B., Mariano, P., & Santana, P. (2020). Game-like 3d visualisation of air quality data. Multimodal Technologies and Interaction, 4 (3). https://doi.org/10.3390/mti4030054

Vicente, K. J. (1999). Cognitive work analysis: Toward safe, productive, and healthy computer-based work. CRC Press.

Vital, D., Mariano, P., Marta Almeida, S., & Santana, P. (2021). A graphical tool for eliciting knowledge of air pollution sources, In 2021 international conference on graphics and interaction (ICGI). https://doi.org/10.1109/ICGI54032.2021.9655276

Weston, J., Bengio, S., & Usunier, N. (2011). Wsabie: Scaling up to large vocabulary image annotation (T. Walsh, Ed.). In T. Walsh (Ed.), Proceedings of the twenty-second international joint conference on artificial intelligence, AAAI Press. https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-460