OpenLUR: Off-the-shelf air pollution modeling with open features and machine learning

Florian Lautenschlager [a],[*], Martin Becker [b], Konstantin Kobs [a], Michael Steininger [a], Padraig Davidson [a], Anna Krause [a], Andreas Hotho [a]

[a] *Chair for Data Science, Computer Science, University of Würzburg, Am Hubland, 97074 Würzburg, Germany*
[b] *Stanford University, USA*

HIGHLIGHTS

- Introduction of globally and openly available features for land use regression (LUR).
- Machine learning featuring automated hyper-parameter tuning for LUR tasks.
- Global features significantly enhance LUR through cross-learning on multiple cities.
- Source code and data available at dmir.org/openlur

ARTICLE INFO

ABSTRACT

To assess the exposure of citizens to pollutants like NO_x or particulate matter in urban areas, land use regression (LUR) models are a well established method. LUR models leverage information about environmental and anthropogenic factors such as cars, heating, or industry to predict air pollution in areas where no measurements have been made. However, existing approaches are often not globally applicable and require tedious hyper-parameter tuning to enable high quality predictions. In this work, we tackle these issues by introducing *Open-LUR*, an off-the-shelf approach for modeling air pollution that (i) works on a set of novel features solely extracted from the globally and openly available data source OpenStreetMap and (ii) is based on state-of-the-art machine learning featuring automated hyper-parameter tuning in order to minimize manual effort. We show that our proposed features are able to outperform their counterparts from local and closed sources, and illustrate how automated hyper parameter tuning can yield competitve results while alleviating the need for expert knowledge in machine learning and manual effort. Importantly, we further demonstrate the potential of the global availability of our features by applying cross-learning across different cities in order to reduce the need for a large amount of training samples. Overall, OpenLUR represents an off-the-shelf approach that facilitates easily reproducible experiments and the development of globally applicable models.

1. Introduction

Epidemiological studies show the negative impact of air pollutants like NO_x or particulate matter (UFP, $PM_{2.5}$ and PM_{10}) on respiratory and cardiovascular health (Pope et al., 1991; Polichetti et al., 2009; Brook et al., 2010). In order to assess the exposure of citizens to such pollutants, many measurement campaigns have been conducted. However, such campaigns are often restricted to very few stationary monitoring sites (Briggs et al., 2000; Carr et al., 2002; Brauer et al., 2003; Sahsuvaroglu et al., 2006; Henderson et al., 2007; Arain et al., 2007; Aguilera et al., 2007; Su et al., 2009; Dons et al., 2013; Ragettli et al., 2014; Montagne et al., 2015; Muttoo et al., 2018; Araki et al., 2018), and even if mobile monitoring devices are used, spatial coverage is limited to road segments or locations that have been chosen for the measurement campaign (Sîrbu et al., 2015; Larson et al., 2009; Zwack et al., 2011; Patton et al., 2014; Hasenfratz et al., 2014; Hankey and Marshall, 2015; Su et al., 2015; Shi et al., 2016; Minet et al., 2017; Basu et al., 2019). To retrieve pollutant concentration in unmeasured locations researchers rely on the correlation of air pollution with environmental and anthropogenic factors such as cars, streets, heating or industry (Jerrett et al.,

2004). In particular, they employ land use regression (LUR) models which leverage features extracted from land use statistics to overcome the limits and predict air quality in a spatially dense manner.
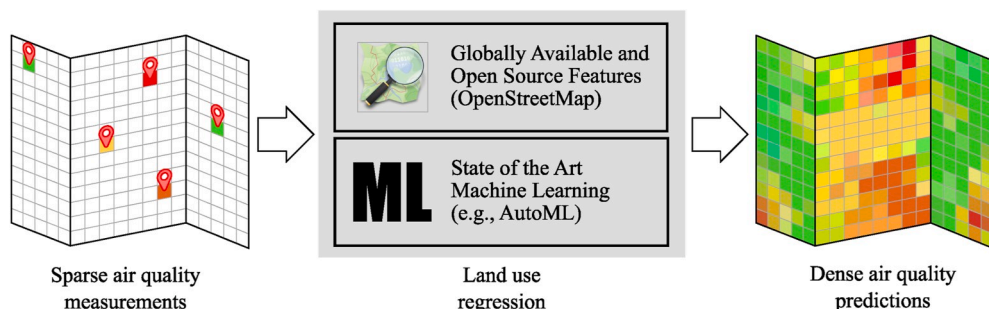
1.1. Problem setting

In previous work, the corresponding features usually stem from very specialized sources like local governments (Brauer et al., 2003; Hoek et al., 2001; Stafoggia et al., 2019), commercial providers (Sahsuvaroglu et al., 2006; Muttoo et al., 2018; Stafoggia et al., 2019), other models (for example traffic or weather models) (Dons et al., 2013; Stafoggia et al., 2019), or custom recordings (Briggs et al., 2000; Carr et al., 2002). For some studies, the source of the underlying land use data is even not easy to access (Montagne et al., 2015; Hankey and Marshall, 2015; Araki et al., 2018). The proposed methods are consequently hard to reproduce and hardly generalize to arbitrary locations.

Additionally, current work is often based on relatively simple models like linear regression (Arain et al., 2007; Aguilera et al., 2007; Muttoo et al., 2018) or generalized additive models (GAM) (Hasenfratz et al., 2014). While some newer work explores more advanced methods (Champendal et al., 2014; Brokamp et al., 2017; Araki et al., 2018; Stafoggia et al., 2019; Basu et al., 2019), state-of-the-art machine learning approaches are still frequently neglected or require tedious hyper-parameter studies.

1.2. Approach

In this work, we address this issue and propose *OpenLUR*, an off-the-shelf solution for air pollution modeling using land use regression (LUR) based on open features and state-of-the-art machine learning (see Fig. 1). First, to ensure reproducible and generalizable models, we derive features solely from openly and globally available data extracted from OpenStreetMap (OSM) (OpenStreetMap contributors, 2017). Second, we apply various state-of-the-art machine learning methods on these features. Besides GAMs and random forests, we specifically focus on methods that feature automated hyper-parameter tuning, for example AutoML (Blum et al., 2015), in order to eliminate the need for tiresome hyper-parameter studies. We evaluate both, our novel features as well as the state-of-the-art methods, on two large scale datasets: mobile air pollution data collected by Hasenfratz et al. (2014) and modelled air pollution data from the London atmospheric emissions inventory (Greater London Authority, 2016). We are able to show (i) that our novel open features outperform previously applied local feature sets on the given data, (ii) that using machine learning with automated hyper-parameter tuning yields high quality, reproducible and spatially generalizable models, (iii) that our features are applicable wherever OpenStreetMap data is available and (iv) that cross-learning on multiple cities can significantly enhance the model performance for small datasets.

1.3. Contribution

Our contributions in this article are: (i) We introduce a set of globally and openly available features for modeling air pollution using land use regression that significantly outperform previously proposed specialized features and show their global applicability. (ii) We evaluate state-of-art machine learning featuring automated hyper-parameter tuning for the application in land use regression tasks. (iii) We assess the enhancements for urban land use regression models achieved by the utilization of data from multiple cities. (iv) We propose OpenLUR as a globally applicable and expendable approach for land use regression and make the source code and our extracted features openly available at https://www.dmir.org/OpenLUR in order to ensure reproducibility and to enable future research.

2. Air quality training data

Our approach is generally applicable to any land use regression scenario. In this work, we train and test our models and features on a year of data collected during the OpenSense project in Zurich starting from April of 2012. Ensuing we show the global applicability of the approach on data extracted from the London Atmospheric Emissions Inventory (LAEI) (Greater London Authority, 2016) and demonstrate the potential of globally available land use features by combining both datasets.

2.1. OpenSense data

In the OpenSense project UFP was continuously measured by sensorboxes fixed to the top of tram cars (Hasenfratz et al., 2014). Hasenfratz et al. (2014) show the good measurement quality through the statistical distribution of measurements, comparison of baseline signals from several measurement devices and evaluation against high-quality datasets. With regard to preprocessing, we follow Hasenfratz et al. (2014): To rule out effects of seasonal variability on air pollution we split the collected data into four seasons of three months each (see Table 1).

To further smooth over smaller temporal and spatial variabilities and outliers, we divided the observation area into squares of $100m \times 100m$ and averaged the measurements for each season and square. Finally, squares with small numbers of samples which are prone to outliers and may negatively impact the model building process were removed. In particular, we kept the 200 squares with the largest amount of measured points (Hasenfratz et al., 2014). The # rows in Table 1 show the mean, min and max amount of measurements in the squares, that were kept in the dataset. The values used for the model training therefore are averages of at least 2000 single measurements which limits the influence of single outliers in the original data.

Table 1 shows statistics of our dataset by season. The mean as well as the standard deviation (SD) tend to be higher for the two later seasons in this dataset.

For a spatial visualization of Season 2, see Fig. 2. The particular

Fig. 1. Abstract-/ToC-Art: Off-the-shelf approach to air pollution modeling using land use regression (LUR) powered by openly available features and state-of-the-art machine learning: On the left this figure shows a set of sparsely collected air quality measurements. To derive a spatially dense map, we train a LUR air quality model using globally and openly available features derived from OpenStreetMap by applying state-of-the-art machine learning featuring automated hyper-parameter tuning. Open-LUR ensures easily reproducible experiments and enables world wide applicable models.

Sparse air quality measurements

Globally Available and Open Source Features (OpenStreetMap)

ML State of the Art Machine Learning (e.g., AutoML)

Land use regression

Dense air quality predictions

Table 1

The four seasonal OpenSense UFP datasets from Zurich and basic statistics. The # rows show mean, min and max count of measurements used for the average in the squares in which we aggregated the air pollution mesurements.

Season	1	2	3	4
From	April 01, 2012	July 01, 2012	October 01, 2012	January 01, 2013
To	June 30, 2012	September 30, 2012	December 31, 2012	March 31, 2013
Mean $[\frac{10^9 particles}{m^3}]$	12.88	13.69	16.08	17.99
SD $[\frac{10^9 particles}{m^3}]$	2.81	2.36	3.72	4.25
Mean #	7292	6111	11712	10986
Min #	2817	2105	3647	3727
Max #	29946	29781	74588	222928

spatial patterns of the measurements are due to the sensor boxes being mounted on tram cars.

2.2. LAEI data

For showcasing the global applicability of our approach and the potential of globally available features, apply our features on annual mean PM_{10} concentrations stemming from the LAEI dataset (Greater London Authority, 2016). The data was obtained from a detailed dispersion model based on a vast number of input factors like road and rail traffic, aviation, agriculture, industry and domestic and commercial fuel burning and fires. For the dataset generation we randomly sampled 3000 datapoints for training and 1500 for testing purposes from the urban central London region. With this comparatively large dataset we are able to provide evaluation scores that are robust against outliers.

Table 2 shows the mean and standard deviation (SD) of both training and testing dataset. Note that the mean concentrations are higher than in Table 1, as PM_{10} includes bigger particles on top of UFP.

3. The OpenLUR approach

In this section, we introduce the main components of OpenLUR, our off-the-shelf approach for building air quality models based on land use regression (LUR): a novel set of open and globally available features derived from OpenStreetMap as well as the concept of automated hyper-

parameter tuning for state-of-the art machine learning methods.

3.1. OpenStreetMap features

In contrast to feature sets used in previous studies (Hasenfratz et al., 2014; Aguilera et al., 2007; Briggs et al., 2000), our features are only based on OSM and thus are openly available and globally applicable.

To assess the validity of our feature set we compare them to a set of features used in previous work. In particular, we focus on the features from Hasenfratz et al. (2014).

Hasenfratz et al. (2014) derived features for each individual grid cell (cf. Section 2), including for example population or industry density, building heights or terrain properties shown in Table 3. While some of these features are derived from OpenStreetMap, most of them stem from data provided by governmental institutions in Switzerland and Zurich. Thus, they are only available in this region, which leads to a model, that is only applicable in Zurich and can not be compared to models designed fo other regions.

Table 2

Statistics for the LAEI PM_{10} dataset from London.

Dataset	Size	Mean $[\frac{10^9 particles}{m^3}]$	SD $[\frac{10^9 particles}{m^3}]$
Training	3000	28.15	2.78
Testing	1500	28.13	2.43

Table 3

Baseline features from Hasenfratz et al. (2014) with their respective source. Except from OSM, none of the features are globally available.

Feature	Source
Population density	Swiss Federal Statistical Office
Industry density	Swiss Federal Statistical Office
Building heights	Swiss Federal Statistical Office
Heating type	Swiss Federal Statistical Office
Terrain elevation	Swiss Federal Statistical Office
Terrain slope	Swiss Federal Statistical Office
Terrain aspect	Swiss Federal Statistical Office
Road type	OSM
Distance to next road	OSM
Distance to next large road	OSM
Distance to next traffic signal	OSM
Average daily traffic volume	Department of Waste, Water, Energy and Air of the Canton of Zürich
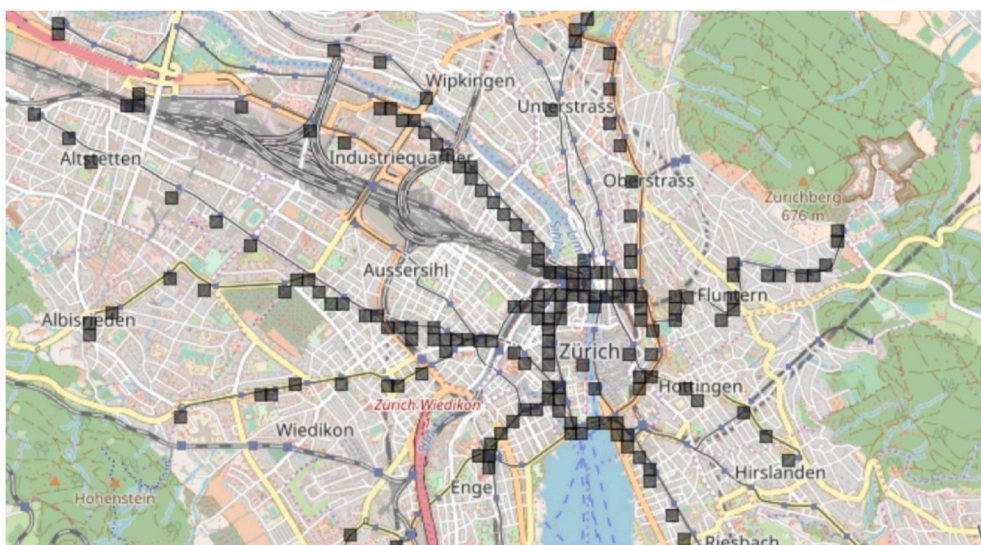
Fig. 2. An excerpt from the air pollution data from Zurich used for training LUR models. The figure shows the spatial distribution of the data from Season 2 of the OpenSense dataset (Hasenfratz et al., 2014) collected via sensor boxes on trams. The individual measurements are aggregated based on $100m \times 100m$ grid cells. Analogously to the experiments in (Hasenfratz et al., 2014), the cells are restricted to those 200 with the most measurements. ©OpenStreetMap contributors (www.openstreetmap.org/copyright).

To derive our novel set of globally available features, we employ OpenStreetMap (OSM) which provides openly and globally available land use data. In this section, we briefly introduce OSM as a data source and describe the features as well as their extraction process.

3.1.1. OpenStreetMap

OSM is an open source map dataset developed and maintained by a large number of volunteers from all around the world (Haklay and Weber, 2008). Many studies confirm the quality of the data provided by OSM (Haklay, 2010; Hecht et al., 2013). Consequently, OSM is a popular data source in a variety of studies ranging from risk management (Schelhorn et al., 2014) and disaster warning (Rahman et al., 2012) to navigation (Hentschel and Wagner, 2010) and routing (Luxen and Vetter, 2011). OSM also contains many variables related to air pollution. For example, it lists *key:value* pairs like *landuse:industry* or *highway: motorway* which can be directly used to derive relevant land usage and land cover statistics (Heymann, 1994; Estima and Painho, 2015; Hasenfratz et al., 2014)..

3.1.2. Feature extraction

To extract air pollution related features, we rely on OSM entities which are stored as polygons, lines or points (such as buildings, streets or traffic lights, respectively). Each entity is associated with a set of *key: value* pairs. In this study, we focus on entities with the keys *landuse* and *highway*. Using entities with these keys, we extracted two types of features: *area/length-based features* and *distance-based features*. These features are generated for each grid cell individually. We provide an overview over the features in Fig. 3.

For the *area/length-based features*, we define a circular zone (buffer) of various sizes around a grid cell's center (see top left of Fig. 3). Within those buffers, we measure the overall area or the overall length covered by those OSM entities relevant to the specific feature. In particular, we derive three area-based features by summing up the areas of entities with the key *landuse* and the values *industrial*, *commercial* and *residential* respectively. For the *length-based features*, we define two categories: roads with heavy traffic and roads with light traffic. For the *heavy traffic features*, we sum up the length of entities with the key *highway* and the values *motorway*, *trunk*, *primary* and *secondary*. For the *light traffic features*, we sum up the length of entities with the key *highway* and the values *tertiary* and *residential*. This provides information about industrial land-use and traffic intensity. The procedure is illustrated in Fig. 3. We varied the buffer radii in 50*m*-steps ranging from 50*m* to 3000*m* for *area-based features* and from 50*m* to 1500*m* for *length-based features* to account for distance-dependencies. The radii were chosen according to their maximum distance of influence (Jerrett et al., 2004; Henderson et al., 2007; Su et al., 2009). Three area-related *key:value* pairs using 60 buffer

radii each and two length-related *key:value* pairs with 30 buffer radii each result in 240 features.

For the *distance-based features*, we focus on the *key:value* pairs *highway:motorway*, *highway:primary*, *highway:traffic_signals* and *landuse:industrial*. For each of these pairs, we calculate the distance between a grid cell's center and the nearest occurrence of an entity with the respective pair as illustrated in Fig. 3. Like for the area-based features, this represents information on the local traffic profile as well as industrial factors which are assumed to negatively influence air quality. Considering the four mentioned *key:value* pairs, this results in 4 features.

Combining both feature classes results in 244 open and publicly available features derived solely from OSM, shown in Table 4. By construction, these features represent land cover and traffic related information and are closely tied to air pollution. We are aware, that this list of land use features is not exhaustive as factors like elevation, population density and other meteorological and environmental covariates can also highly influence air quality. OpenLUR can be extended with additional data sources via an easy to use API. The aim of this study however is to show the capability of OSM to provide land use information that can outperform closed source land use features.

3.2. Automated hyper-parameter tuning

Newer advancements in machine learning often promise better prediction results using the same data. These models however commonly require tedious hyper-parameter tuning and expert knowledge concerning the applied algorithms. In this section we briefly introduce several approaches for automatic hyper-parameter tuning to negate this disadvantage. This is one key feature of our off-the-shelf approach.
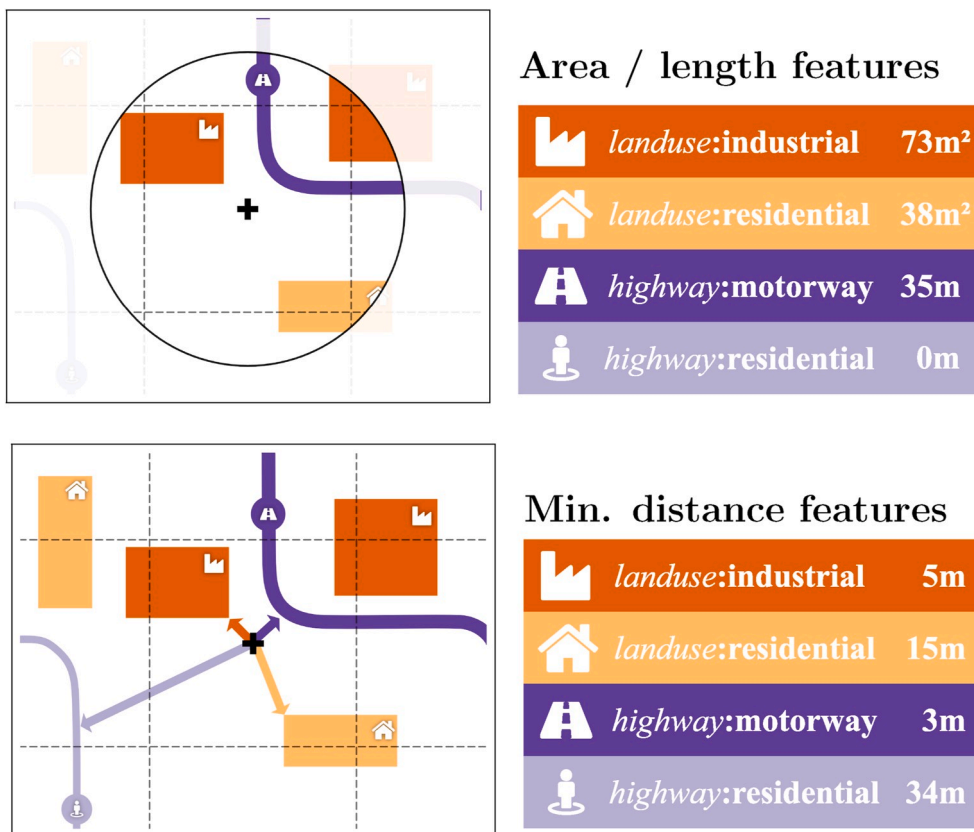
3.2.1. Basic approaches

Most state-of-the-art machine learning methods need to be tailored to specific tasks by selecting an appropriate set of hyper-parameters. For the example of random forests, the number of estimators, the number of features per estimator or the minimal number of samples per leaf have to be tuned. The typical procedure to tune hyper-parameter sets is as follows: The dataset is split in a train, a validation and a test set. Different hyper-parameter sets are trained on the train data and tested on the validation data. The best performing model is used as final model, retrained on train and validation set and tested on the test set. Due to the combinatorial explosion of possible hyper-parameter combinations, this process either requires expert knowledge or has to be automated. In the following, we revisit two commonly used generic methods to automatically optimize hyper-parameters: grid and stochastic search.

3.2.1.1. Grid search.
Grid search is performed by manually choosing a set of candidate values for each hyper-parameter. Then, all possible combinations of these values are evaluated.

3.2.1.2. Stochasic search.
Stochastic search optimizes hyper-parameters by randomly choosing candidate values for each hyper-parameter from a predefined probability distribution (mostly uniform) within a given time budget. This often allows to "find better models by effectively searching a larger, less promising configuration space" (Bergstra and Bengio, 2012) than manual or grid search.

3.2.2. AutoML

Automated Machine Learning (AutoML) (Blum et al., 2015) goes one step further than the standard way to automated parameter tuning. It builds an ensemble learner that exploits the synergy of several weak regressors to produce an improved model. In other words, it simultaneously chooses and combines models from a set of model classes (random forests, support vector machines, naive Bayes, etc.) while *at the same time* optimizing their hyper-parameters. For this, it does not rely on

Table 4
Features derived from OpenStreetMap. The features are divided into two classes: area/length (top part) and distance based features (bottom part), where area/length features use different buffer sizes (50m–3000m/50m–1500m with a step size of 50m). Overall this results in 244 features for each grid cell.

Variable	Unit	key:value pairs in OSM
Industry usage	Area [m^2]	landuse:industrial
Commercial usage	Area [m^2]	landuse:commercial
Residential usage	Area [m^2]	landuse:residential
Heavy traffic	Length [m]	highway:motorway
		highway:trunk
		highway:primary
		highway:secondary
Light traffic	Length [m]	highway:tertiary
		highway:residential
Distance to next motorway	Distance [m]	highway:motorway
Distance to next primary road	Distance [m]	highway:primary
Distance to next traffic signal	Distance [m]	highway:traffic_signals
Distance to next industrial area	Distance [m]	landuse:industrial

Fig. 3. Visualization of our set of open and globally available features. The top picture shows features based on the area/length of land use related entities within a given buffer zone. The bottom picture shows features based on the minimum distance to certain land use related entities.

grid or stochastic search, but utilizes efficient Bayesian optimization methods based on Gaussian processes to intelligently pick the most promising model and hyper-parameter combinations while staying within a given computational budget, such as time or memory usage (Blum et al., 2015).

4. Experimental setup

For OpenLUR we evaluate the two key components introduced in Section 3: our set of OSM features and the concept of hyper-parameter tuning for state-of-the-art machine learning methods. To show the ability of these components to provide an off-the-shelf approach, we compare our globally available OSM features with baseline features from previous work as independent variable and evaluate the competitiveness of machine learning methods featuring automated hyper-parameter tuning.

In this context, our general experimental setup is as follows: We aim to train models to predict the dependent variable UFP concentration at unobserved locations for the four seasons of the OpenSense dataset listed in Section 2. As done in most previous work, R^2 and the root mean squared error (RMSE) are computed as scores to compare their performance. The spatial dependence of the UFP concentration is modelled through the spatial variations of the independent variables. The temporal dependence is ruled out by averaging measurements over seasons (c.f. Section 2) and using only one season for each model building and evaluation process. To account for random outliers of these scores due to the inherently small training sets ($\leq$ 200 labeled samples in the Open-Sense datasets, cf. Section 2), we report the mean of 40 10-fold cross validation scores as the final score for each model (cf. Hasenfratz et al. (2014)): For each of the 40 iterations, the dataset is randomly split into 10 subsets. Ensuingly each subset is used once for the evaluation while the models are built based on the 9 remaining subsets.

As baselines we picked two models that have proven to perform good on state-of-the art land use regression tasks (Hasenfratz et al., 2014; Champendal et al., 2014; Brokamp et al., 2017). To evaluate our apporach we compare them against two machine learning methods featuring automated hyper-parameter tuning. This results in the following list of models:

- GAM: generalized additive model (no hyper-parameter tuning) (Hastie and Tibshirani, 1986)
- RF: random forest (no hyper-parameter tuning) (Breiman, 2001)
- RFOstochastic: random forest (hyper-parameters tuned by stochastic search) (Breiman, 2001; Bergstra and Bengio, 2012)
- AutoML: automated machine learning (automated hyper-parameter tuning) (Blum et al., 2015)

Due to the technical limitations of GAMs, a small set of features needs to be selected. We explain this process in the supplementary material. Beyond evaluating different feature sets, we compare GAMs and untuned random forests, against two state-of-the-art models with automated hyper-parameter tuning. Besides AutoML, we chose to optimize random forests using stochastic hyper-parameter search since (i) random forests are one of the most popular machine learning methods for land use regression (Champendal et al., 2014; Brokamp et al., 2017; Araki et al., 2018; Stafoggia et al., 2019) and (ii) stochastic search is reported to outperform manual or grid search (Bergstra and Bengio, 2012).

Note that, the features we extracted from OpenStreetMap as well as the code used to produce the following results are publicly available at https://www.dmir.org/OpenLUR. A more detailed explanation of the experimental setup can be found in the supplementary material.

5. Results

In this section, we report the results based on the experimental setup described in Section 4. This encompasses (i) results on comparing our novel OSM features against a baseline feature set, (ii) results on comparing machine learning methods with and without hyper-parameter tuning, (iii) the application of OpenLUR on the LAEI dataset to evaluate the number of data samples needed for competitive results and the applicability of cross-learning across different cities, namely Zurich and London, to overcome the limits of small-scale air quality datasets, and (iv) a summary of the results and a recommendation of the overall approach for OpenLUR.

5.1. Feature comparison

In this section, we evaluate the performance of our feature set introduced in Section 3.1 and compare it with specialized — however only locally available — features from previous work (Hasenfratz et al., 2014). For this, we train several air quality models using both sets of features on the four seasons introduced in Section 2.

As a measure of absolute performance gain, we calculated the performance difference of our novel feature set and the OpenSense features for each model and each season with regard to RMSE and R^2, respectively. Table 5 shows the results. For RMSE (R^2) negative values (positive values) indicate a better model performance in favor of our proposed OSM features. Bold values indicate a statistically significant difference using the Wilcoxon signed-rank test ($p < 0.05$ with p values are provided in the supplementary materials).

We observe that in nearly all cases, our novel features yield significantly better results compared to the baseline features. That is, out of the 32 differences, 24 show a significant improvement in model quality. Only in four cases there is a significant tendency towards the baseline features. For the remaining cases our features perform equally well. The latter cases focus on Season 3 pointing towards a very specific data configuration that does not seem to be representative across the evaluated datasets (e.g., due to significant temperature drifts from October to December).

Thus, our novel and open OSM features significantly improve the performance of all studied air quality models compared to specialized and possibly restricted data sources. The globally available features enable air quality models to be trained anywhere on earth where OSM data is available.

5.2. Model comparison

We further evaluate the potential of automated hyper-parameter tuning. We therefore focus on our OSM features since they promise to yield the overall best results. The results for the baseline features are listed in the supplementary material. The results in Table 6 show the performance of each model listed in Section 4 for each season with regard to RMSE and R^2. The models are also ranked from best (1) to worst (4) on their performance in each Season. To facilitate an overall comparison between the models, the table furthermore lists the mean rank across all seasons.

We do not observe statistically significant differences between the regression models. Nevertheless, examining the mean ranks as an alternative evaluation measure, we clearly observe a tendency for models with automated hyper-parameter tuning to perform better than regular models. This holds with regard to both metrics and confirms that hyper-parameter tuning is an essential step for training land use regression models for air quality prediction.

Of those methods featuring automated hyper-parameter tuning, RFOstochastic performs better on both metrics. We assume that due to our rather small dataset the AutoML approach based on Gaussian processes can not exploit its full potential. On top of that Blum et al. (2015) showed, that the AutoML model needs a considerable amount of time to be able to outperform competitors like random forest (Blum et al., 2015). However, we expect the AutoML-based methods to show their advantage on larger datasets, where training models is more expensive and selecting particularly promising sets of hyper-parameters is essential. This needs to be further investigated.

Nevertheless, we have shown that advanced machine learning approaches employing automated hyper-parameter tuning, i.e., AutoML and RFOstochastic, are applicable to air pollution modeling and outperform the baseline methods when considering mean ranks across several experiments. We thus have shown that methods employing hyper-parameter tuning can alleviate manual effort while not compromising on prediction quality.

Fig. 4(a) depicts the spatial distribution of UFP in Zurich predicted by the RFOstochastic trained on OpenSense season 1 with our OSM features. Some patterns are clearly distinguishable: The water area as well as the recreation area in the western central part of the predictions are less polluted than urban Zurich. The higher pollution along some major roads is also visible. Fig. 4(b) shows the standard deviation of multiple resampled OpenLUR runs. The standard deviations are low with values up to 1.6 $[\frac{10^9 particles}{m^3}]$ while absolute predictions are values of up to 17 $[\frac{10^9 particles}{m^3}]$. The deviations are equally low in areas with lower or higher

Table 5

The absolute performance gain of our OSM based features over the OpenSense features with regard to RMSE and R^2 is shown. Negative (positive) RMSE (R^2) values show a better model performance when using our openly available OSM features and are highlighted in gray. In nearly all cases, our OSM features yield significantly better air quality predictions (bold values indicate statistical significance).

Season	AutoML	RFOstochastic	RF	GAM
	RMSE $[\frac{10^9 particles}{m^3}]$			
1	−0.18	−0.24	−0.24	−0.30
2	−0.06	−0.08	−0.07	−0.07
3	−0.05	0.06	**0.06**	**0.10**
4	−0.34	−0.39	−0.40	−0.19
	R^2			
1	**0.11**	**0.16**	**0.16**	**0.18**
2	**0.04**	**0.05**	**0.06**	**0.04**
3	0.03	−0.02	−0.06	−0.06
4	**0.13**	**0.19**	**0.19**	0.03

Table 6

RMSE and R^2 metrics of the models using OSM features. Parenthesis show the rank of the model given a particular season for the corresponding metric. Generally, the model performances do not differ significantly. However, a clear tendency towards models featuring automatic hyper-parameter tuning can be observed judging by their mean rank over all seasons.

Season	AutoML	RFOstochastic	RF	GAM
	RMSE $[\frac{10^9 particles}{m^3}]$ (rank)			
1	2.06 (3)	2.01 (2)	2.12 (4)	2.00 (1)
2	1.75 (2)	1.74 (1)	1.82 (4)	1.75 (2)
3	2.87 (1)	2.91 (2)	3.07 (3)	3.13 (4)
4	3.55 (1)	3.55 (1)	3.69 (3)	3.73 (4)
Mean rank	1.75	1.50	3.50	2.75
	R^2 (rank)			
1	0.40 (3)	0.43 (1)	0.36 (4)	0.42 (2)
2	0.38 (2)	0.39 (1)	0.32 (4)	0.37 (3)
3	0.35 (1)	0.32 (2)	0.25 (3)	0.21 (4)
4	0.19 (1)	0.19 (1)	0.11 (3)	0.08 (4)
Mean rank	1.75	1.25	3.50	3.25

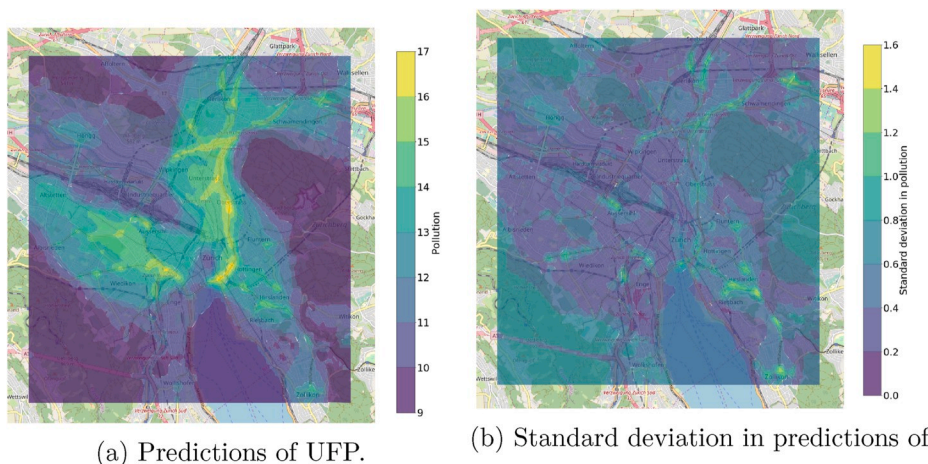(a) Predictions of UFP.   (b) Standard deviation in predictions of

Fig. 4. Predictions and standard deviations of OpenLUR trained on season 1 of the OpenSense data with values in [$\frac{10^9 particles}{m^3}$]. ©OpenStreetMap contributors (www.openstreetmap.org/copyright).

pollution, best visible in the western, higher polluted, region of Zurich. This means that the resampled models have a high agreement on the predictions, independent of the predicted pollution value.

To show that the spatial dependence of the pollutant is modelled by the OSM features, as stated in Section 4, we computed the spatial correlation of the residuals with each independent variable. The correlation coefficient c is between -0.09 and 0.19. The coefficient of determination c^2 describes the percentage of the residuals that can be explained by the independent variable (Taylor, 1990). With $c^2 < 4\%$ the residuals are independent from our OSM features, meaning that the spatial dependence of UFP is indeed modelled by the independent variables.

In conclusion, the predictions achieved with OpenLUR are of a competitive prediction quality (c.f. Tables 5 and 6) and subjectively reasonable (c.f. Fig. 4(a)).

5.3. Global applicability and cross-learning

In this section, we utilize the LAEI dataset to show the applicability and potential of globally available land use features. We will first apply OpenLUR on the LAEI dataset and second demonstrate how the global applicability of our features can be used to improve the performance of LUR for small datasets through cross learning.

We use our OSM features and the best working models from the previous experiment, namely AutoML and RFOstochastic.

To show the global applicability of OpenLUR on datasets of different sizes, we first apply OpenLUR to different subsets (from 10 up to 1000 datapoints) of the LAEI data. The results are shown in Table 7. While the results of the models do not differ significantly, the accuracy rises with the size of the training data with an especially strong increase below 200

training samples. For a high number of training samples (≥ 300), AutoML seems to slightly ($\Delta R^2 \approx 0.01$), though not significantly, outperform RFOstochastic.

Datasets in LUR scenarios however are mostly small (< 200 data points). We can exploit the potential of our globally available OSM features through cross-learning on multiple cities to overcome the limit of small datasets and enhance the model accuracy. For cross-learning we increase the size of our training dataset by adding data samples from another region. In our case, we added 180 OpenSense data points from season 1 to the LAEI training samples.

Since both datasets measure different pollutants (PM_{10} and UFP respectively) the concentration values are in different ranges (see Section 2). To just exploit the similar dependence of both pollutants on land use features, we standardized the measurements of both datasets to a mean of 0 and a standard deviation of 1. To retrieve pollutant predictions, the output of the resulting model has to be transformed back using the mean and standard deviation of the original dataset.

The performance gain is shown in Table 8, where a positive value indicates a better performance through cross-learning. For small datasets, the enhancement of R^2 is significant (up to a LAEI dataset size of 40 for AutoML and 60 for RFOstochastic) with $p < 0.05$ (p-values are shown in the supplementary material). For bigger LAEI subsets, the performance gain is small and not significant. Especially for small datasets, cross-learning on multiple cities provides an opportunity to improve the model performance. Interestingly, for AutoML there is also a significant, however small ($\Delta R^2 \approx 0.01$), improvement for 1000 LAEI samples. The model enhancement through cross-learning on two cities gives a glimpse of the potential of our globally available features.

Table 7
R^2 of both models, RFOstochastic and AutoML, applied on subsets of the LAEI dataset. This shows, that OpenLUR performs well, when a sufficient amount of training data is available.

Number of LAEI samples	AutoML	RFOstochastic
20	-0.08	0.01
40	0.21	0.13
60	0.29	0.27
80	0.39	0.41
100	0.42	0.41
150	0.48	0.49
200	0.51	0.51
300	0.54	0.53
400	0.58	0.57
500	0.59	0.58
1000	0.63	0.63

Table 8
R^2 performance gain (difference of the R^2 of models trained on both datasets and the R^2 of models trained exclusively on LAEI data). Positive values show a performance gain through cross learning. Especially small datasets can significantly take advantage of cross learning.

Number of LAEI samples	AutoML	RFOstochastic
20	0.29	0.21
40	0.10	0.16
60	0.05	**0.10**
80	0.04	0.00
100	0.02	0.02
150	0.02	0.00
200	-0.01	0.00
300	0.01	0.00
400	0.00	0.01
500	0.00	0.01
1000	**0.01**	0.01

5.4. Recommended off-the-shelf approach for predicting air pollution

For our OpenLUR approach, we recommend a combination of our novel OSM features and the AutoML model. With regard to the features, this recommendation is justified by the fact that the OSM features significantly outperformed the baseline features (Hasenfratz et al., 2014) on the given dataset and have the advantages of being openly and globally available. The global availability enables land use regression even for small datasets through cross-learning. With regard to the underlying model, the choice is less clear. While stochastically optimized random forests have a slight advantage with regard to the mean rank, we still recommend AutoML for its more sophisticated hyper-parameter tuning based on Gaussian processes which seems to yield a better prediction performance on larger datasets. The gobal applicability of OpenLUR facilitates through cross-learning LUR studies on small datasets and could furthermore be used for multi-city or even global scale LUR research. Thus, with its novel openly and globally available feature set in combination with the notion of automated hyper-parameter tuning to eliminate tedious parameter studies, OpenLUR provides a reproducible, easily and widely applicable off-the-shelf land-use regression approach for air quality prediction even for small datasets that does not require expert knowledge in machine learning.

6. Future work and implications

In this section, we discuss several directions of future work as well as important implications of OpenLUR. In particular, (i) we discuss further potential features, approaches and models, (ii) we list some limitations of our dataset, and (iii) review the possibilities and limitations of globally available land use regression features.

6.1. Features and machine learning methods

While we have introduced an air quality regression pipeline that outperforms previously proposed methods, our novel features as well as our applied models can be further refined and extended:

First, there is an endless amount of features to incorporate into air quality models: For example, based on OSM data, information about crossings, parks or specific venues like shops or sights has not been explored yet. Also, besides static land usage features, traffic models (Krauβ, 1998; Smith, 1993), open weather data or wind flow models could account for time-dependencies. However, the openly available code of OpenLUR provides the possibility to add custom features without restrictions to type or origin.

Second, other machine learning algorithms than those covered in this work can be considered. In the supplementary material, we present some additional experiments. Concretely, we optimize random forests with auto-sklearn using Bayesian as well as stochastic optimization in combination with ensemble learning, but none was able to consistently outperform AutoML and RFOstochastic. Additionally, we evaluated two more recently applied methods as baselines: geographically weighted regression (Alam and McNabola, 2015) and feed forward neural networks (Hu et al., 2013) (results in the supplementary materials). However, with the limited amount of 200 datapoints, these models did not perform well. Nevertheless, models based on neural networks may be interesting to explore, as they may be able to alleviate the issue of deriving specific features from the OSM attributes by directly providing raw OSM data.

Finally a disadvantage of most nonlinear state-of-the-art machine learning models as used in OpenLUR is the more difficult interpretability: Unlike simple linear regression models, the influence of an independent variable is not measured by a single value — the respective slope — but is hidden in more complex model structures. Research to interpret these models has been conducted (Palczewska et al., 2014; Fabris et al., 2018). This is however out of the scope for this study and will be treated in future work.

6.2. Dataset and measurements

A crucial point for developing air quality regression models is the quality and quantity of measurement data. While traditional studies used stationary monitoring devices that resulted in a small number of datapoints (Montagne et al., 2015; Ragettli et al., 2014), recent studies show the potential of large amounts of mobile measurements. Mobile devices however are usually prone to inaccurate measurements and noise. To counteract short-term disturbances like bypassing trucks or simply wind, the measurements mostly get aggregated temporally or spatially which, analogously to the static case, results in fewer datapoints. Nevertheless, at least for our dataset, the spatial coverage was still a lot larger than using a handful of static devices. In future studies, it may be of interest to directly compare the quality of continuous mobile measurements with static approaches. This will require potentially very expensive, large scale measurement campaigns.

6.3. Towards global land use regression models

Because of the openly and globally available features, the models produced by OpenLUR can be applied in any city with comparable OSM data. With the LAEI data we have shown the global applicability as well as the ability for cross-learning.

But locally differing characteristics of cities or the underlying OSM data (Davidovic et al., 2016), e.g., caused by structural difference of cities in different countries, conceptually differing ways of providing data within local OSM communities, or the general quality of the provided information, can lead to different dependencies of pollutants on the features.

This points to the scientifically highly interesting area of the effects of local air pollution environments as well as the characteristics of area-specific OSM data. Nevertheless, our results are promising and present an important step towards generalized global air pollution models.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Florian Lautenschlager: Conceptualization, Methodology, Software, Formal analysis, Writing - original draft. **Martin Becker:** Writing - review & editing, Conceptualization, Methodology, Funding acquisition, Supervision. **Konstantin Kobs:** Writing - review & editing, Visualization. **Michael Steininger:** Writing - review & editing. **Padraig Davidson:** Writing - review & editing. **Anna Krause:** Writing - review & editing, Supervision. **Andreas Hotho:** Writing - review & editing, Conceptualization, Funding acquisition, Supervision.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.atmosenv.2020.117535.

References

Aguilera, I., Sunyer, J., Fernández-Patier, R., Hoek, G., Aguirre-Alfaro, A., Meliefste, K., Bomboi-Mingarro, M.T., Nieuwenhuijsen, M.J., Herce-Garraleta, D., Brunekreef, B., 2007. Estimation of outdoor no x, no2, and btex exposure in a cohort of pregnant women using land use regression modeling. Environ. Sci. Technol. 42, 815–821.

Alam, M.S., McNabola, A., 2015. Exploring the modeling of spatiotemporal variations in ambient air pollution within the land use regression framework: estimation of pm10 concentrations on a daily basis. J. Air Waste Manag. Assoc. 65, 628–640.

Arain, M., Blair, R., Finkelstein, N., Brook, J., Sahsuvaroglu, T., Beckerman, B., Zhang, L., Jerrett, M., 2007. The use of wind fields in a land use regression model to predict air pollution concentrations for health exposure studies. Atmos. Environ. 41, 3453–3464. https://doi.org/10.1016/j.atmosenv.2006.11.063. http://www.sciencedirect.com/science/article/pii/S1352231006013008.

Araki, S., Shima, M., Yamamoto, K., 2018. Spatiotemporal land use random forest model for estimating metropolitan no2 exposure in Japan. Sci. Total Environ. 634, 1269–1277. https://doi.org/10.1016/j.scitotenv.2018.03.324. http://www.sciencedirect.com/science/article/pii/S0048969718310842.

Basu, B., Alam, M.S., Ghosh, B., Gill, L., McNabola, A., 2019. Augmenting limited background monitoring data for improved performance in land use regression modelling: using support vector regression and mobile monitoring. Atmos. Environ. 201, 310–322. https://doi.org/10.1016/j.atmosenv.2018.12.048. http://www.sciencedirect.com/science/article/pii/S1352231019300159.

Bergstra, J., Bengio, Y., 2012. Random search for hyper-parameter optimization. J. Mach. Learn. Res. 13, 281–305. http://www.jmlr.org/papers/v13/bergstra12a.html.

Blum, M., Feurer, M., Klein, A., Springenberg, J., Hutter, F., Eggensperger, K., 2015. Efficient and Robust Automated Machine Learning. NIPS, pp. 2962–2970. http://papers.nips.cc/paper/5872-efficient-and-robust-automated-machine-learning.

Brauer, M., Hoek, G., van Vliet, P., Meliefste, K., Fischer, P., Gehring, U., Heinrich, J., Cyrys, J., Bellander, T., Lewne, M., Brunekreef, B., 2003. Estimating long-term average particulate air pollution concentrations: application of traffic indicators and geographic information systems. Epidemiology 14, 228–239. https://doi.org/10.1097/01.EDE.0000041910.49046.9B.

Breiman, L., 2001. Random forests. Mach. Learn. 45, 5–32. https://doi.org/10.1023/A:1010933404324 doi:10.1023/A:1010933404324.

Briggs, D.J., de Hoogh, C., Gulliver, J., Wills, J., Elliott, P., Kingham, S., Smallbone, K., 2000. A regression-based method for mapping traffic-related air pollution: application and testing in four contrasting urban environments. Sci. Total Environ. 253, 151–167. https://doi.org/10.1016/S0048-9697(00)00429-0. http://www.sciencedirect.com/science/article/pii/S0048969700004290.

Brokamp, C., Jandarov, R., Rao, M., LeMasters, G., Ryan, P., 2017. Exposure assessment models for elemental components of particulate matter in an urban environment: a comparison of regression and random forest approaches. Atmos. Environ. 151, 1–11.

Brook, R.D., Rajagopalan, S., Pope, C.A., Brook, J.R., Bhatnagar, A., Diez-Roux, A.V., Holguin, F., Hong, Y., Luepker, R.V., Mittleman, M.A., 2010. Particulate matter air pollution and cardiovascular disease. Circulation 121, 2331–2378.

Carr, D., von Ehrenstein, O., Weiland, S., Wagner, C., Wellie, O., Nicolai, T., von Mutius, E., 2002. Modeling annual benzene, toluene, no2, and soot concentrations on the basis of road traffic characteristics. Environ. Res. 90, 111–118. https://doi.org/10.1006/enrs.2002.4393. http://www.sciencedirect.com/science/article/pii/S0013935102943938.

Champendal, A., Kanevski, M., Huguenot, P.E., 2014. Air pollution mapping using nonlinear land use regression models. In: International Conference on Computational Science and its Applications. Springer, pp. 682–690.

Davidovic, N., Mooney, P., Stoimenov, L., Minghini, M., 2016. Tagging in volunteered geographic information: an analysis of tagging practices for cities and urban regions in openstreetmap. ISPRS Int. J. Geo-Inf. 5, 232.

Dons, E., Poppel, M.V., Kochan, B., Wets, G., Panis, L.I., 2013. Modeling temporal and spatial variability of traffic-related air pollution: hourly land use regression models for black carbon. Atmos. Environ. 74, 237–246. https://doi.org/10.1016/j.atmosenv.2013.03.050. http://www.sciencedirect.com/science/article/pii/S1352231013002264.

Estima, J., Painho, M., 2015. Investigating the potential of openstreetmap for land use/land cover production: a case study for continental Portugal. In: OpenStreetMap in GIScience. Springer, pp. 273–293.

Fabris, F., Doherty, A., Palmer, D., De Magalhães, J.P., Freitas, A.A., 2018. A new approach for interpreting random forest models and its application to the biology of ageing. Bioinformatics 34, 2449–2456.

Greater London Authority, A.Q.T., 2016. London atmospheric emissions inventory (laei) 2013. https://data.london.gov.uk/dataset/london-atmospheric-emissions-inventory-2013.

Haklay, M., 2010. How good is volunteered geographic information? a comparative study of openstreetmap and ordnance survey datasets. Environ. Plann. Plann. Des. 37, 682–703.

Haklay, M., Weber, P., 2008. Openstreetmap: user-generated street maps. IEEE Pervasive Computing 7, 12–18.

Hankey, S., Marshall, J.D., 2015. Land use regression models of on-road particulate air pollution (particle number, black carbon, pm2.5, particle size) using mobile monitoring. Environ. Sci. Technol. 49, 9194–9202. https://doi.org/10.1021/acs.est.5b01209 pMID: 26134458.

Hasenfratz, D., Saukh, O., Walser, C., Hueglin, C., Fierz, M., Thiele, L., 2014. Pushing the spatio-temporal resolution limit of urban air pollution maps. In: 2014 IEEE International Conference on Pervasive Computing and Communications. PerCom, pp. 69–77. https://doi.org/10.1109/PerCom.2014.6813946.

Hastie, T., Tibshirani, R., 1986. Generalized additive models. Stat. Sci. 1, 297–310. https://doi.org/10.1214/ss/1177013604 doi:10.1214/ss/1177013604.

Hecht, R., Kunze, C., Hahmann, S., 2013. Measuring completeness of building footprints in openstreetmap over space and time. ISPRS Int. J. Geo-Inf. 2, 1066–1091.

Henderson, S.B., Beckerman, B., Jerrett, M., Brauer, M., 2007. Application of land use regression to estimate long-term concentrations of traffic-related nitrogen oxides and fine particulate matter. Environ. Sci. Technol. 41, 2422–2428. https://doi.org/10.1021/es0606780 pMID: 17438795.

Hentschel, M., Wagner, B., 2010. Autonomous robot navigation based on openstreetmap geodata. In: Intelligent Transportation Systems (ITSC), 2010 13th International IEEE Conference on, IEEE, pp. 1645–1650.

Heymann, Y., 1994. CORINE Land Cover: Technical Guide. Office for Official Publ. of the Europ. Communities.

Hoek, G., Meliefste, K., Brauer, M., van Vliet, P., Brunekreef, B., Fischer, P., Lebret, E., Cyrys, J., Gehring, U., Heinrich, A., 2001. Risk Assessment of Exposure to Traffic-Related Air Pollution for the Development of Inhalant Allergy, Asthma and Other Chronic Respiratory Conditions in Children (Trapca). IRAS University, Utrecht. Final Report.

Hu, X., Waller, L.A., Al-Hamdan, M.Z., Crosson, W.L., Estes Jr., M.G., Estes, S.M., Quattrochi, D.A., Sarnat, J.A., Liu, Y., 2013. Estimating ground-level pm2. 5 concentrations in the southeastern us using geographically weighted regression. Environ. Res. 121, 1–10.

Jerrett, M., Arain, A., Kanaroglou, P., Beckerman, B., Potoglou, D., Sahsuvaroglu, T., Morrison, J., Giovis, C., 2004. A review and evaluation of intraurban air pollution exposure models. J. Expo. Anal. Environ. Epidemiol. 15, 185–204. https://doi.org/10.1038/sj.jea.7500388.

Krauß, S., 1998. Microscopic Modeling of Traffic Flow: Investigation of Collision Free Vehicle Dynamics (Ph.D. thesis).

Larson, T., Henderson, S.B., Brauer, M., 2009. Mobile monitoring of particle light absorption coefficient in an urban area as a basis for land use regression. Environ. Sci. Technol. 43, 4672–4678. https://doi.org/10.1021/es803068e pMID: 19673250.

Luxen, D., Vetter, C., 2011. Real-time routing with openstreetmap data. In: Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems. ACM, pp. 513–516.

Minet, L., Gehr, R., Hatzopoulou, M., 2017. Capturing the sensitivity of land-use regression models to short-term mobile monitoring campaigns using air pollution micro-sensors. Environ. Pollut. 230, 280–290. https://doi.org/10.1016/j.envpol.2017.06.071. http://www.sciencedirect.com/science/article/pii/S0269749117309612.

Montagne, D.R., Hoek, G., Klompmaker, J.O., Wang, M., Meliefste, K., Brunekreef, B., 2015. Land use regression models for ultrafine particles and black carbon based on short-term monitoring predict past spatial variation. Environ. Sci. Technol. 49, 8712–8720. https://doi.org/10.1021/es505791g pMID: 26079151.

Muttoo, S., Ramsay, L., Brunekreef, B., Beelen, R., Meliefste, K., Naidoo, R.N., 2018. Land use regression modelling estimating nitrogen oxides exposure in industrial south durban, South Africa. Sci. Total Environ. https://doi.org/10.1016/j.scitotenv.2017.07.278, 610-611, 1439 – 1447. http://www.sciencedirect.com/science/article/pii/S0048969717319988.

OpenStreetMap contributors, 2017. Planet dump retrieved from. https://planet.osm.org. https://www.openstreetmap.org.

Palczewska, A., Palczewski, J., Marchese Robinson, R., Neagu, D., 2014. Interpreting Random Forest Classification Models Using a Feature Contribution Method. Springer International Publishing, Cham, pp. 193–218. https://doi.org/10.1007/978-3-319-04717-1_9 doi:10.1007/978-3-319-04717-1_9.

Patton, A.P., Collins, C., Naumova, E.N., Zamore, W., Brugge, D., Durant, J.L., 2014. An hourly regression model for ultrafine particles in a near-highway urban area. Environ. Sci. Technol. 48, 3272–3280. https://doi.org/10.1021/es404838k pMID: 24559198.

Polichetti, G., Cocco, S., Spinali, A., Trimarco, V., Nunziata, A., 2009. Effects of particulate matter (pm10, pm2.5 and pm1) on the cardiovascular system. Toxicology 261, 1–8. https://doi.org/10.1016/j.tox.2009.04.035. http://www.sciencedirect.com/science/article/pii/S0300483X09002121.

Pope, C.A., Dockery, D.W., Spengler, J.D., Raizenne, M.E., 1991. Respiratory health and pm10 pollution: a daily time series analysis. Am. Rev. Respir. Dis. 144, 668–674. https://doi.org/10.1164/ajrccm/144.3_Pt_1.668 doi:10.1164/ajrccm/144.3_Pt_1.668.

Ragettli, M.S., Ducret-Stich, R.E., Foraster, M., Morelli, X., Aguilera, I., Basagaña, X., Corradi, E., Ineichen, A., Tsai, M.Y., Probst-Hensch, N., Rivera, M., Slama, R., Künzli, N., Phuleria, H.C., 2014. Spatio-temporal variation of urban ultrafine particle number concentrations. Atmos. Environ. 96, 275–283. https://doi.org/10.1016/j.atmosenv.2014.07.049. http://www.sciencedirect.com/science/article/pii/S1352231014005792.

Rahman, K.M., Alam, T., Chowdhury, M., 2012. Location based early disaster warning and evacuation system on mobile phones using openstreetmap. In: Open Systems (ICOS), 2012 IEEE Conference on, IEEE, pp. 1–6.

Sahsuvaroglu, T., Arain, A., Kanaroglou, P., Finkelstein, N., Newbold, B., Jerrett, M., Beckerman, B., Brook, J., Finkelstein, M., Gilbert, N.L., 2006. A land use regression model for predicting ambient concentrations of nitrogen dioxide in hamilton, ontario, Canada. J. Air Waste Manag. Assoc. 56, 1059–1069. URL:/brokenurl#. https://doi.org/10.1080/10473289.2006.10464542.

Schelhorn, S.J., Herfort, B., Leiner, R., Zipf, A., Albuquerque, J.P.d., 2014. Identifying elements at risk from openstreetmap: the case of flooding. In: International Conference on Information Systems for Crisis Response and Management, vol. 11. The Pennsylvania State University.

Shi, Y., Lau, K.K.L., Ng, E., 2016. Developing street-level pm2.5 and pm10 land use regression models in high-density Hong Kong with urban morphological factors.

Environ. Sci. Technol. 50, 8178–8187. https://doi.org/10.1021/acs.est.6b01807 pMID: 27381187.

Smith, M., 1993. A new dynamic traffic model and the existence and calculation of dynamic user equilibria on congested capacity-constrained road networks. Transp. Res. Part B Methodol. 27, 49–63.

Stafoggia, M., Bellander, T., Bucci, S., Davoli, M., de Hoogh, K., de' Donato, F., Gariazzo, C., Lyapustin, A., Michelozzi, P., Renzi, M., Scortichini, M., Shtein, A., Viegi, G., Kloog, I., Schwartz, J., 2019. Estimation of daily pm10 and pm2.5 concentrations in Italy, 2013–2015, using a spatiotemporal land-use random-forest model. Environ. Int. 124, 170–179. https://doi.org/10.1016/j.envint.2019.01.016. http://www.sciencedirect.com/science/article/pii/S0160412018327685.

Su, J., Jerrett, M., Beckerman, B., 2009. A distance-decay variable selection strategy for land use regression modeling of ambient air pollution exposures. Sci. Total Environ. 407, 3890–3898. https://doi.org/10.1016/j.scitotenv.2009.01.061. thematic Issue - BioMicroWorld Conference. http://www.sciencedirect.com/science/article/pii/S0048969709001442.

Su, J.G., Hopke, P.K., Tian, Y., Baldwin, N., Thurston, S.W., Evans, K., Rich, D.Q., 2015. Modeling particulate matter concentrations measured through mobile monitoring in

a deletion/substitution/addition approach. Atmos. Environ. 122, 477–483. https://doi.org/10.1016/j.atmosenv.2015.10.002. http://www.sciencedirect.com/science/article/pii/S1352231015304210.

Sîrbu, A., Becker, M., Caminiti, S., De Baets, B., Elen, B., Francis, L., Gravino, P., Hotho, A., Ingarra, S., Loreto, V., Molino, A., Mueller, J., Peters, J., Ricchiuti, F., Saracino, F., Servedio, V.D.P., Stumme, G., Theunis, J., Tria, F., Van den Bossche, J., 2015. Participatory patterns in an international air quality monitoring initiative. PloS One 10, 1–19. https://doi.org/10.1371/journal.pone.0136763 doi:10.1371/journal.pone.0136763.

Taylor, R., 1990. Interpretation of the correlation coefficient: a basic review. J. Diagn. Med. Sonogr. 6, 35–39. https://doi.org/10.1177/875647939000600106 doi:10.1177/875647939000600106.

Zwack, L.M., Paciorek, C.J., Spengler, J.D., Levy, J.I., 2011. Characterizing local traffic contributions to particulate air pollution in street canyons using mobile monitoring techniques. Atmos. Environ. 45, 2507–2514. https://doi.org/10.1016/j.atmosenv.2011.02.035. http://www.sciencedirect.com/science/article/pii/S1352231011001695.