A Data-driven System for City-wide Energy Footprinting and Apportionment

PETER WEI and XIAOFAN JIANG, Columbia University

11

Energy footprinting has the potential to raise awareness of energy consumption and lead to energy-saving behavior. However, current methods are largely restricted to single buildings; these methods require energy and occupancy monitoring sensor deployments, which can be expensive and difficult to deploy at scale. Further, current methods for estimating energy consumption and population at scale cannot provide fine enough temporal or spatial granularity for a reasonable personal energy footprint estimate. In this work, we present a data-driven system for city-wide estimation of personal energy footprints. This system takes advantage of existing sensing infrastructure and data sources in urban cities to provide energy and population estimates at the building level, even in built environments that do not have existing or accessible energy or population data.

CCS Concepts: • **Hardware** → **Power and energy**; • **Computer systems organization** → *Sensor networks*; *Real-time systems*;

Additional Key Words and Phrases: Energy footprinting, real-time system, population estimation

ACM Reference format:

Peter Wei and Xiaofan Jiang. 2021. A Data-driven System for City-wide Energy Footprinting and Apportionment. *ACM Trans. Sen. Netw.* 17, 2, Article 11 (January 2021), 24 pages.

<https://doi.org/10.1145/3433639>

1 INTRODUCTION

In urban cities such as New York City, buildings are responsible for up to 75% of total greenhouse gas emissions [43] and over 90% of total benchmarked energy consumption [10]. A significant portion of the energy consumed directly services humans in retail, commercial, and residential buildings. In addition to buildings, transportation is also responsible for large amounts of energy consumption. As sustainability increasingly becomes an important factor in modern society, energy consumption in the built environment is one area where reduction can have a major impact.

Parts of this article have been previously published in the Proceedings of the 6th ACM International Conference on Systems for Energy-efficient Built Environments.

This research was partially supported by the National Science Foundation under Grant Numbers CNS-1704899, CNS-1815274, CNS-11943396, and CNS-1837022. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Columbia University, NSF, or the U.S. Government or any of its agencies.

Authors' addresses: P. Wei and X. Jiang, Department of Electrical Engineering, Columbia University, 500 W. 120th St., Mudd 1310, New York, NY 10027; emails: wei.peter@columbia.edu, jiang@ee.columbia.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1550-4859/2021/01-ART11 \$15.00

<https://doi.org/10.1145/3433639>

Fig. 1. Top: Energy and population are estimated at the building level; personal energy footprints are provided to each person, depending on their location traces (red and blue traces). Bottom: Two real city-wide energy footprints. Labels below denote energy source.

In Reference [12], the authors show that certain feedback mechanisms given to occupants raises awareness and can lead to energy-saving behavior. One such mechanism is by notifying the occupant in real-time of their numerical energy consumption, or energy footprint. However, many areas of the built environment do not have the capabilities to measure personal energy consumption, much less notify people of their personal energy responsibility. Although companies such as Nest and recent research studies have begun to address this challenge through *energy footprinting* in single residential and commercial buildings, there are still many difficulties to scaling these solutions to the city-scale. Most notably, there is a lack of energy and population data with high temporal and spatial granularity. Without sensors to measure this data, or models to estimate this data, energy footprinting is not possible.

In this work, we present *CityEnergy*, a scalable, real-time system for computing personal energy footprint estimates. As shown in Figure 1, *CityEnergy* estimates the energy consumption and population for each building in real-time; these estimates are used to calculate the per capita energy footprint at the building level. *CityEnergy* is able to provide personal energy footprints to people who provide their location data through a mobile application.

An example city-wide personal energy footprint generated by *CityEnergy* is shown in Figure 1. In this example, a *CityEnergy* user, Stephen, (A) begins the day in a residential building. *CityEnergy* uses a detailed model of the particular built environment Stephen is in and computes the energy and occupancy level to estimate the individual energy that Stephen is responsible for.

Stephen's commute consists of (B1) walking to the subway station, (B2) riding the subway downtown, and (B3) walking from the subway station to the office building; for each of the modes of transportation, *CityEnergy* associates the relevant energy consumption to Stephen's personal energy footprint. After arriving at the office, (C) Stephen works until the lunch break. During this period, *CityEnergy* may interface with the local energy footprinting system to determine Stephen's personal footprint, or rely on energy and occupancy models to estimate Stephen's footprint. Finally, Stephen leaves the office and (D) walks to the local deli, where he spends an hour to eat lunch.

In this work, we present the following contributions:

- We create an energy and occupancy *digital twin* of the city, with a focus on major aspects of the built environment including buildings and transportation.
- We present the design and implementation of *CityEnergy*, a **city-scale** energy footprinting system that utilizes the city's digital twin to provide **real-time energy footprints** with a focus on 100% coverage.
- We deploy *CityEnergy* in New York City, utilize local data sources to develop energy and population models, and collect and evaluate the accuracy of real-world personal energy footprint data.
- We have developed a number of tools and applications such as mobile and web applications to provide citizens with insights into their everyday energy consumption and city planners with important information at the city-scale.

CityEnergy is a tool for *estimating* an individual's energy footprint at any location in the city, at any time. Due to the data sources available, the energy footprint estimate may not be accurate for any one person. However, we believe that *CityEnergy* is an important step towards realizing city-wide energy footprinting.

2 RELATED WORKS

There are a number of recent works addressing topics such as building energy consumption estimation, population estimation, transportation detection, and energy footprinting.

Predicting energy consumption of a building normally constitutes one of three approaches: building level regression, software modeling, or city-wide energy estimation. For building level regression, historical energy consumption data is typically collected with fine temporal granularity (minute to hourly frequency). Once the data is collected, a regression model [20, 22, 54, 58] or neural network [20, 23, 54] is trained on the data. After training, the model is validated on a different dataset from the same building. These models tend to produce low error rates (<10 MAE) but require a large amount of data. At a city-scale level, these methods are impractical without data already collected from a major entity, such as a government project.

Another method for predicting energy consumption in a building is through software modeling. EnergyPlus [11] is a popular program for simulating energy consumption in custom buildings under various internal and external conditions. References [33, 46] are two recent works that utilize EnergyPlus to model specific buildings. EnergyPlus is a complex program that requires careful modeling of the building to provide an accurate estimate; for larger number of buildings, this becomes increasingly difficult to scale. Recent efforts have focused on automating the modeling of buildings to enable city-wide modeling such as CityBES [7], which uses GIS data and CBES [18] to automatically generate and run baseline models, and Reference [4], which utilizes UMI [40] to generate models of over 1,200 buildings in Boston. In future works, these models may be used instead, or combined with this work to provide better estimates for energy and population at the building level.

Recently, a number of studies have explored energy estimation from city-wide datasets. In Rotterdam, Reference [30] utilized a Geographical Information System (GIS) to “downscale” the energy consumption in the city to individual residential buildings. A multiple linear regression model was able to achieve a Mean Absolute Percentage Error (MAPE) of 9% for electricity prediction. Similarly, References [26] and [19] utilize city-level energy consumption data in New York City, along with regression models to predict energy consumption at the block level and building level. *CityEnergy* builds on these works by incorporating population estimates to compute personal energy footprints.

There are additionally a number of methods for estimating population in a city. A recent study [29] proposed a model for estimating dynamic populations of communities from subway smart card data. In Reference [45], the authors study flow from multiple transportation modalities using different machine learning models. *CityEnergy* utilizes similar broad ideas to estimate dynamic populations from subway and vehicle transportation modalities, but further develops the ideas towards a more granular population estimate.

Transportation mode recognition using mobile devices is a well-studied problem and primarily relies on GPS, accelerometer, and gyroscope sensor data. Examples of studies using mobile sensor data include References [42, 59], which only utilize GPS data and classifiers such as decision trees to determine the mode of transportation; References [17, 21], which utilize accelerometer and gyroscope data along with support vector machines; and Reference [53], which utilizes GPS, accelerometer, and gyroscope data. While *CityEnergy* incorporates transportation mode detection, we do not claim any novelty in this area.

Recently, the field of personal energy footprinting has begun to grow, with a number of studies combining energy monitoring with different protocols [13], occupant localization, and energy apportionment to estimate real-time personal energy footprints. In residential homes and apartment buildings, References [27] and [44] deployed energy monitoring sensors to detect energy consumption of each occupant. For commercial buildings, References [8, 47, 48] presented a scalable system to apportion energy consumption using different policies; Reference [52] expanded on this system to additionally deliver energy-saving recommendations. However, the main difficulties in scaling these works to multiple buildings is the *cost of deployment*. Most buildings require retrofitting of energy and occupant monitoring systems to enable energy footprinting. Building on our previous work [49], *CityEnergy* is the first city-scale energy footprinting system that does not require such monitoring systems.

3 CHALLENGES

The goal of *CityEnergy* is to provide “full coverage” for a person’s energy footprint. Coverage in this application broadly refers to the percentage of time and locations for which our system can provide a reasonable personal energy footprint estimate. We noted three desirable characteristics to maximize coverage: high spatial granularity, high temporal granularity, and high accuracy. In the design of *CityEnergy*, the two most critical components that must adhere to these characteristics are energy estimation and population estimation.

3.1 Energy Estimation

To enable a real-time energy footprint estimate, a system must be capable of producing an energy consumption estimate for any building at any time of day. The main challenge is in achieving *both high temporal granularity as well as high spatial granularity without sacrificing accuracy*. We demonstrate this challenge by comparing two types of energy estimation methods frequently seen in literature.

Building level regression models, which use historical data of a building to predict energy consumption, can achieve high accuracy for different levels of temporal granularity [20, 22, 54, 58]. However, energy data at sub-daily or sub-hourly time scales is difficult to obtain for even a single building, much less at the city-scale. There are two main obstacles to extending this method to the city-scale. First, fine temporal energy data is often proprietary and cannot be easily accessed through public datasets. Second, this data is often obtained for specific purposes, meaning that most buildings will not have accessible historical energy data. Thus, while building level regression models are a good option for localized energy consumption studies, they are difficult to generalize to the city-scale.

City-wide energy estimation, which uses data containing energy consumption of a sample of buildings, has been used to estimate energy consumption at the city-scale [19, 26, 30]. However, the data is often collected at the yearly or monthly scale and thus does not achieve sufficient temporal granularity. The main obstacle to “downscaling” the monthly or yearly energy consumption data to a finer temporal granularity is the lack of knowledge about the specific behavior of the building for different environmental conditions (weather, time of day, day of week, etc.). City-wide energy estimation is able to achieve the necessary spatial granularity but requires additional information to achieve temporal granularity.

3.2 Population Estimation

Similarly to energy estimation, a system must be capable of producing a population estimate with high spatial and temporal granularity. The most common form of population estimation is from mobility models such as References [24, 31, 55, 56, 59, 60]. However, these models typically require GPS traces or cellular information from base stations, which are often not accessible by the public due to privacy concerns. Another method of population estimation is by studying the different types of transportation and inferring dynamic population.

According to Reference [6], of the 4.8M commuters in New York City, 38.7% ride the subway and 26.9% drive in vehicles. Thus, by estimating these two modalities of transportation, 67% of the dynamic population due to commuting can be accounted for. The remaining challenges are how to estimate each of these modes of transportation and how these estimates translate to dynamic population.

Each of these modes presents unique challenges in both spatial and temporal granularity. For subways in New York City, there exists historical data at four-hour intervals describing the number of people exiting and entering the station through the turnstiles [3]; however, finer temporal granularity, as well as the changes in population of the surrounding area, need to be modeled.

For motor vehicles, New York City traffic cameras can be utilized to determine the density of vehicles [37]. Three challenges are how to determine density of vehicles from traffic cameras, how to infer vehicle density of unobserved roads, and translating vehicle density to dynamic population.

4 CITY-WIDE ENERGY ESTIMATION

4.1 Data Sources

As discussed in Section 3, one major challenge to estimating energy consumption is a lack of data at a sufficiently granular level. In New York City, there are no public datasets of energy consumption at the building level. This poses a problem, as many energy estimation techniques rely on such data.

However, in Reference [26], the authors demonstrate a method for estimating the energy use intensity of any building in New York City by training a machine learning model on the energy benchmarking dataset from New York City’s Local Law 84 [25]; this type of dataset is also collected in other cities with benchmarking laws such as Los Angeles and San Francisco. This dataset, however, only includes buildings greater than 50,000 square feet. One disadvantage of this dataset is that due to the size of the buildings benchmarked, most of the included buildings are commercial; this excludes residential buildings.

We chose to further incorporate an energy dataset of residential buildings from the New York City Housing Authority (NYCHA) [34]. This dataset contains energy consumption data of over 2,400 residential buildings at a monthly scale and is home to 1 in 14 New Yorkers. We reasoned that including a residential dataset would increase the overall model accuracy for energy estimation.

In addition to the Local Law 84 and NYCHA datasets, we also incorporated hourly energy traces from the Department of Energy’s (DOE) reference buildings [14]. This dataset contains baseline

Fig. 2. Diagram of the energy estimation pipeline.

energy consumption values for 16 building types covering commercial and residential buildings. We utilize this dataset to drill down from monthly to hourly energy estimation.

Finally, to account for out-of-sample buildings, we utilize features derived from the Primary Land Use Tax Lot Output (PLUTO) dataset [35]. This dataset contains numerous features of every building in New York City, including building age, gross floor area, and floor area for different use types. We derive many of the input parameters of our energy consumption model from the PLUTO dataset.

4.2 Energy Consumption Model

As discussed in Section 3.1, an energy consumption model should have high spatial and temporal granularity. To achieve this, we developed a two-stage model, as shown in Figure 2. In the first stage, a regression model is trained on monthly energy consumption data; this enables prediction of energy consumption of individual buildings. The second stage is a fitting model to downscale monthly energy consumption to hourly energy consumption.

To enable monthly energy estimation of any building in New York City, we trained a machine learning model on energy consumption data from the Local Law 84 and NYCHA datasets. As inputs to the model, we extracted parameters from the PLUTO dataset that represented important features of each building. As described in Reference [26], features such as gross floor area, year built, office floor area, residential floor area, and borough are the best predictors of energy consumption. In addition to these parameters, we also included weather data, office floor area, retail floor area, garage floor area, and factory floor area.

The first stage outputs a monthly energy consumption estimate for each building; however, *CityEnergy* also requires fine time granularity. To transform monthly energy consumption to hourly energy consumption, we construct a model building based on the Department of Energy Commercial Reference Buildings [14]. The Reference Buildings consist of 16 building types and provide hourly energy simulations in EnergyPlus.

For each building in New York City, the PLUTO dataset provides the floor areas by usage type; these roughly correspond to common building types in the DOE Reference Buildings. The model for each building is constructed via a linear combination of the Reference Building energy traces, with the weights corresponding to the percentage floor area of each type of building. An example is provided in Figure 3(a); a New York City building that is 50% apartment floor area, 30% office floor area, and 20% retail floor area is modeled by combining the respective percentages of the DOE Reference Buildings energy traces.

Finally, to incorporate the monthly energy consumption estimate from the first stage, we scale the model trace to equal the monthly energy consumption estimate. When a client requests an energy footprint of the building, the hourly energy footprint of the building is used in the

(a) One example model building combination from the DOE Reference Buildings using weights corresponding to floor area for this building.

(b) Heatmap of average building energy footprint in Manhattan by block.

Fig. 3. Energy estimation from regression (not shown) and reference building to energy heatmap.

computation of the personal energy footprint. An illustration of the average building energy footprint by block is shown in Figure 3(b). We evaluate the accuracy of this approach in Section 7.1.

5 CITY-WIDE POPULATION ESTIMATION

Another critical source of information for *CityEnergy* is real-time population estimation. An estimate of the number of people in a particular building, coupled with the energy estimate, is enough information to provide a real-time energy footprint. Unfortunately, there are no public datasets detailing population counts at any spatial granularity. In addition, the population varies throughout the day and week due to the large number of commuters; this further complicates modeling of real-time building population.

As stated in Section 3.2, 67% of dynamic population due to commuters can be estimated from subway ridership and driving. By combining dynamic population estimates from different modes of transportation, a real-time population estimate can be built. *CityEnergy* utilizes data from the U.S. Census as a baseline population model and uses real-time and historical data from the New York City Metropolitan Transportation Authority (MTA) and NYC traffic cameras to estimate dynamic population.

5.1 Baseline Population

A preliminary baseline model was constructed using the US Census dataset as initial populations for each block. As population data at the building level is not available from the Census, we down-scaled the block-level population to the building level through a simple model. We assume that people only reside in residential floor area; the residential floor area of each building in the block is collected from the PLUTO dataset. The estimated static population of building x can then be computed by Equation (1), as the ratio of the floor area of the building (FA_x) to the aggregate floor area of the census block times the census block population P_{CB} :

$$P_x = P_{CB} \frac{FA_x}{\sum_{B \in CB} FA_B}. \quad (1)$$

Fig. 4. Illustration of the dynamic population estimation pipeline.

Fig. 5. Illustration of the subway turnstile data and the subway throughput estimation.

5.2 Dynamic Population

To estimate real-time population dynamics, we analyze two of the most common modes of transportation in New York City: subway and motor vehicle. Estimates of the two modes of transportation are combined with Google Places data and population models to estimate dynamic population, as shown in Figure 4.

5.2.1 Subway. Every week, the New York City MTA publishes a new dataset detailing the commuter throughput for turnstiles in NYC subway stations at four-hour intervals [3]. Each subway station has multiple turnstiles, each of which records the number of commuters entering as well as exiting the subway station.

To estimate the number of people entering or exiting from a station, we first construct a station dataset by aggregating the inflow and outflow counts of the turnstiles at each station. We collect this time-series data for six weeks to train a regression model. The parameters used in the training include: historical commuter flow, hour, day of week, and weather conditions.

One consideration to improve the time granularity of inflow and outflow estimates is that commuters enter a subway station at more uniform frequencies than they exit. The reason is that exiting commuters usually correspond to arrivals of subway trains. The regression model can estimate commuter throughput for four-hour intervals; however, by using real-time data, the outflow estimates can be improved. To achieve this, *CityEnergy* utilizes the subway data feeds [2] from the NYC MTA, which provide the real-time locations of all running subway trains.

When a subway reaches a station, the number of commuters exiting the station is calculated by the regression estimate divided by the expected number of trains for the four-hour window. The estimated distribution of these commuters to the surrounding buildings is discussed in Section 5.3. An illustration of the throughput estimates and whole subway estimation pipeline is shown in Figure 5. Evaluation of different regression models for hourly throughput are presented in Section 7.2.1.

Fig. 6. Computer vision pipeline for detecting vehicles. (A) the background is masked from the incoming video stream image, and (B) an object detection network extracts the candidate bounding boxes.

5.2.2 Motor Vehicles. Besides the subway, the second most frequently used mode of transportation by commuters is by motor vehicle. For *CityEnergy*, we estimate motor vehicle commuting by analyzing real-time footage from traffic cameras, building on the work in Reference [50]. Real-time image feeds are publicly available from the New York City Department of Transportation, which provides 752 real-time traffic cameras covering major intersections [37].

The estimate of dynamic population from traffic cameras is set in two stages, as shown in Figure 7. First, the flow of motor vehicles at major intersections is estimated through detection and tracking computer vision methods. Second, the vehicle flow, along with a historical taxi dataset, is used to estimate population change at the block level.

The traffic camera streams present two main challenges: low resolution and low frame-rate. The stream resolution is 352×240 pixels, and are updated at a frequency between 0.3–1 Hz. Thus, to determine motor vehicle flow, we implement a vehicle detection pipeline composed of two parts: background masking and object detection, as shown in Figure 6. Initially, the background of the images of the video stream are masked as in Reference [57]. This ensures that low resolution patterns in the background will not be falsely detected as vehicles.

The masked image is then fed to an object detection network. To increase accuracy, we utilize training data from two sources: the CityCam dataset [57] and a custom hand-labeled dataset comprising 2,000 images and approximately 5,000 vehicles. We use transfer learning to tune a pre-trained SSD-Mobilenet model [28] to better recognize vehicles at low resolution. Transfer learning is achieved by freezing all layers, except for the final layer, and retraining the neural network. Once trained, the network provides bounding boxes of detected vehicles in each image frame.

Due to the low frame-rate, traditional computer vision methods for tracking are unsuitable due to the high displacement of fast-moving vehicles. Instead, specific features are extracted from each bounding box and matched in consecutive frames. We extract the same features as in Reference [50] into a feature vector, including the color histogram and the output from VGG16’s conv3_3 layer. In consecutive frames, we compute correlation scores between pairs of feature vectors. A final pairwise score is produced from a linear combination of the correlation scores. Last, the final matching is determined by solving the maximum weighted bipartite matching problem with the pairwise scores, as described in Reference [50]. This method enables tracking of vehicles across multiple frames and provides the vehicle flow across the major intersections.

Even with the large number of deployed traffic cameras, a majority of the streets remain unobserved. To provide full spatial coverage, we query the Here location framework API¹ (similar to the Google Maps API), which provides indications of traffic density on road segments not covered by traffic cameras.

¹<https://developer.here.com/>.

Fig. 7. Pipeline for determining dynamic population at the block level from traffic cameras.

Fig. 8. Estimated number of taxi trips through each zone (color), estimated percentage of zone source, and destination for taxi trips.

Vehicle flow, however, is not sufficient to predict dynamic population. For example, there are roads with high vehicle flow but low dynamic population, such as highways. Unfortunately, there is no available data describing the relationship between vehicle flow and population dynamics. As a proxy, we utilize the NYC Taxi & Limousine Commission dataset describing millions of taxi trips, pickups, and destinations [9], which is sufficient to give a frequency of transfer to specific areas.

In the TLC dataset, each taxi trip is described by a source and destination zone. However, taxi trips frequently traverse multiple zones between the source and destination zones; these taxis appear in the vehicle flow, but will not contribute to the flow of people in the intermediate zones. To better understand population transfer from vehicles, we estimate three parameters for each taxi zone. First, the total taxi count includes any taxi trip that passes through the zone and is a proxy for the total vehicle flow. Next, based on the sources and destinations of the taxi trips, we determine the percentage of total taxi trips that use the zone as the source or destination. We calculated these parameters for each taxi zone using six months of TLC data, as shown in Figure 8.

After processing these trips, we can estimate the real-time population dynamics of commuters from vehicles in the traffic cameras to the taxi zones by multiplying by the source and destination percentages, as shown in Figure 7.

Fig. 9. Illustration of catchment areas of subway stops in New York City.

Table 1. Distribution of Traveler Destinations by Mode of Transportation, Extracted from the NYC DOT Citywide Mobility Survey

	Office	School	Retail	Errand	Medical
Motor Vehicle	32.5%	4.9%	45.1%	14.1%	4.3%
Subway	61.6%	4.6%	21.9%	7.8%	4.0%

5.2.3 Google Places API. Google’s Places API² provides information about the popularity and the current estimated occupancy (“live” data) of many retail locations in New York City. When current estimated occupancy is available, this value is used as an estimate for the retail location; otherwise, the popularity is used.

5.3 Population Models

Given dynamic population estimates from the subway and motor vehicles, we estimate the dynamic population in buildings by using a population model. First, we define “catchment area” in this application as the spatial regions that are serviced by a transportation hub (a subway station, parking lot, taxi zone). Catchment area is defined by a Voronoi diagram, where distance is defined by street distance. The catchment areas of the subway stops in New York City are shown in Figure 9.

The catchment area determines the buildings that are serviced by either a subway station or parking lot. However, the different buildings receive a disproportionate number of people for different times of the day; for example, the probability that a person is traveling to the office rather than home is higher in the morning, and the reverse for the afternoon. We utilize the Citywide Mobility Survey [39], which provides the destinations of a sample population in New York City for different modes of transportation, including motor vehicles and subway. From this survey, we extracted the following distributions for trip destinations, as shown in Table 1.

Unfortunately, there is no finer granularity data that describes temporal differences (morning/evening) or direction (inflow/outflow). In *CityEnergy*, the percentages are set heuristically as

²<https://developers.google.com/places/web-service/search>.

Fig. 10. Energy footprinting of Manhattan on a typical weekday at 8 AM.

follows: In the morning (7 AM to 1 PM), office and school are set to the mobility survey percentages for inflow and 0% for outflow; in the evening (1 PM to 7 PM), office and school are set to 0% for inflow and set to the mobility survey percentages for outflow. Other destination percentages are scaled up to sum to 100%. In future works, finer granularity data can be incorporated to improve the accuracy of the population models.

Figure 10(a) shows a heatmap of the real-time population in Manhattan at the block level. To our knowledge, this is the highest granularity temporal-spatial mapping of population using dynamic estimation from vehicle modalities. As there are no datasets describing the change in population for individual buildings at different times of day, it is difficult to verify the accuracy of our population models. However, we have provided some high-level evaluations in Section 7.2.

Due to various factors such as model inaccuracies and daily differences, the populations of the different buildings often overshoot or undershoot the census baseline. To correct for this, the census baseline population for each building is reset at 4:00 AM each day.

6 DESIGN AND IMPLEMENTATION

CityEnergy is composed of three main subsystems: energy estimation, population estimation, and energy footprinting. As described in Sections 4 and 5, the energy and population estimation subsystems provide corresponding estimates for each building at fine temporal granularity. The energy footprinting module combines the results from the other two subsystems along with an apportionment policy manager to produce an energy footprint estimate. Sensor data from the user's mobile device is also used to determine mode of transportation and energy consumption when the user is not in a building. Energy footprinting information is relayed to the user via a mobile application. The system architecture is shown in Figure 11. The virtual client layer will be discussed in Section 6.4.1.

CityEnergy is deployed in New York City, which is an ideal testbed for two reasons. First, the city publishes a number of relevant datasets containing information about energy consumption in the built environment and population mobility for various modes of transportation. The abundance of datasets makes possible the study of energy footprinting at a city scale while presenting new challenges in data representation, cleaning, and modeling.

Second, there is a wide variety of buildings and people in New York City. These variations lead to interesting differences in energy usage and mobility. For example, the age of buildings in New York

Fig. 11. System architecture block diagram.

City varies tremendously, which can correlate with energy usage differences due to the available building management systems (BMS) technology. However, the number of people in Manhattan varies greatly between weekdays and weekends due to people commuting from outside of the city, according to NYU Wagner [32]. These variations lead to interesting challenges in the design of a city-scale energy footprinting system.

6.1 Transit Estimation

To complete the idea of “full coverage” for energy consumption, *CityEnergy* includes a transit estimation module. This module is responsible for analyzing location information from the users, determine whether the user is in transit, and provide the corresponding personal energy footprint. If the user is determined not to be in transit, the personal energy footprint is based on apportioned energy, as described in Section 6.2; otherwise, the personal energy footprint is based on the specific mode of transportation of the user. We do not claim any novelty in this section, but it is critical to maintain high temporal granularity of the user’s energy footprint.

6.1.1 Transportation Mode Detection. We focused on classifying the two main modes of transportation estimated in Section 5, as well as walking. Detection of the mode of transportation relies on the location traces collected from the user’s mobile device. A few simple thresholds are used to separate the modes of transportation. If the location change between multiple samples is below a lower threshold, L1, the user is considered to be stationary. Above the lower threshold L1 and below an upper threshold L2, the user is considered to be walking; and above L2, the user is considered to be in a vehicle. Finally, due to the nature of the NYC subway system being primarily underground, the location traces create large jumps, which separate subway transportation from motor vehicle transportation. This section is not the primary focus of this work, and we do not claim any novelty.

6.1.2 Transit Energy Footprint. If an occupant is in transit, *CityEnergy* provides an energy footprint depending on the mode of transportation. Because of the lack of relevant real-time information such as whether a motor vehicle is electric or gasoline, or how many people are sharing a subway car/bus/motor vehicle, we can only offer high-level methods for energy consumption estimates. In future works, it may be possible to obtain better transit energy estimates by incorporating various sensors or user feedback.

An energy footprint estimate for a user on the NYC subway can be computed as the average energy of the entire NYC subway system, divided by the average ridership (weekday or

Table 2. Estimated Energy Consumption for Common Modalities of Transportation in New York City

Mode of Transport	Estimated Power Consumption
Subway (Weekday)	36.7 Wh
Subway (Weekend)	72.2 Wh
Bus	587 Wh/km
Vehicle	874 Wh/km
Walking	0.0 W

weekend) [1]:

$$E_{subway}^p = \frac{E_{subway}}{R_{weekday/weekend}}.$$

An energy estimate per kilometer for a user on a NYC bus is computed as the energy consumption per kilometer over the average number of riders per trip (E_{pkm} is energy per kilometer) [1, 15]:

$$E_{pkm}_{bus}^p = \frac{E_{pkm}_{bus}}{\frac{ridership}{trips}}.$$

Last, to compute the energy estimate per kilometer for a user in a motor vehicle, we divide the approximate energy potential of gasoline by the US average distance per gallon of gasoline ($\overline{FE}$ is average fuel efficiency) [5]:

$$E_{pkm}_{car}^p = \frac{E_{pkm}_{car}}{\overline{FE}}.$$

A summary of energy consumption estimates used in *CityEnergy* is provided in Table 2.

6.2 Apportionment

Once an estimate for a building's energy consumption and population are calculated, a user's personal energy footprint can be estimated by applying an apportionment policy. From Reference [16], different apportionment policies can be applied, depending on the specifics of the building, such as primary building usage or types of occupants. *CityEnergy* includes a policy manager to apportion estimated building energy consumption to the occupants localized to the building. The simplest apportionment policy is uniform apportionment, such that the total energy consumption of the building is distributed uniformly over the number of people. This is the default policy for *CityEnergy*, but additional information can enable more complex policies, as discussed in Section 6.4.2. The energy consumption per capita for a typical morning in Manhattan is shown in Figure 10(b).

6.3 User Applications

Mobile Energy Footprinting: To provide energy footprints and other actionable feedback to everyday users, we developed a mobile application for iOS. The application is responsible for sensing location data, such as GPS and Wi-Fi information, which is then encoded and sent to the server. Once *CityEnergy* has localized the user and determined an energy footprint estimate, corresponding data is returned and displayed to the user in real-time. Screenshots of the mobile application are shown in Figure 12.

City Planner Applications: We also plan to make public web applications displaying city-wide information about energy and population estimates at the block level, such as in Figures 3(b),

Fig. 12. Screenshots of the *CityEnergy* iOS application. Left: User's recent estimated energy footprint and location trace. Right: Settings screen.

10(a), and 10(b). These tools can be useful for a variety of applications, such as transportation planning and energy planning.

6.4 Extensions

CityEnergy is able to provide a personal energy footprint estimate for citizens in the city. However, the high-level datasets utilized by *CityEnergy* provide coverage at the cost of footprint accuracy. To increase the accuracy of the footprint estimates, additional information about the building or occupants can be incorporated. We describe two extensions to *CityEnergy* with the goal of increasing accuracy: support for more complex apportionment policies and integration with bottom-up approaches.

6.4.1 Virtual Client. In the standard implementation of *CityEnergy*, clients query the system for energy footprint estimates at any location in the city. The system does not require knowledge of the number of clients or information about specific clients. The virtual client layer gives *CityEnergy* information that can be used to enable more complex apportionment policies and enable interfacing with building-specific energy footprinting systems.

When a client first queries the system, *CityEnergy* instantiates a virtual client with a number of fields, including: sensor data, parameters, and energy footprints. These fields are updated either by the physical client, the policy manager, or an external energy footprinting system.

6.4.2 Apportionment Policy. In most cases, the personal energy responsibility of different occupants in the building is not equal. The uniform apportionment policy, which estimates the same energy footprint for all occupants, is not sufficiently complex. To motivate the importance of more complex apportionment policies, we present two simplified examples of differing energy responsibilities and the information necessary to enable a relevant policy.

The first example is a mixed-use property consisting of a residential unit and a restaurant of equal size. Because of the high energy use of cooking appliances in the restaurant, the energy consumption may be significantly higher throughout the day than the residential unit. The uniform apportionment policy would estimate the same energy footprint for all occupants, whether in the

residential unit or in the restaurant. A better policy would be able to further divide the building into spaces based on usage (such as residential versus commercial).

The second example is a single retail location, such as a restaurant. Again, there are a number of energy consuming resources including cooking equipment such as refrigerators and stove tops, and resources for comfort such as air conditioning and lighting. Employees and customers in the restaurant may interact with completely different subsets of energy-consuming resources; thus, one group is utilizing more energy-intensive resources. For this reason, it may be reasonable to apportion different energy responsibilities to the employees and to the customers based on resources used.

In both scenarios, personal energy responsibility is dependent on the characteristics of the space (e.g., volume, usage, affiliation), and within each space, on the characteristics or role of the occupants. A two-level apportionment policy can be implemented [48] based on a space-level policy $f(s)$ and an occupant-level policy $g(p, s)$. The apportionment policy can be defined as in Equation (2):

$$E(p) = E \cdot f(s) \cdot g(p, s). \quad (2)$$

CityEnergy only relies on high-level datasets to estimate energy footprint, and thus does not inherently collect the information necessary to implement more complex apportionment policies. However, it is possible to extend *CityEnergy* to incorporate required information for two-level apportionment policies.

One simple method to implement a two-level apportionment policy is to first define the space-level $f(s)$ and occupant-level $g(p, s)$ policies for the building. As an example, if a building manager chooses to define space apportionment based on volume, then the space-level policy for a space s_0 can be defined as a fraction of the total building volume, where $V(\cdot)$ is a volume function:

$$f(s_0) = \frac{V(s_0)}{\sum_s V(s)}. \quad (3)$$

Similarly, the occupant-level policy $g(p, s)$ can be defined. As in the second example, we can define an occupant-level policy based on a subset of the total energy resources (A). If we let P be the set of people of a specific role (such as customer) and the resources that service this set as A_P , then the occupant-level policy can be defined as:

$$g(p_0, s_0) = \frac{1}{|P|} \cdot \frac{\sum_{d \in A_P} E(d)}{\sum_{d \in A} E(d)} \text{ s.t. } p_0 \in P.$$

Both the space- and occupant-level policies can be predefined based on known data (such as square footage of each space). The policies can also be dynamically updated depending on changes within the building (such as changes to a business lease).

Once the space- and occupant-level policies have been defined, we can generate an energy footprint for different types of occupants in each space, as shown in Figure 13. To determine which energy footprint to apportion to a specific user, the *CityEnergy* application can query the user about their specific location within the building and role. With the user's permission, this information can be learned and provided by the virtual client to the system automatically.

Note that the apportionment policy is merely a method for understanding energy consumption at a finer granularity. Thus, a specific building may be apportioned in many different ways, depending on the needs of the occupants. There may be situations where it is not desirable to attribute energy consumption of a specific space to individual occupants. For example, some parts of the building may have the energy consumption attributed to individual occupants, such as office employees, whereas other parts may have the energy consumption directly attributed to the building manager or the business leasing the space. In these cases, any apportionment policy such

Fig. 13. Example apportionment policy based on text examples. Total energy is split into residential and retail by a space-level policy and further split by an occupant-level policy.

as Equation (3) can be used to estimate the energy consumption of the space and allows flexible attribution of energy consumption.

6.4.3 Top-down and Bottom-up Approaches. *CityEnergy* is a top-down approach to energy footprinting, starting with high-level datasets and drilling down to individual footprint estimates. The bottom-up approach, explored in References [47, 52], begins at the physical sensor level and builds up the energy footprint. The two approaches have distinct advantages and disadvantages. The top-down approach is easily scalable to many buildings and can be implemented without any sensors, although the top-down approach suffers from low spatial and temporal accuracy. However, the bottom-up approach can provide an accurate energy footprint but requires extensive deployment.

As an increasing number of buildings become instrumented with sensing capabilities, there is an opportunity to integrate the two methods. Buildings with the appropriate sensing capabilities can leverage the high footprint accuracy of the bottom-up approach, while buildings without the capabilities can utilize the spatial and temporal coverage of the top-down approach.

As an example, Reference [48] demonstrates a bottom-up approach for energy footprinting in a commercial building. To integrate the advantages of this external energy footprinting system into *CityEnergy*, we provide an interface through the virtual client layer. Reference [48] accepts sensor data in the form of Bluetooth RSSI values and returns an energy footprint value. Thus, the virtual client for a user includes a field for sensor data, which can be forwarded to the external energy footprinting system through the web API. The virtual client's energy footprint field is then

updated and forwarded to the physical client, as shown in Figure 11. In this manner, the virtual client layer prevents the need to hand off the client to the appropriate energy footprinting system.

In addition to commercial buildings, most residences are already collecting energy consumption data through a central energy meter. Although this information is currently used by energy companies to determine monthly energy billings, this real-time energy information can be utilized by *CityEnergy* to provide accurate energy footprints for residential properties in the future.

6.5 Future Applications

The primary goal of *CityEnergy* is to provide insight into personal energy consumption throughout a city. However, we envision future applications that can utilize *CityEnergy* to provide intelligent actions for energy savings and to help city planners better understand energy and population distributions for various applications.

6.5.1 Recommendations. Although *CityEnergy* provides an energy footprint estimate to everyday citizens, it does not directly lead to reductions in energy savings. As shown in Reference [52], knowledge about energy consumption, coupled with possible energy-saving actions, can be used to develop a recommender system for intelligent energy-saving recommendations.

For example, a simple recommendation for energy savings can be derived from the method of transportation. Many commuters who take taxi trips in Midtown Manhattan during a typical weekday can save up to 1.2 kWh in fuel energy consumption and save time by instead using a city bike, according to Reference [38]. As each method of transportation has potential advantages and disadvantages for commuters, this information can be included in the recommendation to aid in a commuter's choices.

By collecting estimates of energy footprints through *CityEnergy*, the energy impact of different everyday decisions can be learned. This information, or dataset, can be used to adaptively train a recommender system for intelligent energy-saving recommendations by using a reinforcement learning algorithm such as deep Q-learning [51].

In addition, individual buildings may have energy-saving recommendations that are specific to the building. In Reference [52], recommendations for saving energy within an individual building increased the amount of energy savings over a four-week period. A city-wide recommender system could leverage existing systems at the building level to provide energy-saving recommendations when an occupant is inside of a building.

One important note is that minimizing personal energy footprint is not always desirable. As an example, one method for reducing one's personal energy footprint is by sharing resources, such as riding public transportation during peak hours or dining in crowded restaurants. In many cases, people may not be willing to share resources if the sole benefit is a reduction in personal energy consumption. Thus, future studies on energy-saving recommendations could include additional rewards to help incentivize people to do less desirable energy-saving actions.

6.5.2 City Planning. *CityEnergy* provides a number of interesting results besides the energy footprint, such as energy estimates and population estimates at a finer granularity time scale than existing datasets. This information could be useful for city planners in different applications such as congestion avoidance, air pollution, and load balancing in the power grid.

Although time-series data for energy consumption is collected by energy companies, this information is not readily available to the public. The energy estimates provided by *CityEnergy* could enable researchers and scientists to broadly study energy consumption at the city scale for issues such as load balancing in the power grid.

Population estimation in *CityEnergy* provides estimates for the flow of vehicles through different traffic intersections, the inflow and outflow of people from subway stations, and estimates of the

Table 3. In-sample Mean Absolute Error and Mean Root Squared Error of Different Machine Learning Regressors on the Local Law 84 and NYC Housing Authority Datasets

Model	In-Sample MAE	In-Sample MRSE	R^2
SVR	0.74	1.11	0.28
Linear Regression	0.77	1.06	0.34
Random Forest	0.30	0.48	0.86

As we are comparing the logarithm of energy consumption, units are in $\log(W)$.

Table 4. Comparison of Regression Models for Out-of-sample Energy Estimation of Four Ground Truth Buildings

Model	SVR		Linear Regression		Random Forest	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
Residential	0.66	0.69	1.48	1.49	0.87	0.93
Commercial	1.62	1.65	2.47	2.49	0.47	0.48
Retail	0.93	1.08	2.0	2.07	2.61	2.67

As we are comparing the logarithm of energy consumption, units are in $\log(W)$.

mobility of people. Besides mobile data, which is not available to the public in the United States, there is a lack of population estimation data at any scale, spatial or temporal. This information can be used by city planners for infrastructure studies. As an example, population estimation for mobile vehicle and subway commuters can provide knowledge of congestion and potentially reveal areas for improving population flow.

7 EVALUATION

We conducted the evaluation of *CityEnergy* in three parts: energy estimation, population estimation, and as a complete system.

7.1 Energy Estimation

As described in Section 3.1, energy consumption is estimated by passing building parameters through a trained regression model and downsampled to the hourly scale by using a model building derived from the DOE Reference Buildings. To evaluate the energy estimation pipeline, we evaluated in-sample estimation with different regression models. In addition, we gathered data from a few representative buildings over two months to assess out-of-sample estimation. Note that evaluation uses the logarithm of energy consumption, as in Reference [26].

7.1.1 In-sample Estimation. We tested a few different machine learning models in an effort to improve prediction accuracy. Support vector regression, random forest, and linear regression models were trained on the Local Law 84 and NYCHA datasets using five-fold cross validation. The mean absolute error and mean root squared error for each model are shown in Table 3.

7.1.2 Out-of-sample Estimation. To evaluate the generalizability of the trained energy regression models to the city-scale, we gathered data from one residential building, one commercial building, and one retail building over the course of two months. We generated energy estimation traces for these buildings using *CityEnergy* and calculated the MAE and MRSE against the ground truth in Table 4.

As described in Reference [26], linear regression can produce low MAE for both in-sample and out-of-sample energy estimation; however, in our experiments, support vector regression yields

Table 5. In-sample Mean Absolute Error and Mean Root Squared Error of Different Machine Learning Regressors on Historical Turnstile Data

Model	In-Sample MAE	In-Sample MRSE	R^2
SVR	0.27	0.44	0.80
Linear Regression	0.74	0.93	0.13
Random Forest	0.24	0.43	0.81

As we are comparing the logarithm of outflow and inflow populations, the units are $\log(\text{people})$.

the lowest MAE and MRSE. Although the accuracy is low (especially for commercial buildings), the model can still provide reasonable energy estimates at the city-scale. More complex models can be used to further improve the accuracy of energy estimation in *CityEnergy*.

7.2 Population Estimation

To evaluate the population estimation in *CityEnergy*, we individually evaluate the subway estimates. Instead of evaluating vehicle to population dynamics, we evaluate the total population estimates on five ground truth buildings throughout New York City to assess the correctness of the population models.

7.2.1 Subway. A regression model is trained for each subway station on the previous six weeks of turnstile data. We tested three different regression models to determine the best in-sample MAE and MRSE, as shown in Table 5. From our experiments, random forest provides the highest accuracy. Evaluation uses the logarithm of the outflow and inflow populations.

7.2.2 Population Models. Finally, we evaluated the full dynamic population pipeline. For four different buildings (two residential, one commercial, and one retail), we manually counted the number of occupants in a building as ground truth. We combined the dynamic population estimates from the subway and motor vehicles, as described in Section 5.3, and scaled the total to account for the remaining 33% from other transportation modalities.

Using our population model, we were able to achieve an MAE of 0.7 for the four test buildings. We found that the error varied greatly, depending on the type of building, time of day, and other factors. Although the population models used were sufficient for *CityEnergy*, we believe that covering additional modes of transportation and increasing the existing models can produce a model with much higher accuracy.

7.3 Energy Footprinting Estimation

The most important characteristic for *CityEnergy* is coverage; however, the energy footprint estimate should also be accurate enough to approximate a ground truth energy footprint.

To demonstrate the potential accuracy of *CityEnergy*, we collected a real energy footprint as ground truth. The real energy footprint consisted of three building locations: one residential location, one retail location, and one commercial building. The ground truth energy footprint is collected as follows:

- In the residential location, personal energy footprint is determined using a set of plugmeters and light sensors to measure the energy consumption.
- In the retail location, a camera was deployed to monitor the building's energy meter; the population was recorded manually.
- In the commercial building location, a combination of plugmeters, light sensors, and building management system data was utilized to determine the personal energy footprint.

Fig. 14. Semi-log plot of the measured real-time energy footprint of a user and the energy footprint estimate from *CityEnergy*.

Table 6. *CityEnergy* Yearly Footprint Estimate Compared to DOE and EIA Estimates

	<i>CityEnergy</i>	DOE (New York State)	EIA (New York State)
Estimate	67.9 BTU	106 BTU	188.1 BTU

A comparison of the ground truth energy footprint and the *CityEnergy* energy footprint is shown in Figure 14. *CityEnergy* is able to achieve a MAE of 1.7 kWh. For different buildings and different energy footprints, *CityEnergy* may achieve higher or lower accuracy. However, better energy estimation and population estimation models can help reduce error in future works.

There are a number of estimates of a typical energy footprint over the entire year; however, these estimates vary, depending on what buildings are considered as part of personal energy footprint. We extrapolated the estimated energy consumption from *CityEnergy* to compare with two other yearly estimates from the U.S. Department of Energy [36] and the U.S. Energy Information Administration [41], as shown in Table 6.

While the DOE estimate only considers residential and transportation energy, the EIA estimate considers all buildings, including industrial and commercial. *CityEnergy* considers residential, transportation, and commercial, and thus underestimates the energy footprint compared to the DOE and EIA estimates. Some of the discrepancy may be due to the specific buildings considered in the *CityEnergy* estimate, or due to the differences in energy consumption between New York City and New York State as a whole. This difference may also be due to errors in the *CityEnergy* energy and population models. However, the estimate is within a small factor of the DOE and EIA estimates, which is an encouraging comparison for the accuracy of *CityEnergy*.

Although *CityEnergy* can provide an energy footprint estimate, there is a wide variation in energy consumption between buildings, due to factors such as usage, age, energy efficiency, and occupancy. While the models used to estimate energy consumption are designed to account for these factors, real consumption is likely to be different from the estimated footprint. Likewise, the personal energy footprint may only be accurate in specific situations and may represent a typical energy footprint rather than an individual energy footprint. However, we hope that the models developed in *CityEnergy* and the estimated energy footprint can still be useful for future works.

8 CONCLUSION

In this work, we present *CityEnergy*, a personal energy footprinting system that uses energy and population models to provide coverage throughout an urban environment. *CityEnergy* can be easily extended to other urban cities by interchanging available energy and population data specific to

the urban city. *CityEnergy* is, to our knowledge, the first system to address the problem of personal energy footprinting at the city-scale without relying on building-specific energy and population monitoring deployments.

REFERENCES

- [1] Metropolitan Transportation Authority. 2018. Average Weekday Subway Ridership. Retrieved from http://web.mta.info/nyct/facts/ridership/ridership_sub.htm.
- [2] Metropolitan Transportation Authority. 2018. MTA Real-Time Data Feeds. Retrieved from <https://api.mta.info/>.
- [3] Metropolitan Transportation Authority. 2018. Turnstile Data. Retrieved from <http://web.mta.info/developers/turnstile.html>.
- [4] Edward Barbour, Carlos Cerezo Davila, Siddharth Gupta, Christoph Reinhart, Jasleen Kaur, and Marta C. González. 2019. Planning for sustainable cities by estimating building occupancy with mobile phones. *Nat. Commun.* 10, 1 (2019), 1–10.
- [5] Michael R. Bloomberg. 2007. Inventory of New York City greenhouse gas emissions. *New York City Mayor's Office of Operations, Office of Long-term Planning and Sustainability*. Retrieved from https://www1.nyc.gov/assets/sustainability/downloads/pdf/publications/NYC_GHG_Inventory_2013.pdf.
- [6] U. S. Census Bureau. 2017. U.S. Census Bureau, 2017 American Community Survey 1-Year Estimate. Retrieved from <https://www.census.gov/>.
- [7] Yixing Chen, Tianzhen Hong, and Mary Ann Piette. 2017. Automatic generation and simulation of urban building energy models based on city datasets for city-scale building retrofit analysis. *Appl. Energy* 205 (2017), 323–335.
- [8] Yun Cheng, Kaifei Chen, Ben Zhang, Chieh-Jan Mike Liang, Xiaofan Jiang, and Feng Zhao. 2012. Accurate real-time occupant energy-footprinting in commercial buildings. In *Proceedings of the 4th ACM Workshop on Embedded Sensing Systems for Energy-efficiency in Buildings*. 115–122.
- [9] New York City Taxi & Limousine Commission. 2019. TLC Trip Record Data.
- [10] New York City Taxi and Limousine Commission. 2019. TLC Trip Record Data. Retrieved from <https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page>.
- [11] Drury B. Crawley, Linda K. Lawrie, Frederick C. Winkelmann, Walter F. Buhl, Y. Joe Huang, Curtis O. Pedersen, Richard K. Strand, Richard J. Liesen, Daniel E. Fisher, Michael J. Witte, et al. 2001. EnergyPlus: Creating a new-generation building energy simulation program. *Energy Build.* 33, 4 (2001), 319–331.
- [12] Sarah Darby et al. 2006. The effectiveness of feedback on energy consumption. *Rev. DEFRA Liter. Meter. Bill. Dir. Displ.* 486, 2006 (2006), 26.
- [13] Stephen Dawson-Haggerty, Jorge Ortiz, Xiaofan Jiang, Jeff Hsu, Sushant Shankar, and David Culler. 2010. Enabling green building applications. In *Proceedings of the 6th Workshop on Hot Topics in Embedded Networked Sensors*. 1–5.
- [14] Office of Energy Efficiency Department of Energy and Renewable Energy. 2018. Commercial Reference Buildings. Retrieved from <https://www.energy.gov/eere/buildings/commercial-reference-buildings>.
- [15] Igors Graurs, Aigars Laizans, Peteris Rajeckis, and Aivars Rubenis. 2015. Public bus energy consumption investigation for transition to electric power and semi-dynamic charging. *Eng. Rural. Dev.* 14 (2015), 366–371.
- [16] Simon Hay and Andrew Rice. 2009. The case for apportionment. In *Proceedings of the 1st ACM Workshop on Embedded Sensing Systems for Energy-efficiency in Buildings*. ACM, 13–18.
- [17] Samuli Hemminki, Petteri Nurmi, and Sasu Tarkoma. 2013. Accelerometer-based transportation mode detection on smartphones. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*. ACM, 13.
- [18] Tianzhen Hong, Mary Ann Piette, Yixing Chen, Sang Hoon Lee, Sarah C. Taylor-Lange, Rongpeng Zhang, Kaiyu Sun, and Phillip Price. 2015. Commercial building energy saver: An energy retrofit analysis toolkit. *Appl. Energy* 159 (2015), 298–309.
- [19] B. Howard, L. Parshall, J. Thompson, S. Hammer, J. Dickinson, and V. Modi. 2012. Spatial distribution of urban building energy consumption by end use. *Energy Build.* 45 (2012), 141–151.
- [20] Samuel Humeau, Tri Kurniawan Wijaya, Matteo Vasirani, and Karl Aberer. 2013. Electricity load forecasting for residential customers: Exploiting aggregation and correlation between households. In *Proceedings of the Conference on Sustainable Internet and ICT for Sustainability (SustainIT'13)*. IEEE, 1–6.
- [21] Arash Jahangiri and Hesham A. Rakha. 2015. Applying machine learning techniques to transportation mode recognition using mobile phone sensor data. *IEEE Trans. Intell. Transport. Syst.* 16, 5 (2015), 2406–2417.
- [22] Rishree K. Jain, Kevin M. Smith, Patricia J. Culligan, and John E. Taylor. 2014. Forecasting energy consumption of multi-family residential buildings using support vector regression: Investigating the impact of temporal and spatial monitoring granularity on performance accuracy. *Appl. Energy* 123 (2014), 168–178.
- [23] Radiša Ž. Jovanović, Aleksandra A. Sretenović, and Branislav D. Živković. 2015. Ensemble of various neural networks for prediction of heating energy consumption. *Energy Build.* 94 (2015), 189–199.

- [24] Gijss Klous, Lidwien A. M. Smit, Floor Borlée, Roel A. Coutinho, Mirjam E. E. Kretzschmar, Dick J. J. Heederik, and Anke Huss. 2017. Mobility assessment of a rural population in the Netherlands using GPS measurements. *Int. J. Health Geogr.* 16, 1 (2017), 30.
- [25] Constantine Kontokosta. 2012. Local Law 84 Energy Benchmarking Data: Report to the New York City Mayor's Office of Long-Term Planning and Sustainability. Retrieved from https://www1.nyc.gov/html/gbee/html/plan/ll84_scores.shtml.
- [26] Constantine E. Kontokosta and Christopher Tull. 2017. A data-driven predictive model of city-scale energy use in buildings. *Appl. Energy* 197 (2017), 303–317.
- [27] Seungwoo Lee, Daye Ahn, Sukjun Lee, Rhan Ha, and Hojung Cha. 2014. Personalized energy auditor: Estimating personal electricity usage. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom'14)*. IEEE, 44–49.
- [28] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. 2016. SSD: Single shot multibox detector. In *Proceedings of the European Conference on Computer Vision*. Springer, 21–37.
- [29] Yunjia Ma, Wei Xu, Xiujuan Zhao, and Ying Li. 2017. Modeling the hourly distribution of population at a high spatiotemporal resolution using subway smart card data: A case study in the central area of Beijing. *ISPRS Int. J. Geo-inf.* 6, 5 (2017), 128.
- [30] Alessio Mastrucci, Olivier Baume, Francesca Stazi, and Ulrich Leopold. 2014. Estimating energy savings for the residential building stock of an entire city: A GIS-based statistical downscaling approach applied to Rotterdam. *Energy Build.* 75 (2014), 358–367.
- [31] Satoshi Miyazawa, Xuan Song, Tianqi Xia, Ryosuke Shibasaki, and Hodaka Kaneda. 2019. Integrating GPS trajectory and topics from Twitter stream for human mobility estimation. *Front. Comput. Sci.* 13, 3 (2019), 460–470.
- [32] Mitchell Moss and Carson Qing. 2012. *The Dynamic Population of Manhattan*. Rudin Center for Transportation, NYU Wagner School.
- [33] Alex Nutkiewicz, Zheng Yang, and Rishree K. Jain. 2018. Data-driven urban energy simulation (DUE-S): A framework for integrating engineering simulation and machine learning methods in a multi-scale urban energy modeling workflow. *Appl. Energy* 225 (2018), 1176–1189.
- [34] New York City Housing Authority (NYCHA). 2018. Electric Consumption and Cost (2010–June 2018). Retrieved from <https://opendata.cityofnewyork.us/>.
- [35] Department of City Planning. 2018. Primary Land Use Tax Lot Output. Retrieved from <https://www1.nyc.gov/site/planning/data-maps/open-data/dwn-pluto-mappluto.page>.
- [36] U.S. Department of Energy. 2014. How Much Do You Spend? Retrieved from <https://www.energy.gov/articles/how-much-do-you-spend>.
- [37] New York City Department of Transportation. 2017. NYCDOT - Real Time Traffic Information. Retrieved from <https://webcams.nyctmc.org/>.
- [38] New York City Department of Transportation. 2016. Mobility Report. Retrieved from <https://www1.nyc.gov/html/dot/html/about/mobilityreport.shtml>.
- [39] New York City Department of Transportation. 2017. Citywide Mobility Survey. Retrieved from <https://www1.nyc.gov/html/dot/html/about/citywide-mobility-survey.shtml>.
- [40] Christoph Reinhart, Timur Dogan, J. Alstan Jakubiec, Tarek Rakha, and Andrew Sang. 2013. Umi—An urban simulation environment for building energy use, daylighting and walkability. In *Proceedings of the 13th Conference of International Building Performance Simulation Association*, Vol. 1.
- [41] U.S. Energy Information Administration EIA Independent Statistics and Analysis. 2020. How much energy does a person use in a year? Retrieved from <https://www.eia.gov/tools/faqs/>.
- [42] Leon Stenneth, Ouri Wolfson, Philip S. Yu, and Bo Xu. 2011. Transportation mode detection using mobile phones and GIS information. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, 54–63.
- [43] Jenna Tatum. 2013. PlaNYC: New York City Mayor's carbon challenge progress report. New York: Mayor's Office of Long-Term Planning and Sustainability (2013). Retrieved from http://www.nyc.gov/html/gbee/downloads/pdf/mayors_carbon_challenge_progress_report.pdf.
- [44] Shailja Thakur, Manaswi Saha, Amarjeet Singh, and Yuvraj Agarwal. 2014. WattShare: Detailed energy apportionment in shared living spaces within commercial buildings. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-efficient Buildings*. ACM, 30–39.
- [45] Florian Toqué, Mostepha Khoudja, Etienne Come, Martin Trepanier, and Latifa Oukhellou. 2017. Short & long term forecasting of multimodal transport passenger flows with machine learning methods. In *Proceedings of the IEEE 20th International Conference on Intelligent Transportation Systems (ITSC'17)*. IEEE, 560–566.

- [46] Liping Wang, Robert Kubichek, and Xiaohui Zhou. 2018. Adaptive learning based data-driven models for predicting hourly building energy use. *Energy Build.* 159 (2018), 454–461.
- [47] Peter Wei, Xiaoqi Chen, Rishikanth Chandrasekaran, Fengyi Song, and Xiaofan Jiang. 2016. Adaptive and personalized energy saving suggestions for occupants in smart buildings: Poster abstract. In *Proceedings of the 3rd ACM International Conference on Systems for Energy-efficient Built Environments (BuildSys'16)*. ACM, New York, NY, 247–248. DOI: <https://doi.org/10.1145/2993422.2996412>
- [48] Peter Wei, Xiaoqi Chen, Jordan Vega, Stephen Xia, Rishikanth Chandrasekaran, and Xiaofan Jiang. 2018. A scalable system for apportionment and tracking of energy footprints in commercial buildings. *ACM Trans. Sens. Netw.* 14, 3–4 (2018), 22.
- [49] Peter Wei and Xiaofan Jiang. 2018. A data-driven system for city-scale personal energy footprint estimations. In *Proceedings of the 5th Conference on Systems for Built Environments*. ACM, 194–195.
- [50] Peter Wei, Haocong Shi, Jiaying Yang, Jingyi Qian, Yinan Ji, and Xiaofan Jiang. 2019. City-scale vehicle tracking and traffic flow estimation using low frame-rate traffic cameras. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing and the ACM International Symposium on Wearable Computers*. ACM, 602–610.
- [51] P. Wei, S. Xia, R. Chen, J. Qian, C. Li, and X. Jiang. 2020. A deep-reinforcement-learning-Based recommender system for occupant-driven energy optimization in commercial buildings. *IEEE Internet of Things Journal* 7, 7 (2020), 6402–6413. <https://doi.org/10.1109/JIOT.2020.2974848>
- [52] Peter Wei, Stephen Xia, and Xiaofan Jiang. 2018. Energy saving recommendations and user location modeling in commercial buildings. In *Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization*. ACM, 3–11.
- [53] Peter Widhalm, Philippe Nitsche, and Norbert Brändie. 2012. Transport mode detection with realistic smartphone sensor data. In *Proceedings of the 21st International Conference on Pattern Recognition (ICPR'12)*. IEEE, 573–576.
- [54] Tri Kurniawan Wijaya, S. F. R. J. Humeau, Matteo Vasirani, and Karl Aberer. Residential Electricity Load Forecasting: Evaluation of Individual and Aggregate Forecasts. EPFL, 2014.
- [55] Fengli Xu, Pengyu Zhang, and Yong Li. 2016. Context-aware real-time population estimation for metropolis. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 1064–1075.
- [56] Bailang Yu, Ting Lian, Yixiu Huang, Shenjun Yao, Xinyue Ye, Zuoqi Chen, Chengshu Yang, and Jianping Wu. 2019. Integration of nighttime light remote sensing images and taxi GPS tracking data for population surface enhancement. *Int. J. Geog. Inf. Sci.* 33, 4 (2019), 687–706.
- [57] Shanghang Zhang, Guanhang Wu, Joao P. Costeira, and José M. F. Moura. 2017. Understanding traffic density from large-scale web camera data. *arXiv preprint arXiv:1703.05868* (2017).
- [58] Hai Xiang Zhao and Frédéric Magoulès. 2010. Parallel support vector machines applied to the prediction of multiple buildings energy consumption. *J. Algor. Comput. Technol.* 4, 2 (2010), 231–249.
- [59] Yu Zheng, Quannan Li, Yukun Chen, Xing Xie, and Wei-Ying Ma. 2008. Understanding mobility based on GPS data. In *Proceedings of the 10th International Conference on Ubiquitous Computing*. ACM, 312–321.
- [60] Zimu Zheng, Feng Wang, Dan Wang, and Liang Zhang. 2018. Buildings affect mobile patterns: Developing a new urban mobility model. In *Proceedings of the 5th Conference on Systems for Built Environments*. 83–92.

Received July 2020; revised October 2020; accepted November 2020