

Article information

Article title

A Dataset of Images of Public Streetlights with Operational Monitoring using Computer Vision Techniques

Authors

Ioannis Mavromatis, Aleksandar Stanoev, Pietro Carnelli, Yichao Jin, Mahesh Sooriyabandara, Aftab Khan

Affiliations

Bristol Research and Innovation Laboratory, Toshiba Europe Ltd., 32 Queen Square, Bristol, BS1 4ND, United Kingdom

Corresponding author's email address and Twitter handle

{Ioannis.Mavromatis, Aleksandar.Stanoev, Pietro.Carnelli, Yichao.Jin, Mahesh, Aftab.Khan}@toshiba-bril.com

Keywords

Streetlight, Street Furniture, Light and Emergency Light Maintenance, Lighting Columns, Computer Vision, Machine Learning, Smart Cities

Abstract

A dataset of street light images is presented. Our dataset consists of ~ 350 k images, taken from 140 UMBRELLA nodes installed in the South Gloucestershire region in the UK. Each UMBRELLA node is installed on the pole of a lamppost and is equipped with a Raspberry Pi Camera Module v1 facing upwards towards the sky and lamppost light bulb. Each node collects an image at hourly intervals for 24 h every day. The data collection spans for a period of six months.

Each image taken is logged as a single entry in the dataset along with the Global Positioning System (GPS) coordinates of the lamppost. All entries in the dataset have been post-processed and labelled based on the operation of the lamppost, i.e., whether the lamppost is switched ON or OFF. The dataset can be used to train deep neural networks and generate pre-trained models providing feature representations for smart city CCTV applications, smart weather detection algorithms, or street infrastructure monitoring. The dataset can be found at <https://doi.org/10.5281/zenodo.6046758>.

Keywords: Streetlight, Street Furniture, Light and Emergency Light Maintenance, Lighting Columns, Computer Vision, Machine Learning, Smart Cities

Specifications table

Subject	Computer Science
Specific subject area	Computer Vision and Pattern Recognition
Type of data	The image data for each street light column are provided as JPEG files [1]. Each zipped directory contains all the all images associated to a single light column/lamppost. A CSV file containing the information about all the street lights, their metadata and the status decision outcome from our post-processing algorithm, i.e., if it is switched ON or OFF. We also include another processed CSV file with the number of occurrences per street light column for easier post-processing of the data.
How the data were acquired	Acquired from the UMBRELLA testbed [2], consisting of 140 UMBRELLA nodes, each equipped with a Raspberry Pi Camera Module v1 [3] and an OmniVision OV5647 sensor [4]. The nodes are installed on top of lampposts, around the South Gloucestershire region in the UK, across a ~ 7.2 km stretch of road and the University of West of England Campus. Note, all cameras are facing upwards, towards the light bulbs of the lampposts and the sky.
Data format	Raw data format, as recorded by each UMBRELLA node, pertaining to the UMBRELLA testbed, in the form of JPEG and CSV files. The raw data are analysed to identify the street lights functionality and the results are listed in a CSV file format.
Description of data collection	The data has been recorded on hourly intervals at a best effort basis (for all available nodes on each hourly timeslot), for seven days per week and for a period of six months. A random time delay of 0 s to 600 s was introduced before taking each photo to spread out the analysis resource utilisation on the server side.
Data source location	Roughly $\sim 80\%$ of the nodes installed on a ~ 7.2 km stretch of public road in the Gloucestershire region in the UK, comprising of the following roads: Coldharbour Lane, Filton Road, and A4174 Ring Road. $\sim 20\%$ of the nodes are installed around the University of the West of England Frenchay Campus, in Bristol, UK.
Data accessibility	Data available online at the link below: https://doi.org/10.5281/zenodo.6046758 [5]
Related research article	

Value of the Data

- Our dataset consists of $\sim 350k$ JPEG images of street lights, taken from lighting columns installed on public road infrastructure in the city of Bristol, UK. Each UMBRELLA node installed on a street column provides a unique camera placement, photographic angle, and distance from the street light. Furthermore, several street lights are partially obstructed by vegetation. In addition, several light bulbs are outside of the Field of View (FoV) the Pi Camera module achieves. Finally, the cameras facing the sky are susceptible to weather conditions (e.g., rain, snow, direct sun and moon light, etc.) that can partially or entirely alter the quality of the images taken. The above provides a unique and diverse dataset of images that can be used for training tools and Machine Learning (ML) models for inspection, monitoring and maintenance use-cases.
- By building upon the provided images, experts from the fields of Computer Vision, Internet of Things (IoT), and Smart Cities can develop and train ML models and heuristic tools that assess the status of the street and emergency lights in real-time (e.g., as in [6]). For instance, this can facilitate further research on the provision of Smart City services that can automatically detect whether a street light is ON or OFF, alert a maintenance team, and reduce the person-hours required for on-site monitoring.
- Deep neural network training and the regularisation of the models produced can benefit from pre-training models [7]. Our dataset provides an extensive collection of images taken at different times of the day and night, different weather conditions and different exposure settings. ML models pre-trained on this dataset can provide excellent feature representations in deep neural network-based transfer learning applications, e.g., outdoor smart city CCTV deployments.
- Since the cameras in question are facing the sky, this dataset can also be used for training real-time weather warning and monitoring systems. For instance, building a simple rain detector is straightforward, as very different images are produced under rainy conditions. In this use case, the direction of rain can be estimated, and potentially by using existing infrastructure (e.g., existing traffic cameras), weather warnings can be generated, pin-pointed very precisely and disseminated to drivers approaching specific streets.
- Finally, as this dataset provides a very diverse set of streetlamp images, it can be combined with datasets of images from other street furniture (e.g., traffic lights, street name signs, traffic signs, etc.) or sensors (e.g. LiDARs). Such a dataset can be later used to train object recognition tools (sensor fusion) for autonomous vehicles and drones systems and accelerate the research towards autonomous navigation on public road infrastructures, e.g., as in [8].

Data Description

The raw data files associated with each lamppost (namely the JPEG [1] images and the CSV file), are organised as follows. The dataset is upload as a single Zip file, that, when unzipped, it unfolds to a number of zipped and two CSV files. As our dataset is fairly large, we provide two Zip files, one containing the complete dataset and another one containing an example dataset with a smaller number of JPEG images and UMBRELLA nodes. The naming convention for both is “streetcare-dataset-complete.zip” and “streetcare-dataset-example.zip” respectively. The CSV files within the root directory follow a similar naming convention, i.e., “occurrences-complete.csv”, “streetlights-complete.csv” for the complete dataset and “occurrences-example.csv”, “streetlights-example.csv” for the example one.

Within the root zipped directory, one can find the Zip files containing our JPEG images. These zipped files are named after the “serial ID” of each UMBRELLA node, i.e., a friendly name given to each node. This serial ID is in the form of “*RS[ES]-[A-Z0-9]*”. For example, valid serials are “RSE-A-11-C”, “RSE-A-G-8-C”, “RSS-47-C”, etc. All serial IDs represent the type of the node and the sensors equipped, but this is outside of the scope of this dataset. With what regards this dataset, all nodes are equipped with the same Raspberry Pi Camera Module (i.e., Camera Module V1 [3]). All the “RSE-*” or “RSS-*” subdirectories, when unzipped, contain the JPEG images taken of this particular UMBRELLA node. The file name of each JPEG file is a random string of characters and numbers with a size of 32 digits. This string is assigned at random when the file is created and is always unique. An example of the directory tree of our dataset can be seen below:

Table 2: Definition of the labels of each column of the CSV data files in the dataset.

Field Name	Definition
id	A unique ID given to each node. It increments by one for each image captured.
serial	The UMBRELLA nodes “serial ID” (friendly name).
date	Date and timestamp at the moment an image was captured.
hostname	The UMBRELLA nodes hostname, as seen in the host OS.
lat	Latitude of the position of the current node.
lon	Longitude of the position of the current node.
image_name	The JPEG image file name (32 digit alphanumeric string with “.jpg” extension).
fault_detected	The post-processed result of the streetlights status (operational or not).
confidence	The probability the “fault_detected” result is correct.
daynight	The astrological day or night time cycle.
red	The pixel intensity as perceived under the “red” RGB channel.
green	The pixel intensity as perceived under the “green” RGB channel.
blue	The pixel intensity as perceived under the “blue” RGB channel.

```

streetcare-dataset-complete.zip
├── occurrences-complete.csv
├── streetlights-complete.csv
├── RSE-6-C.zip
│   ├── 9397c82e6dafd06435e13552e2e7d2ca.jpg
│   ├── 9196c26e24a1ed938d8a170939833c79.jpg
│   └── ...
├── RSE-A-12-C.zip
│   ├── d69bb15b397f33b2f706f89926f904a9.jpg
│   ├── f1b60e40d3d295a9eee2971ce5e426bc.jpg
│   └── ...
└── ...

streetcare-dataset-example.zip
├── occurrences-example.csv
├── streetlights-example.csv
├── RSE-L-1-C.zip
│   ├── cf0f229123a6b1ac659a9f51366c469f.jpg
│   ├── b8fb64ece07c34a21791afcff7fad72f.jpg
│   └── ...
├── RSS-12-C.zip
│   ├── e27bd4046e8616906f5daf80ed8b628b.jpg
│   ├── 8aedbdbdf8aab033bc808be3feaeaa34.jpg
│   └── ...
└── ...

```

With what regards to the CSV files, the “streetlights-XXX.csv” contains in a tabular format all the information about the street lights, their metadata and the status decision outcome from our post-processing algorithm, i.e., ON or OFF. The first line of the file shows the label of each column. For the following rows, each row is an “entry”, i.e., it represents a JPEG image taken from a specific UMBRELLA node at a specific time. The “streetlights-XXX.csv” files are structured in the following format:

```

Column_Label_0;Column_Label_1;... Column_Label_n
Value_0_Entry_0;Value_1_Entry_0;... Value_n_Entry_0
Value_0_Entry_1;Value_1_Entry_1;... Value_n_Entry_1
...
Value_0_Entry_m;Value_1_Entry_m;... Value_n_Entry_m

```

A detailed explanation of each column found in the “streetlights-XXX.csv” file can be found in Tab. 2. With regards to the “occurrences-XXX.csv”, it contains a list of serial IDs from all nodes in the dataset, and the number of images taken from this specific node. This file is ordered in descending order based on the number of images. As before, the file is in a tabular format with the first line being the labels of the columns, i.e., “serial” and “occurrences”, and later each line is an entry, i.e., a node serial ID and the number of occurrences in the dataset.

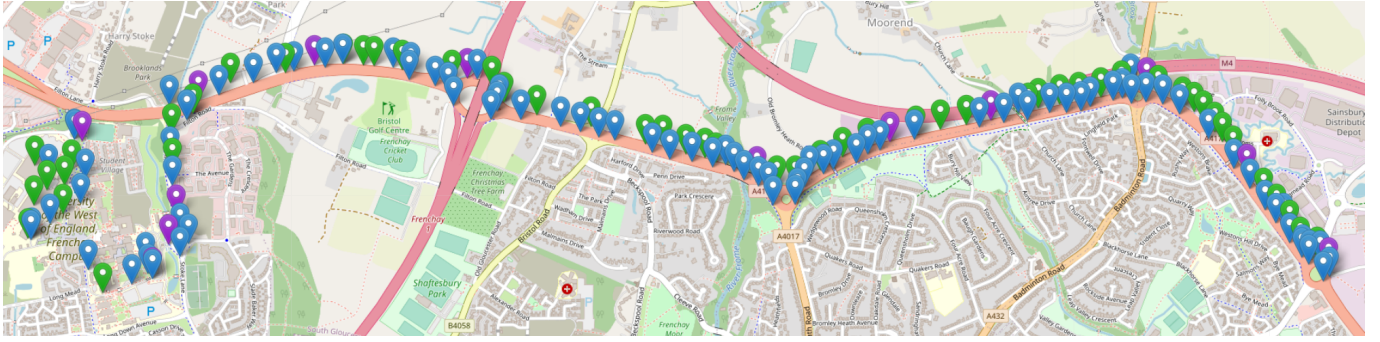


Figure 1: The UMBRELLA network. All nodes are installed on public lampposts across a road of ~ 7.2 km. The colours represent the nodes connectivity, i.e. green is fibre connected, blue is WiFi connected, and purple is fibre connected and can act as a LoRa gateway too.




Figure 2: Umbrella Node on a lamppost, its exploded view, and the camera placement at the top of the node (seen in the red circles).

Experimental design, materials and methods

All UMBRELLA nodes are installed across multiple locations in South Gloucestershire region of the UK, these being a ~ 7.2 km stretch of public road (about $\sim 80\%$ of the nodes) and the University of the West of England (UWE) Frenchay Campus (about $\sim 20\%$ of the nodes). The nodes locations and the region can be seen in Fig. 1. All UMBRELLA nodes are connected to our backend servers via a fibre or a wireless (WiFi) interface. Our servers are used to collect all the images taken and store the information in our database.

Fig. 2 shows an Umbrella node attached to a lamppost. Each rhombus segment contains custom PCBs, designed and manufactured by Toshiba, accommodating ten sensors (e.g., Bosch BME680, accelerometers, microphones, etc.) and seven network interfaces (Bluetooth, fibre, WiFi, LoRa, etc.). At the top of the node, one can find a Raspberry Pi Camera Module v1 [3], used for taking the street light images. The camera module is equipped with an OmniVision OV5647 sensor [4]. All PCBs are connected to a main processing unit (Raspberry Pi 3b+ Compute Module [9]), responsible for controlling the sensors and the camera module, processing the data generated and sending them to the backend servers.

All available nodes periodically take a photo of the lamppost and send it back to our backend for processing and storing. The photos are taken at hourly intervals and are collected for the entire 24 h time period every day. The scheduler works in a best-effort fashion, i.e., if a node is not available on a specific timeslot (e.g., due to downtime), no photos are taken. The photo is scheduled every hour on the hour. A random time delay of 0 s to 600 s is introduced on each node before taking a photo, to ensure the data processing load is spread evenly on the server side. When a photo is saved, an entry is created in our CSV file reporting the “serial ID” of the node (`serial` column - Tab. 2), the date and time the photograph was taken (`date` column - Tab. 2), the hostname of the Raspberry Pi the camera was connected to (`hostname` column - Tab. 2), and the GPS latitude and longitude coordinates of the lamppost (`lat` and `lon` column - Tab. 2). Finally, the JPEG file name is randomly generated using 32 alphanumeric characters and is reported in the CSV file as well (`image_name` column - Tab. 2).

All lampposts operate (are turned ON or OFF) based on the astronomical night time, i.e., they are turned ON 15 min before the astronomical dusk and are turned OFF 15 min after the astronomical dawn. Based on that day/night cycle, we change the configuration of our camera accordingly. The different configurations applied to the camera sensor are listed in Tab. 3. During the “day” time, i.e., the time that




Figure 3: Two examples of the highly exposed photos used for the night time evaluation. On the left, a lamppost with the light directly over the camera. On the right, a lamppost with increased vegetation around the node.

Table 3: Camera Module and Sensor Settings.

Setting	Day	Night	Night - High exposure
Exposure Mode	Automatic	Manual	
ISO	100	400	800
Shutter Speed	Automatic	250000	600000
LED State	Off		
Flash	Off		
Resolution	1024 × 768		
Contrast	0		
Brightness	50%		

a lamppost is expected to be OFF, a camera is left in automatic exposure and shutter speed settings to compensate for the available sunlight under different weather conditions. During the “night” time, i.e., the time that a lamppost is expected to be ON, the exposure, ISO, and shutter speed are increased to compensate for the low-light conditions. During the night time particularly, two images are collected with different configurations. The first one (column “Night” in Tab. 3) is the photo that is saved as part our dataset. The second configuration (column “Night - High exposure” in Tab. 3) is used for our post-processing and the labelling of the dataset. Two examples of the same photo taken with and without the increased exposure can be seen in Fig 6. The astrological day and night cycle reported for each individual photo taken is reported in the column `daynight` in our CSV file, and is calculated using Astral Python library [10].

Each lamppost image is post-processed on the server-side to identify whether the lamppost is ON or OFF based on its expected behaviour and day/night cycle. More specifically, all photos are converted to a 3-channel Red-Green-Blue (RGB) format with a range of values between 0 to 255. The values are reported in our CSV file under the `red`, `green` and `blue` columns. During the nighttime, the results reported under the RGB columns come from the highly exposed photos. When the three channels are reported, the streetlight status is calculated, i.e., operational or not (`fault_detected` column in Tab. 2), as well as a confidence interval (`confidence` column in Tab. 2). As operational is considered a streetlight when it is turned OFF during the day or turned ON during the night cycles. On the other hand, a fault is reported when a streetlight is ON during the day or OFF during the night. An example of how the data are reported and are visualised in our frontend is shown in Fig. 4.

With regards the fault detection mechanism, during the day all images are validated against a ML-trained model. Our model is based on a pre-trained VGG-16 model (using the Imagenet dataset). We initially optimised the final layers of the model for the lamppost ON/OFF classification, followed by some fine tuning for the entire model. The output of the model is a Softmax layer working as a binary classifier (ON/OFF). Our ML model operates with an accuracy of circa $\geq 90\%$. The neural network architecture can be seen in Fig 5. More information about the initial VGG-16 model can be found in [11]. During the night, the RGB values of the node are used for identifying its’ status. For the night use-case we use the highly exposed images and the (raw) RGB values. More specifically, when the median RGB value is ≥ 200 , the lamppost is considered as ON and the confidence is nearly 1. On the other hand, when $RGB_{\text{median}} \leq 100$ the lamppost is considered as OFF and the confidence is 1 again. For the values





Serial number	Date	Lamppost	Section	Location	Image taken	Analysed	Status	Confidence	Daynight
RSE-6-C	06/02/2022 - 08:10	686	M32 - CHL	51.506188581,-2.538194686		Yes	Operational	0.99700540304184	Day
RSS-A-26-C	06/02/2022 - 08:09	438	LGR - FBR	51.499652,-2.480762		Yes	Operational	0.97086906433105	Day
RSS-A-12-C	06/02/2022 - 08:09	496	LGR - FBR	51.505666378,-2.490806003		Yes	Operational	0.98620915412903	Day
RSE-16-C	06/02/2022 - 08:09	530	FBR - BHR	51.504182555,-2.501398829		Yes	Operational	0.8192572593689	Day

Figure 4: An example of a number of images collected and analysed, as seen on our frontend portal. Their status, the confidence of the prediction, the day/night cycle and information about the nodes (serial ID, location, etc.) are reported along with the photo taken.




Figure 5: The neural network architecture used for our “day” time classification.

$100 < RGB_{\text{median}} < 200$ the green channel is used a second step for validating the status of the lamppost. It was shown that the images produced by nodes with increased vegetation around them (as in Fig. 6-right), are more prevalent on the green channel, thus when the green channel is ≥ 200 , then the streetlight is considered as ON. We list these results with a confidence of 0.5.

Acknowledgments

This work is funded in part by Toshiba Europe Ltd. UMBRELLA project is funded in conjunction with South Gloucestershire Council by the West of England Local Enterprise Partnership through the Local Growth Fund, administered by the West of England Combined Authority.

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.
- The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

References

- [1] G. K. Wallace, “The JPEG Still Picture Compression Standard,” *Commun. ACM*, vol. 34, no. 4, p. 30–44, Apr. 1991. [Online]. Available: <https://doi.org/10.1145/103085.103089>
- [2] BRIL Toshiba Europe Ltd., “UMBRELLA Platform and Testbed,” <https://www.umbrellaiot.com>, 2022, Accessed: 2022-1-31.



Figure 6: Examples of correctly or falsely labelled photos taken from our system for both “day” and “night” cases.

- [3] Raspberry Pi, “Raspberry Pi Camera Module ver.1 Datasheet,” https://www.arducam.com/downloads/modules/RaspberryPi_camera/RaspberryPi_Camera_Module_DS_V3.0.pdf, 2015, Accessed: 2022-1-31.
- [4] OmniVision, “Camera Sensor OV5647 Datasheet,” https://cdn.sparkfun.com/datasheets/Dev/RaspberryPi/ov5647_full.pdf, 2009, Accessed: 2022-1-31.
- [5] I. Mavromatis, A. Stanoev, P. Carnelli, Y. Jin, M. Sooriyabandara, and A. Khan, “Images of Public Streetlights with Operational Monitoring using Computer Vision Techniques,” Apr. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.6046759>
- [6] S. Y. Kadirova and D. I. Kajtsanov, “A Real Time Street Lighting Control System,” in *Proc. of International Conference on Electrical Machines, Drives and Power Systems (ELMA) 2017*, Jun. 2017, pp. 174–178. [Online]. Available: <https://doi.org/10.1109/ELMA.2017.7955426>
- [7] Y. Yao, B. Yu, C. Gong, and T. Liu, “Understanding How Pretraining Regularizes Deep Learning Algorithms,” *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–13, Dec. 2021. [Online]. Available: <https://doi.org/10.1109/TNNLS.2021.3131377>
- [8] Z. Li, Q. Zeng, Y. Liu, J. Liu, and L. Li, “An Improved Traffic Lights Recognition Algorithm for Autonomous Driving in Complex Scenarios,” *International Journal of Distributed Sensor Networks*, vol. 17, no. 5, May 2021. [Online]. Available: <https://doi.org/10.1177%2F15501477211018374>
- [9] Raspberry Pi, “Raspberry Pi Compute Module 3b+,” <https://www.raspberrypi.org/products/compute-module-3-plus/>, 2021, Accessed: 2022-1-31.
- [10] Simon Kennedy, “Astral Library - Python calculations for the position of the sun and moon.” <https://github.com/sffjunkie/astral>, 2022, Accessed: 2022-1-31.
- [11] A. Z. Karen Simonyan, “Very Deep Convolutional Networks for Large-Scale Image Recognition,” in *Proc. of International Conference on Learning Representations 2015*, May 2015. [Online]. Available: <https://arxiv.org/abs/1409.1556>