

Multi-Camera Vehicle Counting Using Edge-AI

Revised Manuscript (with changes marked)

Luca Ciampi^a (luca.ciampi@isti.cnr.it), Claudio Gennaro^a
(claudio.gennaro@isti.cnr.it), Fabio Carrara^a (fabio.carrara@isti.cnr.it),
Fabrizio Falchi^a (fabrizio.falchi@isti.cnr.it), Claudio Vairo^a
(claudio.vairo@isti.cnr.it), Giuseppe Amato^a (giuseppe.amato@isti.cnr.it)

^a Institute of Information Science and Technologies of the National Research Council
of Italy (ISTI-CNR), via G. Moruzzi 1 - 56124, Pisa, Italy

Corresponding Author:

Luca Ciampi

Institute of Information Science and Technologies of the National Research Council
of Italy (ISTI-CNR), via G. Moruzzi 1 - 56124, Pisa, Italy

Tel: +39 050 6213054

Email: luca.ciampi@isti.cnr.it

Abstract

This paper presents a novel solution to automatically count vehicles in a parking lot using images captured by smart cameras. Unlike most of the literature on this task, which focuses on the analysis of *single* images, this paper proposes the use of multiple visual sources to monitor a wider parking area from different perspectives. The proposed multi-camera system is capable of automatically estimating the number of cars present in the *entire* parking lot directly on board the edge devices. It comprises an on-device deep learning-based detector that locates and counts the vehicles from the captured images and a decentralized geometric-based approach that can analyze the inter-camera shared areas and merge the data acquired by all the devices. We conducted the experimental evaluation on an extended version of the *CNRPark-EXT* dataset, a collection of images taken from the parking lot on the campus of the National Research Council (CNR) in Pisa, Italy. We show that our system is robust and takes advantage of the redundant information deriving from the different cameras, improving the overall performance without requiring any extra geometrical information of the monitored scene.

Keywords: Smart Parking, Counting Objects, Edge AI, Counting Vehicles, Smart Mobility, Deep Learning

1. Introduction

2 Traffic-related issues are constantly increasing, and tomorrow's cities cannot
3 be considered intelligent if they do not enable smart mobility. Smart mobility
4 applications, such as smart parking and road traffic management, are nowadays
5 widely employed worldwide, making our cities more livable and bringing benefits
6 to the cities and, consequently, to our lives.

7 Images are perhaps the best sensing modality to perceive and assess the flow
8 of vehicles in large areas. Like no other sensing mechanism, city camera net-

9 works can monitor large areas while simultaneously providing visual data to AI
10 systems to extract relevant information from this deluge of data. However, this
11 application is often hampered by the massive flow of data that must be sent to
12 central servers or the cloud for processing. On the other hand, edge computing
13 is a recent paradigm that promotes the decentralization of data processing to
14 the border, i.e., where the data are generated, thus reducing the traffic on the
15 network and the pressure on central servers. No wonder that combination of
16 recent Computer Vision deep learning-based techniques and the edge comput-
17 ing paradigm is an emerging trend, as witnessed, for example, by Khan et al.
18 (2019) that tackles the face recognition task or by Amato et al. (2019b); Ciampi
19 et al. (2020a) that instead can detect people directly onboard surveillance cam-
20 eras. Nonetheless, this promising paradigm brings along with it also some new
21 challenges related to the limited computational resources on the disposable edge
22 devices and also concerning security inside IoT networks (Ujjan et al., 2020).

23 In this work, we tackle the problem of estimating the number of vehicles
24 present in a parking lot using images captured by smart cameras. Whereas
25 classic car counting solutions are sensor-based (e.g., entrance-level photocells,
26 per-space ground sensors), vision-based solutions provide several advantages,
27 such as a) flexibility, as cameras can adapt to more challenging configurations
28 of parking spaces (e.g., undelimited parking lots with non-fixed spaces), b) lower
29 hardware and maintenance cost, as smart cameras can cost few tens of dollars
30 while each monitoring multiple parking spaces, and c) being multi-purpose, as
31 the same hardware can be used to perform additional tasks (e.g., surveillance).
32 However, this vision-based counting task is challenging as the process of un-
33 derstanding the captured images faces many problems, such as shadows, light
34 variation, weather conditions, and inter-object occlusions. Although most of
35 the existing works concerning the vehicle counting task focus on the analysis of
36 *single* images, in many real-world scenarios, one can benefit from using multiple
37 cameras to monitor the same parking lot from different perspectives and view-
38 points. Furthermore, multiple neighboring cameras can also help cover a wider
39 area. At the same time, such an approach introduces issues related to merg-

40 ing the knowledge extracted from the single cameras with partially overlapping
41 fields of views (FOVs), as shown in Figure 1.

42 In this paper, we propose a novel solution to improve car counting when
43 scaled up with multi-camera setups. Specifically, we introduce a multi-camera
44 system that estimates the number of cars present in the *entire* parking lot by
45 combining a state-of-the-art Convolutional Neural Network (CNN), which can
46 locate and count vehicles present in images belonging to individual cameras,
47 along with a decentralized geometry-based approach that is responsible for ag-
48 gregating the data gathered from all the devices. Our solution performs the task
49 directly on the edge devices (i.e., the smart cameras) without using a central
50 server or cloud, consequently reducing the communication overhead. The total
51 count is built exploiting the partial results computed in parallel by the single
52 cameras and propagated through messages. Hence, our system scales better
53 when the number of monitored parking spaces increases. Moreover, our solu-
54 tion does not require any manual intervention or any extra information about
55 the monitored parking area, such as the location of the parking spaces, nor any
56 geometric information about the camera positions in the parking lot. In short,
57 it is a flexible and ready-to-use solution that allows a simple “plug-and-play”
58 insertion of new cameras into the system.

59 To validate our multi-camera solution, we employed the *CNRPark-EXT*
60 dataset (Amato et al., 2017), a collection of images taken from the parking
61 lot on the campus of the National Research Council (CNR) in Pisa, Italy. The
62 pictures are acquired by multiple cameras having partially overlapping fields of
63 view and describing challenging scenarios with different perspectives, illumina-
64 tions, weather conditions, and many occlusions. Since the annotations of this
65 dataset concern single images, we extended it by manually labeling a part of
66 it to be consistent with our algorithm that instead considers the entire parking
67 area. We conducted extensive experiments testing the generalization capabil-
68 ities of the CNN-based technique responsible for detecting vehicles in single
69 images and the effectiveness of our multi-camera algorithm, demonstrating that
70 our system is robust and benefits from the redundant information deriving from



Figure 1: An example of two cameras monitoring the same parking area with partially overlapping fields of view. This redundancy provides robustness and fault-tolerance but also raises the problem of aggregating knowledge extracted from the individual cameras.

71 the different cameras improving the overall performance.

72 To summarize, the main contributions of this work are the followings:

- 73 • We introduce a novel multi-camera system able to automatically estimate
74 the number of cars present in the *entire* monitored parking area. It runs
75 directly on the edge devices and combines a deep learning-based detector
76 together with a decentralized technique that exploits the geometry of the
77 captured images.
- 78 • We specifically extend the *CNRPark-EXT* dataset (Amato et al., 2017),
79 a collection of images acquired by multiple cameras having partially over-
80 lapping fields of views and describing various parking lots. We manually
81 label a subset of it, making it suitable for our considered scenario in which
82 we consider the whole parking area.
- 83 • We conduct an experimental evaluation showing that our system is ro-
84 bust, flexible, and can benefit from redundant information from different

85 cameras while improving overall performance.

86 We organize the rest of the paper as follows. Section 2 reports other works
87 present in the literature related to our topic. Section 3 describes our multi-
88 camera counting algorithm. Section 4 states the experimental setup, describing
89 the dataset, the metrics, and the implementation details. Section 5 presents and
90 discusses the experiments and the obtained results. Finally, Section 6 concludes
91 the paper with some insights on future directions.

92 **2. Related Work**

93 This section overviews some works related to our, organizing them into two
94 categories. The first one concerns the counting task, while the second regards
95 multi-camera parking lot monitoring systems.

96 *2.1. The counting task*

97 The counting task estimates the number of object instances in still images
98 or video frames (Lempitsky & Zisserman, 2010). This topic has recently at-
99 tracted much attention due to its inter-disciplinary and widespread applicability
100 and paramount importance for many real-world applications. Examples include
101 counting bacterial cells from microscopic images (Xie et al., 2016; Ciampi et al.,
102 2022), estimating the number of people present at an event (Boominathan et al.,
103 2016; Benedetto et al., 2022), counting animals in ecological surveys to moni-
104 tor the population of a specific region (Arteta et al., 2016) and evaluating the
105 number of vehicles on a highway or in a car park (Amato et al., 2019a).

106 Several machine learning-based solutions (especially supervised) have been
107 suggested in the last years. Following the taxonomy adopted in Sindagi & Patel
108 (2018), we can broadly classify existing counting approaches into two categories:
109 counting by regression and counting by detection. Counting by *regression* is
110 a supervised method that tries to establish a direct mapping (linear or not)
111 from the image features to the number of objects present in the scene or a
112 corresponding density map (i.e., a continuous-valued function), skipping the

113 challenging task of detecting instances of the objects (Zhang et al., 2016, 2017;
114 Oñoro-Rubio & López-Sastre, 2016; Ciampi et al., 2020b, 2021). Counting by
115 *detection* is, instead, a supervised approach where we localize instances of the
116 objects, and then we count them (Amato et al., 2018; Ciampi et al., 2018). While
117 regression-based techniques work very well in very crowded scenarios where the
118 single object instances are not well defined due to inter-class and intra-class
119 occlusions, they perform poorly in images with a large perspective and oversized
120 objects. Another remarkable drawback of the regression-based approaches is
121 that they cannot precisely localize the objects present in the scene, eventually
122 providing only a coarse position of the area in which they are distributed.

123 In this work, we estimate the number of vehicles present in a park area from
124 images collected by smart cameras having large perspectives. The cars close
125 to the cameras are much larger than those far away from them. Therefore, we
126 employ a detection-based method. Furthermore, another reason which led us to
127 discard counting by regression approaches is that we need to know the precise
128 localization (with boundaries) of the detected vehicles. Most of the existing
129 counting solutions do not directly deal with edge computing devices and the
130 consequent constraints due to the limited available computing resources. They
131 use deep learning-based approaches that typically require the use of a GPU
132 and that are computationally expensive. Moreover, they consider the images
133 as single entities. They do not account for the possible benefits of monitoring
134 the same lots from different perspectives or covering a wider parking area with
135 multiple cameras. Instead, our solution runs directly on the edge devices and
136 can estimate the number of vehicles present in the entire parking lot.

137 *2.2. Multi-camera parking lot monitoring*

138 Only a few works addressed parking lot monitoring considering a multi-
139 camera scenario. In Nieto et al. (2019), the authors applied a homography to
140 project the detected vehicles from the plane of each camera to a common plane,
141 where they performed a perspective correction to correct matching between
142 the vehicle detections and the parking spots. Also, the authors in Vitek &

143 Melničuk (2017) proposed a multi-camera system to classify parking spaces as
144 vacant or occupied. In this solution, the acquired images are processed onboard
145 Raspberry Pi devices. The extracted information about the status of parking
146 spaces is then transmitted to a central server, which evaluates the parking spaces
147 in the overlapping areas. Their algorithm is based on the histogram of oriented
148 gradients (HOG)(Dalal & Triggs, 2005) feature descriptor and support vector
149 machine (SVM) classifier. Since the HOG feature descriptor cannot adequately
150 describe rotated vehicles, the authors have provided a descriptor with additional
151 information about rotation to increase the system accuracy.

152 However, these solutions rely on prior knowledge of the monitored scene, such
153 as the position of the parking spaces or some geometric information concerning
154 the parking area. For instance, the proposed system in Nieto et al. (2019)
155 requires manually annotating the corners of the parking area and the number of
156 spots. In essence, a preliminary annotation of the new areas and a new training
157 phase of the algorithm are often mandatory operations. Consequently, these
158 techniques are not very flexible. On the other hand, we propose a simple yet
159 effective solution that does not need any extra information about the monitored
160 scene. The smart cameras can automatically localize and count the vehicles
161 present in their field of view, propagating the single results to the other edge
162 devices through messages. A decentralized technique, again running directly on
163 the edge devices, is instead in charge of analyzing and merging these results,
164 exploiting the captured images geometry, and automatically outputs the number
165 of cars present in the entire parking area.

166 **3. Proposed approach**

167 *3.1. Overview*

168 In this section, we describe our multi-camera counting algorithm. We based
169 our system on the parallel processing of each of the smart cameras followed by
170 the fusion of their results to estimate the number of vehicles present in the *entire*
171 parking area.

172 Figure 2 shows an example of our multi-camera counting system, together
 173 with its graphical representation. We model our system as a graph G , comprised
 174 of n nodes ν_i and one Sink node S , $V = \{\nu_1, \nu_2, \dots, \nu_n, S\}$. Each node ν_i
 175 represents an independent edge device, i.e., a smart camera in our case. Two
 176 nodes ν_i and ν_j are considered neighbors if their FOVs overlap. In this case,
 177 a directed edge of the graph connects them. Each edge device ν_i can capture
 178 images, localize and count the vehicles present in its FOV exploiting a deep
 179 learning-based detector, and communicate with its neighboring nodes through
 180 messages m_i containing the cars detections. Furthermore, each node ν_i can also
 181 run a local counting algorithm in charge of computing partial counting results
 182 concerning the estimation of the number of vehicles present in overlapped areas
 183 between its FOV and the ones belonging to its neighbors.

184 The fusion of the partial results is performed by the Sink node S , which is
 185 also in charge of providing the final result and synchronizing all the algorithm
 186 steps through synchronization signals headed towards the other nodes ν_i . On
 187 the other hand, the nodes ν_i can also communicate through messages with the
 188 Sink node. Messages can be of two types: i) messages η_i containing the number
 189 of cars captured by the node ν_i in its FOV, and ii) messages $\mu_{j,i}$ representing
 190 the partial counting estimation related to the overlapping area between two
 191 neighboring nodes ν_i and ν_j .

192 In the following sections, we describe all the steps of our algorithm in detail.
 193 First, in Section 3.2, we outline the automatic system initialization performed by
 194 the smart cameras themselves, in which they compute the homographic trans-
 195 formations between the scene they are monitoring and the scene observed by the
 196 neighboring cameras. Then, in Section 3.3, we describe the CNN-based local
 197 counting algorithm that runs on each of the smart cameras and the geometric-
 198 based technique helpful for the overlapped areas. Finally, in Section 3.4, we
 199 depict the global counting algorithm responsible for the fusion of these individ-
 200 ual and partial results, and that finally outputs the number of cars present in
 201 the *entire* parking area.




Figure 2: An example of our multi-camera counting system, with $n = 5$ smart cameras. We model it as a graph G , comprised of n nodes ν_i (one for each camera) and one Sink node S , $V = \{\nu_1, \nu_2, \dots, \nu_n, S\}$. Each node ν_i can capture images, localize and count the vehicles present in its FOV, and communicate with its neighboring nodes through messages m_i containing these detections. Moreover, each node ν_i can run a local counting algorithm in charge of computing partial counting results concerning the overlapped areas between its FOV and the ones belonging to its neighbors, exploiting images geometry. These partial results are sent through messages to the Sink node S , which is responsible for their fusion and provides the final result. Messages to S can be of two types: i) η_i containing the number of cars captured by the node ν_i in its FOV, and ii) $\mu_{j,i}$ representing the partial counting estimation related to the overlapping area between two neighboring nodes ν_i and ν_j .

202 *3.2. Initialization*

203 This step is aimed at *automatically* initializing the system, estimating the
 204 geometric relationship between each node (i.e., each scene monitored by a smart
 205 camera) and its neighbors. The only hypotheses we impose are i) each smart
 206 camera is aware of the IP addresses of its neighbors, i.e., the cameras having
 207 the field of view overlapped with its own; ii) the Sink node S is aware of the IP
 208 addresses of all the smart cameras belonging to the system.

209 The Sink node S starts the initialization phase, sending a synchronization
 210 signal to the other nodes. Once received, each smart camera captures an image
 211 of the scene it monitors and sends it to all its neighbors. Once a smart camera
 212 i receives an image from a neighboring camera j , it computes a homographic
 213 transformation $H_{j,i}$ between the image j and the image i describing its mon-
 214 itored scene. This allows us to establish a correspondence between the points
 215 belonging to the pair of images taken by the two cameras, which will be used
 216 subsequently in the algorithm. We formalized the system initialization for a

217 generic node ν_i in the Algorithm 1.

218 However, finding this homography can be challenging because neighboring
219 cameras can have different angles of view, leading to a perspective distortion be-
220 tween the images captured by them. Given a pair of neighboring nodes ν_i, ν_j , we
221 employ a procedure that starts with finding the SIFT (Lowe, 2004) key-points
222 and feature descriptors of the images i, j captured by the two nodes. Then, we
223 match the two sets of feature descriptors by performing David Lowe’s ratio test
224 (Lowe, 2004), and we further filter the matched feature descriptors by keeping
225 only the pairs whose euclidean distance is below a given threshold. Finally, we
226 obtain the homographic transformation by applying the random sample con-
227 sensus (RANSAC (Fischler & Bolles, 1981)) algorithm to the filtered feature
228 descriptors. All these computations are performed *automatically* without the
229 need of any extra geometric information about the monitored scene, and no
230 manual intervention is needed. Figure 3 shows the concatenation of two neigh-
231 boring images i and j in which we apply the found homographic matrix to the
232 image i , to have the same perspective as the image j .

Algorithm 1 : Initialization

At each Initialization Signal by S , each node ν_i performs the following steps:

- 1: RECEIVEINITSIGNAL() \triangleright waits the initialization signal from S
 - 2: $\text{image}_i \leftarrow \text{CAMERACAPTURE}()$
 - 3: **for each** $j \in J$ **do** $\triangleright J$ is the set of neighboring nodes of node ν_i
 - 4: SENDIMAGE(image_i, ν_j) \triangleright sends image_i to node ν_j
 - 5: $\text{image}_j \leftarrow \text{RECEIVEIMAGE}()$ \triangleright receives image_j from node ν_j
 - 6: $H_{j,i} = \text{COMPUTEHOLOGRAPHY}(\text{image}_j, \text{image}_i)$
-

233 3.3. Local Counting Algorithm

234 This section describes the local counting algorithm that runs directly on-
235 board the edge devices. It combines a CNN-based counting technique in charge
236 of the localization and the estimation of the number of vehicles present in the
237 acquired single images, i.e., the contents of the messages m_i and the quantities



Figure 3: Example of concatenation of two images using a homographic transformation, where it is also visible the overlapping area between them.

238 η_i shown in Figure 2, together with a geometric-based approach responsible of
 239 estimating the number of vehicles present in the overlapping areas between the
 240 nodes and their neighbors, i.e., the quantities $\mu_{j,i}$.

241 *A vehicle counting CNN on the Edge.* Each smart camera needs to indepen-
 242 dently detect and count vehicles from its captured frame. For this step, every
 243 approach providing precise localization of the detected vehicles in the pixel
 244 space is suitable, and the choice of a particular approach should be guided by
 245 resource constraints, e.g., available memory, prediction frequency, or energy con-
 246 sumption, if any. Here, we base our vehicle counting technique on *Mask R-CNN*
 247 (He et al., 2017), a popular deep CNN for instance segmentation that operates
 248 within the ‘recognition using regions paradigm’ (Gu et al., 2009). In particular,
 249 it extends the *Faster R-CNN* detector (Ren et al., 2017) by adding a branch
 250 that outputs a binary mask saying whether or not a given pixel is part of an
 251 object. Briefly, a CNN acts as a backbone in the first stage, extracting the input
 252 image features. Starting from this feature space, another CNN named Region
 253 Proposal Network (RPN) generates region proposals that might contain objects.
 254 RPN slices pre-defined region boxes (called anchors) over this space and ranks
 255 them, suggesting those most likely containing objects. Once RPN produces the
 256 Regions Of Interests (ROIs), they might be of different sizes. Since it is hard
 257 to work on features having different sizes, RPN reduces them into the same di-
 258 mension using the Region of Interest Pooling algorithm. Finally, these fixed-size

259 proposals are processed by two parallel CNN-based branches: one is responsi-
260 ble for classifying and localizing the objects inside them with bounding boxes;
261 the second produces a binary mask that says whether or not a given pixel is
262 part of an object. In the end, given an input image, the network produces per-
263 pixel masks localizing the detected objects together with the associated labels
264 classifying them.

265 To make our counting solution able to run efficiently directly on the edge de-
266 vices, we employ, as a backbone, the *ResNet50* architecture, a lighter version of
267 the popular *ResNet101* (He et al., 2016). This simplification is also justified be-
268 cause the more powerful version of Mask R-CNN based on the ResNet101 model
269 was designed for more complicated visual detection tasks than ours. Originally,
270 Mask R-CNN was trained on the *COCO* dataset (Lin et al., 2014) to detect
271 and recognize 80 different classes of everyday objects. In our case, we have
272 to localize and identify objects belonging to just one category (i.e., the *vehicle*
273 category). To this end, we further simplify the model by reducing the number
274 of the final fully convolutional layers responsible for the classification of the de-
275 tected objects, making the model lighter. Once we have localized the instances
276 of the objects, we count them estimating the number of vehicles present in the
277 scene.

278 *Local counting.* The Sink node S starts this phase, sending a synchronization
279 signal to all the smart cameras belonging to the system. Once received the syn-
280 chronization signal, each node ν_i captures an image belonging to its underlying
281 FOV and feeds it to the previously described CNN-based counting technique
282 obtaining a set of masks masks_i localizing the vehicles present in the scene. The
283 cardinality of this set of masks corresponds to the number of cars present in
284 the image, i.e., the quantity η_i , that is sent with a message to the Sink node S .
285 Then, the node ν_i packs this set of masks masks_i in a message m_i , sends it to
286 all its neighboring nodes ν_j , and receives from them their corresponding set of
287 masks masks_j packed in a message m_j . Once received a message m_j , the node
288 ν_i is responsible for analyzing the potential vehicles present in the overlapped

289 area between its FOV and the one of the node ν_j . To this end, it employs the
 290 homographic transformation $H_{j,i}$ computed during the system initialization, as
 291 described in Section 3.2. Specifically, it projects the masks belonging to the set
 292 masks_j into its image plane, filtering them and discarding the ones that overlap
 293 with the masks belonging to the set masks_i having a value of Intersection over
 294 Union (IoU) greater than a threshold that we empirically found to be optimal
 295 at 0.2. These masks indeed localize vehicles already detected, which should not
 296 be considered a second time. On the other hand, the cars left after this filtering
 297 are vehicles that were not detected in the FOV underlying the node ν_i , but
 298 instead found by the node ν_j , probably because of having a better view of this
 299 object. Referring to our graph modeling the system and reported in Figure 2,
 300 the number of the discarded cars after this filtering operation corresponds to
 301 the message $\mu_{j,i}$, that is sent to the Sink node S . We detail all the described
 302 steps in the Algorithm 2 and in the Procedure 3.

Algorithm 2 : Local Counting

At each Computational Signal by S , each node ν_i performs the following steps:

- 1: RECEIVECOMPUTSIGNAL() ▷ waits the computational signal from S
 - 2: $\text{image}_i \leftarrow \text{CAMERACAPTURE}()$
 - 3: $\text{masks}_i \leftarrow \text{MASKRCNN}(\text{image}_i)$
 - 4: $\eta_i \leftarrow |\text{masks}_i|$
 - 5: SENDMESSAGE(η_i, S) ▷ sends η_i to Sink node S
 - 6: $m_i \leftarrow \text{PACKMESSAGE}(\text{masks}_i)$ ▷ builds message m_i containing masks_i
 - 7: **for each** $j \in J$ **do** ▷ J is the set of neighboring nodes of node ν_i
 - 8: SENDMESSAGE(m_i, ν_j) ▷ sends m_i to node ν_j
 - 9: $m_j \leftarrow \text{RECEIVEMESSAGE}()$ ▷ receives message m_j from node ν_j
 - 10: $\text{masks}_j \leftarrow \text{UNPACKMESSAGE}(m_j)$ ▷ unpacks m_j containing masks_j
 - 11: $\mu_{j,i} \leftarrow \text{COMPUTE_}\mu(\text{masks}_i, \text{masks}_j, H_{j,i})$
 - 12: SENDMESSAGE($\mu_{j,i}, S$) ▷ sends $\mu_{j,i}$ to Sink node S
-

Algorithm 3 : Computation of μ

μ represents the num of cars detected by ν_j and already detected by ν_i

Each node ν_i performs the following procedure:

```
1: procedure COMPUTE_μ(masksi, masksj, Hj,i)
2:   n_cars_already_detected ← 0
3:   for each mask ∈ masksj do
4:     maskh ← PROJECT(Hj,i, mask) ▷ projects mask points on plane i
5:     if maskh falls within imagei then
6:       maskmax ← arg maxm ∈ masksi IoU(maskh, m)
7:       if IoU(maskh, maskmax) > τ then
8:         n_cars_already_detected ++
9:   return n_cars_already_detected
```

303 *3.4. Global Counting Algorithm*

304 In this section, we describe the global counting algorithm that runs on the
305 Sink node S , responsible for the fusion of the partial results coming from all the
306 other nodes, and that finally outputs the number of cars present in the *entire*
307 monitored parking area.

308 This phase starts when S receives all the η_i and the $\mu_{j,i}$ messages, i.e.,
309 the number of vehicles estimated in the single FOVs and the estimation of the
310 number of cars already considered in the overlapping areas between neighbor-
311 ing cameras, from all the nodes belonging to the system. Specifically, for each
312 overlapped area shared between a pair of nodes ν_i, ν_j , the node S receives two
313 messages $\mu_{j,i}$ and $\mu_{i,j}$, the contents of which are computed by the two nodes
314 employing two homographic transformations $H_{j,i}$ and $H_{i,j}$, respectively. These
315 two quantities can be potentially different. We choose the best value by aggre-
316 gating them, choosing between three different functions - max, min and mean,
317 finding that the latter is the best one. Finally, the node S builds the final result,
318 i.e., the estimation of the number of vehicles present in the *entire* parking lot,
319 by summing up the content of all the η_i messages and subtracting the computed
320 aggregated values. We detail all these steps in the Algorithm 4.

Algorithm 4 : Global Counting

The Sink node S performs the following steps:

- 1: **for each** $(\mu_{i,j}, \mu_{j,i})$ **do**
 - 2: $\bar{\mu}_k \leftarrow \text{AGGREGATE}(\mu_{i,j}, \mu_{j,i})$
 - 3: $\text{global_cars_count} \leftarrow \sum_{n=1}^N \eta_n - \sum_{k=1}^K \bar{\mu}_k$
 $\triangleright N$ is the set of nodes, K is the set of aggregations
-

321 4. Experimental Setup

322 In this section, we describe the simulated scenario that we exploited for our
323 experiments. In particular, we extended the *CNRPark-EXT* dataset (Amato
324 et al., 2017), adapting it to be suitable for the counting task so that it was
325 usable for training the vehicles counting CNN running on the smart cameras
326 and applicable to validate our multi-camera algorithm. Furthermore, we briefly
327 describe the *PKLot* dataset (de Almeida et al., 2015), a public dataset compris-
328 ing parking lot scenes that we exploited for further assessing the generalization
329 capabilities of the local vehicles counting network. Then, we illustrate the em-
330 ployed evaluation metrics, and, finally, we report some implementation details.

331 4.1. The *CNRPark-EXT* Dataset

332 In this work, we exploit the *CNRPark-EXT* public dataset introduced in
333 Amato et al. (2017), a collection of annotated images of vacant and occupied
334 parking spaces on the campus of the National Research Council (CNR) in Pisa,
335 Italy. This dataset represents most of the challenging situations that can be
336 found in a real scenario: nine different cameras capture the images under var-
337 ious weather conditions, angles of view, light conditions, and many occlusions.
338 Furthermore, the cameras have their fields of view partially overlapped. Since
339 this dataset is specifically designed for parking lot occupancy detection, it is not
340 directly usable for the counting task. Indeed, each image, called *patch*, contains
341 one parking space labeled according to its occupancy status - 0 for vacant and
342 1 for occupied. Since this work aims at counting the cars present in the parking

343 area, we extended it by considering the full images and adapting the ground
344 truth to our purposes.

345 Specifically, we created a suitable label set to train and evaluate the local ve-
346 hicles counting based on Mask R-CNN. In this case, labels correspond to *binary*
347 *masks*, i.e., binary images identifying the polygons surrounding the vehicles we
348 want to detect. Since mask creation is a very time-consuming operation, dif-
349 ferently from our previous work (Ciampi et al., 2018), we considered the *raw*
350 masks obtained directly from the bounding boxes localizing the occupied park-
351 ing spaces. The idea is that we do not need precise polygons that identify the
352 vehicles we want to detect. Still, we can use the region within the delimiters
353 that identify the occupied parking spaces and the underlying part of the car.

354 On the other hand, to validate our multi-camera algorithm, we built a simu-
355 lated scenario considering some sequences of images belonging to different cam-
356 eras captured simultaneously. In other words, a sequence is defined as the set of
357 images captured by the different smart cameras that are monitoring the parking
358 area at the same moment. Hence, a sequence represents a snapshot of the *entire*
359 parking lot at a given timestamp, and it takes into account all the spaces from
360 the available different views. We manually annotated these sequences to obtain
361 the ground truth car counts. Specifically, we considered the single images com-
362 posing a sequence, counting the vehicles present in the scenes, but taking care of
363 accounting for them just once if they appear in more than one view, i.e., discard-
364 ing the cars from the global count if they were located in the overlapping areas.
365 We labeled six different sequences, two for each weather condition, considering
366 the images belonging from camera₂ to camera₉. We did not consider camera₁
367 since it has small and particularly skewed field-of-view overlaps with the other
368 cameras, hindering the automatic homography estimation and the subsequent
369 projections.

370 4.2. The PKLot Dataset

371 To further validate the generalization capabilities of the CNN-based local
372 vehicles counting algorithm, we exploited an additional public dataset, named

373 *PKLot* (de Almeida et al., 2015). In particular, this dataset is composed by
 374 three different scenarios describing three different parking lot scenes - *UFPR04*,
 375 *UFPR05* and *PUC*. We considered only the first two subsets since the third one
 376 contains images captured from a fixed camera located at the height of the 10th
 377 floor of a building, which provides a slanted view of the parking lot and results
 378 in a different setting without intra-vehicle occlusions. Since also the *PKLot*
 379 dataset, like the *CNRPark-EXT* one, is specifically designed for the parking
 380 lot occupancy detection task, we manually re-labeled the ground truth for our
 381 purposes as already described in Section 4.1, obtaining a simulation scenario
 382 suitable for measure the performance of our solution for the counting task.

383 4.3. Evaluation Metrics

384 Following other counting benchmarks, we exploited Mean Absolute Error
 385 (*MAE*), Mean Square Error (*MSE*), and Mean Relative Error (*MRE*) as the
 386 metrics for the performance evaluation, defined as follows:

$$MAE = \frac{1}{N} \sum_{n=1}^N |c_n^{gt} - c_n^{pred}|, \quad (1)$$

$$MSE = \frac{1}{N} \sum_{n=1}^N (c_n^{gt} - c_n^{pred})^2, \quad (2)$$

$$MRE = \frac{1}{N} \sum_{n=1}^N \frac{|c_n^{gt} - c_n^{pred}|}{num_spaces_n}, \quad (3)$$

387 where N is the total number of the images, c_{gt} , c_{pred} and num_spaces_n are
 388 the actual count, the predicted count, and the total number of parking spaces
 389 of the n -th image, respectively. Note that as a result of the squaring of each
 390 difference, *MSE* effectively penalizes large errors more heavily than small ones
 391 and thus should be more useful when large errors are particularly undesirable.
 392 On the other hand, *MRE* also considers the relation between the error and the
 393 total number of objects present in the image.

394 4.4. Implementation Details

395 We report in this section some implementation details concerning the Mask
396 R-CNN-based algorithm responsible for the prediction of the number of vehi-
397 cles in the single images. In particular, we trained the modified Mask R-CNN
398 initializing the weights of the ResNet50 backbone with the ones of a pre-trained
399 model on *ImageNet* (Deng et al., 2009), a popular dataset for classification
400 tasks, and the remaining ones at random. We froze the backbone for the first
401 10 epochs, and then we trained the whole network for 20 additional epochs.
402 We used Stochastic Gradient Descent (SGD) to perform the CNN parameters
403 update. Concerning the Region Proposal Network, explained in Section 3.3, we
404 exploited a set of five anchors of sizes 16, 32, 64, 128, and 256 pixels. To prevent
405 overfitting, we applied some standard augmentation techniques to the training
406 data: images are horizontally flipped with a 0.5 probability, then their pixels are
407 multiplied by a random value between 0.8 and 1.5, and finally, they are blurred
408 using a Gaussian kernel with a standard deviation of a random value between
409 0 and 5. Then, to support training multiple images per batch, we resized all
410 pictures to the same size. If an image was not square, we padded it with zeros
411 to preserve the aspect ratio. In the end, we obtained images of size 1024×1024 .
412 At inference time, images were resized and padded with zeros to get a square
413 picture of size 1024×1024 , and no other augmentations took place.

414 5. Experiments and Results

415 In this section, we report the experiments and the obtained results. First, we
416 evaluate the performance against other state-of-the-art solutions of the CNN-
417 based technique responsible for estimating the vehicles in the single images
418 directly onboard the smart cameras, also stressing its generalization capabilities.
419 Then, we validate the effectiveness of our multi-camera algorithm by testing it
420 in the simulated scenario previously described. We demonstrate that our system
421 can benefit from the redundant information deriving from the different cameras,
422 obtaining performance improvements in all the considered counting metrics.

423 *5.1. Experiments on the CNN-based counting solution on the edge*

424 *5.1.1. State-of-the-art comparison*

425 We compared our solution with the results obtained in our previous work
426 Ciampi et al. (2018), where we presented a centralized counting approach based
427 on the original version of Mask R-CNN having the ResNet101 model as a fea-
428 tures extractor, which has been fine-tuned on a very small manually annotated
429 subset of the CNRPark-EXT dataset, starting from the model pre-trained on
430 the *COCO* dataset (Lin et al., 2014). We filtered the detections considering
431 only the predictions related to the car class, and we counted them. Although
432 this solution is very computationally expensive and unsuitable for edge devices,
433 it represents a direct comparison in terms of counting on the same dataset.
434 We also compared our technique against the method proposed in Amato et al.
435 (2017), an approach for car parking occupancy detection based on *mAlexNet*,
436 a deep CNN designed explicitly for smart cameras. This work represents an
437 indirect method for counting cars in a parking lot, as the counting problem is
438 cast as a classification problem: if a parking space is occupied, we increment the
439 total number of cars; otherwise, we do not. We illustrate the results in Table 1,
440 where we also report the performance obtained using the Mask R-CNN network
441 without a preliminary fine-tuning on the CNRPark-EXT dataset. Our solution
442 performs better than the other considered methods, considering all three count-
443 ing metrics. In particular, our approach outperforms the solution introduced
444 in Ciampi et al. (2018), despite the latter employing a more deep and powerful
445 CNN, and it is designed to be used as a centralized-server solution. This is ex-
446 plained by the fact that in Ciampi et al. (2018) the authors fine-tuned the CNN
447 using a tiny dataset. Consequently, the algorithm overfits on the training data,
448 and it cannot generalize over the test subset. It is also worthy of notice that our
449 CNN also outperforms the mAlexNet network, even though the latter knows the
450 exact location of the parking spaces. Figure 4 shows some examples of images
451 belonging to different cameras and different weather conditions together with
452 the masks localizing them computed by our counting solution.

| Method | CNRPark-EXT | | | PKLot | | |
|-----------------------|-------------|-------------|-------------|-------------|--------------|-------------|
| | MAE | MSE | MRE | MAE | MSE | MRE |
| (Amato et al., 2017) | 1.34 | 8.00 | 0.04 | | - | |
| (Ciampi et al., 2018) | 1.05 | 4.41 | 0.03 | | - | |
| ResNet50 Mask R-CNN | 11.20 | 247.40 | 0.30 | 16.90 | 522.40 | 0.48 |
| Our solution | 0.49 | 1.04 | 0.01 | 4.56 | 33.88 | 0.13 |

Table 1: Local Counting: Left-side: results obtained using our counting solution on the edge compared with other state-of-the-art approaches; we get the best results on all the three considered counting metrics. Right-side: evaluation of the generalization capabilities on the *PKLot* dataset (de Almeida et al., 2015), using the model trained on the *CNRPark-EXT* dataset; we achieved an error that is approximately four times lower than the one obtained with the COCO pre-trained model.



(a) Image from Camera₂



(b) Image from Camera₈

Figure 4: Two examples of the output of our counting method. Images are taken from the CNRPark-EXT dataset. We report the predictions and the estimate of the number of vehicles present in the scene.

453 *5.1.2. Generalization capabilities*

454 Errors in vehicle detection and counting are due to many reasons, but critical
455 points are different light conditions and diverse perspectives. Weather condi-
456 tions might produce significant illumination changes since puddles and wet floors
457 create a textural pattern that may lead to an error, and sunbeams can create
458 reflections on the car windscreen, covering the majority of the images with sat-
459 urated patterns. When a CNN does not generalize well, it works well only in
460 the conditions where it was trained.

461 To measure the robustness of our approach to these scenarios, we per-
462 formed two types of experiments exploiting the *CNRPark-EXT* dataset: i)
463 *inter-weather* and ii) *inter-camera* experiments. In the former, we trained our
464 CNN with images taken in one particular weather condition, and we computed
465 the performance metrics obtained on images having different weather condi-
466 tions. In particular, we performed three experiments, training respectively on
467 the *Sunny*, *Overcast* and *Rainy* subsets of the CNRPark-EXT dataset. In the
468 latter, we trained our algorithm employing images from one camera, and then
469 we computed the performance metrics on pictures captured by another cam-
470 era. In particular, we performed two experiments, training with images coming
471 respectively from camera₁ and camera₈. We chose these two cameras because
472 they are particularly representative since one has a side view of the parking lot
473 while the other has a pure front view.

474 We report the results of the two experiments in Table 2 and Table 3, respec-
475 tively. We achieve a good generalization in both the considered scenarios. We
476 experienced a larger amount of error when the CNN was trained and tested on
477 two opposite weather conditions, for instance, *Sunny* and *Rainy*, while the more
478 accurate model was the one trained on *Overcast* weather conditions. However,
479 the performance difference is quite small. On the other hand, in *inter-camera*
480 experiments, the model trained on camera₈ is the best, and it has a slight drop
481 in performance only when tested on the camera₁ subset. The model trained on
482 the camera₁ dataset performs in general worse. This is probably due to a bias

| Train Set | Sunny | | | Overcast | | | Rainy | | |
|-----------|-------|------|------|----------|------|-------|-------|------|------|
| | MAE | MSE | MRE | MAE | MSE | MRE | MAE | MSE | MRE |
| Sunny | - | - | - | 0.29 | 0.34 | 0.009 | 0.96 | 2.78 | 0.02 |
| Overcast | 0.62 | 1.09 | 0.02 | - | - | - | 0.56 | 1.26 | 0.01 |
| Rainy | 0.84 | 1.65 | 0.02 | 0.49 | 0.65 | 0.01 | - | - | - |

Table 2: CNRPark-EXT: Results of inter-weather experiments in terms of counting metrics obtained when training on sunny, overcast, or rainy weather.

| Metric | Train Set | Test Set | | | | | | | | |
|--------|-----------|----------|------|------|-------|-------|-------|-------|-------|------|
| | | C1 | C2 | C3 | C4 | C5 | C6 | C7 | C8 | C9 |
| MAE | C1 | - | 0.77 | 1.21 | 2.53 | 3.26 | 2.57 | 2.88 | 2.88 | 1.54 |
| | C8 | 3.87 | 0.85 | 0.76 | 0.45 | 0.48 | 0.71 | 1.07 | - | 0.41 |
| MRE | C1 | - | 0.08 | 0.05 | 0.06 | 0.07 | 0.05 | 0.06 | 0.05 | 0.05 |
| | C8 | 0.11 | 0.09 | 0.03 | 0.01 | 0.01 | 0.01 | 0.02 | - | 0.01 |
| MSE | C1 | - | 1.48 | 2.91 | 10.61 | 20.24 | 13.50 | 19.82 | 17.30 | 7.19 |
| | C8 | 22.60 | 1.78 | 1.36 | 0.57 | 0.74 | 0.95 | 4.97 | - | 2.13 |

Table 3: CNRPark-EXT: Results of inter-camera experiments in terms of counting metrics obtained when training on camera 1 and camera 8.

483 in the CNRPark-EXT dataset, where the majority of the images are captured
484 from a frontal viewpoint.

485 Moreover, to further validate the generalization capabilities of our approach,
486 we considered our counting network trained on the entire training set of the
487 *CNRPark-EXT* dataset, and we tested it over a different dataset, the *PKLot*
488 dataset (de Almeida et al., 2015). Results are shown in Table 1 where we also
489 report the performance obtained using the Mask R-CNN network without a
490 preliminary fine-tuning on the *CNRPark-EXT* dataset. As we can see, using
491 our solution, we achieve an error that is approximately four times lower than
492 the one obtained with the COCO pre-trained model.

493 *5.2. Experiments on the Multi-Camera Scenario*

494 To the best of our knowledge, there are no annotated datasets in the liter-
495 ature suitable for evaluating counting algorithms operating on multiple FOV-
496 overlapping cameras. The most relevant work in this context is Nieto et al.
497 (2019), in which there are only two overlapping cameras facing each other with
498 an extreme perspective transformation between the two; this makes any auto-
499 matic perspective computation nearly impossible without manual intervention,
500 and this is a mandatory assumption for our proposed method. Hence, we per-
501 formed our experiments on the extended version of the CNRPark-EXT dataset
502 created on purpose in this work, which we hope will become a new benchmark
503 for this task. Furthermore, to demonstrate that our algorithm can benefit from
504 the redundant information deriving from the different cameras, we compared the
505 obtained results against a baseline and a simplified version of our algorithm.

506 Specifically, we compared our solution against a system that is not aware
507 of the other cameras' overlapped areas, and so it just sums up all the vehicles
508 detected by all the cameras belonging to a sequence (Naïve Counting **N**). Then,
509 we considered a more conservative approach, where the nodes employ the homo-
510 graphic transformations only with the purpose of black-masking the overlapped
511 areas (Overlap Masking **M**). This latter baseline then loses the ability to take
512 advantage of monitoring the same lots from different views. However, it is still
513 aware of the locations of the overlapping areas, and it considers the vehicles
514 inside them only once.

515 Results are shown in Table 4. Our solution obtains the best results compared
516 to the considered baselines in all the three counting metrics and all the employed
517 scenarios. We report the errors concerning the considered six sequences of the
518 CNRPark-EXT dataset, together with the MAE, MSE, and MRE, which sum-
519 marize the mean results regarding all the scenarios. As an example, in Figure
520 5 we also report the output of our multi-camera algorithm for a pair of images
521 belonging to two different cameras having a shared area in their field of view,
522 where we highlight in red and blue the masks projected from one camera to the
523 other, using the previously computed homographic transformations.

| | Error | | | Absolute Err. | | | Squared Err. | | | Relative Err. (%) | | |
|------------|-------|-------|-------------|---------------|------|------------|--------------|---------|-------------|-------------------|------|------------|
| | N | M | O | N | M | O | N | M | O | N | M | O |
| Overcast-1 | 124 | -33 | 2 | 124 | 33 | 2 | 15,376 | 1,089 | 4 | 71.6 | 19.0 | 1.2 |
| Overcast-2 | 131 | -26 | 1 | 131 | 26 | 1 | 17,161 | 676 | 1 | 76.1 | 15.1 | 0.6 |
| Rainy-1 | 80 | -39 | -5 | 80 | 39 | 5 | 6,400 | 1,521 | 25 | 47.6 | 23.2 | 2.9 |
| Rainy-2 | 105 | -44 | -5 | 105 | 44 | 5 | 11,025 | 1,936 | 25 | 54.4 | 22.8 | 2.6 |
| Sunny-1 | 117 | -38 | 2 | 117 | 38 | 2 | 13,689 | 1,444 | 4 | 68.0 | 22.1 | 1.2 |
| Sunny-2 | 113 | -37 | 2 | 113 | 38 | 2 | 12,769 | 1,444 | 4 | 66.1 | 22.2 | 1.2 |
| Mean | 111.6 | -36.1 | -0.5 | 111.6 | 36.3 | 2.8 | 12,736.6 | 1,351.6 | 10.5 | 63.9 | 20.7 | 1.6 |

N: Naïve Counting; **M**: Overlap Masking; **O**: Ours (mean aggr., IoU Threshold $\tau = 0.2$)

Table 4: Results using our multi-camera counting algorithm, considering the *entire* parking lot. We compare our solution against a baseline and a simplified version of our algorithm. We report the errors obtained on the six considered sequences (two for each weather condition) of the CNRPark-EXT dataset that we extend on purpose.

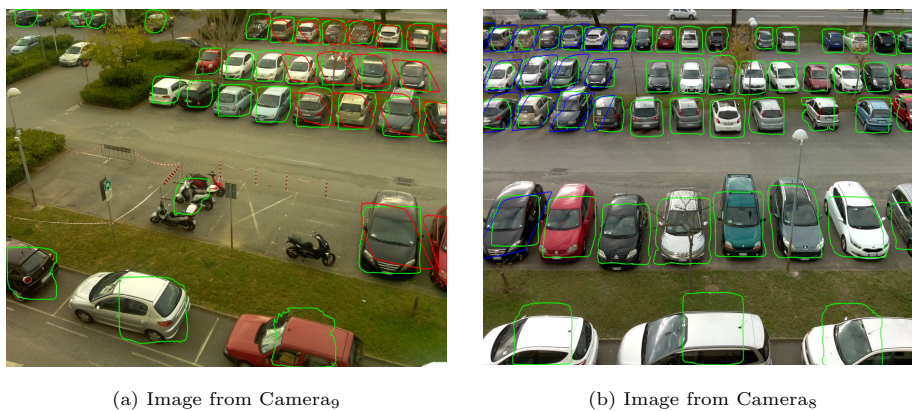


Figure 5: Example of the output of our multi-camera algorithm for a pair of images belonging to two different cameras i, j having a shared area in their FOV. We report in green the masks localizing the vehicles detected by a camera in its own FOV, while in red and blue, the masks projected from camera j to camera i and vice-versa, employing the homographic transformations pre-computed during the system initialization.

524 **6. Conclusion**

525 This paper presented a distributed artificial intelligence-based system that
526 automatically counts the vehicles present in a parking lot using images taken
527 by multiple smart cameras. Unlike most of the works in literature, we intro-
528 duced a multi-camera approach that can estimate the number of cars present in
529 the *entire* parking area and not only in the single captured images. The main
530 peculiarities of this approach are that all the computation is performed in a
531 distributed manner at the edge of the network and that there is no need for
532 any extra information about the monitored parking area, such as the location
533 of the parking spaces, nor any geometric information about the position of the
534 cameras in the parking lot. We modeled our system as a graph. The nodes, i.e.,
535 the smart cameras, are responsible for estimating the number of cars present in
536 their view and merging data from nearby devices with an overlapping field of
537 view. Our solution is simple but effective, combining a deep-learning technique
538 with a distributed geometry-based approach. We evaluated our algorithm on
539 the CNRPark-EXT dataset, which we specifically extended and which we hope
540 will become a new benchmark for counting vehicles in multi-camera parking
541 area scenarios. Through an experimental evaluation, we showed how we bene-
542 fit from redundant information from different cameras while improving overall
543 performance.

544 There are multiple lines of future development that can help improve the
545 proposed system. Although our multi-camera algorithm is flexible, one limita-
546 tion relies on computing the homographic matrix between images captured by
547 cameras placed in completely different locations, such as facing each other. In
548 such cases, the two perspectives are totally different, and manual intervention
549 is required to avoid the generation of an inaccurate geometric transformation.

550 **Acknowledgements**

551 This work was partially supported by H2020 project AI4EU under GA
552 825619, by H2020 project AI4media under GA 951911, and by Tuscany POR

553 FSE 2014-2020 AI-MAP (CNR4C program, CUP B15J19001040004).

554 References

555 de Almeida, P. R., Oliveira, L. S., Britto, A. S., Silva, E. J., & Koerich, A. L.
556 (2015). PKLot – a robust dataset for parking lot classification. *Expert*
557 *Systems with Applications*, 42, 4937–4949. URL: [https://doi.org/10.](https://doi.org/10.1016%2Fj.eswa.2015.02.009)
558 [1016%2Fj.eswa.2015.02.009](https://doi.org/10.1016/j.eswa.2015.02.009). doi:10.1016/j.eswa.2015.02.009.

559 Amato, G., Bolettieri, P., Moroni, D., Carrara, F., Ciampi, L., Pieri, G.,
560 Gennaro, C., Leone, G. R., & Vairo, C. (2018). A wireless smart cam-
561 era network for parking monitoring. In *2018 IEEE Globecom Workshops*
562 *(GC Wkshps)*. IEEE. URL: [https://doi.org/10.1109%2Fglocomw.2018.](https://doi.org/10.1109%2Fglocomw.2018.8644226)
563 [8644226](https://doi.org/10.1109/glocomw.2018.8644226). doi:10.1109/glocomw.2018.8644226.

564 Amato, G., Carrara, F., Falchi, F., Gennaro, C., Meghini, C., & Vairo, C.
565 (2017). Deep learning for decentralized parking lot occupancy detection.
566 *Expert Systems with Applications*, 72, 327–334. URL: [https://doi.org/](https://doi.org/10.1016%2Fj.eswa.2016.10.055)
567 [10.1016%2Fj.eswa.2016.10.055](https://doi.org/10.1016/j.eswa.2016.10.055). doi:10.1016/j.eswa.2016.10.055.

568 Amato, G., Ciampi, L., Falchi, F., & Gennaro, C. (2019a). Count-
569 ing vehicles with deep learning in onboard UAV imagery. In *2019*
570 *IEEE Symposium on Computers and Communications (ISCC)*. IEEE.
571 URL: <https://doi.org/10.1109%2Fiscc47284.2019.8969620>. doi:10.
572 [1109/iscc47284.2019.8969620](https://doi.org/10.1109/iscc47284.2019.8969620).

573 Amato, G., Ciampi, L., Falchi, F., Gennaro, C., & Messina, N. (2019b).
574 Learning pedestrian detection from virtual worlds. In *Lecture Notes*
575 *in Computer Science* (pp. 302–312). Springer International Pub-
576 lishing. URL: https://doi.org/10.1007%2F978-3-030-30642-7_27.
577 doi:10.1007/978-3-030-30642-7_27.

578 Arteta, C., Lempitsky, V., & Zisserman, A. (2016). Counting in the wild.
579 In *Computer Vision – ECCV 2016* (pp. 483–498). Springer International

580 Publishing. URL: https://doi.org/10.1007%2F978-3-319-46478-7_30.
581 doi:10.1007/978-3-319-46478-7_30.

582 Benedetto, M. D., Carrara, F., Ciampi, L., Falchi, F., Gennaro, C., & Am-
583 ato, G. (2022). An embedded toolset for human activity monitoring in
584 critical environments. *Expert Systems with Applications*, 199, 117125.
585 URL: <https://doi.org/10.1016%2Fj.eswa.2022.117125>. doi:10.1016/
586 j.eswa.2022.117125.

587 Boominathan, L., Kruthiventi, S. S. S., & Babu, R. V. (2016). Crowd-
588 Net. In *Proceedings of the 24th ACM international conference on Mul-*
589 *timedia*. ACM. URL: <https://doi.org/10.1145%2F2964284.2967300>.
590 doi:10.1145/2964284.2967300.

591 Ciampi, L., Amato, G., Falchi, F., Gennaro, C., & Rabitti, F. (2018). Counting
592 vehicles with cameras. In S. Bergamaschi, T. D. Noia, & A. Maurino (Eds.),
593 *Proceedings of the 26th Italian Symposium on Advanced Database Systems,*
594 *Castellaneta Marina (Taranto), Italy, June 24-27, 2018*. CEUR-WS.org
595 volume 2161 of *CEUR Workshop Proceedings*. URL: [http://ceur-ws.](http://ceur-ws.org/Vol-2161/paper12.pdf)
596 [org/Vol-2161/paper12.pdf](http://ceur-ws.org/Vol-2161/paper12.pdf).

597 Ciampi, L., Carrara, F., Amato, G., & Gennaro, C. (2022). Counting or localiz-
598 ing? evaluating cell counting and detection in microscopy images. In *Pro-*
599 *ceedings of the 17th International Joint Conference on Computer Vision,*
600 *Imaging and Computer Graphics Theory and Applications*. SCITEPRESS
601 - Science and Technology Publications. URL: <https://doi.org/10.5220%2F0010923000003124>.
602 doi:10.5220/0010923000003124.

603 Ciampi, L., Messina, N., Falchi, F., Gennaro, C., & Amato, G. (2020a). Virtual
604 to real adaptation of pedestrian detectors. *Sensors*, 20, 5250. URL: <https://doi.org/10.3390%2Fs20185250>. doi:10.3390/s20185250.
605 [//doi.org/10.3390%2Fs20185250](https://doi.org/10.3390%2Fs20185250).

606 Ciampi, L., Santiago, C., Costeira, J., Gennaro, C., & Amato, G. (2021). Do-
607 main adaptation for traffic density estimation. In *Proceedings of the 16th In-*
608 *ternational Joint Conference on Computer Vision, Imaging and Computer*

- 609 *Graphics Theory and Applications*. SCITEPRESS - Science and Technology
610 Publications. URL: <https://doi.org/10.5220%2F0010303401850195>.
611 doi:10.5220/0010303401850195.
- 612 Ciampi, L., Santiago, C., Costeira, J. P., Gennaro, C., & Amato, G. (2020b).
613 Unsupervised vehicle counting via multiple camera domain adaptation. In
614 A. Saffiotti, L. Serafini, & P. Lukowicz (Eds.), *Proceedings of the First*
615 *International Workshop on New Foundations for Human-Centered AI (Ne-*
616 *HuAI) co-located with 24th European Conference on Artificial Intelligence*
617 *(ECAI 2020), Santiago de Compostella, Spain, September 4, 2020* (pp. 82–
618 85). CEUR-WS.org volume 2659 of *CEUR Workshop Proceedings*. URL:
619 <http://ceur-ws.org/Vol-2659/ciampi.pdf>.
- 620 Dalal, N., & Triggs, B. (2005). Histograms of oriented gradients for human
621 detection. In *2005 IEEE Computer Society Conference on Computer Vision*
622 *and Pattern Recognition (CVPR'05)*. IEEE. URL: [https://doi.org/10.](https://doi.org/10.1109%2Fcvpr.2005.177)
623 [1109%2Fcvpr.2005.177](https://doi.org/10.1109/cvpr.2005.177). doi:10.1109/cvpr.2005.177.
- 624 Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). ImageNet:
625 A large-scale hierarchical image database. In *2009 IEEE Conference on*
626 *Computer Vision and Pattern Recognition*. IEEE. URL: [https://doi.](https://doi.org/10.1109%2Fcvpr.2009.5206848)
627 [org/10.1109%2Fcvpr.2009.5206848](https://doi.org/10.1109/cvpr.2009.5206848). doi:10.1109/cvpr.2009.5206848.
- 628 Fischler, M. A., & Bolles, R. C. (1981). Random sample consensus. *Com-*
629 *munications of the ACM*, 24, 381–395. URL: [https://doi.org/10.1145%](https://doi.org/10.1145%2F358669.358692)
630 [2F358669.358692](https://doi.org/10.1145/358669.358692). doi:10.1145/358669.358692.
- 631 Gu, C., Lim, J. J., Arbelaez, P., & Malik, J. (2009). Recognition using regions.
632 In *2009 IEEE Conference on Computer Vision and Pattern Recognition*.
633 IEEE. URL: <https://doi.org/10.1109%2Fcvpr.2009.5206727>. doi:10.
634 [1109/cvpr.2009.5206727](https://doi.org/10.1109/cvpr.2009.5206727).
- 635 He, K., Gkioxari, G., Dollar, P., & Girshick, R. (2017). Mask r-CNN. In
636 *2017 IEEE International Conference on Computer Vision (ICCV)*. IEEE.

- 637 URL: <https://doi.org/10.1109%2Ficcv.2017.322>. doi:10.1109/iccv.
638 2017.322.
- 639 He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image
640 recognition. In *2016 IEEE Conference on Computer Vision and Pattern
641 Recognition (CVPR)*. IEEE. URL: [https://doi.org/10.1109%2Fcvpr.
642 2016.90](https://doi.org/10.1109%2Fcvpr.2016.90). doi:10.1109/cvpr.2016.90.
- 643 Khan, M. Z., Harous, S., Hassan, S. U., Khan, M. U. G., Iqbal, R.,
644 & Mumtaz, S. (2019). Deep unified model for face recognition
645 based on convolution neural network and edge computing. *IEEE Ac-
646 cess*, 7, 72622–72633. URL: [https://doi.org/10.1109%2Faccess.2019.
647 2918275](https://doi.org/10.1109%2Faccess.2019.2918275). doi:10.1109/access.2019.2918275.
- 648 Lempitsky, V. S., & Zisserman, A. (2010). Learning to count objects in
649 images. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S.
650 Zemel, & A. Culotta (Eds.), *Advances in Neural Information Process-
651 ing Systems 23: 24th Annual Conference on Neural Information Pro-
652 cessing Systems 2010. Proceedings of a meeting held 6-9 December 2010,
653 Vancouver, British Columbia, Canada* (pp. 1324–1332). Curran Asso-
654 ciates, Inc. URL: [https://proceedings.neurips.cc/paper/2010/hash/
655 fe73f687e5bc5280214e0486b273a5f9-Abstract.html](https://proceedings.neurips.cc/paper/2010/hash/fe73f687e5bc5280214e0486b273a5f9-Abstract.html).
- 656 Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár,
657 P., & Zitnick, C. L. (2014). Microsoft COCO: Common objects in context.
658 In *Computer Vision – ECCV 2014* (pp. 740–755). Springer International
659 Publishing. URL: https://doi.org/10.1007%2F978-3-319-10602-1_48.
660 doi:10.1007/978-3-319-10602-1_48.
- 661 Lowe, D. G. (2004). Distinctive image features from scale-invariant key-
662 points. *International Journal of Computer Vision*, 60, 91–110. URL:
663 <https://doi.org/10.1023%2Fb%3Avisi.0000029664.99615.94>. doi:10.
664 1023/b:visi.0000029664.99615.94.

- 665 Nieto, R. M., Garcia-Martin, A., Hauptmann, A. G., & Martinez, J. M. (2019).
666 Automatic vacant parking places management system using multicam-
667 era vehicle detection. *IEEE Transactions on Intelligent Transportation*
668 *Systems*, *20*, 1069–1080. URL: [https://doi.org/10.1109/2Ftits.2018.](https://doi.org/10.1109/2Ftits.2018.2838128)
669 [2838128](https://doi.org/10.1109/2Ftits.2018.2838128). doi:10.1109/tits.2018.2838128.
- 670 Oñoro-Rubio, D., & López-Sastre, R. J. (2016). Towards perspective-free ob-
671 ject counting with deep learning. In *Computer Vision – ECCV 2016* (pp.
672 615–629). Springer International Publishing. URL: [https://doi.org/10.](https://doi.org/10.1007/978-3-319-46478-7_38)
673 [1007/978-3-319-46478-7_38](https://doi.org/10.1007/978-3-319-46478-7_38). doi:10.1007/978-3-319-46478-7_38.
- 674 Ren, S., He, K., Girshick, R., & Sun, J. (2017). Faster r-CNN: Towards
675 real-time object detection with region proposal networks. *IEEE Trans-*
676 *actions on Pattern Analysis and Machine Intelligence*, *39*, 1137–1149.
677 URL: <https://doi.org/10.1109/2Ftpami.2016.2577031>. doi:10.1109/
678 [tpami.2016.2577031](https://doi.org/10.1109/2Ftpami.2016.2577031).
- 679 Sindagi, V. A., & Patel, V. M. (2018). A survey of recent advances in CNN-based
680 single image crowd counting and density estimation. *Pattern Recognition*
681 *Letters*, *107*, 3–16. URL: [https://doi.org/10.1016/2Fj.patrec.2017.](https://doi.org/10.1016/2Fj.patrec.2017.07.007)
682 [07.007](https://doi.org/10.1016/2Fj.patrec.2017.07.007). doi:10.1016/j.patrec.2017.07.007.
- 683 Ujjan, R. M. A., Pervez, Z., Dahal, K., Bashir, A. K., Mumtaz, R., & González,
684 J. (2020). Towards sFlow and adaptive polling sampling for deep learn-
685 ing based DDoS detection in SDN. *Future Generation Computer Sys-*
686 *tems*, *111*, 763–779. URL: [https://doi.org/10.1016/2Fj.future.2019.](https://doi.org/10.1016/2Fj.future.2019.10.015)
687 [10.015](https://doi.org/10.1016/2Fj.future.2019.10.015). doi:10.1016/j.future.2019.10.015.
- 688 Vitek, S., & Melničuk, P. (2017). A distributed wireless camera system for the
689 management of parking spaces. *Sensors*, *18*, 69. URL: [https://doi.org/](https://doi.org/10.3390/2Fs18010069)
690 [10.3390/2Fs18010069](https://doi.org/10.3390/2Fs18010069). doi:10.3390/s18010069.
- 691 Xie, W., Noble, J. A., & Zisserman, A. (2016). Microscopy cell counting and
692 detection with fully convolutional regression networks. *Computer Meth-*

693 *ods in Biomechanics and Biomedical Engineering: Imaging & Visualiza-*
694 *tion*, 6, 283–292. URL: [https://doi.org/10.1080%2F21681163.2016.](https://doi.org/10.1080%2F21681163.2016.1149104)
695 1149104. doi:10.1080/21681163.2016.1149104.

696 Zhang, S., Wu, G., Costeira, J. P., & Moura, J. M. F. (2017). Understand-
697 ing traffic density from large-scale web camera data. In *2017 IEEE Con-*
698 *ference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
699 URL: <https://doi.org/10.1109%2Fcvpr.2017.454>. doi:10.1109/cvpr.
700 2017.454.

701 Zhang, Y., Zhou, D., Chen, S., Gao, S., & Ma, Y. (2016). Single-image
702 crowd counting via multi-column convolutional neural network. In *2016*
703 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
704 IEEE. URL: <https://doi.org/10.1109%2Fcvpr.2016.70>. doi:10.1109/
705 cvpr.2016.70.