

Towards an Error-free Deep Occupancy Detector for Smart Camera Parking System

Tung-Lam Duong^[0000–0002–9459–4705], Van-Duc Le^[0000–0002–3333–1848],
Tien-Cuong Bui, and Hai-Thien To

Seoul National University, Seoul, Korea
{dtlam26, levanduc, cuongbt91, haithienld}@snu.ac.kr

Abstract. Although the smart camera parking system concept has existed for decades, a few approaches have fully addressed the system’s scalability and reliability. As the cornerstone of a smart parking system is the ability to detect occupancy, traditional methods use the classification backbone to predict spots from a manual labeled grid. This is time-consuming and loses the system’s scalability. Additionally, most of the approaches use deep learning models, making them not error-free and not reliable at scale. Thus, we propose an end-to-end smart camera parking system where we provide an autonomous detecting occupancy by an object detector called OcpDet. Our detector also provides meaningful information from contrastive modules: training and spatial knowledge, which avert false detections during inference. We benchmark OcpDet on the existing PKLot dataset and reach competitive results compared to traditional classification solutions. We also introduce an additional SNU-SPS dataset, in which we estimate the system performance from various views and conduct system evaluation in parking assignment tasks. The result from our dataset shows that our system is promising for real-world applications.

1 Introduction

According to the 2018 UN media, 68% of the world’s population will move to urban areas by 2050[20]. This dense population in towns and cities directly leads to an increase in the number of cars and other vehicles, which raised a major concern on the parking management on their capacities and efficiency. Letting drivers wander in the city to find an appropriate parking slot in a tight city space causes significant air pollution and wastes drivers’ time and energy. It also leaves empty spaces in parking lots and varies statistic measurements on parking occupancy rate, which trouble operators to exploit their facility for revenue. In addition, these factors may worsen, particularly during peak hours when the flow density is at its maximum. For real concrete evidence, a recent report by INRIX [9] shows that on average, a typical American driver spends 17 hours a year looking for a parking space, which can go up to 107 hours when addressing a dense population city like New York. From [24] analysis, the exceeding of CO2 emissions can rise nearly three times due to this problem.

Therefore, a stable future city needs a Smart Parking System (SPS) that can serve as a link between drivers and parking operators and benefit both sides. By suggesting optimal parking places to drivers and managing their destination, a future SPS not only minimizes vehicle emissions (via decremented delays in finding the vacant parking spot [2]), but also provides operators a reliable number of customers to boost revenues through e.g. dynamic pricing [25].

Regardless of potential promises, most SPS functionalities are strictly bounded by the performance of correctly determining the occupancy of a parking lot. Hence, our current parking system relies heavily on sensors as the first layer of the system [2]. However, despite its high precision, this turns out to be pricey when scaling up the parking lot size for future perspective, as each sensor (magnetometer/ultrasonic sensor)[25] is designed to operate solely on a single parking spot. An effective solution for this drawback is applying computer vision(CV) to occupancy detection. A single camera can cover a multiple of parking locations and eliminate the need for a sensor per parking spot[3]. Furthermore, because most of nowadays parking lots have security cameras, it reduces installation and maintenance costs and supports multiple additional tasks for better parking management, such as wrong parking placement, abnormal behavior, and theft detection [29] with which sensors fail to cope.

Although computer vision is a promising approach, no datasets exist for a full CV-SPS intention. Most popular datasets PKLot[10], and CNRPark-Ext[4] and their solutions [29,5,28,22,4] are constrained to a small number of parking lots and treat each parking spot as a binary classification image. Regarding performance, there are three main drawbacks to this type of dataset and their following research. First, it limits solutions to operate in the classification scheme solely. Second, when the number of slots increases, reliable deep-learning classification solutions [28,22] require multiple forward passes and are slow to run in real-time feedback to drivers or to stack more additional tasks. Lastly, a parking operator using solutions from this dataset must reannotation every parking spot for new installation. For example, an operator is in charge of three parking lots with at least 300 parking spots in each facility. He must perform 900 annotations to use the solution, and this procedure will be conducted again when the positions of the cameras change. Therefore, an automatic parking space localization and classification solution for a scalable CV-SPS is needed to deal with the future urban population. Moreover, when we take into the functionalities of SPS in these datasets, there is no information on the parking location or surrounding traffic in this dataset as they only focus on occupancy results. It creates a big gap between existing data and the SPS's scope in the CV paradigm.

Aware of those flaws in current CV approaches and the importance of CV in future SPS, we proposed a complete CV-SPS with a new SPS-based dataset called SNU-SPS. We provide a different solution to the first layer of SPS by treating the parking lot occupancy as an object-detection task and propose a new variant model architecture called OcpDet. Changing the scope to object-detection lifts the system's performance to real-time operation and produces results frame by frame. However, most object detector is not error-free and can

be potentially wrong in a real-world inference. Hence, to maintain a reliable SPS, we provide SPS with a result filter in the second layer by addressing results from two modules of our object detector: the spatial estimator module and the training error module. While the spatial estimator module creates a predicted parking region and compares it with the model output to form the spatial error, the training error module catches the training difference of an inference frame and assigns an error. Incorrect inferences are recorded, marked as unusable information, and collected for fine-tuning and retraining the detection model. Then finally, only correct/believable detection results are stored and analyzed in the middleware layer and transparent to operators and drivers. From this layer, we can support optimal routing for drivers and alert operators about upcoming occupied parking slots at the top layer. SNU-SPS dataset is created to support this idea of the system. It contains parking slots captured from multiple parking lots at various angles, ranges, and positions with different light and contrast settings to train the detectors. Furthermore, we provide parking lot GPS with its occupancy rate and surrounding traffic information for system performance analysis.

We extensively test our proposed system on our dataset for efficiency evaluation and conduct detection measurements with the popular parking datasets PKLot for a detailed benchmark. The results from our experiment raise a competitive performance compared to exhaustive classification methods.

2 Related Work

2.1 Automatic Parking Occupancy Detection

Because most previous work focuses on solving the SPS occupancy task as image classification from datasets [10,4] with manual label mask/grids location of a parking lot, none of available datasets could be found for automatic parking space detection. It leads to a small amount of effort on this topic. To the best of our knowledge, there are currently two main approaches to this topic: a mask-based method and a detector-based method.

A mask-based method aims to provide parking patches directly from captures and perform binary classification for the occupancy. The perspective transformation method [7,8,21] is usually used in this scheme to bring the parking lot to a 2D grid presentation. Therefore, it can save time for self-annotating parking locations and exploit the classification machine-learning and deep-learning backbones. However, since the perspective projection process is highly dependent on the camera setting to the parking lot, classification models need to be retrained for different camera settings, which questions the scalability of those methods. Notice this behavior, [18] has introduced a GAN approach that generates the parking place’s masks from a team of drones, but there is no comprehensive measurement of the correctness of these masks. In addition, this method requires a top-view capture of the lot, making it unrealistic for indoor parking facilities.

In contrast to mask-based solutions, detector-based approaches perform detection and classification tasks in a single process by a CNN architecture instead

of separating them into two processes, which maintains the flexibility and fast inference for CV-SPS infrastructure. The CNN architecture in this realm regresses a parking slot as a foreground or a region of interest and optimizes its classification score. This procedure can be classified into two-stage and one-stage detectors. While the two-stage detector, such as Faster RCNN [26] focuses on the first stage to propose the regions of interest and performs classification on those regions in the second stage, the one-stage detector combines both tasks by grid-anchor regressions. However, because an empty or small parking slot is easily confused as a part of the image’s background, both of these architectures face a lot of flaws [23]. Most recent works [16] only used detectors to find a parked car in the parking lot and determine the occupancy rate by a preowned parking lot’s capacity and location. This approach relaxes the problems into the well-known car detection, but it limits the extension of the SPS for letting drivers know the location of the parking slot. Recent developments in new architectures such as YOLO [6], and RetinaNet [19] have opened some flexibilities in the small object capture. The idea of using a drone’s captures is also used by [14]. The author performs the car detection at top-views by Faster RCNN and YOLO and combines it with the layout proposal. This method faces the same drawback as [18] and is restricted for car detections. For a complete occupancy detection from a detector, only [23] has conducted on a RetinaNet on PKLot [10] dataset. However, the results show much confusion between moving cars and occupied parking slots. While the main reason for this inefficiency is the nature of the PKLot dataset itself (partial area of the parking lot is annotated), the method’s performance can be improved if there is an attention mechanism on the parking lot region. In addition, there is a potential non-optimized model design as there is no information on the grid-anchor feature selection.

2.2 Deep Model Uncertainty

As a model performance is a reflection of the coverage of the training dataset, recent research tends to capture model performance in the wild by measuring its uncertainty/stability or contradiction. While [15,1] aimed to create this uncertainty from the outputs by comparing an image and its noise version, [12,13] measured the stability of the prediction of the same image from different models’ inferences. However, these methods are not designed for online inference and require some delay for a model judgment. Only [30] can achieve the model uncertainty within a single forward pass inference. Instead of replicating the contradiction from inferences, the author trained the model with a different head to predict this uncertainty directly. Even though it sounds naive and simple, his approach has shown promising results on general object detection datasets.

2.3 Qualitative Comparison Between Existing Approaches and the Proposed System

Aware of the lack of discovery in the current approaches for the CV-SPS system and the need for an efficient detector-based method, we provide a compact

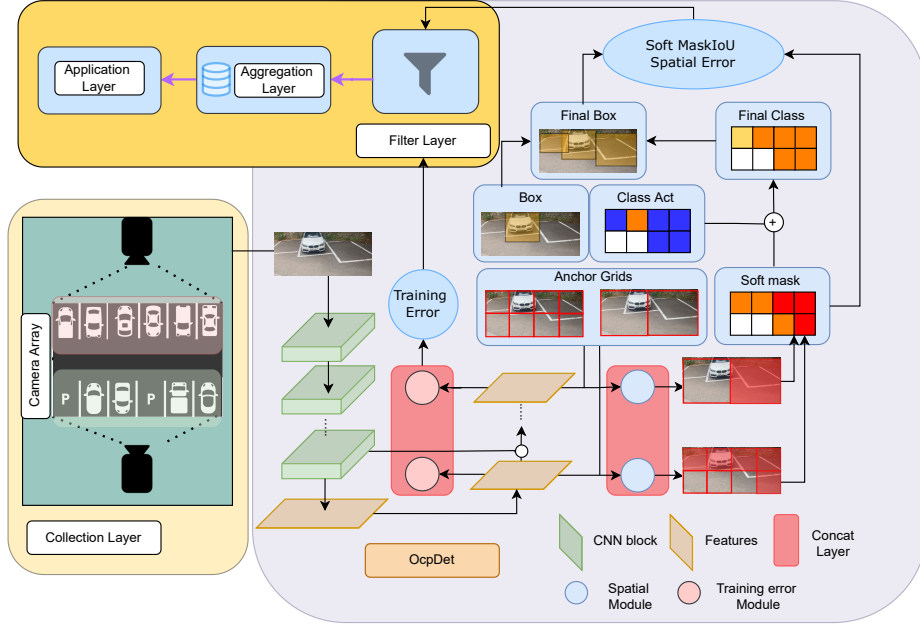


Fig. 1: Our CV-SPS Overall Architecture

detector-based solution with its corresponding dataset. To our knowledge, our CV-SPS is the first end-to-end solution for automatic parking occupancy detection. Our detector-based solution inherits the feature pyramid structure from RetinaNet [19] with two additional module heads: the spatial module and the training error module. While we inherit [30] as the training error module for training error information, we design the spatial estimator module as a parking region proposal to compare the spatial error. It can be understood as a foreground filter, aiming to filter out detection not covered in the predicted foreground region. Hence, this module provides the model with better localization attention, which [23] did not address.

3 Proposed Method and Dataset

3.1 Overall System Architecture

As shown in Fig 1, we divide a parking lot into sectors to create a scalable and efficient CV-SPS. Each sector is controlled and well-observed by a camera and has non-overlapped observing areas among cameras. This constraint reduces the complexity of the problem by duplicating observation or high occlusion. Assuming the parking lot can be set up with this requirement, our overall system architecture consists of 4 layers: a collection layer, a filter layer, an aggregation layer, and an application layer. The collection layer is responsible for gathering

the detection results from distributed cameras in the parking lot as well as their potential error info during inference. Then, these results are propagated to the filter layer to cleanse for reliable results. Non-trusted results are masked out as non-usable spaces. This filtered information is stored in the aggregation layer that acts as middleware of the system. From this layer, the application layer can receive reliable SPS support. Users can have a transparent measurement of the current occupancy capacity of a parking lot, while model engineers can access and inspect poor performance behavior in specific sectors. Especially, optimal routings and parking assignments can execute with high precision by correctly capturing vacant spots.

3.2 Dataset




Fig. 2: SNU-SPS dataset images representation from various indoor and outdoor views. Annotation colors (Red: Occupied, Green: Available, Blue: Restricted and Yellow: Illegal)

We introduced our SNU-SPS dataset containing nearly 3500 images to support our system. Those images are captured from various views, heights (1-3m), and light conditions in indoor and outdoor parking lots. Each parking lot has different parking spot background colors. The total images were manually checked, labeled, and attached by GPS to the corresponding parking slot. The protocol used to construct the SNU-SPS dataset is composed as follows:

Image Acquisition: All images are captured with a full-HD resolution. For the training set, it is captured randomly for one month in 15 parking lots. Meanwhile, the test set is captured consecutively in 6 parking lots from 3-6pm through 5 working days. It should be noted that none of the 6 parking lots are in the training set. Moreover, test samples contains various weather conditions

(sun/rain/cloudy) and has corresponding surrounding traffic measurements from the open government website.

Labeling: For each parking sector, images were labeled as *available/ occupied/ illegal/ restricted* of each parking space. Each annotation is covered by four keypoints that specify for the localization of a parking lot. We formulate the wrapping bounding boxes for the detector from these key points. Especially, we provide optional image masks for the test set to filter out overlapping areas and non-important localization among capture among parking lots. The intention is to maintain the system’s constraints and preserve a better parking assignment benchmark.

Set Type	Total Images	Total Labels	Classes			
			Available	Occupied	Illegal	Restricted
Train	2848	18263	7229	10596	396	42
Test	574	2747	1291	1336	36	84

Table 1: Training and Testing Sets

3.3 Automatic Parking Occupancy Detector

When addressing parking occupancy as an object detector, the most arousing problem is the confusion of the parking slots with the image background information such as moving cars or blank spots. The only meaningful visual information is the thin lines separating spots. However, it is usually missed at the lower level of deep neural networks.

RetinaNet: is a promised solution due to its feature pyramid network (FPN). In short, the FPN backbone combines standard convolutional network lower features with lateral connections of early-level features. Hence, the network can construct rich, multi-scale object features, which maintain the impact of the line features in the network. However, from [23] results, despite capturing good center localization, traditional RetinaNet could not expand the parking space tightly when the mAP dropped dramatically from 63.64 to 4.75 when raising from 0.5 to 0.75 precision. This might be non-optimized anchor grid features and lack of a location attention mechanism. Moreover, the Resnet backbone is quite heavy for computation and may not be able to scale up with other additional SPS tasks, limiting the scope of CV-SPS.

$$L_{size} = \sum_{i=0}^N D_{p,k}^i / D_{p,c}^i \quad (1)$$

OcpDet: Noticing this limitation, we replace the heavy Resnet backbone with a lightweight Mobilenet for faster inference and build up our model from

this called OcpDet. We intend to give OcpDet more reference points rather than just centers and sizes of parking slot boxes. Therefore, instead of solely detecting the bounding boxes by their centers and sizes, our detector also predicts the key points of parking slots in the localization head. From this scope, we design a new loss function L_{size} which aims to maximize the predicted coverage boxes to their corresponding predicted key points. We treat these N keypoints as anchors which pull the box corners closer to them, where N is less than 4. We denoted the distance between a box corner p to the its corresponding keypoint k as $D_{p,k}$ and the distance between a box corner p to the center c as $D_{p,c}$. When we set the keypoints as our dataset keypoints, the loss implicitly guides the model to focus more on the border of a parking slot.

Spatial Estimator Module: Increasing solely reference points is not enough to make our OcpDet robust with background and foreground conflict in a parking lot. Cars vs. parked places or empty spots vs. walls can be highly confused. Because our model is a single-stage object detector, it divides an image into anchor patches and represents them at different scale levels for classification prediction and localization regression. As the parking lot layout is usually aligned, equally separated, and non-overlapped among spots, it is very convenient to ideally wrap a parking slot in one single anchor patch representation. Therefore, we create a soft head estimator for the active anchors that can be considered as parking slots from each level feature generator of FPN.




Fig. 3: Module Architectures for Spatial Estimator Module and Training Error Module

We attach a residual convolution block for each of the N feature levels and average the last channel to get a 2D map of anchor patches. Noted that, we avoid flattening the map to avoid insufficient computations as the fully connected layers can go up to huge connections for early-level features. For example, a 256×256 anchor can cost 65k connection and additional 65 parameters for a dense layer.

With that intuition, as parking slots are in the foreground, we trigger this 2D prediction by a sigmoid activation and treat it as a binary classification (Fig 3). From this activation map, we can create a soft mask for parking slot locations as a reference for the model spatial outputs. To train this activation map, we compare each map M from N feature levels with its corresponding classification target map C in the classification head by a ϵ difference. This loss can act as a regularizer for the model to provide attention to the foreground region of the model. Moreover, as the soft mask head is not directly connected to the localization head prediction, it gives the model another degree of freedom to operate while implicitly improving the localization through top-level features. To deeply understand the role of the spatial module, we investigate the way




Fig. 4: HiRes-GradCam activations following by Image Detections (green: "occupied", blue: "available") and their corresponding soft mask heatmaps. White arrows stand for weak detection area: (a) low confidence score, (b) zero detection. c) is an example of our model inference

of predicting the foreground from multi-level anchor-grid features through the FPN. Regardless of the grid's anchor box size, it will be considered a foreground if the classification activation is switched on. Therefore, we look into the pixel activation map through the class activation map (HiRes-GradCam) [11] to understand which pixel contributes to the final decision. As demonstrated in Fig 4, there is a high density of pixel activations at each middle of parking slots whenever the model predicts the location class, which is reasonable as a class determination belongs to what having inside a parking border. The class confidence is reduced when the number of surrounding activation maps is dimmed Fig 4-a and disappears in Fig 4-b. This means the activation map only produces

a reliable decision when it falls inside a parking spot. We notice that the spatial module’s heatmap can evade this problem by adding the activation of the soft mask prediction. Following this idea, we injected the predicted classification parking-grid information into the classification head by adding its prediction to the ground confidence. We avoid multiplication for overfitting results with the predicted soft-mask.

Training Error Module: as mentioned above, we use [30] method as a model error predictor from training knowledge. In short, [30] represents this error as predicted loss during training. The method selects pairs of training batch samples in each iteration to represent the model inference stability during training. A good pair of samples should have a low loss during training, while a poor inference will raise the loss. Hence, during training, a predictor module will try to estimate this loss by accessing multi-level features, allowing it to choose necessary information between layers to capture the loss behavior correctly. In other words, this method can also act like an abstract regularizer for the model to avoid conflict between batches. Compare to the Spatial Estimator Module in Fig 3, it shares the same feature map, but average out each feature by an average pooling before aggregation for loss prediction. Therefore, this module can operate independently from the model output performance.

3.4 Result Filter

Therefore, after obtaining the detection results with two additional predictions from the spatial and training error modules, the foreground N detections of each sector D are compared with its soft spatial activation map S . We compute the overlap value for each detection and select the strong-belief region by a threshold γ . From an array of overlapping ratios, we can estimate the spatial error $Err_{spatial}$ for each frame capture detection as demonstrated in Eq. 2. Noted that, if $\mathbb{1}(R)$ from a detection reach near the zero bottoms of the overlap ratio, the spatial module will get rid of this detection. We combine this spatial error with the training error $Err_{training}$ from the training error module by averaging aggregation as a final error score Err_{total} for the inference. From this score, we can eliminate whether to believe these detections results or not by thresholding our belief.

$$Err_{spatial} = 1 - \frac{1}{N} \sum_{i=0}^N \mathbb{1}\left(\frac{D_i \cap S}{D_i}\right) \quad \text{s.t.} \quad \mathbb{1}(R) = \begin{cases} 1, & \text{if } R > \gamma \\ R, & \text{otherwise} \end{cases} \quad (2)$$

$$Err_{total} = \alpha Err_{training} + (1 - \alpha) Err_{spatial} \quad (3)$$

4 Experiment

In this section, we demonstrate the benefits of our approach and its efficiency in a real-world scenario. We first study the impact of using the spatial module and reference keypoints for improving parking localization and compare it with the existing solutions. Second, we demonstrate the effectiveness of filtering out

unstable inference to the model’s actual performance. Finally, we measure the impact of OcpDet’s performance with the optimal routing and parking assignment task.

4.1 Experimental Setup

Training settings: For our experiments, with the scope of an efficient and quick response in the first system layer, we only addressed the single-stage object detector: SSD-Mobilenet (denoted MBN), Mobilenet-FPN (our model backbone, denoted MBN-FPN), and our OcpDet. The training engine for these models is Tensorflow Object Detection API. We trained OcpDet for both datasets in 25000 iterations with batchsize 48 by SGD optimizer. To keep the detection robust to the small parking spots, we selected with high-resolution 896-pixel input instead of the traditional 300-pixel or 640-pixel.

Benchmark settings: We used PKLot dataset [10] and our SNU-SPS dataset for the solutions benchmark. We did not use CNRPark-Ext because this dataset’s parking spot border lines are faded and not consistently visible. PKLot dataset contains three sub-datasets capturing from high view: PUCPR, UFPR05, and UFPR04, which leads to a small scale of parking spots. All of these sub-dataset parking locations are partially annotated. Thus, we provide masks for each sub-dataset to clip out non-annotated parking regions to avoid false positives during training and make it suitable for the detection benchmark. We split the dataset from each sub-dataset in half for training and testing, as the PKLot authors suggested. For our SNU-SPS, we have a separate test set. However, due to a small label of *illegal* and *restricted* classes in our dataset, our test set is only addressed with two classes: *occupied* and *available*. In addition, because the scope of our dataset is to detect correctly in a sector, we only address medium and large ground truths in the test set. During the assignment application test, the detection will be filtered out overlapping areas by our provided masks.

4.2 Detection Performance

In this part, we conduct our experiment on the efficiency of our spatial module and the localization improvement of additional reference keypoints. Our main metric for the detection evaluation is mAP(mean average of precision) and recall ranging from 0.5 to 0.95 IoU(intersection over union). Meanwhile, as the classification task is the side task to benchmark with a classification approach, we address the mAP(0.5) from each class for the comparison. As results are summarized in Table 2, the performance of the localization attention leads in all data settings. We did not compare with [23] solution because of his insufficient model’s performance.

In the PKLot dataset, MBN struggles to learn the features because it is lack of top feature generation from the grid and lines to capture the small objects. Meanwhile, thank to the FPN, both MBN-FPN and our model outperform MBN in this dataset. Due to the fixed location of the parking lot’s captures, OcpDet

Method	Test Set	Recall (0.5:0.95)	mAP		Classification Score		
			0.5	0.75	(0.5:0.95)	Occupied	Available
<i>OcpDet</i>	PUCPR	0.88	0.98	0.98	0.84	0.98	0.98
MBN-FPN		0.76	0.86	0.85	0.72	0.83	0.90
MBN[27]		0.41	0.48	0.35	0.31	0.49	0.46
Classifier[28]		-	-	-	-	0.99	0.99
<i>OcpDet</i>	UFPR05	0.98	0.99	0.99	0.97	0.99	0.99
MBN-FPN		0.84	0.93	0.90	0.82	0.93	0.94
MBN[27]		0.42	0.51	0.42	0.37	0.50	0.53
Classifier[28]		-	-	-	-	0.99	0.99
<i>OcpDet</i>	UFPR04	0.96	0.99	0.99	0.93	0.99	0.99
MBN-FPN		0.83	0.95	0.90	0.79	0.95	0.96
MBN[27]		0.43	0.52	0.44	0.36	0.51	0.53
Classifier[28]		-	-	-	-	0.99	0.99
<i>OcpDet</i>	SNU-SPS	0.56	0.81	0.48	0.47	0.83	0.80
<i>OcpDet_{spatial}</i>		0.56	0.83	0.50	0.47	0.85	0.82
<i>OcpDet_{both}</i>		0.56	0.82	0.49	0.47	0.84	0.81
<i>OcpDet_u</i> [30]		0.55	0.82	0.49	0.47	0.83	0.81
MBN-FPN		0.54	0.77	0.46	0.45	0.80	0.74
MBN[27]		0.51	0.71	0.48	0.44	0.73	0.69

Table 2: PKLot and SNU-SPS Detection Benchmark

can strongly overfit the position of each parking space and turn into a grid classifier. As demonstrated in Fig 5, the localization guidance from the soft mask generations helps the anchor patches avoid negative samples that do not belong to the parking area. Our model can boost the performance to near perfection using the soft mask head during training. In addition, by using additional keypoints, our approach improves the localization detection and preserves its tightness among scales. In contrast to the unchanged parking layout of the PKLot dataset, our SNU-SPS creates more challenges for the model to select the correct anchor patch due to its various capture positions. This makes not only our model but also other detectors struggle. Therefore, we omit the classifier method for a fair comparison. In this dataset, both the reference keypoints and the soft mask generator are not much efficient as the model grid is not stationary. Despite the challenge of adaptability, when looking into the heatmap of the soft mask through different scales of the spatial module in Fig 5, the obtained mask on the parking lot still has denser attention than other foreground predictions. Combined with the impact of reference keypoints, *OcpDet* leaves a gap of nearly 5% on mAP(0.5) to original MBN-FPN and 10% to MBN. Moreover, in the first row of Fig 5, it shows that our model is not sensitive with car appearance. The prediction only activates inside in the parking zone where the lines are visible.

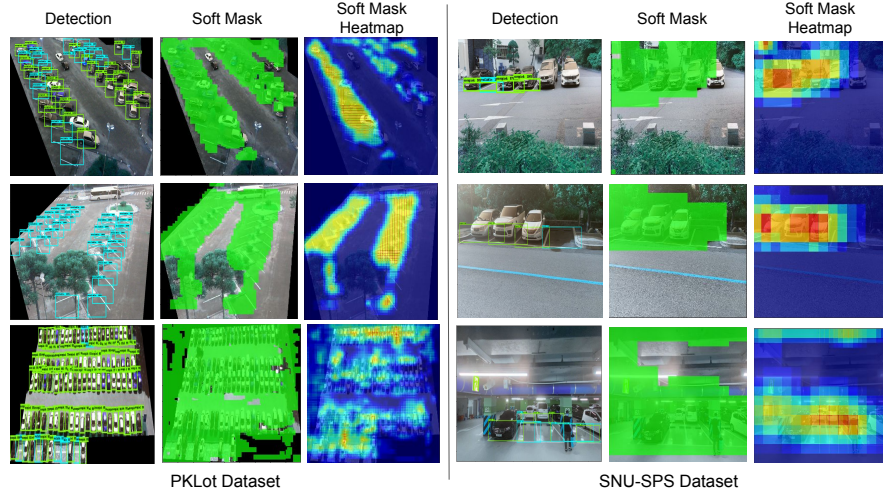


Fig. 5: Visualization occupancy detections, soft mask predictions and soft mask predictions heatmap on PKLot dataset and SNU-SPS dataset

4.3 Result Filter Performance

To evaluate the result filter efficiency, we solely address OcpDet on SNU-SPS as there is still room for model improvement. We calculate the error of an inference sample from the formula Eq 2 and remove a maximum of 100 samples (20% of the test set) from the test set. For the purpose of testing out the benefit from the spatial module and the training error module, we lock the γ by 0.7 and set α from a set of $[0, 0.4, 1]$ with respect to $OcpDet_{ll}$, $OcpDet_{both}$ and $OcpDet_{spatial}$ in Table 2. According to the experiment, the result filter has boosted the model’s overall accuracy, proving that the filter can ensure a better quality from the detector regardless of α assignment. Because the spatial filter can score and remove unreliable results from detections by overlapping with the soft mask, its results are slightly better than the [30] approach on the training filter.

4.4 Optimal Routing and Parking Assignment

To make a comprehensive benchmark on the impact of detection results on the assignment application layer, we collect the traffic information over ten days from the government website and associate that information with the test set to form a close loop simulation. Each day from 3 to 6pm, there will be about 100 requirements for booking a vacant spot to 6 parking lots. The suggested optimal road for each request will be assigned from the MapQuest API. We consider the Hungarian assignment [17] from the masked-out test label as the ground truth for assigning vacant spots.

Then, we compute two evaluation metrics: cost error and assignment error. The cost error is computed by the absolute error of the ground truth assignment cost $C_{g,i}$ and the vacancy detection assignment cost $C_{p,i}$, which is designed by Eq. 4. The assignment cost is computed after getting the assignments. From each assigned parking lot in 6 parking lots, we compute the total number of booked slots N and compare those between the ground truth $N_{g,i,j}$ and the detection $N_{p,i,j}$ by an absolute error. We perform normalization to keep the value from 0 to 1. The simulation will repeat ten times to capture the model’s average performance due to different traffic statuses.



Fig. 6: The cost errors and the assignment errors: averaging simulation for 5 days at 6 parking lots

$$C = \gamma C_{price} + (1 - \gamma)(C_{travel} + C_{distance}) \quad (4)$$

$$Err_{cost} = \sum_{i=3}^6 \frac{|C_{g,i} - C_{p,i}|}{C_{g,i}} \quad Err_{assign} = \sum_{i=3}^6 \sum_{j=1}^6 \frac{|N_{g,i,j} - N_{p,i,j}|}{N_{g,i}} \quad (5)$$

From Fig. 6, OcpDet allows the system to operate at most 40% error for the cost-minimizing budget while maintaining at least 70% correct on assignment. Because the cost error is related to the distance travel, wrong assignments placed on far-distance drives can cause a huge gap to the optimal cost. But this error is not much in terms of matching the number of assignments. From these metrics, we can overview the system’s benefits to users. Operators will benefit the most as their vacant spaces will automatically be assigned with minima error. In contrast, some drivers may get some disadvantages from the system assignment.

5 Conclusion & Future work

This paper proposes a novel end-to-end CV-SPS with a detailed benchmark on both old and new datasets. Even though our dataset is small, it shows challenging factors, and it is the first dataset for computer vision with full CV-SPS scope. Our method has proved its efficiency to some extent and can potentially close the gap to the classifier approach when addressing stationary views. Moreover, we also provide a novel filtering method that can help the system approach an error-free execution. In the future, we will continue building our dataset to a bigger scale with which it can provide functional information for SPS, such as vehicle reidentification or parking type selection.

6 Acknowledgement

Dataset and experiment in this work were supported by the Automotive Industry Building Program (1415177436, Building an open platform ecosystem for future technology innovation in the automotive industry) funded by the Ministry of Trade, Industry Energy (MOTIE, Korea).

References

1. Aghdam, H.H., Gonzalez-Garcia, A., Weijer, J.v.d., López, A.M.: Active learning for deep detection neural networks. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3672–3680 (2019)
2. Al-Turjman, F., Malekloo, A.: Smart parking in iot-enabled cities: A survey. *Sustainable Cities and Society* **49**, 101608 (2019)
3. Lisboa de Almeida, P.R., Honório Alves, J., Stubbs Parpinelli, R., Barddal, J.P.: A systematic review on computer vision-based parking lot management applied on public datasets. *arXiv e-prints* pp. arXiv–2203 (2022)
4. Amato, G., Bolettieri, P., Moroni, D., Carrara, F., Ciampi, L., Pieri, G., Gennaro, C., Leone, G.R., Vairo, C.: A wireless smart camera network for parking monitoring. In: *2018 IEEE Globecom Workshops (GC Wkshps)*. pp. 1–6. IEEE (2018)
5. Amato, G., Carrara, F., Falchi, F., Gennaro, C., Meghini, C., Vairo, C.: Deep learning for decentralized parking lot occupancy detection. *Expert Systems with Applications* **72**, 327–334 (2017)
6. Bochkovskiy, A., Wang, C.Y., Liao, H.Y.M.: Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934* (2020)
7. Bohush, R., Yarashevich, P., Ablameyko, S., Kalganova, T.: Extraction of image parking spaces in intelligent video surveillance systems (2019)
8. Bura, H., Lin, N., Kumar, N., Malekar, S., Nagaraj, S., Liu, K.: An edge based smart parking solution using camera networks and deep learning. In: *2018 IEEE International Conference on Cognitive Computing (ICCC)*. pp. 17–24. IEEE (2018)
9. Cookson, G.: Parking pain—inrix offers a silver bullet. INRIX—INRIX. Online. Available: <http://inrix.com/blog/2017/07/parkingsurvey/>. Accessed: November 21 (2017)
10. De Almeida, P.R., Oliveira, L.S., Britto Jr, A.S., Silva Jr, E.J., Koerich, A.L.: Pklot—a robust dataset for parking lot classification. *Expert Systems with Applications* **42**(11), 4937–4949 (2015)
11. Draelos, R.L., Carin, L.: Use hirescam instead of grad-cam for faithful explanations of convolutional neural networks. *arXiv e-prints* pp. arXiv–2011 (2020)
12. Feng, D., Wei, X., Rosenbaum, L., Maki, A., Dietmayer, K.: Deep active learning for efficient training of a lidar 3d object detector. In: *2019 IEEE Intelligent Vehicles Symposium (IV)*. pp. 667–674. IEEE (2019)
13. Haussmann, E., Fenzi, M., Chitta, K., Ivanecky, J., Xu, H., Roy, D., Mittel, A., Koumchatzky, N., Farabet, C., Alvarez, J.M.: Scalable active learning for object detection. In: *2020 IEEE Intelligent Vehicles Symposium (IV)*. pp. 1430–1435. IEEE (2020)
14. Hsieh, M.R., Lin, Y.L., Hsu, W.H.: Drone-based object counting by spatially regularized regional proposal network. In: *Proceedings of the IEEE international conference on computer vision*. pp. 4145–4153 (2017)
15. Kao, C.C., Lee, T.Y., Sen, P., Liu, M.Y.: Localization-aware active learning for object detection. In: *Asian Conference on Computer Vision*. pp. 506–522. Springer (2018)
16. Kirtibhai Patel, R., Meduri, P.: Faster r-cnn based automatic parking space detection. In: *2020 The 3rd International Conference on Machine Learning and Machine Intelligence*. pp. 105–109 (2020)
17. Kuhn, H.W.: The hungarian method for the assignment problem. *Naval research logistics quarterly* **2**(1-2), 83–97 (1955)

18. Li, X., Chuah, M.C., Bhattacharya, S.: Uav assisted smart parking solution. In: 2017 international conference on unmanned aircraft systems (ICUAS). pp. 1006–1013. IEEE (2017)
19. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
20. Media, U.N.: 68% of the world population projected to live in urban areas by 2050, says un. <https://www.un.org/development/desa/en/news/population/2018-revision-of-world-urbanization-prospects.html> (2018), last Accessed: 2022-06-08
21. Nieto, R.M., García-Martín, Á., Hauptmann, A.G., Martínez, J.M.: Automatic vacant parking places management system using multicamera vehicle detection. *IEEE Transactions on Intelligent Transportation Systems* **20**(3), 1069–1080 (2018)
22. Nyambal, J., Klein, R.: Automated parking space detection using convolutional neural networks. In: 2017 Pattern Recognition Association of South Africa and Robotics and Mechatronics (PRASA-RobMech). pp. 1–6. IEEE (2017)
23. Padmasiri, H., Madurawe, R., Abeysinghe, C., Meedeniya, D.: Automated vehicle parking occupancy detection in real-time. In: 2020 Moratuwa Engineering Research Conference (MERCon). pp. 1–6. IEEE (2020)
24. Paidi, V., Håkansson, J., Fleyeh, H., Nyberg, R.G.: Co2 emissions induced by vehicles cruising for empty parking spaces in an open parking lot. *Sustainability* **14**(7), 3742 (2022)
25. Polycarpou, E., Lambrinos, L., Protopapadakis, E.: Smart parking solutions for urban areas. In: 2013 IEEE 14th International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM). pp. 1–6. IEEE (2013)
26. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* **28** (2015)
27. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4510–4520 (2018)
28. Valipour, S., Siam, M., Stroulia, E., Jagersand, M.: Parking-stall vacancy indicator system, based on deep convolutional neural networks. In: 2016 IEEE 3rd World Forum on Internet of Things (WF-IoT). pp. 655–660. IEEE (2016)
29. Varghese, A., Sreelekha, G.: An efficient algorithm for detection of vacant spaces in delimited and non-delimited parking lots. *IEEE Transactions on Intelligent Transportation Systems* **21**(10), 4052–4062 (2019)
30. Yoo, D., Kweon, I.S.: Learning loss for active learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 93–102 (2019)