



HAL
open science

Few-Shot Object Detection in Real Life: Case Study on Auto-Harvest

Kevin Riou, Jingwen Zhu, Suiyi Ling, Mathis Piquet, Vincent Truffault, Patrick Le Callet

► **To cite this version:**

Kevin Riou, Jingwen Zhu, Suiyi Ling, Mathis Piquet, Vincent Truffault, et al.. Few-Shot Object Detection in Real Life: Case Study on Auto-Harvest. 2020 IEEE 22nd International Workshop on Multimedia Signal Processing (MMSP), Sep 2020, Tampere, Finland. pp.1-6, <10.1109/MMSP48831.2020.9287053>. <hal-03652639>

HAL Id: hal-03652639

<https://hal.science/hal-03652639v1>

Submitted on 19 Feb 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Few-Shot Object Detection in Real Life: Case Study on Auto-Harvest

Kevin Riou¹, Jingwen Zhu¹, Suiyi Ling^{1,2}, Mathis Piquet¹, Vincent Truffault³, Patrick Le Callet¹

¹University of Nantes ²Capacites SAS ³CTIFL

{kevin.riou, jingwen.zhu, mathis.piquet}@etu.univ-nantes.fr;

{suiyi.ling, patrick.lecallet}@univ-nantes.fr Vincent@futuragaia.com

Abstract—Confinement during COVID-19 has caused serious effects on agriculture all over the world. As one of the efficient solutions, mechanical harvest/auto-harvest that is based on object detection and robotic harvester becomes an urgent need. Within the auto-harvest system, robust few-shot object detection model is one of the bottlenecks, since the system is required to deal with new vegetable/fruit categories and the collection of large-scale annotated datasets for all the novel categories is expensive. There are many few-shot object detection models that were developed by the community. Yet whether they could be employed directly for real life agricultural applications is still questionable, as there is a context-gap between the commonly used training datasets and the images collected in real life agricultural scenarios. To this end, in this study, we present a novel cucumber dataset and propose two data augmentation strategies that help to bridge the context-gap. Experimental results show that 1) the state-of-the-art few-shot object detection model performs poorly on the novel ‘cucumber’ category; and 2) the proposed augmentation strategies outperform the commonly used ones.

Index Terms—few-shot object detection, real life few-shot learning, dataset, cucumber detection, auto-harvesting.

I. INTRODUCTION

COVID-19 has brought an extremely painful period for the world. Apart from continuing to take its toll, the pandemic has also spawned a great number of economic losses. Especially in agriculture, due to the confinement, harvesting slowly becomes one of the enormous challenges. Therefore, the development of robust ‘auto-harvesting’ is of far greater urgency. The existing mechanical harvest/auto-harvest system is normally based on 1) auto-detection of the matured vegetables/fruits and 2) the auto-harvest using robots within a limited reachable zoom. In most of the auto-harvest systems, auto-detection of the target regions is usually the bottleneck and thus is more vital.

Recent deep convolutional networks [1], [2] have achieved significant improvement in object detection. Most of them depends heavily on large-scale training dataset like COCO [3] and PASCAL VOC [4], [5]. Such dependency raises questions like 1) how to deal with the new categories that was not involved in the training set; and more importantly, 2) how to employ them in real life, when there are not enough samples for target categories. For auto-detection of vegetables/fruits, it is often required to deal with new categories with limited training samples per category as samples and annotations could be difficult and expensive to collect.

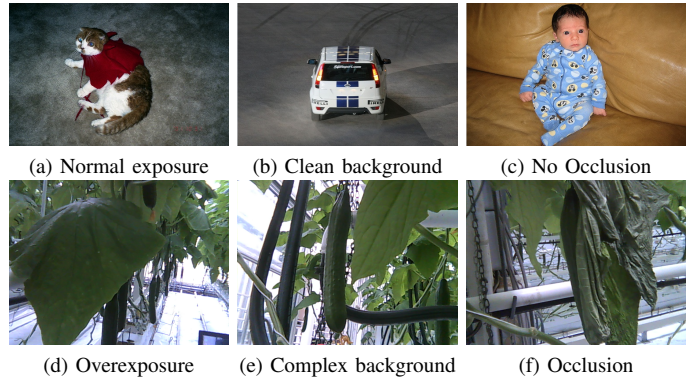


Fig. 1. Examples explaining the ‘Gap’ between the training dataset used in the literature and target real life problem. **First row**: images in commonly used PASCAL dataset; **Second row**: Images from our cucumber dataset.

To tackle such intractable problem, recently, few-shot [6], [7] and meta learning [8] based models were proposed. Most of these models target to obtain a model that is able to address both base and novel categories at test time. In general, there are two training phases/stages within the training scheme of these models, including 1) a base learner training or meta training stage where the base categories are utilized; 2) a fine-tuner/task adaptor training or meta testing stage, where the new categories are used. This type of training-testing schemes target to transfer the task-relevant (recognition/detection) knowledge from the base categories to the novel ones, and more importantly force the model to learn with only a few samples.

Nevertheless, real life few-shot problems are more challenging due to factors like noisy imaging conditions [9]. In the case of auto vegetables/fruits detection, we notice that there is an obvious gap between the image-context of dataset used to train the state-of-the-art few-shot object detection models and the ones of real life vegetable images collected in greenhouse. Examples are depicted in Fig. 1. As shown, the aforementioned challenges include 1) **image conditions**: as images are not taken under controlled camera setting, it is common to see samples collected in real life with worse image conditions, e.g., overexposure as shown in Fig. 1 (d); 2) **not all the objects appear in the image are the target**: since auto harvest relies on the robots to harvest, only objects that are within a reachable area are interesting, e.g. as shown in Fig. 1 (e), only the cucumbers shown in the foreground are the targets;

3) **noisy background and occlusion:** images taken in certain real life scenarios could have a complex and noisy background, and the target objects could be occluded as shown in Fig. 1 (f). Thus, except for transferring the detection knowledge from the base categories to the new ones, the context-knowledge (e.g., image conditions, background *etc.*) of new scenarios should also be transferred for real life applications.

In this study, based on the discussions above, we aim at re-visiting the state-of-the-art few-shot detection models for real life greenhouse applications, and exploring data augmentation strategies to better transfer context-knowledge of the new scenarios. The main contributions of our paper are twofold: 1) A novel cucumber dataset with bounding boxes annotations is released; 2) Two data augmentation strategies are proposed to take the context of the new applications into account.

II. RELATED WORKS

Few-shot object detection, including zero-shot object detection, aims to accurately detect novel category of objects that are not involved in the training procedure using few samples, *i.e.*, shots, of the new category. Compared to the traditional recognition or object detection task, this brand new problem is significantly more challenging due to the ill-posed nature and inherent complexity of detecting absolutely unknown categories. To tackle this challenging problem, many works have been developed. Michaelis *et al.* [10] improved the Mask R-CNN with a Siamese architecture. In [11], the author proposed a zero-shot model based on transductive learning. Recently, natural language descriptions were considered to better address the problem [12]. An R-CNN based re-weighting network was presented in [13], which disentangles multi-object information and turns Faster/Mask R-CNN into a meta-learner for few-shot object detection. Following a similar recipe, Kang *et al.* proposed a region proposal free re-weighting network based on YOLOv2.

Nevertheless, most of the aforementioned models are trained and tested on the same type of datasets, *e.g.* the PASCAL dataset. The distributions of these datasets could be significantly different compared to the images collected in real life scenarios. More specifically, there will be a large domain-shift between the source (datasets used to train existing models) and the target (real life applications) domains. Therefore, they are prone to fail when applied in typical real life scenarios. An example is depicted in Fig. 2, where the model proposed in [10] incorrectly segments the background as the target ‘cucumber’ regions.

III. OUR CUCUMBER DATASET

To remedy the lack of real life dataset for automatic harvesting, we collected a new dataset that contains images of cucumbers in the greenhouse. The dataset is publicly available at <https://github.com/KevinRiou22/Labeled-cucumber-dataset>. Details of the dataset is summarized in Table I.

Data collection: The collection of the dataset was organized by *Centre Technique Inter-professionnel des Fruits et Légumes*




Fig. 2. Result of Siamese Mask-R-CNN model on zero-shot cucumber segmentation.

(CTIFL) [14] and the images were taken within their greenhouse ($47^{\circ}17'11.6''N$ $1^{\circ}27'34.3''W$). CTIFL is an ‘inter-professional center for the fruits and vegetables’. Especially, it is a french agency that aims at developing the knowledge and expertise of all professions around fruits and vegetables sectors.

TABLE I
SUMMARY OF OUR CUCUMBER DATASET.

| | |
|--------------------------------|-----------|
| Number of images | 555 |
| Resolution | 640 x 480 |
| Number of annotation per image | 1 - 4 |

In their greenhouses, the plants, *i.e.*, the cucumbers, were planted in rows. Between every two rows, rails were set up to allow trolleys or automatic harvesting robots to forward and harvest the ripe vegetables/fruits automatically. To facilitate the automation of harvesting, cameras were mounted on the trolley, around 80cm away from the plants as shown in Fig. 3 (a). The angles of the cameras are adjustable. To avoid ‘motion blur’, the trolleys stopped constantly along their moving trajectory to take images from different angles. In general, images of plants were taken in front of the rows, *e.g.*, Fig. 3 (b-d). Additionally, we also enriched the dataset by varying the cameras’ angles, as presented in Fig. 3 (e-f). It is worth emphasizing that only the row closest to the trolley is interesting, as the automatic robot targets only to harvest the plants (cucumbers) within a constraint surrounding range.

Annotations: When annotating the locations of target cucumbers, we followed the instructions provided by [4], and thus followed a ‘PASCAL-VOC’ format. Only the cucumbers from the closest row to the trolley, *i.e.*, those in the foreground, were labeled, as they are within the accessible range of the robotic-arm of the trolley.

IV. RE-VISIT STATE-OF-THE-ART FEW-SHOT OBJECT DETECTION MODELS FOR REAL LIFE SCENARIO

As summed up in Section II, the few-shot object detection model proposed in [15], namely FS-FRW, is one of the state-of-the-art models and achieves the best performances. Therefore, in this study, we re-visit and adapt this approach to few-shot cucumber detection in real life scenarios.




Fig. 3. Examples from our dataset, where the green boxes are the labels indicating the location of the cucumbers: (a) shows the trolley that carried the camera along the rails in the greenhouse; (b,c) are examples taken facing directly against the ‘cucumber rows’ as depicted in (a); (d, e) are examples collected with different camera angles.

The overall adapted FS-FRW is summarized in Fig. 4. The framework is composed of two main parts, including (1) a generalized feature learner D that is trained to detect objects of novel categories using large-scale training samples per class, as shown in the upper part of Fig. 4; and (2) a feature re-weighter M that reveals the vital meta-features, *i.e.*, vital components of the feature learner, and benefits the accurate identification/detection of objects belonging to a new category with only a few samples of this new category, as shown in the lower part of Fig. 4. They are trained together in an end-to-end and few-shot fashion so that novel categories of objects could be well detected.

In this study, ‘cucumber’ is considered as the novel category, and the goal is to accurately detect them utilizing FS-FRW under a complex, noisy greenhouse environment. Aiming at narrowing the gap between the common context and the context of real life greenhouse conditions, we consider different data augmentation techniques and verify their effectiveness for the adaptation of the state-of-the-art few-shot learning models. Fine-grained descriptions are shown below.

Concretely, taking an image I (size of $w \times h$) as input, the generalized feature learner D yields a set of meta features $F = D(I)$, $F \in \mathbb{R}^{w \times h \times m}$. The feature re-weighter M embeds images from the support images into a set of re-weighting vectors, *i.e.*, coefficients, w_i , and returns the category-relevant feature F_i by

$$F_i = F \otimes w_i \quad (1)$$

where i indicates a certain category, and $i = 1, \dots, N$, and \otimes is the channel-wise multiplication.

With the obtained F_i , it is then fed into a prediction model P , which returns the likelihood of object existence o_i , the predicted location of the object bounded by the offset of (x, y, h, w) , and a corresponding classification score c_i :

$$\{o_i, x_i, y_i, h_i, w_i, c_i\} = P(F_i) \quad (2)$$

A novel two-stage learning scheme was proposed [15] to guarantee the generalization of the model when dealing with

new categories with few samples.

The first stage is the base learning stage. During this stage, each base category is of abundant training samples with location labels. To ensure the model’s capability of detecting target objects by making full use of a good re-weighting vector, D, M, P are trained jointly.

Formally, the base categories from the base training set were split into a set of few-shot detection task T_j . Each task was defined as $T_j = S_j \cup Q_j = \{(I_1^j, M_1^j), \dots, (I_N^j, M_N^j)\} \cup \{(I_N^q, M_N^q)\}$, where S is the support set containing N samples from a different base category, and Q_j is the Query set for evaluating the performance. Finally, the three modules are optimized jointly by minimizing the loss defined below:

$$\min_{\theta_D, \theta_M, \theta_P} \sum_j \mathcal{L}_j = \sum_j \mathcal{L}_{det}(P_{\theta_P}(D_{\theta_D}(I_j^q) \otimes M_{\theta_M}(S_j)), M_j^q), \quad (3)$$

where θ_D , θ_M , and θ_P are the parameters of D, M and P respectively. \mathcal{L}_{det} is an adopted detection loss defined as the sum of (1) the cross-entropy loss over the calibrated category scores [15] (2) the bounding box regression loss [16], and (3) the objectiveness regression loss [16].

The second stage is the few-shot fine-tuning stage. During this stage, the model is trained on both base and novel categories, including ‘cucumber’, but only k bounding boxes are available for each novel category. Similarly, only k bounding boxes are included for the base categories. The training process is the same as the one in the first stage, but with less iterations and novel categories are involved.

As pointed out in Section I that there could be significant domain-shift between the categories of base categories and the target novel category in real life applications. As thus, we should transfer not only the detection knowledge from the base categories to the new ones, but also the context-knowledge (*e.g.*, noisy image conditions, occlusions *etc.*) within real life scenarios.

Therefore, in this study, we explore different simple data augmentation techniques to narrow the gap between the distributions of base categories and target novel category in real life cases, *i.e.*, the ‘cucumber’.

The first stage of FS-FRW is crucial for the final performance, since the tuning stage relies heavily on the generalized features. Intuitively, data augmentation is one of the straightforward and simplest way to bridge the gap of the distributions between the base categories and the new categories that come from real life greenhouse scenarios. To this end, we propose two data augmentation strategies that are dedicated for few-shot object recognition:

- **Background replacement of base training images:** The backgrounds and image conditions of real life greenhouse images are relatively more complex compared to the ones in commonly used benchmark datasets like PASCAL VOC and COCO, as demonstrated in Fig 1. Hence, we collected a set of greenhouse background images and randomly replace the backgrounds of the base training

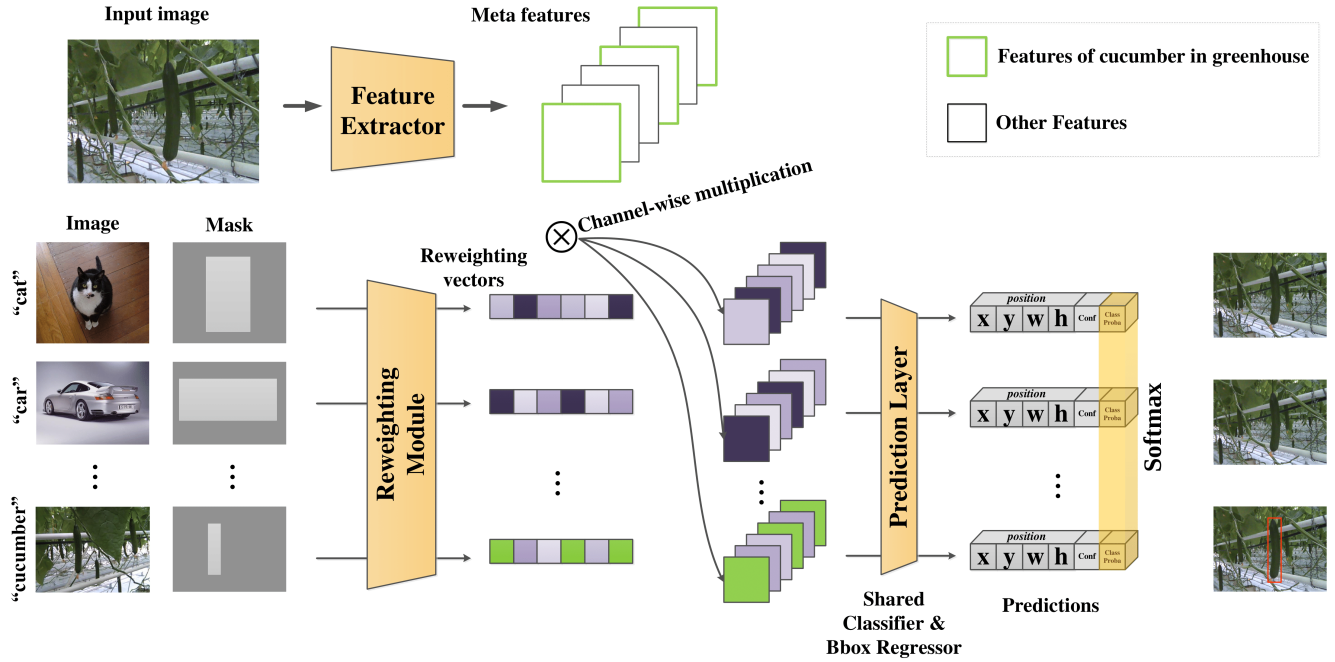


Fig. 4. The overall framework of the adapted few-shot object detection model based on feature re-weighting in greenhouse scenario [15].



Fig. 5. Examples of generated samples for base categories with replaced greenhouse background.

set with the collected ones. Examples of the new base training samples are depicted in Fig. 5. As shown, only the foreground base objects were kept. By doing so, the model is forced to take into account the image conditions/real life environment statistics. Some feature-dimension highlighted by the generalized feature extractor could therefore be dedicated to the greenhouse contexts. In parallel, it is then possible for the feature re-weighter to focus on these relevant features and yield more robust re-weighted vectors.

- **Adding ‘target background’ as one of the base categories:** During the base training stage, an extra category ‘background’ is added to the base dataset. This category could be adapted according to the specific task. In this study, it is the background images taken from the greenhouse. It is worth mentioning that the target objects do not show up in the images, and possible background

regions of the target objects, e.g. regions of branches and leaves without cucumber in the greenhouse, were annotated as the ground truth. By including the task-relevant ‘background’ as one of the base category, the model is then able to obtain important meta features that is related to the image conditions/environment of the target task and the feature re-weighter is able to associate these relevant-features to the target context.

In a nutshell, both of these two augmentation approaches enrich the meta-training procedure so that it has access to the context (e.g. the image conditions, environment) without exposing the novel object itself. More implementation details of the two strategies are provided in Section V.

V. EXPERIMENTS

In this section, (1) the performances of the state-of-the-art few-shot object detection model FS-FRW in real life greenhouse applications are reported; (2) different data augmentation strategies were compared.

A. Experimental Setup

Similar to the setup in [15], The PASCAL VOC 07 [4], and 12 train/val [4] set was utilized for training while the VOC 07 test set was employed for testing. Following the same evaluation setting, among the 20 categories in the dataset, 15 of them were randomly selected as the base categories in the first stage for training generalized feature extractor, and the remained were considered as the novel ones.

In the second stage for training the feature re-weighter, a very small set of training images were kept so that each category of objects only contain k annotated bounding boxes for k -shot object detection problem. It has to be emphasized that we replaced one of the novel categories considered in [15] by our target ‘cucumber’. Details of base and novel category are summarized in Table II.

TABLE II
SUMMARY OF BASE & NOVEL CATEGORIES IN THE EXPERIMENTS.

| | |
|----------------|---|
| Base category | aeroplane, bicycle, boat, bottle, car, cat, chair, dining-table, dog, horse, person, potted-plant, sheep, train, tv-monitor |
| Novel category | bird, bus, cow, motobike, cucumber |

2) *Data Augmentation*: To bridge the gap between the context-knowledge of the scenarios used for training, and the ones of the target new scenarios in real life, in this study, we explored and examined not only the proposed data augmentation strategies presented in Section IV, but also the commonly used strategies including:

- **Background Replacement of base training images (BR)**: Inspired by the data augmentation method proposed in [17], the pixel-wise segmentation masks provided by PASCAL VOC were utilized to extract the foreground semantic objects and the extracted objects were further posed on the greenhouse backgrounds.
- **Adding ‘Target Background’ as one of the base categories (ATB)**: To implement this strategy, we simply replaced one of the base category ‘airplane’ with the manually collected ‘greenhouse background’ as described in Section IV.
- **Illuminance Adjustment (IA)**: The illuminance of all the images of categories except for the ‘cucumber’ were simply enhanced using gamma correction by a factor of 1.5.
- **Contrast Adjustment (CA)**: For this strategy, we increased the contrast of images belong to all categories except for the ‘cucumber’ via contrast stretching with a factor of 2.

B. Experimental Results

The Mean Average Precision (mAP) with IoU threshold of 50% was calculated to evaluate the performances of considered models¹. According to [15], FS-FRW achieves best performances among the compared models, it is thus considered as the baseline model. The 10-shots performance on the novel categories is reported in the TABLE III. It is obvious that the performance of FS-FRW on the target ‘cucumber’ is worse than most of the other novel categories, better solutions are required to narrow the gaps. Note that ‘bird’ class achieves

¹Different batch size was set in this study, therefore, the reported results are lower compared to the ones in [15].

even lower performances than ‘cucumber’, which was already the case as reported in [15] on PASCAL VOC. Most of bird images may be captured in wild and complex backgrounds. Such contexts also present a significant gap between most of the base classes, such as ‘aeroplane’, ‘bike’, etc., which hinders the few shots tuning on this class.

TABLE III
10-SHOT PERFORMANCE (MAP %) OF THE 5 NOVEL CATEGORIES.

| | | | | |
|-------|-------|-------|----------|-----------------|
| bird | bus | cow | motobike | cucumber |
| 15.66 | 41.06 | 27.24 | 35.59 | 27.00 |

As highlighted in Section I, one of the goals of this study is to explore different data augmentation methods, so that the state-of-the-art few-shot detection models could be better adjusted for real life cases. To this end, we compare the performance of FS-FRW equipped with different data augmentation approaches. Results are shown in TABLE IV.

TABLE IV
PERFORMANCE COMPARISON USING DIFFERENT DATA AUGMENTATION STRATEGIES.

| Approach | mAP for ‘cucumber’ (%) |
|--------------|------------------------|
| FS-FRW | 27.00 |
| FS-FRW + BR | 29.33 |
| FS-FRW + ATB | 33.12 |
| FS-FRW + IA | 14.04 |
| FS-FRW + CA | 23.69 |

It could be observed that the two proposed strategies, *i.e.* BR and ATB, outperform the others. We also note that ATB performs slightly better than BR, one of the possible reasons is that by taking the task-relevant context into account earlier in an early stage (base learning training stage) is more efficient. Because the re-weighter fine-tuned in the second stage focus more on the characteristics of the object instead of the context of the real life applications.

Examples are shown in Fig. 5 to better explain why better improvements could be achieved via ATB. Ideally, the model should detect only cucumbers in the target/first row so that the machine could harvest them within a reachable distance. Therefore, any other cucumbers that appear elsewhere were not labeled in our dataset as mentioned in Section 3c. By taking backgrounds as one of the base categories, the base model to some extent captures the fore/background information of the context, and thus able to better detect objects in the foreground other than the unrelated ones in the background.

VI. CONCLUSION AND FUTURE WORK

In this study, in order to improve the robustness of few-shot object detection models for real life auto-harvest scenarios, 1) a ‘cucumber’ dataset is released for the community; 2) the state-of-the art few-shot models is re-visited and employed for



Fig. 6. Examples explaining why performance could be improved by using ATB. (a) is the output of the feature re-weighter trained with ATB approach; (b) is the output of the feature re-weighter model trained without using special data augmentation method.

cucumber detection; 3) two novel data augmentation methods dedicated to real life few-shot applications are presented. Throughout experiments, it is verified that tested few-shot object detection model still need to be further improved to be applied in agricultural scenarios. The proposed data augmentation methods are proven to be effective.

For future work, 1) higher shot setting, e.g., 15-shot results could be considered to further verify the impacts of the proposed strategies in low-shot regime; 2) more traditional data augmentation methods, e.g., flipping, resizing, blurring etc., could be compared and combined with the proposed ones to strengthen few-shot models in real life scenarios; 3) performance significant test, e.g. the one employed in [18], could be further utilized to ensure whether the performances are statistically improved.

REFERENCES

- [1] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama *et al.*, “Speed/accuracy trade-offs for modern convolutional object detectors,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 7310–7311.
- [2] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.
- [3] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C. L. Zitnick *et al.*, “Common objects in context,” in *European Conference on Computer Vision*, 2014.
- [4] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, “The pascal visual object classes (voc) challenge,” *International journal of computer vision*, vol. 88, no. 2, pp. 303–338, 2010.
- [5] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, “The pascal visual object classes challenge: A retrospective,” *International journal of computer vision*, vol. 111, no. 1, pp. 98–136, 2015.
- [6] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra *et al.*, “Matching networks for one shot learning,” in *Advances in neural information processing systems*, 2016, pp. 3630–3638.
- [7] J. Snell, K. Swersky, and R. Zemel, “Prototypical networks for few-shot learning,” in *Advances in neural information processing systems*, 2017, pp. 4077–4087.
- [8] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR.org, 2017, pp. 1126–1135.
- [9] S. Ling, A. Pastor, J. Li, Z. Che, J. Wang, J. Kim, and P. L. Callet, “Few-shot pill recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9789–9798.
- [10] C. Michaelis, I. Ustyuzhaninov, M. Bethge, and A. S. Ecker, “One-shot instance segmentation,” *arXiv preprint arXiv:1811.11507*, 2018.
- [11] S. Rahman, S. Khan, and N. Barnes, “Transductive learning for zero-shot object detection,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 6082–6091.
- [12] Z. Li, L. Yao, X. Zhang, X. Wang, S. Kanhere, and H. Zhang, “Zero-shot object detection with textual descriptions,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 8690–8697.
- [13] X. Yan, Z. Chen, A. Xu, X. Wang, X. Liang, and L. Lin, “Meta r-cnn: Towards general solver for instance-level low-shot learning,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 9577–9586.
- [14] Ctifl website. [Online]. Available: <http://ctifl.fr/>
- [15] B. Kang, Z. Liu, X. Wang, F. Yu, J. Feng, and T. Darrell, “Few-shot object detection via feature reweighting,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 8420–8429.
- [16] J. Redmon and A. Farhadi, “Yolo9000: better, faster, stronger,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 7263–7271.
- [17] S. Ling, J. Li, J. Wang, and P. L. Callet, “Gans-nqm: A generative adversarial networks based no reference quality assessment metric for rgb-d synthesized views,” *arXiv preprint arXiv:1903.12088*, 2019.
- [18] S. Ling, J. Gutiérrez, K. Gu, and P. Le Callet, “Prediction of the influence of navigation scan-path on perceived quality of free-viewpoint videos,” *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 9, no. 1, pp. 204–216, 2019.