Weakly Supervised Faster-RCNN+FPN to classify small animals in camera trap images

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

23-11-2021 / 06-12-2021

CITATION

Pochelu, Pierrick; Erard, Clara; Cordier, Philippe; Petiton, Serge G.; Conche, Bruno (2021): Weakly Supervised Faster-RCNN+FPN to classify small animals in camera trap images. TechRxiv. Preprint. <https://doi.org/10.36227/techrxiv.17068454.v1>

DOI

[10.36227/techrxiv.17068454.v1](https://doi.org/10.36227/techrxiv.17068454.v1)

Weakly Supervised Faster-RCNN+FPN to classify small animals in camera trap images

Pierrick Pochelu
TotalEnergies SE
Pau, France
pierrick.pochelu@totalenergies.com

Clara Erard
MODIS
Pau, France
clara.erard95@gmail.com

Philippe Cordier
TotalEnergies SE
Paris, France
philippe.cordier@totalenergies.com

Serge G. Petiton
Univ. Lille, CNRS, UMR 9189 CRISTAL
F-59000 Lille, France
serge.petiton@univ-lille.fr

Bruno Conche
TotalEnergies SE
Pau, France
bruno.conche@totalenergies.com

Abstract—Camera traps have revolutionized animal research of many species that were previously nearly impossible to observe due to their habitat or behavior.

Deep learning has the potential to overcome the workload to the class automatically those images according to taxon or empty images. However, a standard deep neural network classifier fails because animals often represent a small portion of the high-definition images. Therefore, we propose a workflow named Weakly Object Detection Faster-RCNN+FPN which suits this challenge. The model is weakly supervised because it requires only the animal taxon label per image but doesn't require any manual bounding box annotations. First, it automatically performs the weakly supervised bounding box annotation using the motion from multiple frames. Then, it trains a Faster-RCNN+FPN model using this weak supervision. Experimental results have been obtained on two datasets and an easily reproducible testbed.

I. INTRODUCTION

Through their variability of living organisms, ecosystems yield a flow of vital services ranging from production (food, water, oxygen...) to regulation (soil purification, climate regulation...) through cultural and recreational benefits [1]. Habitat loss and fragmentation, pollution, climate change, species introduction, and disturbance are the current drivers of biodiversity decline. Biodiversity does not only provide services but is also linked to human health. Human health and well-being are influenced by the health of local plant and animal communities and the integrity of the local ecosystems that they form. The COVID-19 pandemic is a reminder that 60.3% [2] of all emerging infectious diseases affecting humans are zoonotic i.e., transferred from animals to humans. Therefore, ecosystem functions and services are now being altered, raising vexing issues about how to best monitor species and ecosystems over time - and if needed - correct the impact of these developments on them.

Camera traps due to their non-invasive nature, affordable technology, high mobility, and battery autonomy have revolutionized animal research of many species that were previously nearly impossible to observe due to their habitat or behavior. Several camera traps take thousands of pictures of animals

in specific areas to conduct reliable surveys by identifying them on each. While being one of the main advantages of the technique, this large amount of collected data also proves to be highly time-consuming for ecologists to annotate and count animals.

Taking advantage of the big data era, Deep Learning has had a breakthrough impact on many domains. Since its popularity, deep learning was applied to many colored image datasets often containing several animal classes. Furthermore, deep learning has already shown human-like performance on animal detection. That is why it is a natural choice to help ecologists to leverage the workload.

Faster-RCNN+FPN [3] is specifically suitable to detect small regions in high-definition images but requires high annotation cost including animal localization and animal taxa. By contrast, classification neural networks require only the taxa annotation but they generally fail when images contain small objects to classify and cluttered backgrounds. Based on those previous works, we propose Weakly Supervised Faster-RCNN+FPN to combine those two advantages: image-level annotation cost (Weakly Supervised) and recognition ability similar to the full supervised Faster-RCNN+FPN.

This paper follows this structure: (2) First we introduce the relevant literature to tackle the challenge of the huge data labeling needs for successful Deep Learning applications. (3) Presentation of both datasets to experiment with our workflow and their inherent challenges. (4) The proposed workflow is introduced and discussed. (5) The experimental results of our workflow are compared to baselines and the performances are analyzed.

II. RELATED WORKS

The last 10 years saw the conduct of several projects aiming to classify animals on camera traps [4], [5], [6] but they are often applied to large Savannah animals or livestock.

Faster-RCNN [7] is a popular neural network method to recognize objects (here animals) in images automatically but they require important annotation effort because the user not

only give an image-level label (here, its taxon or 'empty') but also its localization into the image (i.e., bounding box). Some upgrades allow increasing the recognition ability in the case of small objects thanks to an additional module named Features Pyramid Network [3].

Previous studies [8] success to classify images according to fixed size small objects using a bag of local features[8]. However, the good performance in practice with varied objects size have not been shown yet.

The noise affecting real-world applications make image recognition task more or less challenging for CNNs. The lower the signal-to-noise ratio is such cluttered background, small objects to detect (measured with object-to-image), the more labeled data is required [8]. However, collecting and labeling data is time taking. Indeed, an ideal system would require only a few thousand classified images to save ecologists' time and therefore be able to detect small signal signatures in a messy background.

A. Tackle annotation cost

To reduce the amount of annotation work, three main tracks exist: extract knowledge from similarly labeled datasets; using crowd-sourced data labeling software and collaborating with many annotators; or applying weak supervision.

Open dataset [9] contains terabytes of camera traps images in a few areas of the world which can be unlabelled, partially labeled, or containing multiple annotation errors. And more, these data are not suited for the supervised learning of specific biodiversity in a wild region where animals have been rarely photographed. A popular way to extract useful knowledge from other datasets is referred to as transfer learning, under this umbrella, pre-training weights is a common method [10].

Another natural method is the usage of crowd-sourced labeling data software [5] where many annotators can work together to divide the huge workload out between them. Unfortunately, only a few ecologists worldwide know to distinguish and annotate some wild animal taxa on a given area of the world.

B. Weakly Supervised Object Detection in images and videos

Weakly Supervised Object Detection is a method able to provide localization of objects with only class annotation. The first approaches to WSOD formulate this task as a workflow of 3 consecutive steps: *region extraction* to extract candidate regions from the image; *region representation* to compute features representations of regions; *region final decision* such a regressor allows to refine the bounding boxes or a classifier to decide if a bounding box contains an object.

Over the last 5 years, the landscape of weakly object detection pushed forward the adoption of efficient end-to-end learning frameworks. Some methods extract the saliency map (i.e., object localization probability map) from a CNN classifier trained with an image-level annotation [11]. More advanced methods with the same label requirement consist in combining a CNN classifier and, with a saliency map extracted from it, training another semantic segmentation CNN [12].

However, they share in common that the majority of works focus on VOC2007 or VOC2012 datasets where objects are rather centered and occupy a large portion of the image.

Another weak supervision method consists in using a standard strongly supervised model but a labelling rule-based function to automatically build labels with unknown accuracy. For example, JFT-300M dataset was built automatically gathering 375 million images from the web associated with one of the 19 thousand categories depending on web signals. It can be noisy (estimated to 20% error) in terms of label confusion and incorrect labels which are not cleaned by humans. This weak supervision machine learning method is popular when it comes to entity extraction where a label function can be found with some relevant hypothesis.

This label function is domain-specific which is why we propose one to localize animals on camera traps and evaluate its performance. In the 'experimental results' section, we compare the complete workflow compared to some relevant baselines.

Weak supervision is especially useful in video recognizing [13] [14] containing tens of frames per seconds. They often rely on motion cues on images and propagated the objectness information over the neighboring pixels (spatial) and neighboring frames (temporal). Generally, the Weakly Supervised Object Detection in videos assumes that relevant objects are similar from frame to frame which is a not relevant assumption in camera traps data. We are facing a rather different problem: the animal can be present in only one frame or present on multiple frames but with very different positions.

In summary, this paper is built on previous research on weak supervision methods and supervised object detection frameworks. The goal is to save the annotation effort on camera traps and to benefit from the state-of-the-art supervised object detection method to spot tiny objects (animals). While most recent initiatives are focusing on how to apply WSOD to localize objects with an end-to-end CNN classifier with some reasonable size objects (VOC2007, VOC2012) or video with multiple similar frames, we focus on building a labeling function applied to small moving objects on a short sequence of images.

III. THE TWO DATASETS AND GOAL

A. Papua New Guinea biodiversity

Images are collected from a monitoring campaign that took place in Papua New Guinea. 8 motion-triggered wildlife camera traps have been used for 15 months with the same settings. They have been operated tied to a tree from 43 to 101 days and took between 474 and 1,272 photographs. When triggered, they shot a burst of three photographs spaced by 1 second, increasing the chances to have at least one good image of the animal. As a majority of mammals are nocturnal, camera traps are also equipped with a flash mode. To decrease potential errors made by our localization algorithm, tuning the settings of cameras is an axis of improvement including the number of shots per sequence and the delay between shots in one sequence.

25 taxa have been captured (dorcopsis luctuosa, bandicoot, scrubfowl, emerald dove...) into 5400 images with 8 cameras. Each one is fixed to a tree and captures similar background when it is triggered, even if the background may slowly evolving due to leaves fall, weather, flowers blooming ... 6 cameras contains 4,320 labeled pictures, 95% of them serving as training and the remaining ones being the validation dataset. The testing dataset contains 1,080 pictures taken by the 2 remaining cameras. It is aimed to label and train on a few cameras and that the system generalizes well on new cameras. All pictures are colored with a resolution of 3,776x2,124 pixels. The labeling is edited with a mapping file between the identifier number of the picture and the identifier number of the class. If a picture does not contain any animal, the class identifier is set to '0'. Otherwise, it is the identifier number of the taxon.

B. Missouri biodiversity

This dataset comes from Missouri campaign [15] [16]. It contains approximately 24,673 camera trap images. There are about 1277 fully labeled images with bounding boxes and classes (squirrel, European hare, red deer, red fox, ...). Among them, a statistic we made on 100 randomly drawn images shows 10% false positive errors (bounding box and animal taxon whereas we do not see animal into), no false-negative errors, and a majority of badly framed bounding box annotations. The tiny portion of the labeled dataset and the fact that annotations contain errors make the task particularly challenging. We checked and corrected 1000 images to create the test dataset.

C. Computer vision challenges

Wildlife camera trap automatic recognition present some challenging situation for deep learning methods. First, the animals are small in a cluttered background making it difficult to recognize them. We measure on the Papua New Guinea dataset that the surface of animals can occupy from the entire image to $< 0.2\%$ of pixels of the photo depending on its actual size and its proximity to the camera, with a median of 4%. Additionally, both datasets contain unbalanced classes.

In addition, we summarize the difficulty to recognize an animal in five challenging situations: when the animal is hidden, when the animal is badly framed, when the animal appear blur (e.g., motion blur when it is running or jumping), hidden behind foliage, unusually small ($< 0.2\%$ of pixels of the image), when the animal is unknown to the training dataset. They are illustrated in supplementary material section A.

IV. THE METHOD

To measure the suitability of the methods we use the percentage of well-classified images: the animal in taxa or the reject class "empty" when no animal is present in the photo. Indeed, the exact position of the bounding boxes of Faster-RCNN does not count in our target metric that is why we use the percentage of well-classified images as relevant metrics.

When several instances of the same species are in the same image, the taxon must be correctly identified but we do not count them. Due to the scarcity of inter-species social behavior, we witnessed only once different animal taxa on the same photograph in the Papua New Guinea dataset. In this case, in the evaluation phase, we count 2 possible correct predictions and in the training phase, we only label the biggest animal (pixel squared).

A. The workflow usage

Before applying deep learning, the ecologists manually check all images and they input a mapping file that maps the name of the photograph with a taxon number or '0' if there is no animal in the photo.

Now, after labeling a few thousand images, the ecologist runs the training phase before running the deep learning inference model on the other images.

Because no deep learning model is perfect, to save ecologist's time and reduce errors, all predictions of the systems are sorted in ascending order to their Softmax posterior probability. Therefore, the first images have more chance to contain challenging animals absent from the training dataset. After that, the ecologist goes through the predictions and associated images and corrects the mapping file when needed.

B. Deep Learning

Many object detectors including R-FCN, Faster-RCNN are 'meta-architectures' and features extractors based on 'architecture' such as VGG or ResNet50. The meta-architecture defines the different modules and how they work together. The feature extractor is a deep learning architecture only made of convolution layers to extract spatial features.

To do object detection, researchers have developed fully integrated meta-architectures like Faster-RCNN [7] which both yield object localization and their classes. The meta-architecture defines the different modules and how they work together. To find an acceptable trade-off between accuracy and computational cost, researchers have proposed SSD or YOLO which are especially popular to recognize objects in videos, that is to say, millions of low-resolution images or usable for real-time applications. Our need is to recognize a few thousand highly detailed images containing often challenging small objects. That is why we give priority to accuracy over inference speed in this work.

Faster-RCNN is a meta-architecture made of three trainable modules trained end-to-end. The CNN extractor is a standard deep learning architecture without any classification layer which extracts spatial features shared with the Region Proposal Network (RPN). The RPN is trained to localize Region of Interests (ROIs) such as the position of the animal in the picture and it sends them to the "classifier and regressor module". This last module classifies each ROI and is trained to classify the localized animal and also to refine the ROIs coordinates.

In theory, standard CNN is adapted to detect different levels of abstraction in the input image. In practice, it fails to preserve tiny signals in the input image of the greatest importance on

the final prediction. FPN [3] is trained to extract and make decisions on different levels of abstraction and resolution (the features pyramid). Faster-RCNN with FPN shows a significant improvement in several applications compared to basic Faster-RCNN models.

Our model is the fruit of those previous works. It uses a meta-architecture Faster-RCNN [7] with Features Pyramid Network [3] and ResNet50 features extractor [17]. It is pre-trained [10] with ImageNet dataset containing 1,000 classes of which 398 are animal classes. Pre-training allows to reduce by about 5 times the convergence speed but it does not make the training converges significantly higher.

C. Weakly Supervised Object Detection proposed workflow

The main criticism against deep learning is the long labeling time needed for the algorithms to be accurate. The advantage of the proposed classification Deep Learning method is that only one piece of information is needed by an animal in a picture - its taxon - which makes it faster and easier to label training images.

Our weakly supervised Faster-RCNN+FPN method is winning on both sides: it can accurately localize and classify a small signal in a cluttered background thanks to the FPN module while requiring classification labeling effort (taxa) and not the usual object detection effort (localization annotations and taxa). The overall proposed method is shown in figure 1 and described in the following paragraphs.

The localization algorithm. The presence of burst mode on modern camera traps makes it possible to capture the same animal in a short sequence of frames in a few seconds when the camera is triggered. Therefore, the animal motion in those frames is possible to distinguish from the background. A motion-based localization allows the ecologists to handle the workload induced by automatically computing bounding boxes of the animals and it enables to feed the object detection neural network.

Our localization algorithm proceeds following those 6 steps and they are illustrated in figure 2:

- 1) Input a short sequence of images $I_1, I_2, \dots, I_n$.
- 2) It computes a background B computed with median filtering of all the n images of the burst.
- 3) Motion map $M_1, M_2, \dots, M_n$ are euclidean distance between each pixel into $I_1, I_2, \dots, I_n$ and each associated pixel in B
- 4) A binary threshold t is applied on M with $t=12\%$. This new map named T may contain salt and pepper noise.
- 5) A morphological opening operation is applied on those previous binary map T . First an erosion operation with kernel 3×3 allows us to erase noisy connected components. Then, a dilation operation with kernel 151×151 ensures all animal parts are connected. It yields D standing for “denoising”.
- 6) The bounding boxes of the largest connected components (if any) is computed from D and returned.

The accuracy of the localization algorithm is 77.6% on the Papua New Guinea dataset test dataset, with 14.7% of false-

positive error (motion from raindrop, flying bugs, or wind in the vegetation) and 7.7% of false-negative error (The animal motion is missed).

The FP box correction algorithm. All the FP errors are automatically corrected based on a simple if-then-else code structure. It compares the presence of a detected box and the presence of an animal in the class label to avoid FP errors. Thus, Faster-RCNN+FPN is trained on 92.3% of correct labels.

The time performance of the workflow depends mostly on the hardware, the resolution of images, and the implementation. Our $3,776 \times 2,124$ resolution images are predicted with a throughput of 1.78 images/sec on an NVIDIA Tesla V100 GPU inside an NVIDIA DGX1 computer. Regarding previous studies on mammalian species in the African Savannah [5], it took over 28,000 registered citizen scientists between 2 and 3 months to classify 3.2 million images. Our code would take 21 days with one single modern GPU with equivalent resolution. An actual production code would be faster than the current inference throughput due to software optimizations.

V. EXPERIMENTAL RESULTS

In this section, we compare different workflows in terms of accuracy. Then we make an in-depth analysis of its errors on the Papua New Guinea dataset. Finally, we assess different classification methods on a testbed by varying the object-to-image rate (O2I).

A. Comparison

The accuracy comparison of different workflows is shown in table I. Those workflows, their settings, and their characteristics are discussed in the following text.

CNN Classifier. It is a trained classification deep learning model ResNet50 [17]. It takes as an input the overall photograph and classifies among 26 classes, 25 animal taxa, or the empty class. We tested also VGG16, InceptionV3, EfficientNet-B4 with similar settings: 50 epochs and batch size of 32. ResNet50 is more accurate than VGG16 and performs the same as InceptionV3. The neural network is trained with SGD and the learning rate starts at 10^{-3} and is divided by 10 after the 20th and the 40th epochs. Despite we spend time assessing many neural network architecture and optimizer settings, we have not obtained satisfactory results. It takes only 1 hour to be trained, thus this workflow and RP+Classifier are the fastest to converge.

RP+Classifier. (Region proposals+Classifier) This method uses also a ResNet50 classifier with the same training set as the CNN Classifier workflow. The difference is that it is trained on the region of interests yielded by the localization algorithm presented in section 2 and the overall image such as Faster-RCNN and classifier models. An empty class is added because the localization algorithm provides frequent (14.7%) false-positive error motions (again, wind, flying bugs, ...). We explain the performance gained compared to the standard classifier by the fact that the localization algorithm crops most of the cluttered background, thus the classifier part is trained and predicts with a better focus on the moving object.

Fully Supervised Object Detection

Weakly Supervised Object Detection

Fig. 1. Comparison between the usual supervised object detection and the proposed Weakly Supervised Object Detection workflow.

Auto-Keras. (version 1.0.12) It is an AutoML strategy named Auto-Keras Image Classifier [18]. It is a Bayesian optimization algorithm searching the best neural network architecture using validation metrics and calibrating the weights of the candidate on the training dataset. We observe a small improvement compared to the ResNet50 classifier after tuning 100 models trained a maximum of 20 epochs but the training computing cost is about multiplied by 100 (4 days on 1 NVIDIA Tesla GPU). It shows that the famous ResNet50 neural architecture is relevant in our datasets and spend days to search a specific neural architecture may only improve the results a little bit.

Weakly Supervised Faster-RCNN. We use ResNet50 inside Faster-RCNN and Faster-RCNN+FPN meta-architecture. The neural network converges after 200 epochs during 4 hours on Tesla GPU with a batch size of 32. The model is trained with SGD and the learning rate starts at 10^{-3} , then divided by 10 after the 100th, 170th and 190th epochs. In the case of the Papua New Guinea dataset, we compare our workflow with and without manually correcting the bounding box annotations. And in the case of the Missouri dataset, we compare with our workflow and the original bounding boxes downloaded with the Missouri dataset.

The downloaded dataset annotation of the Missouri dataset

Data	Annotation	Method	Presence err.		Presence OK	
			FN	FP	Taxa error	Acc.
PNG	Supervis. classif.	CNN classifier	31.4	6.7	34.3	27.6
	Supervis. classif.	Auto-Keras	28.7	5.1	37.6	28.6
	Weak Obj. Det.(2)	RP+Classifier	7.7	11	10.3	71
	Weak Obj. Det.(2)	F-RCNN	27.7	18.1	12	42.2
	Weak Obj. Det.(2)	F-RCNN+FPN(ours)	2.8	1.7	13.8	81.7
	Supervised O.D.(1)	F-RCNN+FPN	1.8	1.6	10.3	86.3
MIS	Supervis. classif.	CNN classifier	24.2	12.5	30.9	32.4
	Supervis. classif.	Auto-Keras	20.2	11.7	34.4	33.7
	Weak Obj. Det.(2)	RP+Classifier	19.9	8.5	22.5	49.1
	Weak Obj. Det.(2)	F-RCNN	26.4	10.4	22.6	40.6
	Weak Obj. Det.(2)	F-RCNN+FPN(ours)	15.9	4.2	20.5	59.4
	Weak Obj. Det.(3)	F-RCNN+FPN	18.1	4.6	27.3	50

TABLE I

SUMMARY OF THE DIFFERENT DEEP LEARNING METHODS TESTED AND THEIR PERFORMANCE ON TWO DATASETS MISSOURI (MIS) AND PAPUA NEW GUINEA (PNG). REGARDING THE METHOD, THE LOCALIZATION REFERS EITHER TO CLASSIFY AN ANIMAL IS PRESENT IN THE PICTURE (IN THE CASE OF CLASSIFIERS) OR TO LOCALIZE ITS BOUNDING BOXES (IN THE CASE OF FASTER-RCNN MODELS). (1) OUR LOCALIZATION ALGORITHM WITH HANDCRAFT CORRECTIONS, (2) OUR LOCALIZATION ALGORITHM MOTION-BASED (3) LOCALIZATION ALGORITHM BY [16]

((3) in the figure I) [16] contains inaccurate bounding box annotations. It suffers from about 10% images which contain wrong animal labels or boxes without a visible animal in them. Our proposed workflow can cancel a bounding box when the

Fig. 2. Our localization algorithm proceeds running step-by-step. The pacific emerald dove is detected despite its camouflage. After training on images like the sixth one, Faster-RCNN+FPN can be accurate on similar images.

class of the image is “empty” but it cannot handle when a wrong animal class is given and thus the biggest move (e.g., the wind on the grass) is used as a bounding box associated to a false taxon. We run our Localization based on the motion on this same dataset and provide better annotations ((2) in the figure I).

RP+Classifier works better than classifier thanks to its ability to focus on the region of interest but is still inferior to Faster-RCNN+FPN. RP+Classifier loses contextual information: the relative size of the region of interest is lost because it is resized to the fixed size 256x256 to feed the classifier. And more, the surrounding is also lost it may be a major issue when the localization poorly frames the animal, thus some animal parts are cropped and not given to the neural network.

Finally, we observe on the Papua New Guinea dataset that fully supervised deep learning performs +4.6% on the test compared to the weakly supervised counterpart (again, 7.7% of the training dataset is affected by false-negative error boxes).

B. In-depth analysis of challenges

Table II shows that our method displays an overall accuracy of 81.7%. In practice, when an error is made, other images in the sequence often contain at least once the right animal identification.

We do our best to break down all the test datasets into those types of challenging images but choosing for each image is sometimes subject to interpretation. However, we show that

	Data test photograph characteristic						sum
	good image	bad framed	blur	hidden	tiny	unknown	
correct	62.5%	9.3%	5.3%	3.8%	0.7%	0.0%	81.7%
good localization, badly classified	2.2%	3.0%	1.2%	0.2%	0.5%	6.8%	13.8%
False negative localization	0.5%	0.3%	0.2%	0.8%	0.8%	0.2%	2.8%
False positive localization	1.7%	-	-	-	-	-	1.7%
sum	66.8%	12.7%	6.7%	4.8%	2.0%	7.0%	100,00%

TABLE II

DISTRIBUTION OF THE DIFFERENT CHALLENGES AND ERRORS ON PAPUA NEW GUINEA DATASET. THE “TINY” TAG IS FORMALLY DEFINED BY $O2I < 0.2\%$ WHICH ARE UNUSUALLY SMALL REGIONS.

the majority of errors are caused by those challenging images and 7% of errors are unavoidable due to the discovery of new species. We also show that 1/3 of images are challenging.

C. Tiny object recognition

We observe in the previous section that on two datasets Faster-RCNN+FPN was superior to other approaches. To better understanding this accuracy, we use the nMNIST testbed [8] consisting in to classify if the image contains or not the digit ‘3’ in a randomly generated cloud of digits. The experiment is repeated with four different O2I (object-to-image): {19.1%, 4.8%, 1.2%, 0.3%} and the corresponding number of digits in

each image generated is $\{3, 6, 26, 101\}$. For all values of O2I ratio, 11276, 1972, 4040 of training, validation, and testing images, out of which 50% are negative and 50% are positive images.

We report results of the workflow previously proposed [8] and meta-architectures. We observe roughly similar results when changing the architecture (such replace ResNet50 with EfficientNet-B4) so we decide do not to show them to keep the figure readable.

Fig. 3. We apply 5 workflows on the nMNIST testbed. They all be ran at least 3 times and the standard deviation plotted, at the exception of Auto-Keras which is run only once because of prohibitive run time.

We observe that when O2I=19.1%, all workflows perform around 95% but [Pawlowski et al 2020] which is around 92%. When O2I=0.3%, Faster-RCNN+FPN outperforms all other workflows with a 76.83% accuracy, and [Pawlowski et al 2020] shows very unstable results with different run times. More surprisingly, when decreasing O2I, Faster-RCNN without FPN accuracy drops faster than all the other workflows. However, on our biodiversity datasets, Faster-RCNN performs better than classification neural networks.

Those experiments confirm that Faster-RCNN+FPN is the more stable to recognize needle-in-haystacks and performs also very well when the object appears bigger. In camera trap applications, due to the accuracy of Faster-RCNN+FPN, we conclude it deserves to use an additional localization algorithm to compute bounding box annotations to feed it.

VI. FUTURE WORKS

An important line of research consists in evaluating uncertainty estimates allowing to focus ecologists' attention on the most uncertain predictions (badly framed, blurred, unknown species, ...). Ensemble Deep Learning has already been shown to be not only useful to boost the accuracy but also to produce qualitative uncertainty estimates [19]. However, those methods multiply the computing cost over one single neural network. It seems only experiments and throughout analysis may study the balance between costs and benefits of ensembling.

CONCLUSION

Deep learning has led to progress in image recognition and has been tested extensively to identify animals and yet its application to biodiversity monitoring is still in its infancy. The limitation of the application to a real need comes from the need to have a dataset previously labeled for a given region. The successful application of our weakly supervised Faster-RCNN with Feature Pyramid Network addresses the need to recognize small objects with a few thousand labeled images. Compared to the fully supervised counterpart we divide by factor 5 the amount of information (the taxon and 4 box coordinates) and the accuracy drops by less than 5%. Now the ecologists can check if the predictions are correct in priority where the model gives a high uncertainty estimate. This allows focusing on the most challenging images or animals absent from the training dataset. It is by this means that in our campaign we discover the Palm Cockatoo taxon which was absent from our initial training dataset. This method could boost camera traps adoption by tackling the inherent challenges. Additionally, we also hope it will shed light on the benefits of weakly supervised deep learning methods for all disciplinary communities.

Acknowledgement: We would like to thank TotalEnergies SE and its subsidiaries for allowing us to share this material and make available the needed resources. We also thank Aurelien Puiseux for labelling the Papua New Guinea dataset.

REFERENCES

- [1] M. Finlayson, R. Cruz, N. Davidson, J. Alder, S. Cork, Groot, C. Lvque, G. Milton, G. Peterson, D. Pritchard, B. Ratner, W. Reid, C. Revenga, M. Rivera, F. Schutyser, M. Siebentritt, M. Stuij, R. Tharme, S. Butchard, and D. Taylor, "Millennium ecosystem assessment: Ecosystems and human well-being: wetlands and water synthesis," 01 2005.
- [2] K. E. Jones, N. G. Patel, M. A. Levy, A. Storeygard, D. Balk, J. L. Gittleman, and P. Daszak, "Global trends in emerging infectious diseases," *Nature*, vol. 451, no. 7181, pp. 990–993, Feb 2008. [Online]. Available: <https://doi.org/10.1038/nature06536>
- [3] T. Lin, P. Dollr, R. B. Girshick, K. He, B. Hariharan, and S. J. Belongie, "Feature pyramid networks for object detection," *CoRR*, vol. abs/1612.03144, 2016. [Online]. Available: <http://arxiv.org/abs/1612.03144>
- [4] M. Norouzzadeh, A. Nguyen, M. Kosmala, A. Swanson, M. Palmer, C. Packer, and J. Clune, "Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 115, no. 25, pp. E5716–E5725, june 2018.
- [5] A. Swanson, M. Kosmala, C. Lintott, R. Simpson, A. Smith, and C. Packer, "Snapshot serengeti, high-frequency annotated camera trap images of 40 mammalian species in an african savanna," *Scientific Data*, vol. 2, no. 1, p. 150026, 2015. [Online]. Available: <https://doi.org/10.1038/sdata.2015.26>
- [6] M. Willi, R. T. Pitman, A. W. Cardoso, C. Locke, A. Swanson, A. Boyer, M. Veldhuis, and L. Fortson, "Identifying animal species in camera trap images using deep learning and citizen science," *Methods in Ecology and Evolution*, vol. 10, no. 1, pp. 80–91, 2019. [Online]. Available: <https://besjournals.onlinelibrary.wiley.com/doi/abs/10.1111/2041-210X.13099>
- [7] S. Ren, K. He, R. B. Girshick, and J. Sun, "Faster R-CNN: towards real-time object detection with region proposal networks," *CoRR*, vol. abs/1506.01497, 2015. [Online]. Available: <http://arxiv.org/abs/1506.01497>
- [8] N. Pawlowski, S. Bhooshan, N. Ballas, F. Ciompi, B. Glocker, and M. Drozdal, "Needles in haystacks: On classifying tiny objects in large images," 2019.
- [9] "LILA BC labeled information library of alexandria: Biology and conservation," 2021, accessed: 2021-06-03. [Online]. Available: <http://lila.science/datasets>
- [10] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, "A survey on deep transfer learning," *CoRR*, vol. abs/1808.01974, 2018. [Online]. Available: <http://arxiv.org/abs/1808.01974>
- [11] Y. Wei, J. Feng, X. Liang, M. Cheng, Y. Zhao, and S. Yan, "Object region mining with adversarial erasing: A simple classification to semantic segmentation approach," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 6488–6496.
- [12] Z. Yu, Y. Zhuge, H. Lu, and L. Zhang, "Joint learning of saliency detection and weakly supervised semantic segmentation," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 7222–7232.
- [13] F. Xiao and Y. J. Lee, "Track and segment: An iterative unsupervised approach for video object proposals," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 933–942.
- [14] K. K. Singh and Y. J. Lee, "You reap what you sow: Using videos to generate high precision object proposals for weakly-supervised object detection," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 9406–9414.
- [15] "LILA BC labeled information library of alexandria: Biology and conservation," 2021, accessed: 2021-06-03. [Online]. Available: <http://lila.science/datasets/missouricameratraps>
- [16] Z. Zhang, Z. He, G. Cao, and W. Cao, "Animal detection from highly cluttered natural scenes using spatiotemporal object region proposals and patch verification," *IEEE Transactions on Multimedia*, vol. 18, no. 10, pp. 2079–2092, 2016.
- [17] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015. [Online]. Available: <http://arxiv.org/abs/1512.03385>
- [18] H. Jin, Q. Song, and X. Hu, "Auto-keras: An efficient neural architecture search system," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 2019, pp. 1946–1956.
- [19] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in *NIPS*, 2017.

APPENDIX

CHALLENGING IMAGE TYPES

Fig. 1. **Badly framed mammals** (right) are difficult to distinguish because many mammals have brown or grey fur in this area. To compare, a good image of “dorcopsis luctuosa” taxon is shown (left). The two predictions are correct.

Fig. 2. **Motion blur** effect (right). To compare, a good image of “bandicoot” taxon is shown (left). The two predictions are correct.

Fig. 3. **Occlusions** of leaves and branches (right) can hide important parts of an animal that are necessary to identify it. To compare, a good image of “feral pig sus scrofa” taxon is shown (left). The two predictions are correct.

Fig. 4. **In absence** of a taxon in our training dataset, the classifier will classify them among known classes leading to an avoidable error. The “palm cockatoo” taxon (right) was not photographed during our previous campaigns.

Fig. 5. **Tiny** pixel representation of a “rat” taxon. The rat seen from the back is running and jumping in the opposite direction of the camera. The right image rat representation takes only 122x110 pixels. It makes 0.17% of the surface of the overall image because not only is it a small mammal but it is also far from the camera. Contrast enhancement and magnification allow us to see it: its head is in the right bottom corner, its tail above its head and back paws do not touch the ground. We can enumerate three reasons for this high level of accuracy: Imagenet pre-training already contains rat images, we collect also many rat images in our training dataset, finally, Faster+FPN is specially adapted for this “needle-in-haystack” situation.