

# Ensemble Learning techniques for object detection in high-resolution satellite images

Arthur Vilhelm<sup>1</sup>, Matthieu Limbert<sup>1</sup>, Clement Audebert<sup>1</sup>, and Tugdual Ceillier<sup>1</sup>

Earthcube, 75009 Paris, France  
[www.earthcube.eu](http://www.earthcube.eu)

**Abstract.** Ensembling is a method that aims to maximize the detection performance by fusing individual detectors. While rarely mentioned in deep-learning papers applied to remote sensing, ensembling methods have been widely used to achieve high scores in recent data science competitions, such as Kaggle. The few remote sensing articles mentioning ensembling mainly focus on mid resolution images and earth observation applications such as land use classification, but never on VHR images for defense-related applications or object detection. This study aims at reviewing the most relevant ensembling techniques to be used for object detection on very high resolution imagery and shows an example of the value of such technique on a relevant operational use case (vehicle detection in desert areas).

**Keywords:** Deep learning · Convolutional neural networks · earth observation · ensembling · data augmentation · object detection.

## 1 Introduction

### 1.1 Object detection

Recent increases in the quality and volume of Very High Resolution (VHR) Earth observation images from commercially-available satellites has allowed breakthroughs in tools that can be used for mapping and processing objects and areas of interest to the defence and security communities.

Photo interprets can only analyse a small part of all this data and may take hours to find a vehicle in a desert area. Computer vision algorithms and deep learning solutions applied to satellite images allow the development of automatic tools to monitor large areas in strategic regions. Large quantities of commercial satellite imagery of very high quality can be obtained to label and train algorithms to detect, classify and identify infrastructures such as airplanes, vehicles, roads or buildings. At a time when it has become easy to classify cats and dogs in various situations, Earth observations raise new challenges by the diversity of scenes and by the expectations of end-users.

Ensembling is an approach using a combination of multiple algorithms (from 3 to 1000s) to achieve better performances than could be obtained from any

of the constituent machine learning algorithms alone. Several academic papers describe the theoretical basis of ensemble learning with first implementations starting in the 1990's. While rarely mentioned in deep-learning papers applied to remote sensing, ensembling methods have been widely used to achieve high scores in recent Kaggle competitions. Remote sensing papers mentioning ensembling mainly focus on mid resolution images and earth observation applications such as landuse classification, but never on VHR images for high resolution segmentation.

### 1.2 Ensemble of neural networks

Convolutional neural networks (CNN) is the most promising machine learning algorithm for object detection and classification in images. It is a class of supervised algorithms that takes advantage of a large dataset of labeled objects to learn a generalized model of this data. At prediction time, the association of convolutional kernels with efficient GPU make CNN an efficient tool with the potential for human-level accuracy.

Ensemble techniques are used in a wide variety of problems of machine learning. By combining several algorithms, it is known to reduce the bias of a single detector or its variance. Those problems of bias and variance of prediction are well known and can be caused by using too big a model with too small a dataset or by the fact that Earth observation data is too heterogeneous for the model to capture the full picture. In other words, there is not a single model that can properly model the full set of data to process. For instance, stacking is known to reduce the bias of the algorithm whereas bagging or Test-Time Augmentation decrease its variance.

Most ensembling methods increase the computing time (training time and/or prediction time) but provide an efficient way to push forward the state-of-the-art. This study aims at reviewing the most relevant ensembling methods to be used for object detection based on VHR satellite imagery. The selected methods have been implemented in a Python-keras framework and the performance gain of a combination of two methods is detailed in the last section.

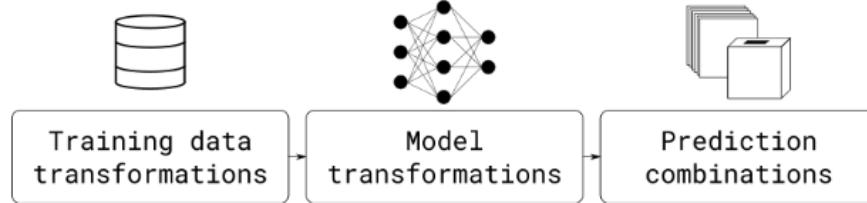
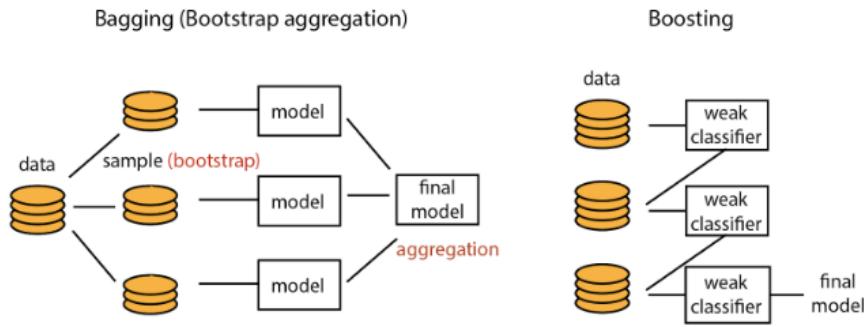
## 2 State of the art of Ensembling techniques

### 2.1 Introduction

All ensemble learning methods relies on combinations of variations in the following spaces:

1. Training data space: Transformations of the data used to train the models
2. Model space: Transformations of the models trained on the dataset
3. Prediction space: Transformations and combinations of the prediction

While research articles usually focus on single approaches in one of those categories, challenge-winning solutions use a combination of methods to achieve

**Fig. 1. Ensembling pipeline architecture****Fig. 2. Comparison between Bagging and Boosting.**

**Details:** The main dataset is divided into 3 subsets that are used to train 3 replicas of a model. This is done in parallel for Bagging and iteratively for Boosting.

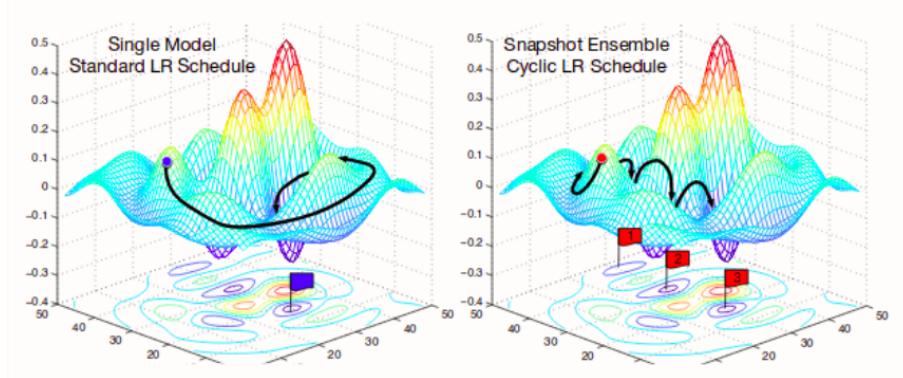
the highest performances. In the following section, the ensemble learning methods relevant to object detection based on VHR satellite imagery will be listed according to the 3 categories previously introduced.

## 2.2 Training data variation-based methods

**Bagging [1] [2]:** Bagging, or bootstrap aggregating, generates new data-sets by extracting random sub-samples from the original one. Training the model on those random sub-samples reduces the variance of the training data and improves the accuracy of final ensembling model. Compared to boosting, which samples the original dataset in order to correct the errors of the previous models iteratively, bagging can be used to train models in parallel. Another way of performing bagging without splitting the dataset, known as synthetic bagging, is to use data augmentation : models are then trained on the same dataset with different kinds of augmentations.

## 2.3 Model variation-based methods

In remote sensing, different CNN architectures are known to have different strong-points: some models might be better at separating adjacent objects, de-

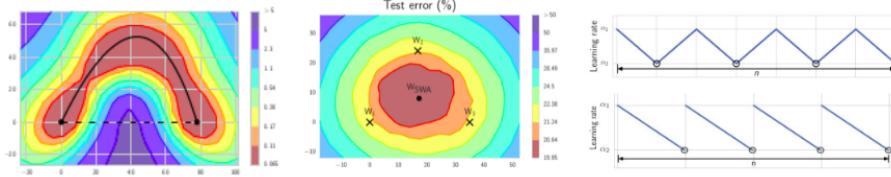


**Fig. 3. classical SGD optimization vs Snapshot Ensembling.** source: [5]

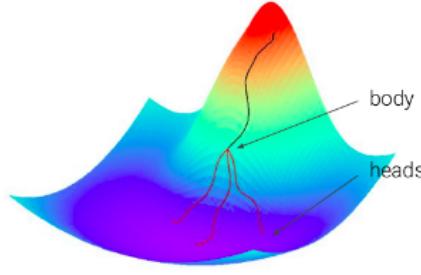
tecting smaller elements or providing more precise segmentation. Varying hyperparameters, such as loss and optimizer, also provides interesting models to combine.

**Snapshot Ensembling methods:** The eponym sub-method Snapshot Ensembling (SSE) [5] is a method providing cheaply a large number of complementary CNN that could be combined to improve performances. While those CNN share the same architecture, the use of a clever “cyclical scheduling” of the learning rate allows the networks to reach different wide or narrow local optima along training. The learning rate is gradually decreased to allow the network to converge to some local minima. The algorithm then takes a Snapshot of the weights, before increasing the learning rate again in order to jump to another optimum. This cycle can be repeated many times to produce a set of snapshots of the desired size. Fast Geometric Ensembling (FGE) [6] is a variation of snapshot ensembling using a different scheduling scheme: rapidly cycling linear piecewise variations of the learning rate instead of the original smooth cosine one. This approach is justified by the existence, according to the authors of [10], of low-loss paths between interesting minimas. FGE is reportedly faster to train than Snapshot ensembling while showing performances improvements.

**Stochastic Weight Averaging (SWA)** [7] [8]: Stochastic Weight Averaging is an optimization method approximating Fast Geometric Ensembling while reducing the computational cost. It uses a combination of two models: the first keeps a running average of the weights at each iteration, while the second model explores the weight space following a cycling learning rate scheduler. At inference, Stochastic Weight Averaging provides a single model with averaged weights.



**Fig. 4. Illustration of FGE and SWA methods.** **Left:** example of a low error rate path between 2 local minimas (FGE method). **Middle:** example of 3 weights networks  $W_1$ ,  $W_2$  and  $W_3$ , and their weights average  $W_{\text{swa}}$  that is on a wider local minima with lower error rate (SWA method). **Right:** the corresponding scheduling schemes of the learning rate (FGE on the top, SWA on the bottom). Sources: [6] [7]



**Fig. 5. Illustration of the optimization process in the Hydra framework.** Source: [9]

**Hydra** [9]: Hydra is a CNN ensembling method that consists of training an initial “coarse” model, called the “body”: a snapshot of the training with average performances. From this body, a set of “Heads” are fine-tuned by applying various data augmentations. Just like snapshot ensembling, this approach provides a set of local minima to combine from a single neural network model.

## 2.4 Prediction variation-based methods

**Voting-based methods** [10]: This very simple fusion method combines prediction by counting the agreeing detectors. This vote could be counted either on the raw pixel outputs of the detector or object-wised from the vectorised shapefile.

**Test time augmentation (TTA)** [11]: CNN are not intrinsically invariant to translation, rotation, scale and other transformations. Applying such transformations to the input image before feeding it to the network can correct this variability and increase the performances of the model. While test time augmentations dont need any new training, it increases linearly the prediction time.

**Stacked generalization [12] [13]:** Stacking relies on a “Meta-learner”, which take the predictions of the individual detectors as its input and is trained to output a refined output. While stacking has the potential of learning the optimal fusion of detectors, it is quite a rigid method as it involves additional training and is very specific to the input detectors.

**Dropout as a Bayesian Approximation [14]:** Dropout is the process of switching off random neurons within the network during training. During training, this method helps avoiding overfitting and can be seen as an ensemble of network sub-sampled in the model space. [14] introduces the idea of using dropout not only during training but also in prediction phase. Several predictions are run by removing different random neurons from the network. Those predictions can then be averaged to achieve better generalisation.

### 3 Application

In this article, we present the results of the tests of three ensembling techniques on the vehicle detection tasks in Very High Resolution images (VHR).

**Use case:** In object detection for military purposes, a use case of interest is detecting vehicles in the desert. On the operational side, it is very difficult and time consuming to manually find those small objects in the middle of very wide areas, as it can be seen on Fig. 6. On the machine learning side, the task of detecting vehicles in such environment is very challenging, even for high performance detectors, since a lot of objects look like vehicles and are very difficult to distinguish (small trees and rocks, very small buildings, etc.), which leads to an increase of false positives. An example of such objects is shown on fig. 7.

**Original model:** We performed the tests of the ensembling techniques on a very generic and high performance model owned by Earthcube (92% of F1-score on a diverse environment testing set detailed in fig. 6). It is a custom ResNet-Unet architecture model with 8.5M parameters. It has been trained for the segmentation task on a very large training dataset (330K civil vehicles labeled at a pixel scale) comprising VHR satellite images of different types of environments (urban, vegetal, desertic, coastal) and conditions (aerosol, summer, snow, etc.) from sites located all around the world.

**Testing set:** In order to assess and compare the tested ensembling techniques on the use case, we used a very challenging testing set composed by images of desert environment that is particularly difficult for vehicle detection (detailed in fig. 6). We used a very large testing set (composed of 36 different sites for a total area of 623.7 km) in order to make the assessment of the ensembling techniques relevant.

	Diverse environment	Desert environment
Environment type	diverse (coastal, desert, urban, dense vegetation, etc.)	desert
# of images	29	36
Satellite	WV3	WV3
Pixel size	0.3m	0.3m
size (km <sup>2</sup> )	25	580
# vehicle / km <sup>2</sup>	755	4

**Fig. 6. Information on the testing sets.** **Left:** Information on the diverse environment testing set (original model has 92% of F1-score on this testing set). **Right:** Information on the desert environment testing set (very challenging for vehicle detection since objects look similar to vehicles).

**Tested ensembling techniques:** Based on the state of the art survey, the ensembling techniques that are suitable to the considered use case have been selected. Since the training dataset is very large, the additional number of epochs to train the model on is a key constraint.

Three ensembling techniques have been tested: TTA combined with Bayesian Dropout, FGE and SWA. TTA and Bayesian Dropout do not need any additional training since are applied at prediction time. Based on previous experiments, the performance is increased when these two techniques are combined. On the other hand, FGE and SWA are promising and only need fine-tuning of the original model so they are not time-consuming to implement.

Other ensembling techniques that need a large number of epochs to train, such as Bagging, SSE, and Hydra, have been discarded (and may need a full article by themselves). Bagging is both time consuming and complex to train since the training dataset size is modified so new tests have to be performed to optimize the model size. SSE needs to train the model from scratch several times in a row and provides lower performance increase than FGE and SWA according to preliminary study. Finally, Hydra needs to ensure convergence of very different models to diversify augmentations and vary fundamental parameters which is both complex and time-consuming.

**Result of the tests:** The results of the tests are gathered on the table in fig. 8. The selected ensembling techniques have been tested on two high recall operating point.

TTA combined with Bayesian Dropout achieves a very large performance increase on such challenging desert environment: +0.86% of recall while decreasing the number of false positives by 25% for an operating point with 92% of recall. This result means that this technique has operational value since it can



**Fig. 7.** Hard examples in the desert environment testing set where many objects can be confused with vehicles. Yellow points: vehicles in the ground truth. **Top:** small trees can look similar to vehicles. **Bottom:** Very small buildings can look similar to vehicles.

boost very significantly on challenging environment a detector that has already operational performance.

FGE achieves a very similar performance boost to TTA / Bayesian Dropout: -31% of false positives while achieving the same recall for the 92% recall operating point. So it provides similar operational value than TTA / Bayesian.

	Ensembling techniques	Computation time		Detection performance			
		# predictions	# Add. epochs to train	FP/km <sup>2</sup>	Recall	FP/km <sup>2</sup> relative change	Recall absolute change
Operating point 1	original model	1	0	15.04	91.85%	-	-
	TTA / Bayesian Dropout	7	0	11.16	92.71%	-25.80%	+0.86%
	FGE	5	20	10.32	91.85%	-31.38%	+0.00%
	SWA	1	40	14.76	92.05%	-1.86%	+0.20%
Operating point 2	original model	1	0	6.43	86.15%	-	-
	TTA / Bayesian Dropout	7	0	3.40	85.46%	-47.12%	-0.69%
	FGE	5	20	4.77	85.87%	-25.81%	-0.26%
	SWA	1	40	5.72	86.23%	-11.04%	+0.08%

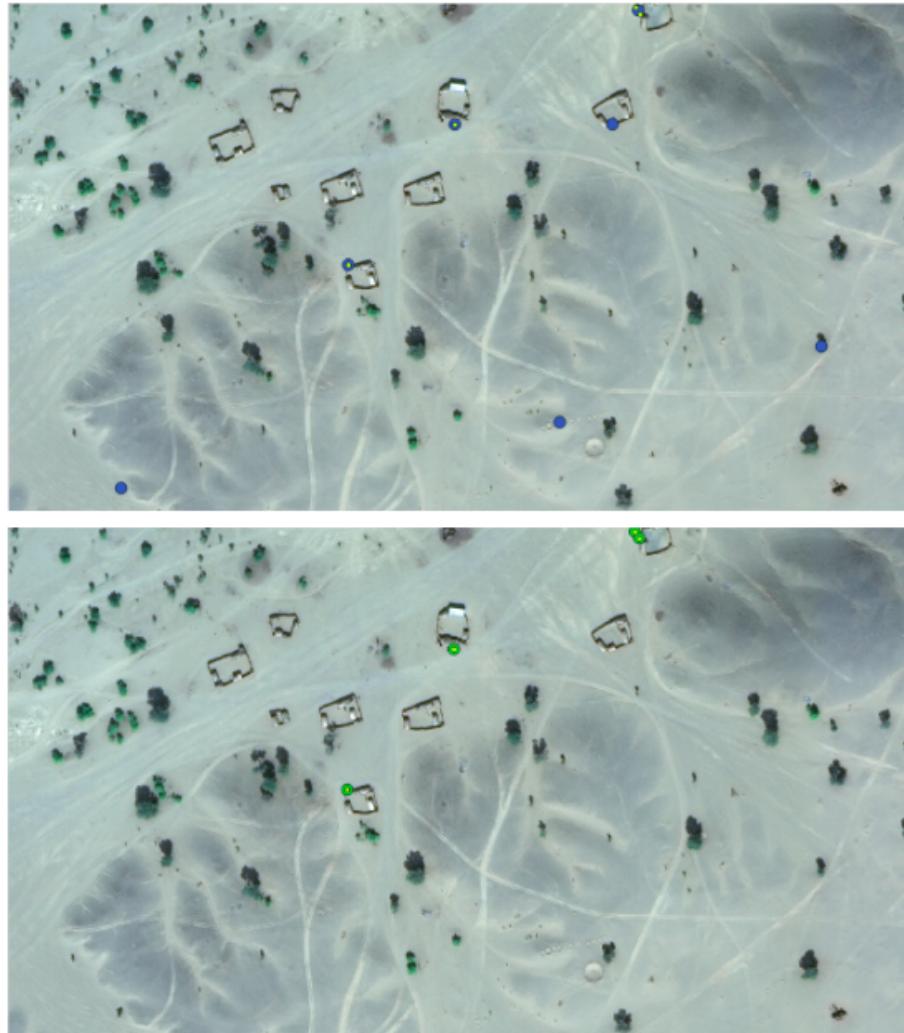
**Fig. 8. Results of the tests on ensembling techniques on the desert environment testing set.** **Interpretation:** Those techniques are compared on two operating points. The first one has a very high recall close to 92%, the second one has a lower recall close to 86% with less false detections. For each technique and each operating point, we assessed both the detection performance and the computation time. **Detection performance metric:** For the detection performance, we assessed the performance gain by using both the absolute increase in recall, and the relative decrease of false positives per km<sup>2</sup>. We have not used the F1-score metric since it is not suitable for desert areas: there are very few vehicles in very large areas, which make the F1-score very unstable. **Computation time metric:** To assess the prediction time, we use the number of predictions to perform so that it does not depend on the available hardware or the image size. For training time, we used the number of additional epochs to train the model.

Finally, SWA achieves a significant performance boost even though it is smaller than TTA / Bayesian / Dropout and FGE: -11% of false positives on the 85% recall operating point, and small performance increase on the 92% recall operating point. However, this technique only needs one prediction (rather than respectively 7 and 5 for TTA / Bayesian Dropout and FGE). So this technique is highly recommended for applications that cannot use additional time during prediction. One could conclude that for a very high level of recall (+90%), substantial performance boosts can only be achieved through voting-based techniques (such as TTA / Bayesian and FGE) and that using one unique model might not be sufficient.

## 4 Conclusion

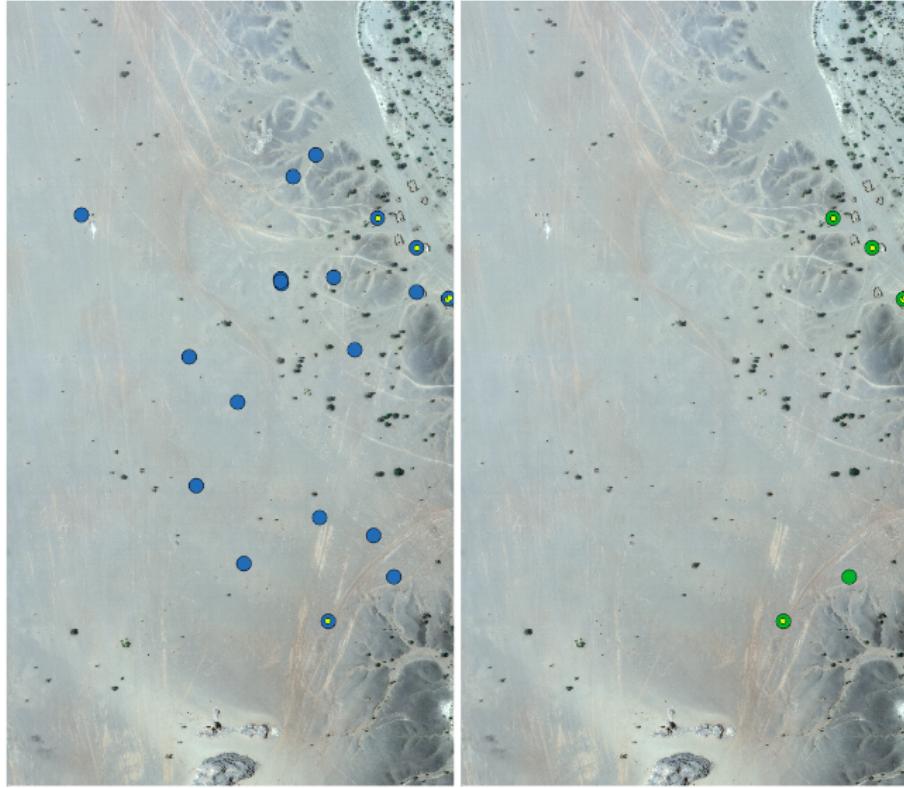
Ensembling techniques provide a toolbox to increase performances of a given model with a given training dataset, at each level of the detection pipeline (training data, model, prediction).

After detailing the most relevant ensembling techniques for object detection in VHR images, three techniques have been selected for the use case of vehicle



**Fig. 9. Example of detections on the desert environment testing set.** **Top:** Detection of the original model in blue (operating point 1) and ground truth in yellow. **Bottom:** Detection of the model with TTA and Bayesian Dropout combined in green (operating point 1) and ground truth in yellow. **Interpretation:** We observe that the original model has a high recall and provides some false detections on small trees and buildings that look similar to vehicles. The combination of TTA with Bayesian Dropout removes these false detections and increases the recall by detecting an additional vehicle on the top.

detection on VHR images with an original very high performance model trained on a very large dataset. These techniques have been assessed and compared and



**Fig. 10. Example of detections on the desert environment testing set.** **Left:** Detection of the original model in blue (operating point 1) and ground truth in yellow. **Right:** Detection of the model with TTA and Bayesian Dropout combined in green (operating point 1) and ground truth in yellow. **Interpretation:** We observe on the left side that the original model achieves a very good recall with some false detections. The detection using ensembling removes most of these false detections without affecting the recall and increase the overall performance.

provide substantial performance increases on a very challenging desert environment testing set.

At the price of extra training/prediction time, ensembling levels up state of the art algorithms and brings automatic vehicle detection at the level of a human expert. In general, the work presented here allows for better performances and higher reliability for any object detection solution. This can give the armed forces better tools to perform their missions, enabling them to automatically extract relevant information from satellite images.

This can be used for specific tasks such as identifying isolated vehicles in deserts or measuring activity at a given location but also for more broadly defined tasks such as strategic sites monitoring. It can be used for any type of

object (vehicles, vessels, buildings, roads, etc.) and can either be deployed alone without any human looking at the images or be applied as a tool to help image analysts in their missions.

## References

1. Breiman, Leo. “Bagging predictors.” *Machine learning* 24.2 (1996): 123-140.
2. Li, Hanxi Li, Yi Porikli, Fatih Mingwen, Wang. (2016). “Convolutional Neural Net Bagging for Online Visual Tracking.” *Computer Vision and Image Understanding*. 153. 10.1016/j.cviu.2016.07.002.
3. Perez, L., Wang, J. (2017). The effectiveness of data augmentation in image classification using deep learning. *arXiv preprint arXiv:1712.04621*.
4. Luke Taylor and Geoff Nitschke. Improving deep learning using generic data augmentation, 2017.
5. Gao Huang, Yixuan Li, Geoff Pleiss, Zhuang Liu, John E. Hopcroft, and Kilian Q. Weinberger. Snapshot ensembles: Train 1, get M for free. *CoRR*, abs/1704.00109, 2017.
6. Timur Garipov, Pavel Izmailov, Dmitrii Podoprikhin, Dmitry Vetrov, and Andrew Gordon Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns, 2018.
7. Pavel Izmailov, Dmitrii Podoprikhin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization, 2018.
8. Ben Athiwaratkun, Marc Finzi, Pavel Izmailov, and Andrew Gordon Wilson. Improving consistency-based semi-supervised learning with weight averaging, 2018.
9. Rodrigo Minetto, Mauricio Pamplona Segundo, and Sudeep Sarkar. Hydra: an ensemble of convolutional neural networks for geospatial land classification, 2018.
10. Roberto Battiti and Anna Maria Colla. Colla, m.: Democracy in neural nets: Voting schemes for classification. *neural networks* 7(4), 691-707. *Neural Networks*, 7:691707, 12 1994.
11. Murat Seckin Ayhan and Philipp Berens. Test-time data augmentation for estimation of heteroscedastic aleatoric uncertainty in deep neural networks. 2018.
12. David H. Wolpert. Stacked generalization. *Neural Networks*, 5:241259, 1992.
13. Milad Mohammadi and Subhasis Das. SNN: stacked neural networks. *CoRR*, abs/1605.08512, 2016.
14. Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning, 2015.