

# INTRODUCING TOPOGRAPHY IN CONVOLUTIONAL NEURAL NETWORKS

Maxime Poli\*    Emmanuel Dupoux\*<sup>†</sup>    Rachid Riad\*

\*ENS, PSL Research University, EHESS, CNRS, INRIA

<sup>†</sup>Meta AI Research

## ABSTRACT

Parts of the brain that carry sensory tasks are organized topographically: nearby neurons are responsive to the same properties of input signals. Thus, in this work, inspired by the neuroscience literature, we proposed a new topographic inductive bias in Convolutional Neural Networks (CNNs). To achieve this, we introduced a new topographic loss and an efficient implementation to topographically organize each convolutional layer of any CNN. We benchmarked our new method on 4 datasets and 3 models in vision and audio tasks and showed equivalent performance to all benchmarks. Besides, we also showcased the generalizability of our topographic loss with how it can be used with different topographic organizations in CNNs. Finally, we demonstrated that adding the topographic inductive bias made CNNs more resistant to pruning. Our approach provides a new avenue to obtain models that are more memory efficient while maintaining better accuracy.

**Index Terms**— Convolutional neural networks, topography, inductive bias, pruning

## 1. INTRODUCTION

Convolutional neural networks (CNNs) reached state-of-the-art performances in a number machine learning tasks dealing with naturalistic signals, such as image classification [1, 2], audio pattern recognition [3], or speech recognition [4].

CNNs combine a number of basic operations (convolutions, non-linearity, and downsampling) to build useful hierarchical representations, and discard useless information. The introduction of these computations was in some way inspired by the brain, and in return these models can yield insights for neuroscience [5]. Indeed, in the earlier stages of CNNs, learned filters only detect colors, edges, or orientations; whereas higher level and specialized representations are extracted later in the model [6]. This is *reminiscent* of what is observed experimentally in the visual cortex [7].

---


This work was performed using HPC resources from GENCI-IDRIS (Grant 2022-AD011013620). This work was supported in part by the Agence Nationale pour la Recherche (ANR-17-EURE-0017 Frontcog, ANR-10-IDEX-0001-02 PSL\*, ANR19-P3IA-0001 PRAIRIE 3IA Institute) and grants from CIFAR (Learning in Machines and Brains) and Meta AI Research (Research Grant).

However, there is no organization within a certain layer and no notion of proximity between each filter, which differs from spatial structures observed in sensory streams in the brain. Indeed, neurons in the visual [8, 9] and in the audio pathways [10, 11] are organized in "maps", i.e. nearby neurons have similar responses. In brains, these topographic organizations of responses are observed both for (1) low-level computation such as the frequency tuning in the temporal lobe of the auditory cortex (tonotopy) [12], but also for (2) higher-level representations such as those observed for face identification [13] (faciotopy). Therefore, in this study, we ask the following question: Is the topographic organization of computations a useful inductive bias for CNNs?

In machine learning, topography as an inductive bias has been tested only in later stages of computations. It was introduced in generative models [14], and in variational autoencoders [15], as a prior over the latent variables, but was only used in these higher representations and only in vision. In CNNs, Lee et al. induced topography thanks to an auxiliary loss in the last stages of their Topographic Deep Artificial Neural Network [16]. They mapped the positions of biological neurons in the brain to the fully connected layers of an AlexNet [17]. This loss is derived by considering the pairwise similarities of the activations of neurons as a function of their distance. A limitation of this study is that it relies on invasive brain recordings to assign positions to artificial neurons.

Thus, in this study, we propose a new methodology that introduces topography in convolutional layers of CNNs and which do not rely on any external data. We enforced topography with an auxiliary loss that can be used with any topographic organization. We evaluated our approach on four classification tasks, both on images and audio with three different architectures.

Finally, the over-parametrization of CNNs is often considered as a core ingredient of their predictive power [18]. Yet, from an *engineering* point of view, the search of memory efficient models is crucial for their deployment. On the other hand, biological neural networks need to balance between robustness and accuracy with a limited number of connections. This phenomenon, referred as *synaptic pruning* in neuroscience, promotes specialized and efficient pathways. Indeed, the number of connections between biological neurons (axons and dendrites) decreases greatly between early



**Fig. 1:** Our approach to introduce topography in convolutional layers with the topographic loss.

childhood and puberty [19]. From these two perspectives, we also evaluated our topographic models for their resistance to magnitude-based pruning [20], and found out that they were more robust than baselines.

The source code to reproduce our methods and experiments is available at [github.com/bootphon/topography](https://github.com/bootphon/topography).

## 2. METHODS

### 2.1. Notations

We consider a CNN with  $L$  convolutional layers. The input of the network is  $x$ , and each convolutional layer  $\ell$  receives a tensor of intermediate representations of shape  $(C_{in}^\ell, H_{in}^\ell, W_{in}^\ell)$  and outputs a tensor of representations  $y^\ell$  with a shape  $(C_{out}^\ell, H_{out}^\ell, W_{out}^\ell)$ .  $\tilde{y}^\ell$  of size  $(C_{out}^\ell, H_{out}^\ell, W_{out}^\ell)$  is a flattened view of  $y^\ell$ .

### 2.2. Topographic loss

We modified existing CNNs by assigning positions to each channel of the convolutional layers. We propose that the topography in CNNs can be expressed as: nearby channels in the topography must have similar activations. In order to enforce this, we introduce an auxiliary loss, that is to be minimized in addition to the classification loss (See fig. 1).

We define a distance function  $d^\ell : \{1, \dots, C_{out}^\ell\}^2 \rightarrow \mathbb{R}^+$  where  $d^\ell(i, j)$  is the euclidean distance between the positions associated to the output channels at indices  $i$  and  $j$  of the layer  $\ell$ . We denote by  $S_{i,j}^\ell$  the cosine similarity function between those outputs channels:

$$S_{i,j}^\ell : y^\ell \mapsto \frac{1}{\|\tilde{y}_i^\ell\|_2 \|\tilde{y}_j^\ell\|_2} \sum_{\substack{1 \leq h \leq H_{out}^\ell \\ 1 \leq w \leq W_{out}^\ell}} y_{i,h,w}^\ell y_{j,h,w}^\ell. \quad (1)$$

For a given  $y^\ell$ , the cosine similarity matrix  $S^\ell(y^\ell)$  is efficiently computed by matrix product. All operations considered here remain differentiable and can be added to the computational graph. The associated topographic loss  $\mathcal{L}_{topo}^\ell$  at layer  $\ell$ , and the full loss to be minimized  $\mathcal{L}$  are defined as:


$$\mathcal{L}_{topo}^\ell = \frac{2}{C_{out}^\ell(C_{out}^\ell - 1)} \sum_{1 \leq i < j \leq C_{out}^\ell} \left( S_{i,j}^\ell - \frac{1}{d^\ell(i, j) + 1} \right)^2, \quad (2)$$

$$\mathcal{L} = \mathcal{L}_{classif} + \lambda \cdot \underbrace{\frac{1}{L} \sum_{1 \leq \ell \leq L} \mathcal{L}_{topo}^\ell}_{\mathcal{L}_{topo}}. \quad (3)$$

$\lambda$  is the weight given to the topographic loss. We chose the same target similarity profile in our loss as what was found out by [16]. The outputs  $y^\ell$  considered in the topographic loss are taken before any regularization or activation function.

### 2.3. Topographic organization

In the formulation of the loss, the topographic organization, *i.e.* the positions associated to the channels and used in  $d^\ell$ , was left undefined. Here, we tested and enforced different topographic organizations between the channels in convolutional layers. The schemes we used are illustrated in fig. 2, using 256 channels as an example.



**Fig. 2:** Positions of 256 channels for the different topographic organizations considered.

First, we used a 2D grid with regularly distributed positions. Coordinates are between 0 and 1 as the impact of this loss is weighted by the topographic weight  $\lambda$  (See fig. 2a). Let  $N^\ell = \lceil \sqrt{C_{\text{out}}^\ell} \rceil$  be the number of positions per axis. For each channel  $1 \leq i \leq C_{\text{out}}^\ell$ , the associated position is  $(\frac{r}{N^\ell}, \frac{q}{N^\ell})$ , where  $q$  and  $r$  are the quotient and remainder in the euclidean division of  $i - 1$  by  $N^\ell$ .

Second, we investigated the influence of the density of positions with a variant of the regular grid: the "nested grid" (See fig. 2b). We tested such grid as density of neurons vary across and within brain regions [21]. It is made of nested squares, with positions at each vertex and in the middle of each edge. Thus, there is a higher density in the center.

Then, as [22, 23] found the helix topography to be useful for audio, we also used positions taken regularly from a circle (See fig. 2c). For each channel at index  $1 \leq i \leq C_{\text{out}}^\ell$ , the associated position is  $(\cos(\frac{2\pi}{C_{\text{out}}^\ell} i), \sin(\frac{2\pi}{C_{\text{out}}^\ell} i))$ . Finally, we compared these flat 2D schemes to their equivalent in 3D (See figs. 2d to 2f).

### 3. EXPERIMENTS

We provide the full experimental setup and the procedure to reproduce our results including all hyperparameters.

#### 3.1. Experimental setup

We used the ResNet-18 [1], VGG-16 [24] and DenseNet-121 [2] architectures. Models were trained with stochastic gradient descent with a learning rate of 0.01, a batch size of 256, a momentum of 0.9, and a weight decay of 0.01. We used a learning rate scheduler with cosine annealing [25] without restart and linear warmup of 30% of the number of epochs.

With the exception of the number of epochs and the data augmentation which are set according to the dataset, we kept the training procedure identical for all tasks and comparisons. We used the introduced parameters in [1] for the first layer of the ResNet-18 and DenseNet-121 with a kernel size of 3, a stride of 1, and a padding of 1. For ResNet-18 (resp. DenseNet) we added topography only to convolutional layers at the end of each building block (resp. dense block) to account for residual connectivity (resp. concatenation).

For each approach, we selected the best model checkpoint and  $\lambda$  based on the validation set. We report the mean accuracy and the standard deviation on the test set across 5 seeds. Our baselines on both modalities replicated [1, 2, 26, 27, 28].

#### 3.2. Image classification

We benchmarked our models on CIFAR10 and CIFAR100 [29]. CIFAR10 (resp. CIFAR100) is made of  $32 \times 32$  color images split into 10 (resp. 100) classes, each containing 6000 images. We used the official split with 50000 images for the development set and 10000 for the test set. We then split the development set into training and validation sets, with 5000

images for the validation set. This split is the same for each run. We applied image normalization, random cropping with a padding of 4, and horizontal random flipping with a 50% probability; and we train for 100 epochs on both datasets.

#### 3.3. Audio classification

We evaluated our models on audio classification with two different tasks: bird sound detection (BirdDCASE) [32] and speech command classification (SpeechCommands) [33]. SpeechCommands is a dataset of spoken English words, while BirdDCASE is made of audio recordings with either the presence or the absence of bird sounds of any kind. BirdDCASE has the particularity that it is originally made from three development sets, each recorded under differing conditions. Those development sets have different balances of positive/negative cases, different bird species, different background sounds, and different recording equipment.

Therefore, we used the official split for SpeechCommands, whereas for BirdDCASE we split the three development datasets with the same proportion, using 80% for the training, 10% for the validation set, and 10% for testing. Audio is sampled at 16 kHz. The inputs of the networks are log mel spectrograms with a window of 25ms computed every 10ms, normalized by the training data statistics.

We trained for 36 epochs for BirdDCASE and 12 epochs for SpeechCommands to see 1M examples during training. On BirdDCASE, as the samples are longer than 1 second, we train with random cropped windows of 1 second. We evaluated by splitting the samples into non-overlapping 1-second windows and averaging the logits over windows, as [28].

#### 3.4. Pruning experiments

We compared the resistance to pruning with the  $L^2$  weight-based pruning [20]. For this experiment, we considered a ResNet-18 trained on CIFAR10 with a 2D grid topography and  $\lambda = 1$ . For each topographic layer, we pruned a fixed proportion of channels that have the weights with the lowest  $L^2$ -norm. We performed similarly on the corresponding layers of the baseline counterpart. We did not retrain the models after pruning.

## 4. RESULTS AND DISCUSSION

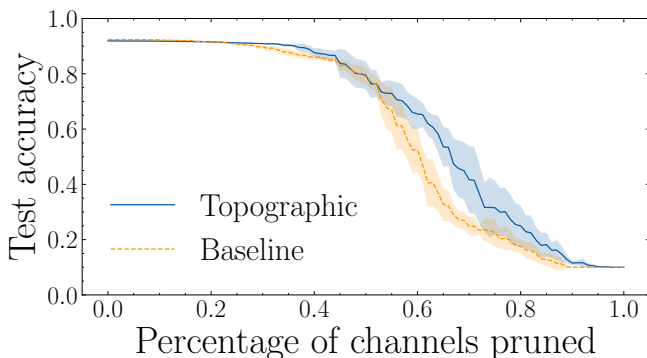
Introducing topography leads to results similar to baselines overall (tables 1 and 2), although there are configurations with topography that provide slight improvements for the different architectures, datasets, and modalities. In addition, for comparison, we also showed the trade-off between accuracy and proportion of pruned channels for the baseline and one topographic model (See Figure 3). Without pruning, the mean accuracy for this given topographic model (ResNet-18, 2D grid,  $\lambda = 1$ ) is 91.9% (92.45% for the baseline). As pruning increases, our topographic model maintains better performance

	CIFAR10			CIFAR100		
	DenseNet121	ResNet18	VGG16	DenseNet121	ResNet18	VGG16
2D grid	<b>93.58</b> $\pm$ 0.21	92.55 $\pm$ 0.32	93.55 $\pm$ 0.10	73.98 $\pm$ 0.38	70.86 $\pm$ 0.31	72.53 $\pm$ 0.28
Circle	93.46 $\pm$ 0.39	<b>92.68</b> $\pm$ 0.26	93.51 $\pm$ 0.15	74.05 $\pm$ 0.26	70.76 $\pm$ 0.48	72.78 $\pm$ 0.40
2D nested	93.42 $\pm$ 0.18	92.53 $\pm$ 0.14	93.45 $\pm$ 0.17	73.84 $\pm$ 0.38	70.93 $\pm$ 0.19	<b>72.87</b> $\pm$ 0.34
3D grid	93.47 $\pm$ 0.18	92.38 $\pm$ 0.15	<b>93.61</b> $\pm$ 0.19	74.02 $\pm$ 0.46	70.43 $\pm$ 0.25	72.77 $\pm$ 0.28
Sphere	93.55 $\pm$ 0.19	92.50 $\pm$ 0.09	93.55 $\pm$ 0.13	73.97 $\pm$ 0.28	<b>71.04</b> $\pm$ 0.41	72.44 $\pm$ 0.25
3D nested	93.55 $\pm$ 0.17	92.34 $\pm$ 0.31	93.53 $\pm$ 0.26	73.94 $\pm$ 0.50	70.91 $\pm$ 0.26	72.59 $\pm$ 0.09
Base	93.43 $\pm$ 0.13	92.45 $\pm$ 0.17	93.57 $\pm$ 0.18	<b>74.10</b> $\pm$ 0.22	70.76 $\pm$ 0.19	72.82 $\pm$ 0.26

**Table 1:** Mean test accuracy in % ( $\pm$  std.) over 5 runs on CIFAR10 and CIFAR100. For reference, the state-of-the-art on CIFAR10 (resp. CIFAR100) is [30] (resp. [31]) with an accuracy of 99.5% (resp. 96.1%).

	SpeechCommands			BirdDCASE		
	DenseNet121	ResNet18	VGG16	DenseNet121	ResNet18	VGG16
2D grid	97.21 $\pm$ 0.07	97.35 $\pm$ 0.04	97.18 $\pm$ 0.09	82.12 $\pm$ 0.29	82.36 $\pm$ 0.44	82.51 $\pm$ 0.22
Circle	97.22 $\pm$ 0.03	97.35 $\pm$ 0.11	97.14 $\pm$ 0.08	82.17 $\pm$ 0.36	<b>82.44</b> $\pm$ 0.28	<b>82.80</b> $\pm$ 0.49
2D nested	<b>97.27</b> $\pm$ 0.07	<b>97.38</b> $\pm$ 0.09	97.17 $\pm$ 0.03	82.48 $\pm$ 0.39	82.30 $\pm$ 0.37	82.68 $\pm$ 0.34
3D grid	<b>97.27</b> $\pm$ 0.11	97.33 $\pm$ 0.15	<b>97.28</b> $\pm$ 0.06	<b>82.58</b> $\pm$ 0.08	82.27 $\pm$ 0.43	82.50 $\pm$ 0.44
Sphere	97.19 $\pm$ 0.13	97.34 $\pm$ 0.03	97.21 $\pm$ 0.06	82.15 $\pm$ 0.49	82.40 $\pm$ 0.22	82.62 $\pm$ 0.39
3D nested	97.19 $\pm$ 0.10	97.33 $\pm$ 0.07	97.17 $\pm$ 0.08	82.27 $\pm$ 0.36	82.28 $\pm$ 0.16	82.67 $\pm$ 0.49
Base	97.20 $\pm$ 0.08	97.36 $\pm$ 0.07	97.19 $\pm$ 0.05	82.46 $\pm$ 0.31	<b>82.44</b> $\pm$ 0.21	82.49 $\pm$ 0.60

**Table 2:** Mean test accuracy in % ( $\pm$  std.) over 5 runs on SpeechCommands and BirdDCASE.



**Fig. 3:** Mean test accuracy ( $\pm$  std.) over 5 runs with  $L^2$ -pruning of a ResNet-18 trained on CIFAR10 with and without topography (2D grid,  $\lambda = 1$ ).

than the baseline. When pruning 62% of channels of each topographic layer, the topographic model has 62.0% test accuracy while the baseline drops to 40.6%. We observed similar results for the other architectures. This indicates that our topographic methods can obtain models with a better memory footprint for the same accuracy.

We did not find a better topographic organization that yields the best performance for audio or vision. This represents a potential limitation of our work as no topography stood out and is left as a modeling choice. This could be

alleviated in future work if topography could be learned or selected during training. From a neuroscience perspective, topographic organization in the brain had been proposed as the solution to the "wiring length" minimization problem [34], mainly for biological constraints and not performance. However, our pruning results potentially suggest that topography in CNNs could also be a solution to find more efficient pathways.

## 5. CONCLUSION

In this work, we introduce topography in CNNs by assigning positions to the output channels of convolutional layers and by enforcing, via an auxiliary loss, that nearby channels have similar activations. We show that, although the topographic loss is not related to the classification itself, topographic models perform similarly to the baselines. With  $L^2$  weight pruning, we evidence that topography has led to a finer selection of useful channels.

In future work, instead of using predefined position schemes, we plan to also learn the topographic organization itself. Such approach could yield scientific hypotheses for a given sensory task concerning the topography itself. The learned topographies through our machine learning method could be put in parallel to the different brain regions.

## 6. REFERENCES

- [1] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep Residual Learning for Image Recognition,” in *CVPR*, 2016. 1, 3
- [2] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger, “Densely Connected Convolutional Networks,” in *CVPR*, 2017. 1, 3
- [3] Qiuqiang Kong, Yin Cao, Turab Iqbal, Yuxuan Wang, Wenwu Wang, and Mark D. Plumbley, “Panns: Large-scale pre-trained audio neural networks for audio pattern recognition,” *IEEE/ACM TASLP*, 2020. 1
- [4] Tom Sercu, Christian Puhersch, Brian Kingsbury, and Yann Lecun, “Very deep multilingual convolutional neural networks for lvcstr,” in *ICASSP*, 2016. 1
- [5] Demis Hassabis, Dharshan Kumaran, Christopher Summerfield, and Matthew Botvinick, “Neuroscience-inspired artificial intelligence,” *Neuron*, 2017. 1
- [6] Matthew D Zeiler and Rob Fergus, “Visualizing and understanding convolutional networks,” in *ECCV*, 2014. 1
- [7] Daniel J Felleman and David C Van Essen, “Distributed hierarchical processing in the primate cerebral cortex,” *Cerebral cortex*, 1991. 1
- [8] Benjamin W. Zeff, Brian R. White, Hamid Dehghani, Bradley L. Schlaggar, and Joseph P. Culver, “Retinotopic mapping of adult human visual cortex with high-density diffuse optical tomography,” *PNAS*, 2007. 1
- [9] Tomas Knäpen, “Topographic connectivity reveals task-dependent retinotopic processing throughout the human brain,” *PNAS*, 2021. 1
- [10] Brian Barton, Jonathan H. Venezia, Kourosh Saberi, Gregory Hickok, and Alyssa A. Brewer, “Orthogonal acoustic dimensions define auditory field maps in human cortex,” *PNAS*, 2012. 1
- [11] Patrick W. Hullett, Liberty S. Hamilton, Nima Mesgarani, Christoph E. Schreiner, and Edward F. Chang, “Human Superior Temporal Gyrus Organization of Spectrotemporal Modulation Tuning Derived from Speech Stimuli,” *The Journal of Neuroscience*, 2016. 1
- [12] Bahar Khalighinejad, Prachi Patel, Jose L. Herrero, Stephan Bickel, Ashesh D. Mehta, and Nima Mesgarani, “Functional characterization of human Heschl’s gyrus in response to natural speech,” *NeuroImage*, 2021. 1
- [13] Linda Henriksson, Marieke Mur, and Nikolaus Kriegeskorte, “Faciotopy—a face-feature map with face-like topology in the human occipital face area,” *Cortex*, 2015. 1
- [14] Max Welling, Simon Osindero, and Geoffrey E Hinton, “Learning sparse topographic representations with products of student-t distributions,” in *NeurIPS*, 2002. 1
- [15] T. Anderson Keller and Max Welling, “Topographic VAEs learn equivariant capsules,” in *NeurIPS*, 2021. 1
- [16] Hyodong Lee, Eshed Margalit, Kamila M. Jozwik, Michael A. Cohen, Nancy Kanwisher, Daniel L. K. Yamins, and James J. DiCarlo, “Topographic deep artificial neural networks reproduce the hallmarks of the primate inferior temporal cortex face processing network,” *bioRxiv*, 2020. 1, 2
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, “Imagenet classification with deep convolutional neural networks,” in *NeurIPS*, 2012. 1
- [18] Kenji Kawaguchi, Leslie Pack Kaelbling, and Yoshua Bengio, “Generalization in deep learning,” *arXiv*, 2017. 1
- [19] Peter R Huttenlocher, “Synaptic density in human frontal cortex—developmental changes and effects of aging,” *Brain Res*, 1979. 2
- [20] Song Han, Jeff Pool, John Tran, and William Dally, “Learning both weights and connections for efficient neural network,” in *NeurIPS*, 2015. 2, 3
- [21] Christine E Collins, David C Airey, Nicole A Young, Duncan B Leitch, and Jon H Kaas, “Neuron densities vary across and within cortical areas in primates,” *PNAS*, 2010. 3
- [22] Vincent Lostanlen and Carmine-Emanuele Cella, “Deep convolutional networks on the pitch spiral for musical instrument recognition,” *ISMIR*, 2016. 3
- [23] Vincent Lostanlen, Sripathi Sridhar, Brian McFee, Andrew Farnsworth, and Juan Pablo Bello, “Learning the Helix Topology of Musical Pitch,” in *ICASSP*, 2020. 3
- [24] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *ICLR*, 2015. 3
- [25] Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li, “Bag of tricks for image classification with convolutional neural networks,” in *CVPR*, 2019. 3
- [26] Shuying Liu and Weihong Deng, “Very deep convolutional neural network based image classification using small training sample size,” in *ACPR*, 2015. 3
- [27] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger, “Deep networks with stochastic depth,” in *ECCV*, 2016. 3
- [28] Rachid Riad, Olivier Teboul, David Grangier, and Neil Zeghidour, “Learning strides in convolutional neural networks,” in *ICLR*, 2022. 3
- [29] Alex Krizhevsky et al., “Learning multiple layers of features from tiny images,” 2009. 3
- [30] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby, “An image is worth 16x16 words: Transformers for image recognition at scale,” in *ICLR*, 2021. 4
- [31] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur, “Sharpness-aware minimization for efficiently improving generalization,” in *ICLR*, 2021. 4
- [32] D. Stowell, Y. Stylianou, M. Wood, H. Pamula, and H. Glotin, “Automatic acoustic detection of birds through deep learning: the first bird audio detection challenge,” *Methods in Ecology and Evolution*, 2018. 3
- [33] Pete Warden, “Speech Commands: A Dataset for Limited-Vocabulary Speech Recognition,” *arXiv*, 2018. 3
- [34] Alexei A Koulakov and Dmitri B Chklovskii, “Orientation preference patterns in mammalian visual cortex: a wire length minimization approach,” *Neuron*, 2001. 4