



Fast threshold optimization for multi-label audio tagging using Surrogate gradient learning

Thomas Pellegrini, Timothée Masquelier

► To cite this version:

Thomas Pellegrini, Timothée Masquelier. Fast threshold optimization for multi-label audio tagging using Surrogate gradient learning. IEEE International Conference on Acoustics, Speech and Signal Processing, Jun 2021, Toronto, Canada. pp. 651-655, 10.1109/ICASSP39728.2021.9414091. hal-03153644

HAL Id: hal-03153644

<https://hal.science/hal-03153644>

Submitted on 26 Feb 2021

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.

FAST THRESHOLD OPTIMIZATION FOR MULTI-LABEL AUDIO TAGGING USING SURROGATE GRADIENT LEARNING

Thomas Pellegrini¹, Timothée Masquelier²

¹IRIT, Université Paul Sabatier, CNRS, Toulouse, France

²CERCO UMR 5549, CNRS – Université Toulouse 3, Toulouse, France

ABSTRACT

Multi-label audio tagging consists of assigning sets of tags to audio recordings. At inference time, thresholds are applied on the confidence scores outputted by a probabilistic classifier, in order to decide which classes are detected active. In this work, we consider having at disposal a trained classifier and we seek to automatically optimize the decision thresholds according to a performance metric of interest, in our case F-measure (micro-F1). We propose a new method, called SGLThresh for Surrogate Gradient Learning of Thresholds, that makes use of gradient descent. Since F1 is not differentiable, we propose to approximate the thresholding operation gradients with the gradients of a sigmoid function. We report experiments on three datasets, using state-of-the-art pre-trained deep neural networks. In all cases, SGLThresh outperformed three other approaches: a default threshold value (defThresh), an heuristic search algorithm and a method estimating F1 gradients numerically. It reached 54.9% F1 on AudioSet eval, compared to 50.7% with defThresh. SGLThresh is very fast and scalable to a large number of tags¹.

Index Terms— Audio tagging, Surrogate Gradient Learning, Automatic threshold optimization.

1. INTRODUCTION

In this work, we are interested in sound event detection (SED), more precisely in audio tagging (AT). AT consists in automatically assigning sets of binary labels, called tags, to real-life recordings. Audio events are often polyphonic, which means that several types of events may occur simultaneously, *e.g.*, cat meowing and speech. Thus, AT is a multi-label classification (MLC) task. AT and SED in general — SED consists in localizing the onset and offset of events — have gained a lot of interest within the audio content analysis community this last decade, because of deep neural

networks (DNNs) and the availability of large-scale datasets such as Audioset [1]. Stimulating challenges have also been organized, *e.g.*, the Detection and Classification of Acoustic Scenes and Events (DCASE) challenges², in which AT is key in several tasks.

State-of-the-art AT systems rely on DNNs, typically convolutional neural networks (CNN) and convolutional recurrent neural networks [2, 3]. Some of these models are open-source and available to the research community, both their architectures and their trained weights [3]. The models are trained by minimizing a loss function suitable to MLC, typically binary cross-entropy (BCE, [4]), whereby the predictions of the output neurons are independent. At inference time, a model outputs confidence scores for each class to be active, and thresholds are used to derive final binary predictions. One of the performance metrics most used is F-score, also called F1, which summarizes precision and recall values in a single number. In previous work, we compared parametric and non-parametric algorithms to search for optimal threshold values according to F1 [5]. In [6], the authors proposed to directly optimize F1 using gradient descent by computing F1 gradients numerically. These methods do bring consistent gains in F1 compared to non-optimized thresholds, but they have a strong limitation in terms of scalability to a large number of target classes.

In this work, we propose a gradient descent based method that is fast and scalable thanks to the use of automatic differentiation tools. F1 is not differentiable because the inherent operation of thresholding is not. To circumvent this issue, instead of computing the gradients numerically as in [6], we propose to use Surrogate Gradient Learning (SGL, [7]), that gives the name of our proposed approach *SGLThresh*. To our knowledge, this method has never been used to optimize thresholds in multi-label classification, and for AT in particular. We report F1 results on three datasets, including AudioSet to assess scalability since AudioSet involves 527 target classes and about 20k examples.

¹To facilitate reproducibility, data and source code in Pytorch are available online: <https://github.com/topel/SGLThresh>

²<http://dcase.community/>

2. PROBLEM STATEMENT

In this work, we aim at optimizing F1, a common metric to characterize performance with a single number. F1 is defined as the harmonic mean between precision and recall, and the larger, the better. More precisely, we use the micro-averaged label F1 score, estimated at label level. The ground-truth and predictions are flattened out before computing micro-F1, as if there were a single class in the task at hand:

$$\text{label micro-F1} = \frac{2 \sum_{n,l} y_l^{(n)} \hat{y}_l^{(n)}}{\sum_{n,l} y_l^{(n)} + \sum_{n,l} \hat{y}_l^{(n)}}$$

where $y_l^{(n)}$ and $\hat{y}_l^{(n)}$ are binary values corresponding to the ground-truth and predictions for instance n and label l . We chose to optimize micro-F1 because it has been used in Task 4 (SED task) of DCASE 2017 and 2019, two datasets on which we report results in this paper.

The best threshold values depend on the target class of interest. We will illustrate this in the result section with a few examples. The central question of this work is how we could optimize the threshold values in order to obtain optimal F1 score values.

3. BASELINE THRESHOLD ESTIMATION METHODS

We used two baseline methods for comparison purposes: `defThresh`, standing for “default threshold value”, and `dichoThresh`, a dichotomic search algorithm.

3.1. `defThresh`

This method is the very baseline one, using a single default threshold value for all the classes, without any optimization procedure involved. We simply choose a value that gives the best compromise for all the classes.

3.2. `dichoThresh`: dichotomic search

We proposed this method in previous work [5]. This method optimizes class-dependent threshold values through a dichotomic search algorithm. The algorithm searches the best threshold values first within a coarse interval and coarse precision. It picks the ones that yield the best F1 score. From this first guess, small threshold decay values are randomly sampled from a Gaussian distribution centered on the preceding best threshold values. The complete process is repeated with an increased precision and a reduced search space. It stops when a maximum number of steps is reached.

The dichotomic search algorithm, when compared to an exhaustive search of all the possible combinations, considerably reduces the time needed to reach efficient thresholds. However, the execution time is still dependent on the number of thresholds to tune and the amount of iterations for every step.

4. GRADIENT DESCENT BASED METHODS

In this section, we describe the two gradient-descent based methods used in this work: `numThresh` [6] and `SGLThresh` (ours). `NumThresh` and `SGLThresh` share the same threshold optimization algorithm. Similarly to `dichoThresh`, we optimize class-dependent threshold values on a validation set (valid) and use them on an evaluation subset (eval) afterwards.

The thresholds are first initialized to a default value. Then, for a given number of epochs, the following steps are repeated: the network predictions (scores) are thresholded with the current threshold values with the Heaviside step function. The F1 objective function and its gradients with respect to the thresholds are calculated. Then, the threshold values are updated with some optimization rule. In our experiments we used Adam [8]. At each epoch, the whole validation subset is used once to obtain robust gradient estimations (*batch training*). Using minibatches led to worse performance in our experiments.

The Heaviside step function is not differentiable at the origin and has a zero-valued gradient elsewhere. `NumThresh` and `SGLThresh` differ in the way the gradient is approximated, as we shall describe below.

4.1. `numThresh`: estimating F1 gradients numerically

To handle F1 optimization in audio tagging, Kong and colleagues [9] approximate the gradients numerically:

$$\nabla_t \text{F1}(t) = \frac{\text{F1}(t + \Delta t) - \text{F1}(t)}{\Delta t} \quad (1)$$

where t is a threshold, Δt a small number (0.01 in practice). This computation is done for all the thresholds, which are then updated with the Adam rule [8]. The implementation provided by [6] is very slow for large numbers of classes and examples, since all the directions are tested to maximize F1: for each threshold, a change in global F1 is tested if we increment t by Δt up to $10\Delta t$. Complexity is $\mathcal{O}(nC^2)$, where C is the number of classes and n is the number of examples.

4.2. `SGLThresh`: Surrogate Gradient Learning (ours)

SGL is a method to reconcile step functions, *e.g.* the sign function and the Heaviside step function (HSF),

with backpropagation. It has been proposed to train “binarized neural networks”, which are deep neural networks with weights and activations constrained to +1 or -1 [10], and spiking neural networks, whose firing threshold is a HSF [7, 11]. In SGLThresh, we use SGL to estimate the gradient $\nabla_t \text{F1}$ with the help of backpropagation and automatic differentiation.

The idea behind of SGL is very simple: while the forward pass uses the HSF, the backward pass uses a “surrogate gradient”, *i.e.* the derivative of a differentiable function, a sigmoid function in our work:

$$\Theta'(x) \approx \text{sig}'_a(x) = a \text{sig}_a(x) \text{sig}_a(-x) \quad (2)$$

where $x = p - t$, p being the network score predictions, and t the class-dependent thresholds to be optimized. This approximation is used when computing the F1 gradients with respect to the thresholds:

$$\frac{\partial \text{F1}(t)}{\partial t} = \frac{\partial \text{F1}}{\partial \hat{y}} \frac{\partial \hat{y}}{\partial t} \approx \frac{\partial \text{F1}}{\partial \hat{y}} \text{sig}'_a(p - t) \quad (3)$$

In practice, SGLThresh can be seen and is implemented as an activation function. In its forward mode, SGLThresh is the HSF, and in its backward mode the partial derivatives of a sigmoid function. Our implementation in PyTorch is very fast since it benefits from PyTorch’s vectorization and automatic differentiation capabilities to compute $\frac{\partial \text{F1}}{\partial \hat{y}}$.

5. AUDIO TAGGING EXPERIMENTS

In this section, we report audio tagging experiments on three datasets: DCASE 2017 [12] and DCASE 2019 [13] task 4 and AudioSet eval [1]. DCASE refers to the Detection and Classification of Acoustic Scenes and Events yearly challenges, and its task 4 aims at detecting sound events in domestic environments. Regarding AudioSet, we used the official eval balanced subset.

5.1. Datasets

Table 1 shows the dataset characteristics: the number of target tags (classes), the number of samples used to optimize the thresholds (the *val* subset), and the number of samples of the *eval* subset, on which we apply the optimized thresholds for the final evaluation of the methods. We did not use the subsets on which the neural networks were trained. We hypothesize that optimizing thresholds on samples held out from the model training subset leads to more robust thresholds.

DCASE2017 and DCASE2019 are similar in terms of number of target classes: 17 and 10, respectively. AudioSet has many more classes (527 different tags), which is particularly interesting to test the scalability

of the threshold optimizing methods regarding this aspect. For DCASE2017 and DCASE2019, we used the official val/eval splits, whereas for AudioSet, there is only an eval subset available. For that reason, in the case of AudioSet only, we ran a 3-fold cross-validation.

Table 1. Dataset characteristics: number of target tags (classes), number of samples in val and eval subsets.

name	classes	val	eval
DCASE2017	17	488	1103
DCASE2019	10	1122	692
AudioSet eval	527	15278	5093

5.2. Pre-trained neural networks

Different state-of-the-art pretrained networks have been used for each of the three datasets. The acoustic features in all cases are log-Mel spectrograms, with slight differences in their extraction process though.

For DCASE2017, we used the best performing model open-sourced from [6], called “CNN-transformer-avg” by the authors. It is comprised of four convolution blocks, each with two convolutional layers with 3×3 kernel size, each layer followed by batch-normalization, Rectifier Linear Unit (ReLU) and 2×2 average pooling. The number of filters of the four blocks increases by a factor two: 64 filters for the first block, 128, 256 and 512 filters for the following blocks. A self-supervision module is placed after the last convolution block, hence the name invoking transformers [14]. The total number of learnable parameters is 5.75 million parameters.

For DCASE2019, we used the convolutional recurrent neural network of our own best submission to the challenge [15], called “model 1” in [15], with which we ranked fourth out of nineteen teams³. It is a small network of 165 thousand parameters, comprised of three convolution blocks followed by a bi-directional Gated Recurrent Unit layer (64 cells), a fully-connected layer of 10 units for localization and an audio tagging head with a 1024-unit fully-connected layer with ReLU and a 10-unit output layer with sigmoid activation functions.

Finally, for the AudioSet experiments, we used the predictions made by the open-sourced model, with its available pretrained weights on AudioSet, called “Cnn14_mAP=0.431” [3]. This model is comprised of six convolution blocks, each followed by a 20% dropout layer, followed by the sum of a max- and a mean-pooling, a 2048-unit fully-connected layer with ReLU,

³<http://dcase.community/challenge2019/task-sound-event-detection-in-domestic-environments-results>

Table 2. Micro-averaged F1 (%) results in terms of on the evaluation subsets of DCASE2017, DCASE2019, AudioSet, between parentheses: on the valid subsets on which F1 was optimized. Optimization time is given for AudioSet.

	DCASE2017	DCASE2019	AudioSet (3-fold CV)	Optimization time on AudioSet
defThresh	59.0 (53.8)	71.3 (70.1)	50.7 (50.7)	-
dichoThresh	61.4 (58.1)	71.4 (71.7)	53.5 (53.6)	9 hours (CPU)
numThresh	60.3 (58.7)	72.2 (71.4)	52.1 (53.6)	5 days (CPU)
SGLThresh (ours)	62.0 (61.8)	73.0 (72.9)	54.9 (56.7)	1.5 s (GPU), 3 s (CPU)

and a 527-unit output layer with sigmoid activations. It totals 80.8 Million parameters.

5.3. Results

NumThresh thresholds were optimized with a learning rate of $1e-2$, with the Adam update rule, for 100 epochs, following [6]. We used the same setting for SGLThresh, but with a $1e-3$ learning rate. The initial thresholds were set to 0.3 for DCASE2017 and AudioSet, and 0.5 for DCASE2019. In SGLThresh, we set the sigmoid slope a to 50 in all the experiments. A large range of values for a , between 20 and 100, led to the same results. We ran the experiments on a machine with an Intel Xeon CPU E5-2623 and a GeForce GTX 1080 GPU.

In principle, the three methods should converge to the same threshold values, the ones that give the best F1 value on the valid subsets. Nevertheless, in practice, it is not the case, probably due to convergence issues. Hyperparameters, such as the threshold decay for dichoThresh and the learning rate for numThresh and SGLThresh, may not be optimal for each target class individually. Table 2 compiles all the results in terms of micro-F1 values on the evaluation subset and within parentheses on the valid subset, on which the thresholds were optimized. The time needed to converge on AudioSet is given in the last column. As a first comment, we observe that our method, SGLThresh, is by far the fastest method: only a few seconds to reach convergence either on GPU or CPU. DichoThresh needed about 9 hours and numThresh several days to converge, showing that the available implementation of numThresh is very inefficient when dealing with a large number of target classes and training examples. Globally speaking, the three optimization methods brought significant gains compared to using a default threshold (defThresh), and our proposed method, SGLThresh, led to the best results on the eval subsets. On AudioSet in particular, the SGLThresh gains are the largest ones, with 56.7% and 54.9% F1 values on the valid and eval subsets respectively. We ran 3-fold cross-validation on AudioSet, the F1 values are averaged on the three folds.

Standard deviation was about 0.2% in all cases.

DCASE2017 is the only corpus for which micro-F1 measures are available for reference in the literature. In 2017, the system best ranked in the challenge reached 55.6%. Our model, which architecture is taken from [6] and trained by ourselves, reached 59.0% with defThresh and 60.3% with numThresh, to be compared with 62.9% and 64.6% respectively, as reported by the authors. Although we did not reach these performance when training their model, we observe a similar relative gain in F1 brought by numThresh: 2.2% relative.

We reported results on micro-F1, but we also ran experiments optimizing macro-F1 scores, which we do not report for the sake of space. The results are similar to those of micro-F1: SGLThresh outperformed the other methods in macro-F1 scores: 62.5% compared to 53.4% with defThresh on DCASE2017, for instance.

6. CONCLUSIONS

We proposed a new method to optimize F1 score, by searching for threshold values using surrogate gradient learning. We ran polyphonic audio tagging experiments on three benchmark datasets, and our method, called SGLThresh, outperformed the use of a default threshold value by a large margin, but also two other optimization algorithms. Besides being accurate, SGLThresh is also very fast thanks to automatic differentiation and vectorization. It scaled well on AudioSet, unlike the two other tested methods, and reached 54.9% micro-F1 value on the eval subset. Future work will consist of end-to-end training models to optimize F1 instead of cross-entropy.

7. ACKNOWLEDGEMENTS

This work was partially supported by the French ANR agency within the LUDAU project (ANR-18-CE23-0005-01) and the French "Investing for the Future — PIA3" AI Interdisciplinary Institute ANITI (Grant agreement ANR-19-PI3A-0004). We used HPC resources from CALMIP (Grant 2020-p20022).

8. REFERENCES

- [1] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *Proc. ICASSP*, New Orleans, 2017, IEEE, pp. 776–780.
- [2] Tuomas Virtanen, Mark D Plumbley, and Dan Ellis, *Computational analysis of sound scenes and events*, Springer, 2018.
- [3] Qiuqiang Kong, Yin Cao, Turab Iqbal, Yuxuan Wang, Wenwu Wang, and Mark D Plumbley, "PANNS: Large-scale pretrained audio neural networks for audio pattern recognition," *arXiv preprint arXiv:1912.10211*, 2019.
- [4] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep learning*, MIT press, 2016.
- [5] Léo Cances, Patrice Guyot, and Thomas Pellegrini, "Evaluation of post-processing algorithms for polyphonic sound event detection," in *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2019, pp. 318–322.
- [6] Qiuqiang Kong, Yong Xu, Wenwu Wang, and Mark D Plumbley, "Sound event detection of weakly labelled data with cnn-transformer and automatic threshold optimization," *arXiv preprint arXiv:1912.04761*, 2019.
- [7] Emre O. Neftci, Hesham Mostafa, and Friedemann Zenke, "Surrogate Gradient Learning in Spiking Neural Networks," *IEEE Signal Processing Magazine*, vol. 36, no. October, pp. 51–63, 2019.
- [8] Diederik P Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [9] Qiuqiang Kong, Changsong Yu, Yong Xu, Turab Iqbal, Wenwu Wang, and Mark D Plumbley, "Weakly labelled audioset tagging with attention neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 11, pp. 1791–1802, 2019.
- [10] Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio, "Binarized Neural Networks: Training Deep Neural Networks with Weights and Activations Constrained to +1 or -1," *arXiv preprint arXiv:1602.02830*, 2016.
- [11] Romain Zimmer, Thomas Pellegrini, Srisht Fateh Singh, and Timothée Masquelier, "Technical report: supervised training of convolutional spiking neural networks with pytorch," *arXiv preprint arXiv:1911.10124*, 2019.
- [12] A. Mesaros, T. Heittola, A. Diment, B. Elizalde, A. Shah, E. Vincent, B. Raj, and T. Virtanen, "DCASE 2017 challenge setup: Tasks, datasets and baseline system," in *Proc. of the Detection and Classification of Acoustic Scenes and Events 2017 Workshop (DCASE2017)*, November 2017, pp. 85–92.
- [13] Romain Serizel, Nicolas Turpault, Hamid Eghbal-Zadeh, and Ankit Parag Shah, "Large-scale weakly labeled semi-supervised sound event detection in domestic environments," in *Proc. of the Detection and Classification of Acoustic Scenes and Events 2018 Workshop (DCASE2018)*, November 2018, pp. 19–23.
- [14] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [15] Léo Cances, Thomas Pellegrini, and Patrice Guyot, "Multi-task learning and post processing optimization for sound event detection," Tech. Rep., IRT, Université de Toulouse, CNRS, Toulouse, France, June 2019.