

Parameter Selection: Why We Should Pay More Attention to It

Jie-Jyun Liu¹, Tsung-Han Yang¹, Si-An Chen^{1,2}, and Chih-Jen Lin²

¹ASUS Intelligent Cloud Services

²National Taiwan University

{eleven1_liu, henry1_yang}@asus.com

{d09922007, cjlin}@csie.ntu.edu.tw

Abstract

The importance of parameter selection in supervised learning is well known. However, due to the many parameter combinations, an incomplete or an insufficient procedure is often applied. This situation may cause misleading or confusing conclusions. In this opinion paper, through an intriguing example we point out that the seriousness goes beyond what is generally recognized. In the topic of multi-label classification for medical code prediction, one influential paper conducted a proper parameter selection on a set, but when moving to a subset of frequently occurring labels, the authors used the same parameters without a separate tuning. The set of frequent labels became a popular benchmark in subsequent studies, which kept pushing the state of the art. However, we discovered that most of the results in these studies cannot surpass the approach in the original paper if a parameter tuning had been conducted at the time. Thus it is unclear how much progress the subsequent developments have actually brought. The lesson clearly indicates that without enough attention on parameter selection, the research progress in our field can be uncertain or even illusive.

1 Introduction

The importance of parameter selection in supervised learning is well known. While parameter tuning has been a common practice in machine learning and natural language processing applications, the process remains challenging due to the huge number of parameter combinations. The recent trend of applying complicated neural networks makes the situation more acute. In many situations, an incomplete or an insufficient procedure for parameter selection is applied, so misleading or confusing conclusions sometimes occur. In this opinion paper, we present a very intriguing example showing that, without enough attention on parameter selection, the research progress in our field can be uncertain or even illusive.

In the topic of multi-label classification for medical code prediction, Mullenbach et al. (2018) is an early work applying deep learning. The evaluation was conducted on MIMIC-III and MIMIC-II (Johnson et al., 2016), which may be the most widely used open medical records. For MIMIC-III, besides using all 8,922 labels, they follow Shi et al. (2017) to check the 50 most frequently occurring labels. We refer to these two sets respectively as

MIMIC-III-full and MIMIC-III-50.

We will specifically investigate MIMIC-III-50. Based on Mullenbach et al. (2018), many subsequent works made improvements to push the state of the art. Examples include (Wang et al., 2018; Sadoughi et al., 2018; Xie et al., 2019; Tsai et al., 2019; Cao et al., 2020a,b; Ji et al., 2020; Teng et al., 2020; Chen, 2020; Vu et al., 2020; Dong et al., 2021).

For the data set MIMIC-III-full, Mullenbach et al. (2018) tuned parameters to find the model that achieves the best validation performance. However, when moving to check the set MIMIC-III-50, they *applied the same parameters without a separate tuning*. We will show that this decision had a profound effect. Many works directly copied values from Mullenbach et al. (2018) for comparison and presented superior results. However, as demonstrated in this paper, if parameters for MIMIC-III-50 had been separately tuned, the approach in Mullenbach et al. (2018) easily surpasses most subsequent developments. The results fully indicate that parameter selection is more important than what is generally recognized.

This paper is organized as follows. In Section 2, we analyze past results. The main investigation is in Section 3, while Section 4 provides some discussion. Some implementation details are in the appendix. Code and supplementary materials can be found at http://www.csie.ntu.edu.tw/~cjlin/papers/parameter_selection.

Table 1: Experimental results from Mullenbach et al. (2018). Macro-F1 and Micro-F1 are Macro-averaged and Micro-averaged F1 values, respectively. P@ n is the precision at n , the fraction of the n highest-scored labels that are truly associated with the test instance.

(a) MIMIC-III-full: 8,922 labels			
	Macro-F1	Micro-F1	P@8
CNN	0.042	0.419	0.581
CAML	0.088	0.539	0.709
(b) MIMIC-III-50: 50 labels.			
	Macro-F1	Micro-F1	P@5
CNN	0.576	0.625	0.620
CAML	0.532	0.614	0.609

2 Analysis of Works that Compared with Mullenbach et al. (2018)

The task considered in Mullenbach et al. (2018) is to predict the associated ICD (International Classification of Diseases) codes of each medical document. Here an ICD code is referred to as a label. The neural network considered is

$$\begin{aligned} &\text{document} \rightarrow \text{word embeddings} \\ &\rightarrow \text{convolution} \rightarrow \text{attention} \rightarrow \text{linear layer}, \end{aligned} \quad (1)$$

where the convolutional operation was based on Kim (2014). A focus in Mullenbach et al. (2018) was on the use of attention, so they detailedly compared the two settings¹

$$\begin{aligned} \text{CNN:} & \quad (1) \text{ without attention,} \\ \text{CAML:} & \quad (1). \end{aligned}$$

For the data set MIMIC-III-full, CAML, which includes an attention layer, was shown to be significantly better than CNN on all criteria; see Table 1a. However, for MIMIC-III-50, the subset of the 50 most frequent labels, the authors reported in Table 1b that CAML is not better than CNN.

The paper (Mullenbach et al., 2018) has been highly influential. By exactly using their training, validation, and test sets for experiments, many subsequent studies have proposed new and better approaches; see references listed in Section 1. Most of them copied the CNN and CAML results from (Mullenbach et al., 2018) as the baseline for comparison. Table 2 summarizes their superior results on MIMIC-III-50.²

¹After convolution, each word is still associated with a short vector and attention is a way to obtain a single vector for the whole document. For CNN where attention is not used, Mullenbach et al. (2018) followed Kim (2014) to select the

While using the same MIMIC-III-50 set, these subsequent studies differ from Mullenbach et al. (2018) in various ways. They proposed sophisticated networks and may incorporate additional information (e.g., label description, knowledge graph of words, etc.). Further, they may change settings not considered as parameters for tuning in Mullenbach et al. (2018). For example, Mullenbach et al. (2018) truncated each document to have at most 2,500 tokens, but Vu et al. (2020) used 4,000.

3 Investigation

We investigate the performance of the CNN and CAML approaches in Mullenbach et al. (2018) for the set MIMIC-III-50. Some implementation details are left in supplementary materials.

3.1 Parameter Selection in Mullenbach et al. (2018)

Mullenbach et al. (2018) conducted parameter tuning on a validation set of MIMIC-III-full. By considering parameter ranges shown in Table 3, they applied Bayesian optimization (Snoek et al., 2012) to choose parameters achieving the highest precision@8 on the validation set; see the selected values in Table 3 and the definition of precision in Table 1. However, the following settings are fixed instead of being treated as parameters for tuning.

- Each document is truncated to have at most 2,500 tokens. Word embeddings are from the CBOW method (Mikolov et al., 2013) with the embedding size 100.
- The stochastic gradient method Adam implemented in PyTorch is used with its default setting. However, the batch size is fixed to be 16 and the learning rate is considered as a parameter. Binary cross-entropy loss is considered.
- The Adam method is terminated if the precision@8 does not improve for 10 epochs. The model achieving the highest validation precision@8 is used to predict the test set for obtaining results in Table 1a.

Interestingly, for the 50-label subset of MIMIC-III, Mullenbach et al. (2018) did not conduct a parameter-selection procedure. Instead, a decision was to *use the same parameters selected for the*

maximal value across all words.

² We exclude papers that used the same MIMIC-III-50 set but did not list values in Mullenbach et al. (2018) for comparison. Anyway, results in these papers are not better than what we obtained in Section 3.

Table 2: MIMIC-III-50 results from past works that have *directly listed values* in Mullenbach et al. (2018) for comparison. For Macro-F1, please see a note on its definition in the appendix.

	Macro-F1	Micro-F1	P@5	Code	Notes
Baseline considered					
CNN (Mullenbach et al., 2018)	0.576	0.625	0.620	Y	
CAML (Mullenbach et al., 2018)	0.532	0.614	0.609	Y	
New network architectures					
MVC-LDA (Sadoughi et al., 2018)	0.597	0.668	0.644	N	multi-view convolutional layers
DACNM (Cao et al., 2020b)	0.579	0.641	0.616	N	dilated convolution
BERT-Large (Chen, 2020)	0.531	0.605	-	N	BERT model
MultiResCNN (Li and Yu, 2020)	0.606	0.670	0.641	Y	multi-filter convolution and residual convolution
DCAN (Ji et al., 2020)	0.615	0.671	0.642	Y	dilated convolution, residual connections
G-Coder without additional information (Teng et al., 2020)	-	0.670	0.637	N	multiple convolutional layers
LAAT (Vu et al., 2020)	0.666	0.715	0.675	Y	LSTM before attention
New network architectures + additional information (e.g., label description, label co-occurrence, label embeddings, knowledge graph, adversarial learning, etc.)					
LEAM (Wang et al., 2018)	0.540	0.619	0.612	Y	label embeddings used
MVC-RLDA (Sadoughi et al., 2018)	0.615	0.674	0.641	N	label description used
MSATT-KG (Xie et al., 2019)	0.638	0.684	0.644	N	knowledge graph
HyperCore (Cao et al., 2020a)	0.609	0.663	0.632	N	label co-occurrence and hierarchy used
G-Coder with additional information (Teng et al., 2020)	-	0.692	0.653	N	knowledge graph, adversarial learning
Results of our investigation in Section 3 are listed below for comparison (values averaged from Table 4)					
CNN	0.606	0.659	0.634	Y	parameter selection applied
CAML	0.635	0.684	0.651	Y	

Table 3: Parameter ranges considered in Mullenbach et al. (2018) and the values used.

Parameter	Range	Values used	
		CNN	CAML
d_c : # filters	50-500	500	50
k : filter size	2-10	4	10
q : dropout prob.	0.2-0.8	0.2	0.2
η : learning rate	0.0003, 0.0001, 0.003, 0.001	0.003	0.0001

full-label set. Further they switch to present precision@5 instead of precision@8 because on average each instance is now associated with fewer labels.

The decision of not separately tuning parameters for MIMIC-III-50, as we will see, has a profound effect. In fact, because in Table 1b CAML is slightly worse than CNN, Mullenbach et al. (2018) have suspected that a parameter tuning may be

needed. They stated that “we hypothesize that this³ is because the relatively large value of $k = 10$ for CAML leads to a larger network that is more suited to larger datasets; tuning CAML’s hyperparameters on this dataset would be expected to improve performance on all metrics.” However, it seems no subsequent works tried to tune parameters of CNN or CAML on MIMIC-III-50.

3.2 Reproducing Results in Mullenbach et al.

To ensure the correctness of our implementation, first we reproduce the results in Mullenbach et al. (2018) by considering the following two programs.

- The public code by Mullenbach et al. (2018) at github.com/jamesmullenbach/caml-mimic.
- Our implementation of CNN/CAML by following the description in Mullenbach et al. (2018). The code is part of our development

³Here “this” means that CAML is not better than CNN.

Table 4: MIMIC-III-50 results after parameter selection. We consider three random seeds, where 1,337 was used in Mullenbach et al. (2018). Under each seed, we select the five models achieving the best validation precision@5, use them to predict the test set, and report mean/variance.

	Seed	Macro-F1	Micro-F1	P@5
CNN	1337	0.608 ± 0.006	0.659 ± 0.005	0.634 ± 0.002
	1331	0.601 ± 0.013	0.660 ± 0.007	0.634 ± 0.003
	42	0.608 ± 0.007	0.658 ± 0.006	0.633 ± 0.003
CAML	1337	0.640 ± 0.004	0.686 ± 0.004	0.650 ± 0.002
	1331	0.631 ± 0.004	0.682 ± 0.003	0.651 ± 0.001
	42	0.634 ± 0.009	0.684 ± 0.004	0.651 ± 0.002

on a general multi-label text classification package `LibMultiLabel`.⁴

Parameters and the random seed used in Mullenbach et al. (2018) are considered; see Table 3.

After some tweaks, on one GPU machine both programs give *exactly the same results* in the following table

	Macro-F1	Micro-F1	P@5
CNN	0.585	0.626	0.617
CAML	0.532	0.610	0.609

Values are very close to those in Table 1b. The small difference might be due to that our GPUs or PyTorch versions are not the same as theirs.

We conclude that results in Mullenbach et al. (2018) are reproducible.

3.3 Parameter Selection for MIMIC-III-50

We apply the parameter-selection procedure in Mullenbach et al. (2018) for MIMIC-III-full to MIMIC-III-50; see details in Section 3.1. A difference is that, because training MIMIC-III-50 is faster than MIMIC-III-full, instead of using Bayesian optimization, we directly check a grid of parameters that are roughly within the ranges given in Table 3. Specifically, we consider

$$d_c = 50, 150, 250, 350, 450, 550$$

$$k = 2, 4, 6, 8, 10$$

$$q = 0.2, 0.4, 0.6, 0.8$$

Because Mullenbach et al. (2018) switched to report test precision@5 for MIMIC-III-50, for validation we also use precision@5.

To see the effect of random seeds, besides the one used in Mullenbach et al. (2018), we checked two other seeds 1,331 and 42, selected solely because they are the lucky numbers of an author.

⁴<https://github.com/ASUS-AICS/LibMultiLabel>

3.4 Results and Analysis

Table 4 shows CNN/CAML results after parameter selection and we have the following observations.

- Both CNN and CAML achieve better results than those reported in Table 1b by Mullenbach et al. (2018). The improvement of CAML is so significant that it becomes better than CNN.
- From details in supplementary materials, for some parameters (e.g., d_c and q for CAML), the selected values are very different from those used by Mullenbach et al. (2018). Thus parameters selected for MIMIC-III-full are not transferable to MIMIC-III-50 and a separate tuning is essential.
- Results are not sensitive to the random seeds.⁵
- A comparison with Table 2 shows that most subsequent developments cannot surpass our CAML results. Some are even inferior to CNN, which is the baseline of all these studies.
- We checked if subsequent developments conducted parameter selection. A summary is in the supplementary materials.

Based on our results, how much progress past works have made is therefore unclear.

4 Discussion and Conclusions

The intention of this paper is to provide constructive critiques of past works rather than place blame on their authors. For the many parameter combinations, it is extremely difficult to check them. However, what our investigation showed is that if resources or time are available, more attention should be paid to the parameter selection. For Mullenbach et al. (2018), as they have done a comprehensive selection on a super-set MIMIC-III-full, the same procedure on the simpler MIMIC-III-50 is

⁵CNN is slight more sensitive to seeds than CAML. More investigation is needed.

entirely feasible. The decision of not doing so leads to a weak baseline in the subsequent developments.

In conclusion, besides proposing new techniques such as sophisticated networks, more attention should be placed on the parameter selection. In the future this helps to ensure that strong baselines are utilized to check the progress.

References

- Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao, Shengping Liu, and Weifeng Chong. 2020a. [HyperCore: Hyperbolic and co-graph representation for automatic ICD coding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 3105–3114.
- Pengfei Cao, Chenwei Yan, Xiangling Fu, Yubo Chen, Kang Liu, Jun Zhao, Shengping Liu, and Weifeng Chong. 2020b. [Clinical-coder: Assigning interpretable ICD-10 codes to Chinese clinical notes](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 294–301.
- Yiyun Chen. 2020. [Predicting ICD-9 codes from medical notes—does the magic of BERT applies here?](#)
- Hang Dong, Vctor Suárez-Paniagua, William Whiteley, and Honghan Wu. 2021. [Explainable automated coding of clinical notes using hierarchical label-wise attention networks and label embedding initialisation](#). *Journal of Biomedical Informatics*, 116:103728.
- Shaoxiong Ji, Erik Cambria, and Pekka Marttinen. 2020. [Dilated convolutional attention network for medical code assignment from clinical text](#). In *Proceedings of the Third Clinical Natural Language Processing Workshop*, pages 73–78.
- Alistair E.W. Johnson, Tom J. Pollard, Lu Shen, Li-Wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. 2016. [MIMIC-III, a freely accessible critical care database](#). *Scientific Data*, 3(1):1–9.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751.
- Fei Li and Hong Yu. 2020. [ICD coding from clinical text using multi-filter residual convolutional neural network](#). In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*, pages 8180–8187.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems (NIPS)*, pages 3111–3119.
- James Mullenbach, Sarah Wiegrefe, Jon Duke, Jimeng Sun, and Jacob Eisenstein. 2018. [Explainable prediction of medical codes from clinical text](#). In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, pages 1101–1111.
- Juri Opitz and Sebastian Burst. 2021. [Macro F1 and Macro F1](#). ArXiv preprint arXiv:1911.03347.
- Najmeh Sadoughi, Greg P. Finley, James Fone, Vignesh Murali, Maxim Korenevski, Slava Baryshnikov, Nico Axtmann, Mark Miller, and David Suendermann-Oeft. 2018. [Medical code prediction with multi-view convolution and description-regularized label-dependent attention](#). ArXiv preprint arXiv:1811.01468.
- Haoran Shi, Pengtao Xie, Zhiting Hu, Ming Zhang, and Eric P. Xing. 2017. [Towards automated ICD coding using deep learning](#). ArXiv preprint arXiv:1711.04075.
- Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. [Practical Bayesian optimization of machine learning algorithms](#). In *Advances in Neural Information Processing Systems 25 (NIPS)*, pages 2951–2959.
- Fei Teng, Wei Yang, Li Chen, LuFei Huang, and Qiang Xu. 2020. [Explainable prediction of medical codes with knowledge graphs](#). *Frontiers in Bioengineering and Biotechnology*, 8:867.
- Shang-Chi Tsai, Ting-Yun Chang, and Yun-Nung Chen. 2019. [Leveraging hierarchical category knowledge for data-imbalanced multi-label diagnostic text understanding](#). In *Proceedings of the Tenth International Workshop on Health Text Mining and Information Analysis*, pages 39–43.
- Thanh Vu, Dat Quoc Nguyen, and Anthony Nguyen. 2020. [A label attention model for ICD coding from clinical text](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3335–3341.
- Guoyin Wang, Chunyuan Li, Wenlin Wang, Yizhe Zhang, Dinghan Shen, Xinyuan Zhang, Ricardo Henao, and Lawrence Carin. 2018. [Joint embedding of words and labels for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association of Computational Linguistics (ACL)*.
- Xiancheng Xie, Yun Xiong, Philip S. Yu, and Yangyong Zhu. 2019. [EHR coding with multi-scale feature attention and structured knowledge graph propagation](#). In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM)*, pages 649–658.

A Additional Implementation and Experimental Details

Before a stochastic gradient step on a batch of data, Mullenbach et al. (2018) pad sequences with zeros

so that all documents in this batch have the same number of tokens. Thus results of the forward operation depend on the batch size. This setting causes issues in validation because a result independent of the batch size is needed. Further, for many applications one instance appears at a time in the prediction stage. Thus we follow [Mullenbach et al. \(2018\)](#) to use

$$\text{batch size} = 1$$

in validation and prediction.

After the convolutional layer, [Mullenbach et al. \(2018\)](#) consider the tanh activation function. For both convolutional and linear layers, a bias term is included.

Before the training process, [Mullenbach et al. \(2018\)](#) sort the data according to their lengths. In the stochastic gradient procedure, data are not reshuffled. Therefore, instances considered in each batch are the same across epochs. While this setting is less used in other works, we follow suit to ensure the reproducibility of their results.

In the stochastic gradient procedure, we follow ([Mullenbach et al., 2018](#)) to set 200 as the maximal number of epochs. This setting is different from the default 100 epochs in the software `LibMultiLabel` employed for our experiments. In most situations, the program does not reach the maximal number of epochs. Instead, it terminates after the validation P@5 does not improve in 10 epochs. This criterion also follows from [Mullenbach et al. \(2018\)](#).

All models were trained on one NVIDIA Tesla P40 GPU compatible with the CUDA 10.2 platform and cuDNN 7.6. Note that results may slightly vary if experiments are run on different architectures.

B A Note on Macro-F1

[Mullenbach et al. \(2018\)](#) report macro-F1 defined as

F1 value of macro-precision and macro-recall, where macro-precision and macro-recall are respectively the mean of precision and recall over all classes. This definition is different from the macro-F1 used in most other works. Specifically, F1 values are obtained for each class first and their mean is considered as Macro-F1; see the discussion of the Macro-F1 definitions in [Opitz and Burst \(2021\)](#). Because works mentioned in [Table 2](#) may not indicate if they use the same Macro-F1 formula as [Mullenbach et al. \(2018\)](#), readers should exercise caution in interpreting Macro-F1 results in [Table 2](#).

However, based on Micro-F1 and P@5 results the main point of this paper still stands.