Mathematical Biosciences and Engineering

http://www.aimspress.com/journal/mbe

MBE, 20(2): 3624–3637.
DOI: 10.3934/mbe.2023169
Received: 22 August 2022
Revised: 28 November 2022
Accepted: 04 December 2022
Published: 08 December 2022

Letter

Fairness-aware genetic-algorithm-based few-shot classification

Depei Wang[1], **Lianglun Cheng**[2] **and Tao Wang**[1,*]

[1] School of Automation, Guangdong University of Technology, Guangzhou 510006, China

[2] School of Computer Science and Technology, Guangdong University of Technology, Guangzhou 510006, China

* **Correspondence:** Email: wangtao_cps@gdut.edu.cn.

Abstract: Artificial-intelligence-assisted decision-making is appearing increasingly more frequently in our daily lives; however, it has been shown that biased data can cause unfairness in decision-making. In light of this, computational techniques are needed to limit the inequities in algorithmic decision-making. In this letter, we present a framework to join fair feature selection and fair meta-learning to do few-shot classification, which contains three parts: (1) a pre-processing component acts as an intermediate bridge between fair genetic algorithm (FairGA) and fair few-shot (FairFS) to generate the feature pool; (2) the FairGA module considers the presence or absence of words as gene expression, and filters out key features by a fairness clustering genetic algorithm; (3) the FairFS part carries out the task of representation and fairness constraint classification. Meanwhile, we propose a combinatorial loss function to cope with fairness constraints and hard samples. Experiments show that the proposed method achieves strong competitive performance on three public benchmarks.

Keywords: fairness; few-shot; genetic algorithm; meta-learning; feature selection

1. Introduction

Machine-learning models are impacting most of the automatic tasks in our daily lives, and have achieved significant social impact. Data-driven models meet new challenges when faced with less domain labeled data. Furthermore, these models suffer from data bias due to data imbalance, leading to unequal treatment of individuals and small groups in decision-making. In actual processing, feature selection plays the role of data pre-processing for various learning tasks [**?**, **?**, **?**]. Regarding few-shot (FS) learning's success in multiple domains [**?**, **?**, **?**, **?**], the models are moving away from being reliant on large-scale data. The area of algorithm fairness [**?**, **?**, **?**] is precisely devoted to guaranteeing that sensitive attributes are treated equally.

In text classification tasks, common feature-selection methods can be divided into four categories:

Filter, wrapper, embedded, and hybrid [?]. Among these, the wrapper method utilizes some search strategy to mine a subset and sends the subset to the classifier to measure the classifier's performance. These two steps are repeated until the subset meets pre-defined criteria. The genetic algorithm (GA) is a popular swapper method used in feature-selection tasks. Ming Fan et al. [?] proposed ExpGA to detect discriminatory individuals, which is an explanation-guided fairness testing approach through GA. In addition, they used ConceptNet to complete semantic word level crossover and mutation on text dataset. Since fairness algorithms have been devised, some researchers have started to think about fairness in their models. Xing et al. [?] proposed fairness-aware unsupervised feature selection, which introduces fairness constraints into the feature-selection task. The method preserves the information in the original feature space as much as possible while minimizing the correlation with protection properties. Zhao et al. [?] introduced fairness into meta-learning, reducing the dependence of predictions on protected variables. In the process of fairness algorithm research, the definition of balance must be mentioned. The index used to measure fairness was first proposed by Chierichetti et al. [?] for the case with two protected groups. Kleindessner et al. [?] and Ziko et al. [?] used a balance penalty in clustering tasks. However, there is still room for improvement. Different from the above ideas, we propose a two-stage framework named FairGAFS under fair penalty to discover efficient feature subsets for few-shot learning. To protect group features in GA selection processing, a gene expression method using 1 or 0, and a fair constraint fitness function are proposed. In addition, we find that hard samples, which get big loss value in training processing, have a negative effect on the classifier training process. Thus, we use focal loss [?] to weaken the effect of hard samples on classification results. Furthermore, to avoid unfair categories from occurring in meta-learning tasks, we add a fairness constraint to the loss function.

The main contributions of this letter are the following.

- FairGA protects individual sensitive genes and avoids trait bias through fairness constraints.
- FairFS improves classifier performance by adding a fairness measure to a regular loss function.
- FairGAFS reduces the input dimension, solves the problem of hard samples, and improves classification accuracy through fairness constraints.
- Three open datasets are used to evaluate the effectiveness of FairGAFS in regulating two protection attributes. Results indicate that the strategy under consideration is capable of resolving the few-shot categorization problem.

2. Methods

The motivation for the proposed model is twofold. First, input texts usually contain numerous useless features, resulting in wasted computing power and biased classification results. Evolutionary algorithms tend to select excellent individuals during the evolution process while ignoring protected genes. Therefore, we seek a solution to retain important features while maintaining the fairness of the population during evolution. Furthermore, we propose the focal loss take harder mining situations into account.

Second, the hard samples and sensitive attributes may mislead classifiers during few-shot learning; thus, we designed a meta-learning framework (FairFS) in which fairness notions are considered as constraints added onto focal loss.

The overall structure of the proposed schema is depicted in Figure **??**, which includes three phrases:

preprocessing, FairGA, and FairFS. The preprocessing component is the bridge between FairGAFS.

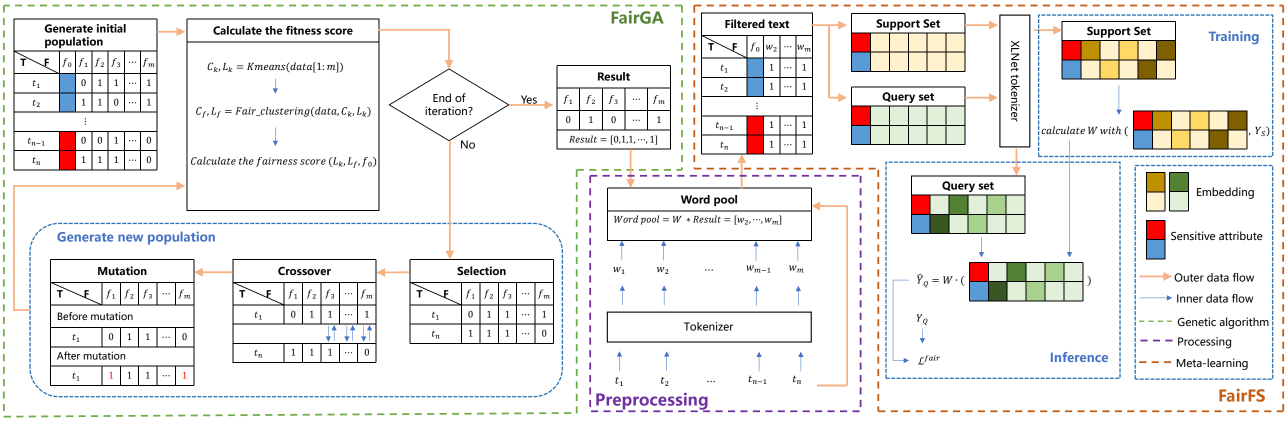

Figure 1. Overall FairGAFS architecture. The network consists of three stages: preprocessing, FairGA, and FairFS. Preprocessing plays the role of a transit station for data flow.

2.1. *Fairness-aware objective*

The fairness restrictions in fair clustering urge clusters to contain balanced demographic groups according to particular sensitive attributes (e.g., sex, gender, and race), so that any type of data-inherent bias can be counteracted with. The balance index is a popular index used to value fairness in many fairness algorithms. We propose a novel fitness object with balance. Our goal is to capture fairness in feature-selection processing.

Taking the dataset V that contains h groups v_s such that $V = \cup_{s\in[h]} v_s$, the balance [**?**] index requires that every cluster contains approximately the same number of protected elements from each group. For a clustering $V = c_1 \cup c_2 \ldots \cup c_k, k = h$. The balance of each cluster c_k is defined as

$$balance(c_k) = \min_{s\neq s'\in[h]} \frac{|v_s \cap c_k|}{|v_{s'} \cap c_k|} \in [0,1]. \tag{1}$$

The higher the balance of each cluster, the fairer the clustering. To investigate various degrees of trade-off between fairness and clustering objectives, Ziko et al. [**?**] formulated the fairness clustering problem in the following form:

$$\min_V \left\{ \mathcal{E}(V) := \underbrace{\mathcal{F}(V)}_{\text{clustering}} + \underbrace{\gamma \sum_k \sum_s -\mu_s \log P(s \mid k)}_{\text{fairness}} \right\}, \tag{2}$$

where $\mathcal{F}(V)$ means K-means clustering, γ is a constant coefficient, $\mu_s \in [0,1]$ is some given proportion of demographic within v_s, $P(s \mid k) = \frac{|v_s \cap c_k|}{|c_k|} \forall s$.

The fitness function of FairGA is written as follows:

$$Maximize \left(\frac{1}{k} \sum_{l=1}^{k} balance(c_l) \right). \tag{3}$$

2.2. Genetic algorithm

Typically, a document has hundreds or even thousands of unique words that are considered features, but many of these words may be obtrusive, redundant, or otherwise unhelpful when compared to the document's class designation [?]. The GA is a popular technology for non-linear optimization problems and is a parallel heuristic intelligent method. Owing to its advantages, the GA has been frequently utilized as an effective approach for feature selection in text classification [?]. Different from previous works [?, ?] applied to tabular dataset for binary classification, we apply FairGA to multi-category text classification.

During data pre-processing, we adapt the pre-trained XLNet model to tokenize the text set $T = \{t_1, t_2, ..., t_n\}$. Meanwhile, we delete high-frequency stop words that contain no information useful for determining the class label. For text representation, We treat each word as a feature. The filtered feature set $F = \{f_1, f_2, ..., f_m\}$ is utilized , such as $t_n = \{f_1, f_2, f_3, ..., f_m\} = \{1, 1, 1, ..., 0\}$, where $f_1 = 1$ denotes that the feature f_1 is in the text t_n and otherwise is 0.

The GA method starts with a randomly initialized population $P_{n \times m}$, which is an array of length proportional to the number of items and characteristics. For feature selection in evolutionary algorithms, we use the same method as for text representation. The selected feature is set to one, and the others to zero. The subset population $P_{t \times m}$ was chosen at random from $P_{n \times m}$, where t is the number of individuals selected. Then, we employ Eq (??) to determine each individual's fitness.

The new population generation is combined with selection, crossover, and mutation. In the selection operation, individuals with high fitness are more inclined to be selected. The crossover operation occurs when the parental chromosomes are evenly swapped. The mutation operation changes between 0 and 1 according to a pre-set probability value. It is important to note that the crossover and mutation operators are executed based on two changeable parameters called the crossover probability (pc) and mutation probability (pm).

The main steps of the FairGA algorithm can be summarized as follows.

Step 1. Give the population size $PSize$, pc, pm, and the number of iterations $iters$. Let G_m and i denote the final generation and number of evolutions, respectively. Then generate at random the first population $P(i)$, and assign $i = 0$.

Step 2. Assess the fitness of $P(i)$ by using the objective function.

Step 3. Pick $\frac{PSize}{2}$ persons with greater fitness relative to the following generation: $P(i + 1)$.

Step 4. Execute the crossover procedure on $P(i)$ to obtain $C(i)$.

Step 5. Conduct the mutation operation on $C(i)$, $i = i + 1$, and construct $P(i + 1)$.

Step 6. If $iters$ is reached, the process terminates, and the best candidate is produced in G_m; else, return to *Step 2*.

2.3. Fair meta-learning

The goal of meta-learning for few-shot learning is to develop a meta-learner that can learn from a limited quantity of data on a variety of different tasks. One well-known statistically based meta-learning framework is R2D2 [?]. We suggest a novel statistic-based strategy in which we apply sensitive characteristic constraints to each task to generalize fairness in a classification issue with few samples. Our goal is to learn representations that enhance the feature's generalization and add fair constraints that strengthen the classifier.

In the present work, we conducted a splitting of a dataset $\mathcal{D} = \{\mathcal{D}^{S}, \mathcal{D}^{V}, \mathcal{D}^{Q}\}$, which are a training set, validation set, and testing set, respectively. We emphasize that the post-partitioned dataset is not repeated, i.e., $\mathcal{D}^{S} \cap \mathcal{D}^{V} \cap \mathcal{D}^{Q} = \emptyset$. Then, we constructed a meta-training task $\mathcal{T}_j = \{(\mathcal{D}^{S}_j, \mathcal{D}^{V}_j)\}_{j=1}^{T}$ by sampling $\mathcal{N}$ classes with $\mathcal{K}$ samples from $\mathcal{D}^{S}$ and $\mathcal{D}^{V}$, respectively. The subscript j represents the j-th task, which is sampled from $p(\mathcal{T})$. Through the meta-training processing, $\mathcal{D}^{S}_j$ is used to update the model parameters in the outer loop, while $\mathcal{D}^{V}_j$ is used to fine-tune the model parameters in the inner loop. The goal of meta-learning in a supervised-learning environment is to learn a model f_θ with optimum parameters θ. Meta-training entails iteratively updating θ with the minimized loss $l_{\mathcal{T}_j}$. The processing is indicated as follows:

$$\theta^* = \arg\min_{\theta} \mathbf{E}_{\mathcal{T}_j} l_{\mathcal{T}_j}(f_\theta). \tag{4}$$

The classifier was also trained during the learning process. We chose the *ridge regression* classifier to adapt to minimal sample data while avoiding over-fitting. Before entering the classifier, we used an embedding function to encode the data:

$$emb(x) := BiLSTM(f_{w2v}(x_i)), \tag{5}$$

$$y = emb(x_S)W + b, \tag{6}$$

where $x_i \in x$ represents a sentence sequence, f_{w2v} is the pre-trained word vector, y is the predicted category probabilities, x_S is the support set, and W is the weight matrix.

As there are many hard samples in the dataset, we utilized focal loss to fade their impact. Specifically, we denote the loss function as follows:

$$\mathcal{L}^{focal}(\theta) = \alpha \cdot (1 - p_t)^{\sigma} \cdot \mathcal{L}^{CE}, \tag{7}$$

$$p_t = e^{-\mathcal{L}^{CE}}, \tag{8}$$

$$\mathcal{L}^{CE} = \hat{y} log(y), \tag{9}$$

where α and σ are adjustable constants, θ is the model's parameters, $\hat{y}$ is the true label of each sample, and $\mathcal{L}^{CE}$ is the cross-entropy loss function.

In addition, we focus on binary sensitive attributes, taking fairness into account during the training process. We added the fairness constraint* to $\mathcal{L}^{focal}$. Then, the loss function will be rewritten as follows:

$$\mathcal{L}^{fair} = \mathcal{L}^{focal} + \lambda \frac{\sum_{i=0}^{m} \omega_i f_i(y, \hat{y})}{\min\limits_{\forall i \in [0,m]} f_i(y, \hat{y})}, \tag{10}$$

where m is the number of values of protect attributes, λ is a constant, ω_i is a weight parameter, y is the true labels, $\hat{y}$ is the predicted labels, and f_i is the fairness score function, is based on accuracy with protected features.

Finally, the model's performance testing is based on $\mathcal{D}^{Q}$. We utilized the distribution $p(\mathcal{T})$ to generate meta-testing tasks and evaluated the average classification accuracy across all testing episodes.

*http://vi.le.gitlab.io/fair-loss/

3. Experiments

3.1. Experimental setup

Dataset: We validated the proposed method on three datasets: THUCnews [?], 20Newsgroups [?], and Amazon product data [?]. For each dataset, 500 samples are included in each category. Furthermore, we used blue and red to denote protected features. Table **??** displays the statistics for these datasets. Table **??** shows some examples of datasets. In FairGA, we marked 250 samples with blue and the others with red in each class to build a balanced dataset for fair clustering. Similarly, we labeled half of the classes with blue and the others with red in each meta task to build a balanced task for fair classification in FairFS.

Table 1. Statistics of three benchmark datasets.

Parameters	THUCnews	20Newsgroups	Amazon
Classes	14	15	15
Samples	7000	7500	7500
Avg. text length	324	362	141
Vocab. size	9707	7151	3249
Train/validate/test classes	6/4/4	5/5/5	5/5/5

Implementation details: In fair feature-selection processing, pc was set to 0.05, pm was set to 0.2, $PSize$ was fixed to 100, and $iters$ is adjusted to 12.

We employed a BiLSTM with 64 hidden units for meta knowledge representation and a pre-trained word vector for word embedding following Bao et al. [?]. All parameters were tuned with RMSProp at a learning rate of 0.01.

We ran 100 training episodes per epoch during meta training. When the accuracy on the validation set did not improve after 20 epochs, we applied early halting. We assessed the model's performance using over 1,000 testing episodes. All the experiments were conducted on a PC with a Core i9 processor and an NVIDIA RTX2080Ti graphical processing unit.

Metrics: We conducted the accuracy measurement to evaluate the category validity by comparing the predicted and ground-truth labels. In addition, we utilized the selection_rate metric to evaluate the fairness of the FS model. The measurements are calculated as follows:

$$Accuracy(f_\theta; \mathcal{D}^Q) = \frac{1}{n}\sum_{i=1}^{n}(f_\theta(x_i) = label_i), \tag{11}$$

$$selection\_rate = \frac{TP + FP}{P + N}, \tag{12}$$

where $x_i \in \mathcal{D}^Q$, the $label_i$ is the ground-truth label of item x_i, n is the count of test dataset $\mathcal{D}^Q$, and $f_\theta(x_i)$ is the model's predicted label. TP means true positive (P), FP denotes false positive, and N represents negative.

Table 2. The dataset examples.

Dataset	Label	Content
THUCnews	sports	Hornets vs Lakers Kobe with injury battle Paul Gasol redemption battle Sina Sports April 27 Beijing time, the first round of the NBA playoffs Los Angeles Lakers...
	property	The value of real estate in the financial crisis is reflected in the quality of life (group photo) Moderator: The following is a professor of the Department...
	education	Reading sprint before the CET4/6 exam: how to play a normal level English CET4/6 researcher Miao Jinhua Reading is the main event in the CET4/6 exam...
20Newsgroups	alt.athesim	... First, I'll make the assumption that you agree that a murderer is onewho has commited murder.Well, I'd say ...
	comp.graphics	... OK, with all the discussion about observed playback speeds with QuickTime,the effects of scaling and so ...
	sci.med	... I am writing this to find out the following: 1.) Any information on surgery to prevent reflux esophagitis. 2.)The name(s) of a doctor(s) who specialize in such surgery ...
Amazon	Arts	I really enjoy these scissors for my inspiration books that I am making (like collage, but in books) and using these different textures these give is just wonderful, makes a great statement with the pictures and sayings. Want more, perfect for any need you have even for gifts as well. Pretty cool!...
	Jewelry	I bought this necklace on a whim; I love butterflies and it looked so dainty and sweet. It was actually a little more weighty than I expected, although it's not a solid piece. The chain is shiny and nicer than I expected...
	Cell phone	Great product- tried others and this is a ten compared to them. Real easy to use and sync's easily. Definite recommended buy to transfer data to and from your Cell.

3.2. Comparison with state-of-the-art methods

To validate the proposed model's performance, we conducted tests on the three aforementioned datasets and compared the results to those of other state-of-the-art approaches [?, ?, ?, ?, ?]; see Figure **??**. The proposed algorithm produces better results than the baseline models (MAML [?] and PROTO [?]). The proposed approach yields the best performance on the Amazon dataset, i.e., 8.42% higher than HATT [?] in 5-shot classification. However, the proposed model's performance is 10.1% lower than HATT [?] on 20Newsgroups in 1-shot classification. Meanwhile, the proposed approach achieved better competitiveness on THUCnews. This is due to the weak encoder capability in our model. It can be inferred that the classification results are improved by feature selection on datasets with shorter length. In addition, our proposed method maintains a better performance in 5-shot, although the results are weak in 1-shot due to feature reduction.

(a) THUCnews **(b)** 20Newsgroups **(c)** Amazon

Figure 2. Comparisons of proposed method (FairGAFS) with other state-of-the-art methods.

3.3. Parameter tuning

We designed model parameter tuning experiments to observe the effects of adjustable parameters on model outcomes.

First, we observed the metrics that measure the balance of the feature-selection process in FairGA, and the results are shown in Figure **??**. The fitness value on the THUCnews dataset is approximately 0.7, which is smaller than that of the other two datasets. The maximum values all appear in the fifth iteration using 60% features. In addition, the maximum value on the Amazon dataset is continuously postponed with an increasing number of features. The results signify that original feature has a weak impact on balance index, and the iteration is increased with a modest number of features.

Second, we verified how the number of features affects network performance, and the results are shown in Figure **??**. Better classification progress can be achieved using 70% of the features, especially on the THUCnews and Amazon datasets. The performance is undesirable on the 20Newsgroups dataset, which is lower by 4.08% in 1-shot and by 0.34% in 5-shot classification. The results indicate that the method's performance can be improved by removing 30% of features.

Third, we confirmed how hyper-parameters (α & σ) in focal loss affect model performance. As shown in Figure **??**, one may find that optimal classification accuracy is achieved with high probability when σ is set to 5. Moreover, the maximum value is obtained in 1 shot when α is set to 0.75 and σ is set to 5. In addition, all classification results differ within 1%, and the gains obtained by adjusting the hyper-parameters are insignificant. It can be seen that the reduced features cause larger fluctuation in few-shot loss. Therefore, a large penalty is needed to suppress the loss function, and choosing a larger

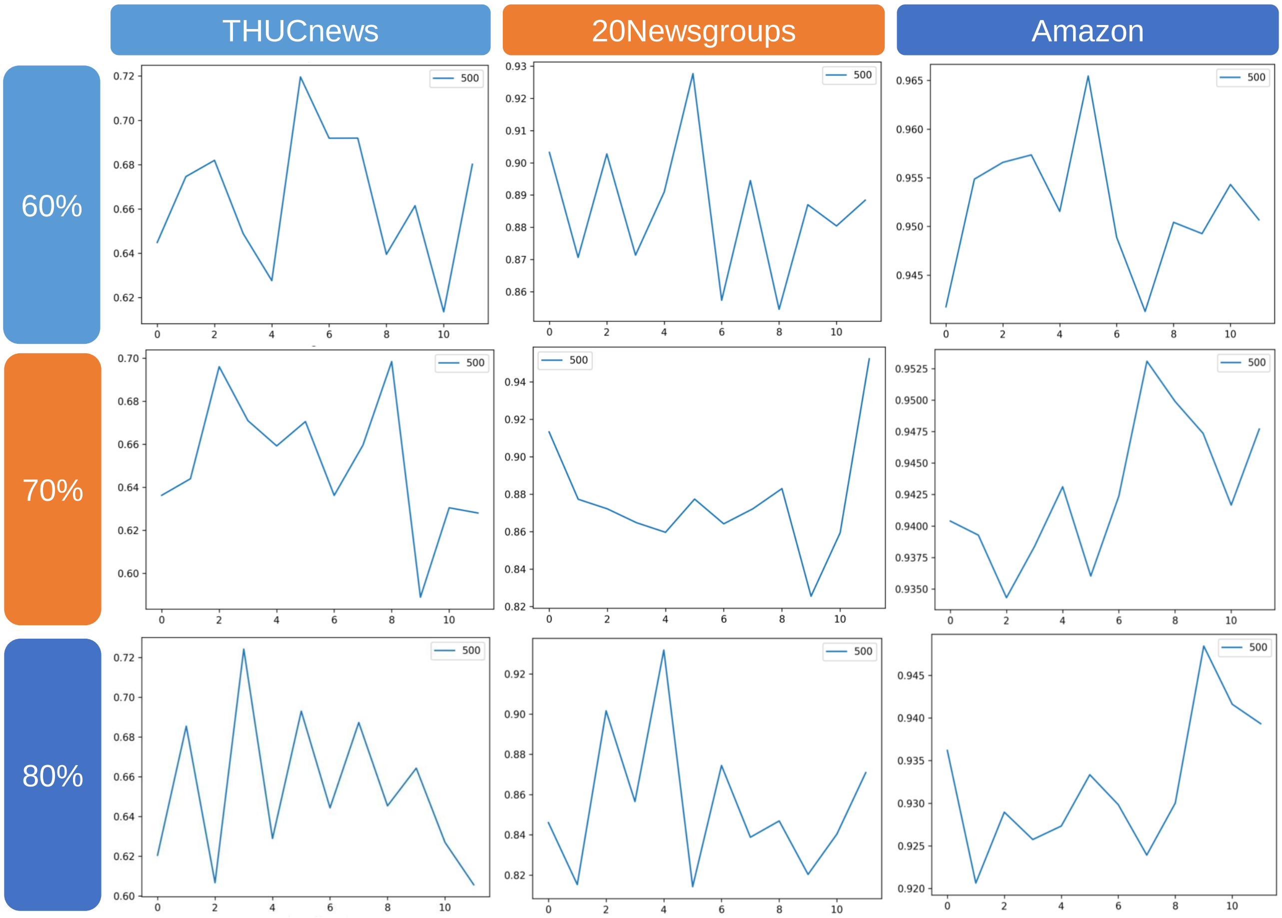

Features(%)	THUCnews	20Newsgroups	Amazon
60	0.72/5	0.928/5	0.965/5
70	0.7/8	0.955/11	0.953/7
80	0.725/3	0.935/4	0.948/9

Figure 3. Change in fitness as iteration proceeds under different feature amounts. Vertical axis shows fitness values, while horizontal axis shows number of iterations. 60%, 70% and 80% indicate the different numbers of features extracted by FairGA. X/Y denotes value of fitness (X) and number of iterations (Y).

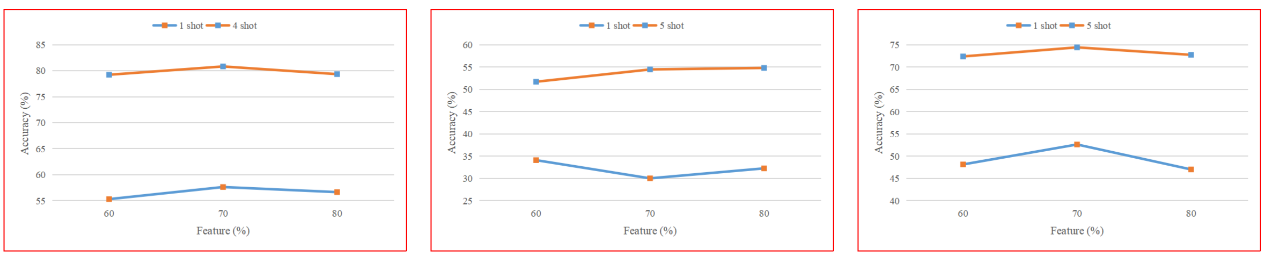

(a) THUCnews **(b)** 20Newsgroups **(c)** Amazon

Figure 4. Accuracy corresponding to different numbers of features.

value of σ acts more than α.

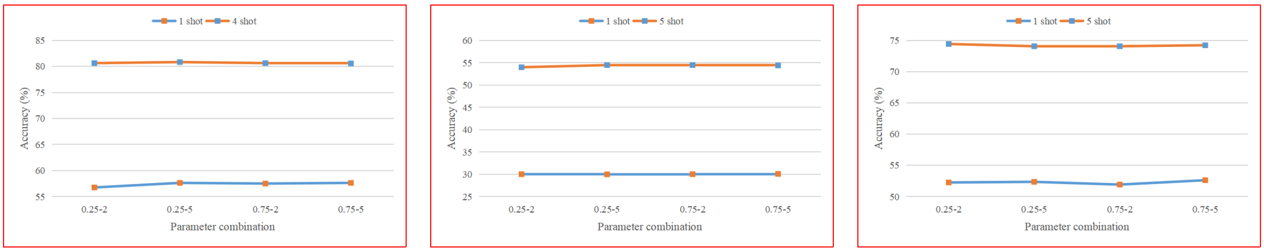

(a) THUCnews **(b)** 20Newsgroups **(c)** Amazon

Figure 5. Accuracy corresponding to adjustable parameters. X-Y denotes α (X) and σ (Y).

Finally, we investigated how the loss function affects how well the proposed model performs. From the results in Figure **??**, it can be seen that the proposed loss function can achieve optimal classification accuracy on most tasks, but does not perform as well as the other two loss functions on the 20Newsgroups dataset. The results show that adding a fairness score penalty to the loss function can lead to performance improvements.

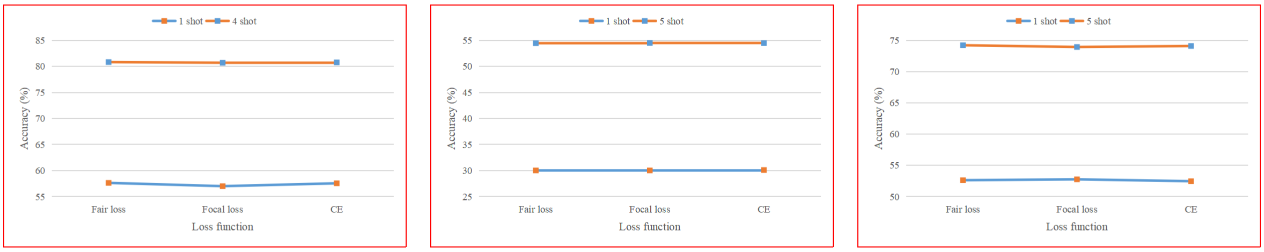

(a) THUCnews **(b)** 20Newsgroups **(c)** Amazon

Figure 6. Accuracy corresponding to different loss functions. The CE is the cross-entropy loss function. The fair loss is proposed by us, shown in Eq (**??**).

3.4. Fairness efficiency

We provided fairness evaluation of our model under different components in Figure **??**. The fairness metrics are accuracy and selection rate. After pre-processing with FairGA, the classification accuracy of FairGAFS decreases, with a particularly significant decrease on Amazon. However, the classification accuracy is improved in terms of equalization over each sensitive feature. In addition, the selection rate decrease, and the value gap on sensitive characteristics is reduced. The results indicate that our model improves the fairness of the FS classification model by introducing FairGA.

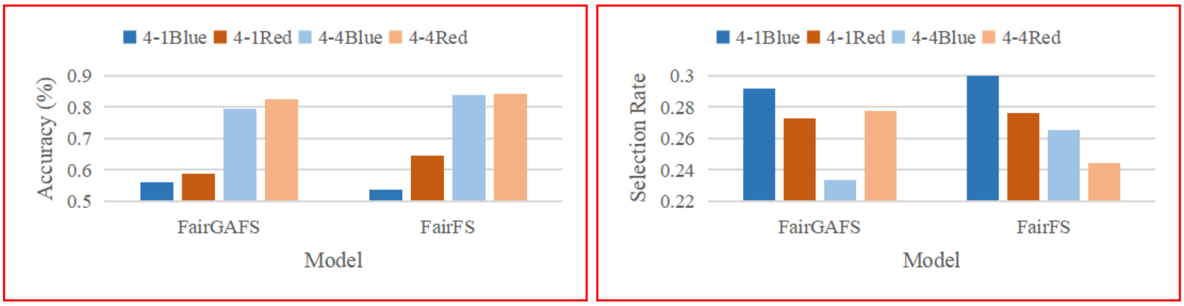

(a) THUCnews

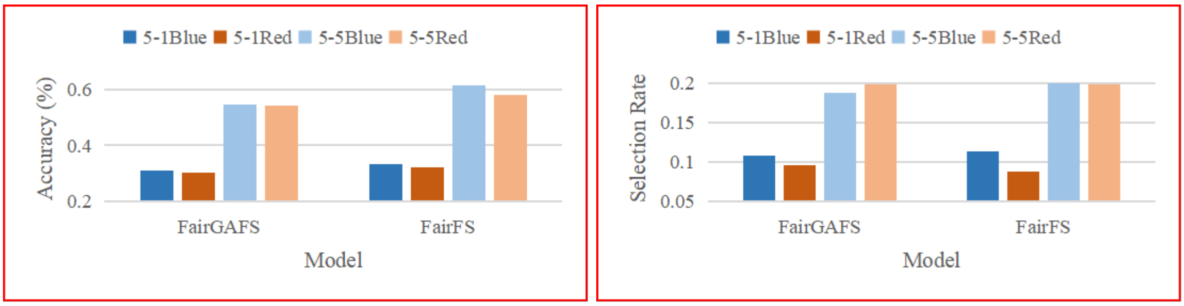

(b) 20Newsgroups

(c) Amazon

Figure 7. The fairness metrics evaluation under different models. Blue and red are sensitive attributes. M-K means M way K shot. The ideal value of selection rate is zero [**?**].

4. Conclusions

In this letter, we proposed a novel network that considered fairness, selected key features, and learned hard samples to achieve high classification accuracy. To a certain extent, FairGA and FairFS reduced the number of input features and improved the performance of the model, respectively. The proposed method achieved strong competitive performance on three public datasets. Our future work will involve developing a computationally efficient FairGA to reduce the computational consumption and improve the fairness calculation strategy in FairFS to further enhance the network.

Acknowledgments

This work was supported in part by the National Key R&D project under Grant 2019YFB1705503, in part by the Key Program of NSFC-Guangdong Joint Funds under Grant U1801263, in part by the

Guangdong Provincial Key Laboratory of Cyber-Physical Systems under Grant 2020B1212060069 and in part by the Science and Technology Projects of Guangzhou under Grant 202007040006.

Conflict of interest

The authors declare that there is no conflict of interest.

References

1. Y. Li, J. Yang, J. Wen, Entropy-based redundancy analysis and information screening, *Digital Commun. Networks*, **2021** (2021). https://doi.org/10.1016/j.dcan.2021.12.001
2. Y. Lin, J. Yang, W. Lu, Q. Meng, Z. Lv, H. Song, Quality index for stereoscopic images by jointly evaluating cyclopean amplitude and cyclopean phase, *IEEE J. Sel. Top. Signal Process.*, **11** (2017), 89–101. https://doi.org/10.1109/JSTSP.2016.2632422
3. Y. Li, J. Yang, Z. Zhang, J. Wen, P. Kumar, Healthcare data quality assessment for cybersecurity intelligence, *IEEE Trans. Ind. Inf.*, **19** (2023), 841–848. https://doi.org/10.1109/TII.2022.3190405
4. S. Yang, L. Liu, M. Xu, Free lunch for few-shot learning: Distribution calibration, preprint, arXiv: 2101.06395. https://doi.org/10.48550/arXiv.2101.06395
5. Y. Wang, Q. Yao, J. T. Kwok, L. M. Ni, Generalizing from a few examples: A survey on few-shot learning, **53** (2020), 1–34. https://doi.org/10.1145/3386252
6. D. Chen, Y. Chen, Y. Li, F. Mao, Y. He, H. Xue, Self-supervised learning for few-shot image classification, in *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2021)*, (2021), 1745–1749. https://doi.org/10.1109/ICASSP39728.2021.9413783
7. J. Xu, Q. Du, Learning transferable features in meta-learning for few-shot text classification, **135** (2020), 271–278. https://doi.org/10.1016/j.patrec.2020.05.007
8. M. Morik, A. Singh, J. Hong, T. Joachims, Controlling fairness and bias in dynamic learning-to-rank, in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, (2020), 429–438. https://doi.org/10.1145/3397271.3401100
9. P. Li, H. Zhao, H. Liu, Deep fair clustering for visual learning, in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, (2020), 9070–9079.
10. A. Chhabra, K. Masalkovaitė, P. Mohapatra, An overview of fairness in clustering, *IEEE Access*, **9** (2021), 130698–130720. https://doi.org/10.1109/ACCESS.2021.3114099
11. X. Deng, Y. Li, J. Weng, J. Zhang, Feature selection for text classification: A review, *Multimedia Tools Appl.*, **78** (2019), 3797–3816. https://doi.org/10.1007/s11042-018-6083-5
12. M. Fan, W. Wei, W. Jin, Z. Yang, T. Liu, Explanation-guided fairness testing through genetic algorithm, in *2022 IEEE/ACM 44th International Conference on Software Engineering (ICSE)*, (2022), 21843296. https://doi.org/10.1145/3510003.3510137
13. X. Xing, H. Liu, C. Chen, J. Li, Fairness-aware unsupervised feature selection, in *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, (2021), 3548–3552. https://doi.org/10.1145/3459637.3482106 .

14. C. Zhao, C. Li, J. Li , F. Chen, Fair meta-learning for few-shot classification, in */2020 IEEE International Conference on Knowledge Graph (ICKG)*, (2020), 275–282. http://arxiv.org/abs/2009.13516 .

15. F. Chierichetti, R. Kumar, S. Lattanzi, S. Vassilvitskii, Fair clustering through fairlets, *Adv. Neural Inf. Process. Syst.*, **30** (2017).

16. M. Kleindessner, S. Samadi, P. Awasthi, J. Morgenstern, Guarantees for spectral clustering with fairness constraints, in *International Conference on Machine Learning*, (2019), 3458–3467.

17. I. M. Ziko, E. Granger, J. Yuan, I. B. Ayed, Clustering with fairness constraints: A flexible and scalable approach, preprint, arXiv: 1906.08207. http://arxiv.org/abs/1906.08207 .

18. T. Y. Lin, P. Goyal, R. B. Girshick, K. He, P. Dollár, Focal loss for dense object detection, *IEEE Trans. Pattern Anal. Mach. Intell.*, **42** (2020), 318–327. https://doi.org/10.1109/tpami.2018.2858826

19. H. Chen, W. Jiang, C. Li, R. Li, A heuristic feature selection approach for text categorization by using chaos optimization and genetic algorithm, *Math. Problems Eng.*, **2013** (2013), e524017. https://doi.org/10.1155/2013/524017

20. S. Tizpaz-Niari, A. Kumar, G. Tan, A. Trivedi, Fairness-aware configuration of machine learning libraries, in *2022 IEEE/ACM 44th International Conference on Software Engineering (ICSE)*, 2022. https://doi.org/10.1145/3510003.3510202

21. A. U. Rehman, A. Nadeem, M. Z. Malik, Fair feature subset selection using multiobjective genetic algorithm, in *GECCO '22: Proceedings of the Genetic and Evolutionary Computation Conference Companion*, (2022), 360–363. https://doi.org/10.1145/3520304.3529061

22. Y. Bao, M. Wu, S. Chang, R. Barzilay, Few-shot text classification with distributional signatures, preprint, arXiv:1908.06039. http://arxiv.org/abs/1908.06039

23. J. Li, M. Sun, Scalable term selection for text categorization, in *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, (2007), 774–782. https://aclanthology.org/D07-1081

24. K. Lang, Newsweeder: Learning to filter netnews, in *Proceedings of the Twelfth International Conference on International Conference on Machine Learning*, (1995), 331–339. https://doi.org/10.1016/B978-1-55860-377-6.50048-7

25. R. He, J. McAuley, Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering, in *Proceedings of the 25th International Conference on World Wide Web*, (2016), 507–517. https://doi.org/10.1145/2872427.2883037

26. C. Finn, P. Abbeel, S. Levine, Model-agnostic meta-learning for fast adaptation of deep networks, in *International Conference on Machine Learning*, (2017), 1126–1135.

27. J. Snell, K. Swersky, R. Zemel, Prototypical networks for few-shot learning, in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, (2017), 4080–4090.

28. R. Geng, B. Li, Y. Li, X. Zhu, P. Jian, J. Sun, Induction networks for few-shot text classification, in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, (2019), 3904–3913. http://dx.doi.org/10.18653/v1/D19-1403

29. T. Gao, X. Han, Z. Liu, M. Sun, Hybrid attention-based prototypical networks for noisy few-shot relation classification, in *Proceedings of the AAAI Conference on Artificial Intelligence*, (2019), 6407–6414. https://doi.org/10.1609/aaai.v33i01.33016407

30. D. Wang, Z. Wang, L. Cheng, W. Zhang, Few-shot text classification with global-local feature information, *Sensors*, **22** (2022), 4420. https://doi.org/10.3390/s22124420

31. S. Majumder, J. Chakraborty, G. R. Bai, K. T. Stolee, T. Menzies, Fair enough: Searching for sufficient measures of fairness, preprint, arXiv:2110.13029. https://doi.org/10.48550/arXiv.2110.13029