arXiv:2012.01675v2 [cs.CL] 6 Apr 2022

Federated Learning for Personalized Humor Recognition

XU GUO, School of Computer Science and Engineering, Nanyang Technological University, Singapore
HAN YU, School of Computer Science and Engineering, Nanyang Technological University, Singapore
BOYANG LI*, School of Computer Science and Engineering, Nanyang Technological University, Singapore
HAO WANG, Alibaba Group, China
PENGWEI XING, School of Computer Science and Engineering, Nanyang Technological University, Singapore
SIWEI FENG, School of Computer Science & Technology, Soochow University, China
ZAIQING NIE, Institute for AI Industry Research, Tsinghua University, China
CHUNYAN MIAO†, School of Computer Science and Engineering, Nanyang Technological University, Singapore

Computational understanding of humor is an important topic under creative language understanding and modeling. It can play a key role in complex human-AI interactions. The challenge here is that human perception of humorous content is highly subjective. The same joke may receive different funniness ratings from different readers. This makes it highly challenging for humor recognition models to achieve personalization in practical scenarios. Existing approaches are generally designed based on the assumption that users have a consensus on whether a given text is humorous or not. Thus, they cannot handle diverse humor preferences well. In this paper, we propose the FedHumor approach for the recognition of humorous content in a personalized manner through Federated Learning (FL). Extending a pre-trained language model, FedHumor guides the fine-tuning process by considering diverse distributions of humor preferences from individuals. It incorporates a diversity adaptation strategy into the FL paradigm to train a personalized humor recognition model. To the best of our knowledge, FedHumor is the first text-based personalized humor recognition model through federated learning. Extensive experiments demonstrate the advantage of FedHumor in recognizing humorous texts compared to nine state-of-the-art humor recognition approaches with superior capability for handling the diversity in humor labels produced by users with diverse preferences.

CCS Concepts: • **Computing methodologies** → **Distributed artificial intelligence**.

Additional Key Words and Phrases: subjectivity, personalization, natural language understanding

*Corresponding Author
†Corresponding Author

Authors' addresses: Xu Guo, School of Computer Science and Engineering, Nanyang Technological University, Singapore, xu008@ntu.edu.sg; Han Yu, School of Computer Science and Engineering, Nanyang Technological University, Singapore, han.yu@ntu.edu.sg; Boyang Li, School of Computer Science and Engineering, Nanyang Technological University, Singapore, boyang.li@ntu.edu.sg; Hao Wang, Alibaba Group, China; Pengwei Xing, School of Computer Science and Engineering, Nanyang Technological University, Singapore; Siwei Feng, School of Computer Science & Technology, Soochow University, China; Zaiqing Nie, Institute for AI Industry Research, Tsinghua University, China; Chunyan Miao, School of Computer Science and Engineering, Nanyang Technological University, Singapore, ascymiao@ntu.edu.sg.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.
XXXX-XXXX/2018/4-ART $15.00
https://doi.org/10.1145/1122445.1122456

ACM Reference Format:
Xu Guo, Han Yu, Boyang Li, Hao Wang, Pengwei Xing, Siwei Feng, Zaiqing Nie, and Chunyan Miao. 2018. Federated Learning for Personalized Humor Recognition. 1, 1 (April 2018), 18 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

In the past decade, we have witnessed broad and far-reaching success of artificial intelligence (AI) in novel application areas such as conversational characters, which provide increasingly human-like interactions to users. One particularly challenging issue in the understanding of human language is the modeling of humor. Humor is a prominent example of linguistic creativity [37] and plays a vital role in human communication. Unlike AI tasks with objective ground truth, the task of humor recognition is characterized by the subjectivity in humor understanding. Due to individual differences in cognitive processes, the same joke can be perceived differently by its audience [3, 15, 25]. This is empirically proved via data analysis on a real-world dataset [16] in which each joke is rated by five persons in terms of its funniness, and the results show that the variance of human's perceived funniness on the same joke is nontrivial (Figure 1).

Computational approaches for recognizing humor generally regard the task as a binary classification problem [7, 39]. Traditional research mainly focuses on extracting expressive features as inputs to a classifier in order to improve classification performance. Neural networks such as Convolutional Neural Networks (CNNs) [7] and transformer networks [24] have shifted the focus from feature engineering to automatic feature extraction. Existing humor recognition models are generally based on the assumption that users have a consensus about whether or not a given text is humorous (as illustrated in Figure 2(a)) [7, 21, 24, 39, 43]. However, in reality, humor preference is highly subjective. As a result, these approaches perform well recommending humorous contents with popular appeal, but cannot achieve personalized humor recognition.

In this paper, we bridge this important gap in the humor recognition literature by proposing a personalized humor recognition approach - FedHumor. We adopt the assumption that humor labels from different users regarding the same text contents can be diverse, and formalize the problem as a conditional binary classification task. The federated learning (FL) architecture is used as the basis for our proposed personalized humor recognition model.

FedHumor extends the popular Federated Averaging (FedAvg) algorithm [26] with a diversity adaptation strategy to enhance the handling of disparate user preferences in humor recognition. As a result, it can adjust the order of humor contents (e.g., jokes) recommended to different users according to their individual preferences (as illustrated in Figure 2(b)). The federated training process allows multiple views over diverse label distributions, thus enhancing generalization of the humor recognition model. Meanwhile, the distributed learning process allows personalized adaptations to different users' preferences to be learned locally, thus enhancing the personalization performance of the model. Moreover, this approach does not require sensitive data concerning each user's personal humor preference to be exposed, complying with the General Data Protection Regulation (GDPR) requirement on data privacy preservation.

To the best of our knowledge, FedHumor is the first federated learning based humor recognition model that explicitly considers the diversity of humor preferences. We conducted the personalized humor recognition by transforming existing humor recognition datasets labeled with explicit ratings into the dataset labeled with implicit binary labels reflecting different humor preferences. Extensive experiments against nine state-of-the-art approaches show that FedHumor achieves the best performance, surpassing the best performing existing approach by 2.5% in terms of F_1 score. FedHumor represents a promising direction for personalized recommendation of humorous content

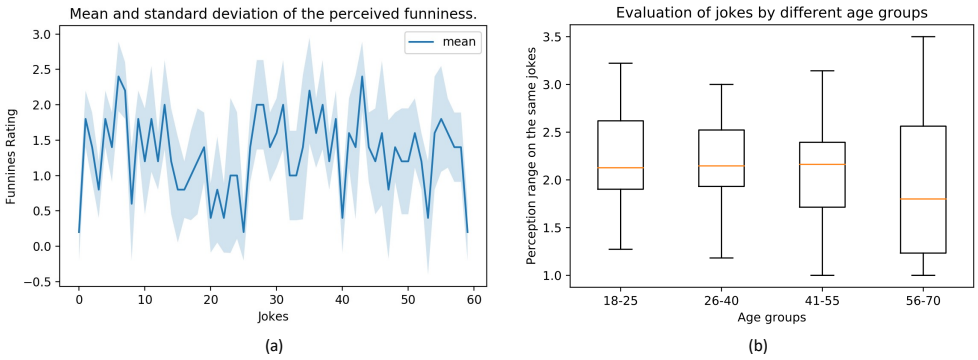

Fig. 1. Empirical analysis of a random set of 60 jokes from a real-world humor rating dataset reflecting non-trivial subjectivity in human perception. (a) shows that users' perceived funniness on the same jokes vary from person to person and the variance differs from joke to joke (shaded area). (b) shows that the effect in (a) is consistent across different age groups, albeit at different levels of variance.

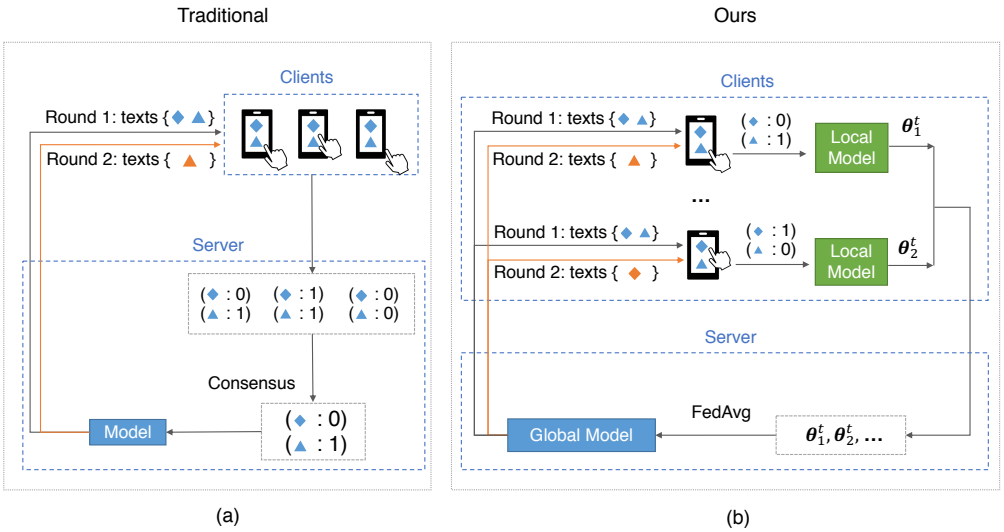

Fig. 2. Traditional setting versus personalized federated learning setting for training and applying a humor recognition model (*Best viewed in color*). In (a), a humor recognition model is trained on a centralized dataset whose labels are determined by majority voting (consensus) from round 1; the trained model is used to recommend funny texts to all clients at round 2, without distinction. In (b), a humor recognition model is trained on distributed datasets where the individual labels and distributions are preserved on local devices at round 1; the trained model will recommend texts to each client at round 2, possibly in different sequences.

under tightened data privacy protection regulations [41], thereby enabling innovative forms of human-AI interaction to emerge.

Our contributions are summarized as follows:

- We relaxed the common assumption in the humor recognition literature that users have a consensus about whether or not a given text is humorous.
- We proposed an FL-based solution with a personalized adaptation strategy to enable personalized humor recognition with good generalization, while not exposing private humor preference data.
- We conducted extensive experiments to evaluate our proposed approach. Results show that our approach is significantly superior to existing approaches in terms of personalizing humor recognition.

2 RELATED WORK

2.1 Humor Recognition

Based on well-established computational humor theory [4, 13, 34], early research works on text-based humor recognition mainly rely on heuristic rules to manually extract features for a classifier to detect humor [21, 36, 39]. Humor-related features can be broadly divided into stylistic features and linguistic structures. Stylistic features such as alliteration, antonyms, and slang [28] are prevalent in stereotypical one-liners in online repositories. Linguistic structures are employed to quantitatively analyze the existence of humor patterns in a sentence [21, 36]. These patterns are exploited to reflect humor in aspects such as incongruity, ambiguity, interpersonal effects, and phonetic style [39]. Nevertheless, these manual feature engineering methods are labor-intensive.

With the advance in computational power and the explosion of data, deep learning has been widely adopted to assist or substitute traditional feature engineering in many areas. Compared with other fields, the development of deep learning for humor recognition is more recent. The existing line of research can be divided according to the milestones of deep learning for natural language processing (NLP). The first category is to design neural architectures, such as CNNs, to learn sentence representations from pretrained word embeddings (e.g., GloVe) [6, 7]. Another category is to fine-tune pretrained language models (e.g., BERT) on humor recognition datasets [24]. Nevertheless, these techniques did not bring new perspectives on humor studies. Different from previous research, this paper is concerned with the issue of label diversity in existing humor recognition datasets, as the existing datasets are all labeled by a few annotators. For example, One Liners [28], Pun of the Day [39], The Onion News [27], SemEval-2017 shared task 7 [30] and [43] were collected for binary classification. The SemEval-2017 shared task 6 [32] is a tweets dataset annotated with three classes determined via ranking. These datasets generally overlooked the differences between individual preferences. Moreover, the non-humorous samples they choose are simply selected from other normal texts, which puts the model at a risk of learning shortcut features [12] rather than understanding the humor in texts. The recent SemEval-2020 shared task 7 [16] avoided these issues by creating a dataset in which the jokes are produced by some users and the funniness scores for those jokes are the average scores rated by other crowd-sourced users. However, this dataset is not dedicated for personalization. In this paper, we will reuse this dataset for our studies.

In psychological literature, it has been shown that the gap between the intended humor and the perceived humor is non-trivial [38]. The lack of label diversity in existing humor datasets precludes existing models from learning such personalized preferences. This paper attempts to address this problem by leveraging federated learning to enhance personalized humor recognition.

2.2 Federated Learning and Personalization

FL aims to protect data privacy by facilitating joint training of a global model across multiple data silos without exposing local data [17, 40]. FL can be divided into two major categories: horizontal FL and vertical FL. Horizontal FL involves data owners with local datasets in the same feature space, while vertical FL involves data owners with local datasets in different feature spaces. Our work belongs to the category of horizontal FL. The mainstream research works in horizontal FL generally assume an even partitioning on data samples and that $P(y|\boldsymbol{x})$ holds the same across devices as in FedAvg [26], FedProx [20], FedOpt [2] and many other FL problems such as incentive mechanism research for fairness [23, 42]. Therefore, these approaches only develop common outputs for all users. In many scenarios, the assumption of independent and identically distributed (i.i.d.) samples does not hold across different data silos.

Existing works on heterogeneous horizontal FL can be divided into two categories. The first category of research aims to eliminate statistical heterogeneity to ensure robust convergence of the global model. For example, FedProx [20] added a weight approximation term to restrict the solutions of the local models to be closer to that of the global model. In [8], a credit-weighted weights averaging approach is proposed to handle label disparities arising from healthcare practitioners' varying skill levels. The second category aims to adapt to the naturally occurring data heterogeneity under FL settings and personalize the global model for each data owner.

There are mainly two ways to achieve personalization. Firstly, we can adapt the global FL model weights for each data owner. For example, in Federated Multi-task Learning [35], a global model is optimized on all data silos using FedAvg while each device learns a separate model with the whole model's weights [19] or with only the final layers' weights [1] being optimized on the local dataset. On top of this, meta-learning has also been used to find better local solutions for each data owner [10]. Secondly, the model predictions or output logits can be adapted for each data owner. For example, [33] proposed a diversification mechanism to tune the diversity of the predicted emoji hints made by FedAvg during users' typing. Our work belongs to this category.

A common default assumption made in horizontal FL is that the datasets are partially owned by each user and their features are heterogeneously distributed while the FL model is forced to learn across each user's device to build a robust machine learning model. Different from this assumption, we focus on the problem in which the same input data (i.e., humorous texts) are labeled differently by different users when viewed on their own devices.

3 PRELIMINARIES

3.1 Notations

In this paper, we use bold letters for vectors and matrices, plain capital letters for sets and tuples, and plain lowercase letters for scalars. The probability of an event is denoted by P and the probability density function of a random variable is denoted by p. $|\cdot|$ indicates set cardinality while $||\cdot||$ indicates the Euclidean norm of a vector or matrix. We use Greek letters in the subscript to indicate function parameters. For example, $f_{\boldsymbol{\theta}}$ denotes a function parameterized by $\boldsymbol{\theta}$. $\mathbb{E}_{p(\boldsymbol{x})}[f\boldsymbol{x})]$ denotes taking the expectation of $f(\boldsymbol{x})$ over the distribution $p(\boldsymbol{x})$. The nabla $\triangledown$ indicates the gradient. For example, $\triangledown_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{\theta})$ denotes the gradient of $\mathcal{L}(\boldsymbol{\theta})$ with respect to the parameters $\boldsymbol{\theta}$.

3.2 Federated Gradient-based Optimization

We find the optimal parameters $\boldsymbol{\theta}^*$ for the function $f_{\boldsymbol{\theta}}$ by optimizing a differentiable loss function $\mathcal{L}(\boldsymbol{\theta})$ on a training set $\mathcal{D}$, where $\boldsymbol{x} \sim p(\boldsymbol{x}), \forall \boldsymbol{x} \in \mathcal{D}$.

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \mathcal{L}(\mathcal{D}; \boldsymbol{\theta}) \tag{1}$$

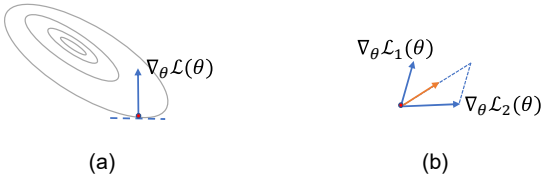

Fig. 3. Model parameters updated following standard gradient descent (a), and averaged gradient descent in federated learning (b). *Best viewed in color.*

The optimal $\boldsymbol{\theta}^*$ is found by iterated gradient descent. At each step t, we update $\boldsymbol{\theta}$ with learning rate η controlling the step size in the direction of the gradient $\nabla_{\boldsymbol{\theta}^t} \mathcal{L}(\mathcal{D}; \boldsymbol{\theta}^t)$:

$$\boldsymbol{\theta}^{t+1} \leftarrow \boldsymbol{\theta}^t - \eta \nabla_{\boldsymbol{\theta}^t} \mathcal{L}(\mathcal{D}; \boldsymbol{\theta}^t) \tag{2}$$

The above update is repeated until convergence or a predefined number of iterations is reached.

Under the Federated Learning setting, m participants in the federation wish to jointly optimize a model without sharing data with each other. Federated Averaging (FedAvg) [26] provides a general method to perform this optimization across m data silos, $\mathcal{D} = \bigcup_{i=1}^{m} \mathcal{D}_i$ and $\boldsymbol{x} \sim p^{(i)}(\boldsymbol{x})$, $\forall \boldsymbol{x} \in \mathcal{D}_i$. The local model parameters are adapted by gradient descent from global parameters $\boldsymbol{\theta}^t$ at each step. To distinguish local parameters and global parameters, we let $\boldsymbol{\theta}_i^{t,k}$ denote the local parameters at the i^{th} participant after applying k local updates to the global parameter $\boldsymbol{\theta}^t$.

$$\begin{aligned} \boldsymbol{\theta}_i^{t,0} &\leftarrow \boldsymbol{\theta}^t \\ \boldsymbol{\theta}_i^{t,k+1} &\leftarrow \boldsymbol{\theta}_i^{t,k} - \eta_t \nabla_{\boldsymbol{\theta}_i^{t,k}} \mathcal{L}(\mathcal{D}_i; \boldsymbol{\theta}_i^{t,k}) \end{aligned} \tag{3}$$

After every K local iterations, we synchronize the parameters across participants by averaging local updates. The global parameter update can be written as

$$\boldsymbol{\theta}^{t+1} \leftarrow \frac{1}{m} \sum_{i=1}^{m} \boldsymbol{\theta}_i^{t,K} \tag{4}$$

This implies that the model parameters are updated following the direction of the sum of the gradient vectors as illustrated in Figure 3. Traditionally, the parameters are updated iteratively on batches of labeled samples randomly sampled from the training set. In this case, the gradient descent direction is guided by the loss over a batch of observed samples (Figure 3(a)). In federated learning, multiple batches of samples are observed by the global model during every training iteration, while only a compromised step (i.e. $\frac{1}{2}(\nabla_{\boldsymbol{\theta}} \mathcal{L}_1(\boldsymbol{\theta}) + \nabla_{\boldsymbol{\theta}} \mathcal{L}_2(\boldsymbol{\theta}))$) is taken to update parameters $\boldsymbol{\theta}$ (Figure 3(b)).

4 PROBLEM FORMULATION

Consider a practical scenario in which a conversational character wants to tell a user a joke and hence needs to predict whether the user will perceive the joke as humorous or not. The conversational character has a set of historical texts that have binary feedback from a set of users, who may disagree with each other on if a joke is funny. Instead of attempting to reconcile the differences of the user-generated labels, in this paper, we study a personalized humor recognition problem in which the predicted classes of jokes are conditioned on the personal preferences of users.

Formally, $\boldsymbol{X} = (\boldsymbol{x}_1, ..., \boldsymbol{x}_n)$ denotes a total of n input jokes, where $\boldsymbol{x}_j \sim p(\boldsymbol{x})$ is determined by a content publisher. $Y_i = \{y_1^{(i)}, ..., y_n^{(i)}\}$, where $y_j^{(i)} \in \{0, 1\}$, denotes the target labels produced by user i on inputs $\boldsymbol{X}$. We use α, a quantified funniness threshold lying between the liked and disliked jokes in the historical data of a user, to denote personal humor preference. For example, the distribution $p(y|\boldsymbol{x}, \alpha_i)$ over the binary labels of user i is determined by α_i. Given m users' historical data, the task here is to predict Y for a user having $\alpha \sim p(\alpha)$. We treat this task as a conditional binary classification task.

In a standard supervised learning setting, the probability of input $p(\boldsymbol{x})$ is assumed to be unchanged. A traditional binary classification task tries to train a model $f_{\boldsymbol{\theta}}$ to estimate the true $p(y|\boldsymbol{x})$ based on sufficient observations in order to construct the relationship between data and labels $p(\boldsymbol{x}, y)$:

$$p(\boldsymbol{x}, y) = p(y|\boldsymbol{x})p(\boldsymbol{x}). \tag{5}$$

Note that the standard learning task assumes that the input data $\boldsymbol{x} \sim p(\boldsymbol{x})$ while the label $y \sim p(y|\boldsymbol{x})$ only depends on the input data $\boldsymbol{x}$. When applied to existing humor recognition datasets, it involves an underlying assumption that different users follow the same label distribution $p(y|\boldsymbol{x})$. However, this assumption may not hold for humor perception in practice.

Here, we assume that $p(y|\boldsymbol{x})$ relates to users' preference α. The problem thus becomes training the model $f_{\boldsymbol{\theta}}$ to estimate the conditional probability $p(y|\boldsymbol{x}, \alpha)$ based on given $\alpha \sim p(\alpha)$ and observed $\boldsymbol{x} \sim p(\boldsymbol{x})$:

$$p(\boldsymbol{x}, y, \alpha) = p(y|\boldsymbol{x}, \alpha)p(\boldsymbol{x}, \alpha) = p(y|\boldsymbol{x}, \alpha)p(\boldsymbol{x})p(\alpha). \tag{6}$$

Here, $p(\boldsymbol{x})$ is determined by the content publisher and $p(\alpha)$ is determined by users. So they are independent and satisfy: $p(\boldsymbol{x}, \alpha) = p(\boldsymbol{x})p(\alpha)$.

5 THE PROPOSED MODEL

In this section, we describe the detailed design of the proposed FedHumor model for the task of personalized humor recognition. We first introduce the contextualized text encoder model, then we present the federated learning-based model training, followed by the diversity adaptation design for different users, and finally we introduce the model validation under FL setting. The pseudo-code for FedHumor is given in Algorithm 1.

5.1 FedHumor Network

We employed the pretrained language model - BERT [9] - to capture contextualized sentence representations as the input features to a classification layer, and fine-tune the pretrained weights together with classifier parameters on our task. Briefly, BERT tokenizes a piece of text into a sequence of word IDs indexed by its stationary vocabulary. The IDs are used to fetch their corresponding word embeddings from an embedding table: $\boldsymbol{W} = (\boldsymbol{W}_1, ..., \boldsymbol{W}_l)$, $\boldsymbol{W} \in \mathbb{R}^{l \times d}$. l is the fixed maximum length of a sentence and d is the word embedding size. The interactions between words and order information are captured through the self-attention modules in twelve transformer layers, denoted as $\boldsymbol{T}$. The first vector ([CLS]) of the final hidden states at the last transformer layer is used to represent the final contextualized sentence representation, $\boldsymbol{x}$, for the given sentence. A non-linear classification layer, denoted as $\boldsymbol{C}$, followed by a $tanh$ activation function, accepts the sentence representations and makes predictions. The trainable parameters in the FedHumor model are $\boldsymbol{\theta} = [\boldsymbol{W}; \boldsymbol{T}; \boldsymbol{C}]$ and time complexity for each round is $O(|\boldsymbol{W}||\boldsymbol{T}||\boldsymbol{C}|K)$. K is the local training iterations on each user's device while all users perform local adaptation synchronously.

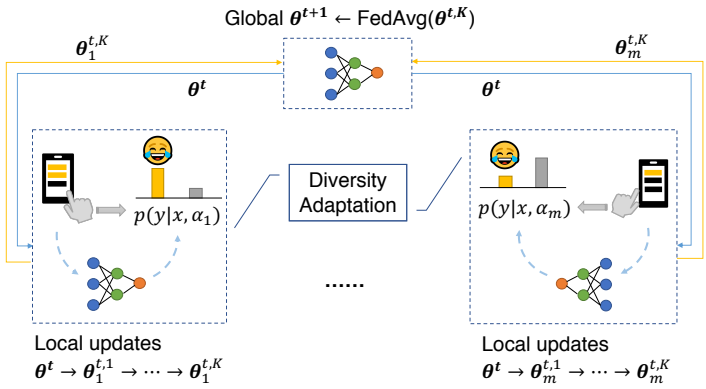

Fig. 4. The training of FedHumor involves three steps: 1) server sends the global model to clients; 2) the clients train the model locally based on their own labels, and send their updated parameters to server; 3) server aggregates local updates to produce a new global model.

5.2 Weight-tying Federated Training

We follow the FedAvg algorithm [26] to aggregate the model parameters updated on each user's data silo $\mathcal{D}_i = \{\boldsymbol{X}, Y_i\}$ after each round of federated training to produce a global model, as illustrated in Figure 4. In every training round, the current global FL model with parameters $\boldsymbol{\theta}$ is sent to m users randomly selected from the device population, and the adapted local parameters are denoted as $\boldsymbol{\theta}_i$. For each user whose humor preference is α_i, we use cross entropy as our loss function, and the instance loss for user i is computed as follows:

$$\ell(\boldsymbol{\theta}_i; \alpha_i) = -y\log(\tilde{p}(y|\boldsymbol{x}, \alpha_i)) - (1-y)\log(1-\tilde{p}(y|\boldsymbol{x}, \alpha_i)) + \lambda||\boldsymbol{\theta}_i||_2^2. \tag{7}$$

Here, $\tilde{p}(y|\boldsymbol{x}, \alpha_i))$ is the adapted estimated probability computed using Equations (10) and (11). λ is the hyperparameter controlling L2 norm regularization over local model parameters $\boldsymbol{\theta}_i$ (a.k.a. weight decay). $\boldsymbol{\theta}_i$ is determined by gradient descent to minimize the following training loss on user i's data silo:

$$\mathcal{L}(\boldsymbol{\theta}_i; \alpha_i) = \frac{1}{|\mathcal{D}_i|} \sum_{\mathcal{D}_i} \ell(\boldsymbol{\theta}_i; \alpha_i) \tag{8}$$

In each federated training round t, the local model parameters are adapted from $\boldsymbol{\theta}^t$ to $\boldsymbol{\theta}_i^{t,K}$ for K local iterations. The global model parameters are adapted from $\boldsymbol{\theta}^t$ to $\boldsymbol{\theta}^{t+1}$ by averaging the aggregated parameters:

$$\boldsymbol{\theta}^{t+1} \leftarrow \boldsymbol{\theta}^t - \eta_t \cdot \frac{1}{m} \sum_{i=1}^{m} \nabla_{\boldsymbol{\theta}_i^{t,K}} \mathcal{L}(\boldsymbol{\theta}_i^{t,K}; \alpha_i) \tag{9}$$

where m is the number of users joining the federated learning and η_t is the learning rate at each round and its value depends on the learning rate scheduler. Throughout the process, users' personal humor preference information (i.e. $p(y|\boldsymbol{x}, \alpha)$) are kept on their own devices.

5.3 Diversity Adaptation

The purpose of applying federated learning to humor recognition is to enhance the generalizability of traditional humor recognition models when dealing with diverse user preferences while preserving data privacy. FedHumor is designed similar to Google's GBoard scenario [33] based on the FL paradigm. Inspired by the mechanism employed in GBoard to prevent the federated model being dominated by highly frequent emojis, we adapt the model predictions, $\widehat{p}(y|\boldsymbol{x})$, on each class $c = 1, .., C$, for user i as follows:

$$v_c = \frac{\widehat{P}(y = c|\boldsymbol{x})}{P(y = c|\boldsymbol{x}, \alpha \in [\alpha_i, \alpha_i + \tau))^{\beta_i}} \tag{10}$$

where $P(y|\boldsymbol{x}, \alpha \in [\alpha_i, \alpha_i+\tau))$ is the empirical marginal label distribution from users whose threshold α falls in the range $[\alpha_i, \alpha_i + \tau)$. τ is a small real number interval that reflects a significant change in label distribution when the preference varies. For example, $\forall\alpha \in [1.0, 1.1)$ is deemed the same preference, and differs from $\forall\alpha \in [1.1, 1.2)$. For brevity, we use $\alpha = \alpha_i$ to denote an established humor preference. β_i is a user-specific scaling factor tuned on the validation set. Although this introduces one tunable hyperparameter for every user, the sensitivity analysis in Section 7.2 shows that it is necessary to tune β_i only when α_i takes on extreme values, which result in very imbalanced label distributions. The new probability is computed using the softmax function:

$$\tilde{P}(y = c|\boldsymbol{x}, \alpha = \alpha_i) = \frac{\exp(v_c)}{\sum_{j=1}^{C} \exp(v_j)}. \tag{11}$$

Intuitively, this approach can penalize predictions that are too confident about a certain class due to its dominant frequency in the observed samples. It is based on the heuristic that users may possess different levels of arousal in response to the same humorous content. Those having *easy-to-amuse* or *hard-to-amuse* personalities can result in a drastically unbalanced dataset. For example, given $\alpha = \alpha_i$, if the predicted probability is $\widehat{P}(y = 1|\boldsymbol{x}) = 0.9$ while the empirical probability is $P(y = 1|\boldsymbol{x}) = 0.7$, the denominator serves as a punishment to adjust the prediction to be $\tilde{P}(y = 1|\boldsymbol{x}) \approx 0.72$ for $\beta = 1$ (computed using Equations (10) and (11)), which is closer to the real probability.

5.4 Federated Model Selection

After completing each round of FL training, the global FL model will be tested on the validation set on each user's device, $\mathcal{D}^{val} = \{\boldsymbol{X}^{val}, Y^{val}|\alpha\}$. In this stage, $\widehat{p}(y^{val}|\boldsymbol{x}^{val}, \alpha)$ is the estimated probability distribution over the implicit labels in the validation set by the model. The overall validation loss for each user is calculated as follows:

$$\mathcal{L}(\boldsymbol{\theta}; \alpha_i) = -\frac{1}{|\mathcal{D}^{val}|} \sum_{(\boldsymbol{x},y)\in\mathcal{D}^{val}} [y\log(\widehat{p}(y|\boldsymbol{x}, \alpha_i)) + (1 - y)\log(1 - \widehat{p}(y|\boldsymbol{x}, \alpha_i))]. \tag{12}$$

The validation performance across all users is calculated as follows:

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{m} \sum_{i=1}^{m} \mathcal{L}(\boldsymbol{\theta}; \alpha = \alpha_i) \tag{13}$$

The best global model $\boldsymbol{\theta}$ is selected based on the lowest validation loss, $\mathcal{L}(\boldsymbol{\theta})$, among all federated rounds.

6 EXPERIMENT SETUP

In this section, we present the details of our experimental setting. We first introduce a real-world humor recognition dataset and the preparation for personalized humor recognition. Then, we

Algorithm 1: FedHumor. The m users are indexed by i; D_{train} denotes training set and D_{val} denotes validation set; α_i is the humor preference of user i; β_i indicates diversity adaptation for user i; $\mathcal{B}$ denotes a mini-batch of D_{train}; E is local training epochs and T is total federated training rounds; η_t is the learning rate at each round t.

Initialize global model parameters $\boldsymbol{\theta}$;
for $t = 1, 2, ..., T$ **do**
  for *each user* $i = 1, 2, ..., m$ **in parallel do**
    $\boldsymbol{\theta}_i^{t+1} \leftarrow$ ClientUpdate $(i, \boldsymbol{\theta}^t, \beta_i)$;
  $\boldsymbol{\theta}^{t+1} \leftarrow \frac{1}{m} \sum_{i=1}^{m} \boldsymbol{\theta}_i^{t+1}$: global parameters update (Eq. (9));
  for *each user* $i = 1, 2, ..., m$ **in parallel do**
    $\mathcal{L}(\boldsymbol{\theta}^{t+1}; \alpha_i, t) \leftarrow$ Inference$(i, \boldsymbol{\theta}^{t+1})$;
  $\mathcal{L}(\boldsymbol{\theta}^{t+1}, t) = \frac{1}{m} \sum_{i=1}^{m} \mathcal{L}(\boldsymbol{\theta}^{t+1}; \alpha_i, t)$ (Eq. (13));
Output: $\boldsymbol{\theta}^* \leftarrow \min(\mathcal{L}(\boldsymbol{\theta}^{t+1}, t))$;

ClientUpdate *(i,θ,β)* :
  Prepare $D_{train} = \{\boldsymbol{X}, Y_i | \alpha_i\}$;
  for $j = 1, 2, ..., E$ **do**
    for $\mathcal{B}$ *in* D_{train} **do**
      $\widehat{p} = f_{\boldsymbol{\theta}}(\boldsymbol{x}; \alpha_i)$: make predictions;
      $\widetilde{p} = \text{softmax}(\widehat{p} / p(y|\alpha_i)^{\beta_i})$: adapt predictions (Eq. (10), (11));
      $\mathcal{L}(\boldsymbol{\theta}; \alpha_i) = \frac{1}{|\mathcal{B}|} \sum_{\mathcal{B}} \ell(\widetilde{p}, y; \boldsymbol{\theta})$: compute training loss (Eq. (8));
      $\boldsymbol{\theta}_i \leftarrow \boldsymbol{\theta} - \eta_t \bigtriangledown_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}; \alpha_i)$: local parameters update;
  return $\boldsymbol{\theta}_i$

Inference *(i, θ)*
  Prepare $D_{val} = \{\boldsymbol{X}, Y_i | \alpha_i\}$;
  $\widehat{p} = f_{\boldsymbol{\theta}}(\boldsymbol{x}; \alpha_i)$: make predictions;
  $\mathcal{L}(\boldsymbol{\theta}; \alpha_i) = \frac{1}{|\mathcal{D}_{val}|} \sum_{\mathcal{D}_{val}} \ell(\widehat{p}, y; \boldsymbol{\theta})$: validation loss (Eq. (12));
  return $\mathcal{L}(\boldsymbol{\theta}; \alpha_i)$

introduce the evaluation metrics used for performance comparison. Finally, we describe our model setting.

6.1 Dataset Description

In our experiments, we use a newly published dataset from SemEval-2020 shared Task 7[1] - assessing the funniness of edited news headlines [16], which contains humorous headlines created based on the incongruity theory. Each created headline was sent to five crowdsourced annotators through Amazon Mechanical Turk and each annotator was required to rate its funniness using a score from the integer interval [0, 1, 2, 3]. For example, a news headline *"Royal wedding: Meghan's dress in detail"* was micro-edited by replacing *dress* with *elbow* to produce a funny version *"Royal wedding: Meghan's elbow in detail"*, which received five ratings: 0, 1, 3, 3 and 3. The original dataset made a consensus assumption on the perceived funniness of the edited headline. They reassigned an average of the five ratings, 2.0, to be the funniness rating for this joke. This example shows a

[1]SemEval 2020 Task 7 (Sub-task 1) dataset downloaded from: https://github.com/n-hossain/semeval-2020-task-7-humicroedit

common assumption of an average user preference does not reflect the diversity among different users. The statistics of the dataset are summarized in Table 1.

Table 1. Statistics of the public dataset

	Train	Validation	Test
Number of samples	9,652	2,419	3,024
Average Rating	0.936	0.935	0.940
Minimum Rating	0.000	0.000	0.000
Maximum Rating	3.000	3.000	2.800

6.2 Implicit Label Generation

Jokes can be perceived differently and thus, receive different funniness ratings from users. Unfortunately, there are no public datasets for training a NLP model under the problem of label distribution shift across people. It is also difficult to collect true feedback from a diverse population in a privacy protection setting. To simulate the diversity in the human perception of jokes, we generate a synthetic binary humor recognition dataset to study this problem from the SemEval ratings dataset.

(1) First, sort the jokes by their original average funniness ratings. Specifically, given a set of n jokes rated by the content publisher with funniness ratings drawn from a non-negative ordinal interval $I = [s_{min}, s_{max}]$, e.g., $I = [0, 3]$. $\mathcal{D} = \{(\boldsymbol{x}_j, s_j))\}_{j=1}^{n}$, where $s_j \in I$ is the original dataset with ratings sorted in ascending order (Figure 5(a)). Users receiving this dataset would agree on the jokes with funniness ratings that are close to the boundaries of the rating interval (i.e. not funny (s_{min}) and very funny (s_{max})), and may disagree on the jokes with ratings lying in between (i.e. slightly funny and moderately funny ($s_{min} < s_j < s_{max}$)).
(2) Second, assume that every user can have only one unique humor preference, denoted as α_i. Based on a user's historical implicit feedback, the user's humor preference, α_i, is reflected as a specific threshold within the funniness rating interval, i.e., $s_{min} < \alpha_i < s_{max}$. The user dislikes jokes below this threshold rating, $s_j < \alpha_i$, and likes jokes above it, $s_j \geq \alpha_i$ (Figure 5(b)). In effect, due to the limited number of training samples observed for every funniness rating, $\forall \alpha_i \in [\alpha_i, \tau)$ is treated as the same humor preference when τ is small enough.

Based on the first assumption, we create a set of implicit labels, $Y_i = \{y^j\}_{j=1}^{n}$ for users with α_i on the same jokes $\boldsymbol{X}$. Based on the second assumption, we can create multiple sets of implicit binary labels of diverse distributions $p(y|\boldsymbol{x}, \alpha_i)$, where $y_j = 0$ for $s_j < \alpha_i$ and $y_j = 1$ for $s_j \geq \alpha_i$.

To simulate the real-world diversity of perceived funniness, we generate a diverse population with humor preferences α ranging from $\alpha = 0.2$ to $\alpha = 2.0$ with a step size of 0.1 (i.e., $\tau = 0.1$), as illustrated in Figure 6. The higher the α value, the more non-humorous labels will be generated. In our experiments, a valid user is regarded as having $\alpha_i \in [0.2, 2.0]$. Specifically, the funniness ratings of the dataset are transformed into implicit binary labels generated by a given user with the preference α on his own device. The distribution over the generated implicit labels is $p(y|\boldsymbol{x}, \alpha)$. The boundary values (e.g., $\alpha = 0$ and $\alpha = 3$) that result in a one-class dataset are not considered in our experiments.

6.3 Evaluation Metrics

To evaluate the classification performance of FedHumor and baseline approaches, we use the *macro-averaged* F_1 score as our main metric, which is the average of the F_1 scores of two classes.

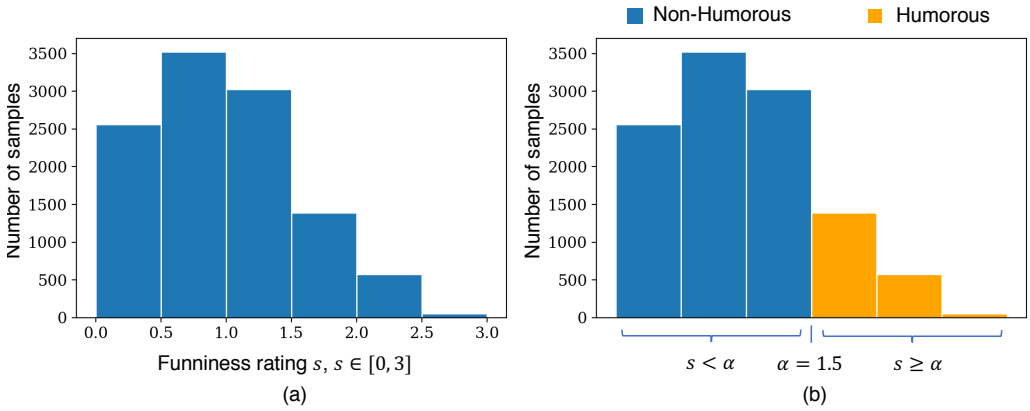

Fig. 5. Transform explicit ratings into binary labels. The distribution of explicit funniness ratings on a set of jokes rated by the content publisher is shown in discrete intervals (a). An example in which a user's humor preference is quantified by $\alpha = 1.5$ (b). *Best viewed in color.*

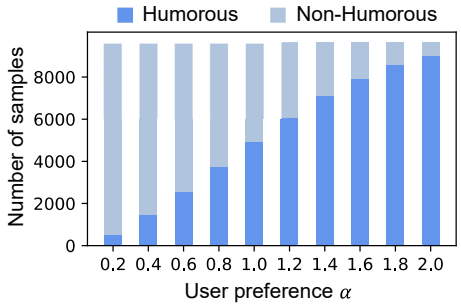

Fig. 6. Generated binary labels with different distributions from funniness ratings when diverse humor preferences are considered. *Best viewed in color.*

We additionally report the *macro-averaged* precision and recall scores. Note that a *macro-averaged* metric implies that the generalization performance on all classes are equally important. This is essential in personalized humor recognition, because some users with very low or very high threshold of humor preference may produce very imbalanced labels. In this case, a model can achieve an artificially high accuracy by predicting all the samples to be the dominant class. Thus we do not use the accuracy metric for evaluation. All the methods are compared on the hold-out test set and their predicted implicit labels for each class are counted to calculate the *macro-averaged* metrics.

6.4 Model Setting

FedHumor is built based on top of the base version of BERT [9]. The pretrained weights we utilize include its word embedding layer and twelve transformer layers[2]. A pooling layer nonlinearized

[2]PyTorch implementation of the pretrained BERT model : https://huggingface.co/transformers/pretrained_models.html

by *tanh* activation function[3] is applied on the first ([CLS]) token representation from the last transformer layer to be the final contextualized sentence representation. A dropout of 0.1 is applied on the pooled output and is further sent to a classifier, which is a fully connected layer randomly initialized from a uniform distribution, $\mathcal{U}(-\sqrt{1/d}, \sqrt{1/d})$, where d is the size of the input features[4]. In our model, $d = 768$, the length of the hidden vector corresponding to the [CLS] token. We follow the recommended hyperparameter setting introduced in the paper [9] to tune learning rate, batch size and weight decay for fine-tuning BERT on downstream tasks based on Federated Averaging setting[5].

7 PERFORMANCE EVALUATION

In this section, we conduct experiments to evaluate the performance of FedHumor. We compare it against nine state-of-the-art humor recognition methods on the hold-out test set. We also study the properties of FedHumor through experiments and compare it with other learning approaches to show the advantages of following the federated learning paradigm.

7.1 Comparison of Different Training Strategies

Given that the property of our data is Differently and independently distributed, we first compare different training strategies against the federated weight-tying method.

(1) **INDV**: A simple approach to this problem is to fine-tune a BERT model for each *individual* user based on their local ground truth labels. This method represents the traditional way of providing personalized humor recognition, which creates a unique model for each user. Yet, it's computationally expensive and might easily overfit to small imbalanced datasets in practice. Moreover, when new users registered, server has to establish new model to be trained for them.
(2) **AGG**: The second approach is to first *aggregate* all the labeled data from each user into a central database, and treat all of them as ground truth labels to train a BERT-based humor recognition model. This follows the traditional way of training a single-task machine learning model. However, this approach needs to have access to user's data, which may cause privacy breach.
(3) **FED**: he third approach is to tie the weights of each individual BERT model through *federated learning*. In this setting, users can share learned BERT parameters while keeping their personal adaptation module.

To test these approaches under different levels of preference diversity, we prepare two groups of user preferences as shown below. Due to the limit of dataset size and the funniness score range, the maximum number of different user preferences we can generate from the given dataset is 18.

- Group 1: a group of three users with different preferences: $\alpha = 0.3$ represents the easy-to-amuse personality, $\alpha = 0.9$ represents the neutral personality, and $\alpha = 1.8$ represents the hard-to-amuse personality;
- Group 2: a group of 18 users with different preferences: α ranging from 0.2 to 1.9 in increments of 0.1, representing a more diverse population.

We compare the average test performance of a group between the three models. Comparison Results are shown in Table 2. Take the column of F score as an example. Aggregated training approach works the worst, meaning that training data with conflict labels can degrade the performance. Individual training comes after, meaning that such issue can be alleviated by separating

[3]PyTorch implementation of the tanh activation function: https://pytorch.org/docs/stable/generated/torch.nn.Tanh.html
[4]PyTorch initialization of linear layer: https://pytorch.org/docs/stable/generated/torch.nn.Linear.html
[5]Code downloaded from: https://github.com/shaoxiongji/federated-learning

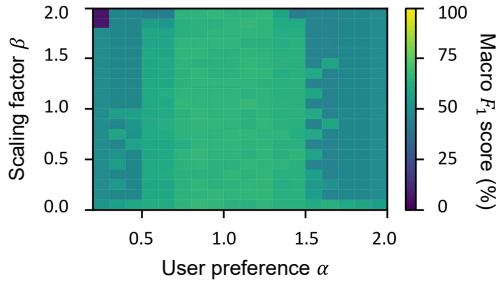

Fig. 7. Tune hyper-parameter β w.r.t. α for single user. *Best viewed in color.*

them. Federated training further improves the performance, meaning that allowing weight-tying between different users in the federated setting can bring performance gains. Comparing the two groups, the more diverse the users, the better the Federated Learning approach performs compared to the rest.

Table 2. Average test performance (in %) of three learning strategies on two groups of users. Values in bold denote the best results. Underlined values indicate the second best results.

		Precision	Recall	F_1 score
Group 1	AGG	<u>58.59</u>	54.89	41.66
	INDV	56.30	<u>55.32</u>	<u>53.52</u>
	FED	**60.03**	**65.57**	**55.61**
Group 2	AGG	57.40	51.25	33.05
	INDV	<u>58.14</u>	<u>55.61</u>	<u>53.03</u>
	FED	**61.67**	**66.62**	**57.48**

7.2 Hyperparameter Sensitivity Analysis

As introduced in Section 5.3, we allow a hyperparameter β on each user to control how much we want to scale the model's estimated probabilities $P(y|x, \alpha_i)$ w.r.t. the marginal local distribution $P(y|x, \alpha_i)$. As such, we conduct a sensitivity analysis through grid search to reveal whether it affects the performance of FedHumor and how to set β for each user. To do so, we vary the values of α from low ($\alpha = 0.2$) to high ($\alpha = 2.0$) and at the same time, increase the scale factor β from small ($\beta = 0$) to large ($\beta = 2.0$), all in increments of 0.1. This results in 399 experiments in total. The implicit binary labels are generated each time the α value changes. We train the BERT-based classification model for each combination of α and β and the best corresponding model is selected on the validation set. We report the macro-averaged F_1 score on the hold-out test set.

The results on the test set are shown in Figure 7. In general, a model can achieve better generalization performance on the preference range of $[0.5, 1.5]$. From Figure 6, we can see that this preference interval leads to the range of empirical probability from $P(y = 1|\boldsymbol{x}, \alpha = 0.5) \approx 0.2$ to $P(y = 1|\boldsymbol{x}, \alpha = 1.5) \approx 0.75$. In this case, for every α value, there are multiple β values that can achieve the best performance. This means that our model is not very sensitive to β and a reasonable value would suffice. For users who have very low ($\alpha < 0.5$) or very high ($\alpha > 1.5$) humor preference

values, their labels are too imbalanced for the model to achieve good generalization performance. In this case, β is recommended to set to a low value (e.g., around 0.1).

7.3 Comparison of Different Humor Recognition Models

Finally, we compare our model with three kinds of baselines. The first three models are the representative humor recognition models in literature, which do not use pretrained language model. The second six models are different versions of pretrained language models with fine-tuning. Our model use both a pretrained language model and federated training on diverse humor preferences.

(1) **DV-LR**: We trained Doc2Vec using the distributed bag of words approach [29] for sentence representation and applied logistic regression classifier on the features for classification.
(2) **WV-RF**: We reproduced the humor recognition model in [39], which used a pretrained word2vec[6] model for sentence representation and a Random Forest classifier for classification.
(3) **WV-CNN-HN**: We reproduced the deep learning-based humor recognition model in [7], which augmented a CNN with a Highway Network for end-to-end humor recognition[7].
(4) **BERT-FZ**: We use the BERT model pretrained on lower-cased English text with 768 hidden size and 12 transformer layers. We *freeze* the pretrained model parameters and only update the parameters of a fully connected classification layer.
(5) **BERT-FT**: Different from the previous method, we *fine-tuned* the pretrained parameters together with the classifier parameters. This is a strong baseline as many downstream NLP tasks have shown improved performance through fine-tuning BERT on the task-specific datasets [9]. It is also the same base model adopted by FedHumor.
(6) **BERT-L/C/M**: Advanced by huge model capacity trained on large corpora, pretrained language models can provide generally effective representations across domains. As such, we adopt other versions of BERT. BERT-L is pretrained with a much *larger* hidden size (1024) and deeper (24) layers. BERT-C means the true *case* and accent markers are preserved when training BERT. BERT-M means a *multilingual* BERT model that was pretrained on lower-cased text in the top 102 languages with the largest Wikipedia. The three models are all *fine-tuned* together with the classifier to form three baseline approaches.
(7) **ALBERT**: There are also many strong pretrained language models after BERT. We adopt ALBERT [18], which had been evaluated to have better scalability on downstream tasks than BERT. We fine-tune its pretrained parameters together with a classification layer on the humor recognition task for comparison.

Test Results are shown in Table 3. The metrics are the macro average values over all classes. Generally, pretrained language models with domain adaptive fine-tuning outperform all traditional models. Our model trained based on the relaxed assumption, exceeds all baselines across metrics.

The first three models (i.e., DV-LR, WV-RF, WV-CNN-HN) use static word representations which were widely adopted in humor recognition tasks before the advent of large pretrained language models. Their performance are generally not good on this dataset compared to pretrained language model-based methods. The following six models (i.e., BERT-FZ, BERT-FT, BERT-L/C/M, ALBERT) are BERT variants. The results achieved by BERT-FZ are much worse than the rest. This shows that the pretrained language models, whose learning phase does not take advantage of labeled data, can only provide general representations. They should be fine-tuned towards a particular task by further learning from domain-specific label information. BERT-L is close to but lower than BERT-FZ showing that larger model parameters do not bring performance gain. BERT-C and BERT-M are lower than BERT-FZ indicating the preservation of the true case and knowledge from

[6]Google word2vec downloaded from: https://code.google.com/p/word2vec/
[7]PyTorch implementation of Highway Network: https://gist.github.com/dpressel/3b4780bafcef14377085544f44183353

Table 3. Test performance (in %) on the user with $\alpha = 1.0$ achieved by FedHumor and all baseline approaches. Values in bold indicate the best results. Underlined values indicate the second best results.

	Precision	Recall	F_1 score
DV-LR	53.69	53.67	53.64
WV-RF	56.70	56.10	55.20
WV-CNN-HN	56.20	54.70	51.90
BERT-FZ	54.15	53.71	52.53
BERT-FT	<u>64.91</u>	<u>64.88</u>	<u>64.87</u>
BERT-L	64.48	64.48	64.47
BERT-C	62.69	62.65	62.62
BERT-M	62.11	62.08	62.06
ALBERT	61.06	61.05	61.04
FedHumor	**66.60**	**66.56**	**66.53**

other languages did not benefit this task. ALBERT, though tested to be more scalable to downstream tasks than BERT on some datasets, also failed to outperform BERT on this task. FedHumor achieved the best results across all the three evaluation metrics and 2.5% relative improvement on F_1 score than the second best result. This reveals that it is not necessary to reconcile the label difference between users. By leveraging the diversity and training the model in a distributed manner, we can improve the model's performance in personalized humor recognition.

8 CONCLUSIONS AND FUTURE WORK

In this paper, we propose FedHumor - a humorous text recognition model following the federated learning paradigm - to perform personalized humor recognition based on labels from distributed data sources. FedHumor is able to account for the diversity in each person's humor preference in the prediction of perceived humorous content. Through extensive experiments, we compared FedHumor against nine state-of-the-art approaches and evaluated the impact of the federated training approach. The results show that it is able to achieve better generalization performance in humor recognition. To the best of our knowledge, FedHumor is the first federated learning-based personalized humor recognition model.

The conceptualization and evaluation of personalized humor recognition under the FL paradigm has many practical values. In FL, data heterogeneity can sometimes reflect users' preferences in highly subjective tasks such as humor recognition and sarcasm detection [14]. Through FedHumor, we show that preserving and adapting to data heterogeneity can bring performance gains in the face of diverse user preferences. Through joint training of the text encoder and adapting the predictions for each user, FedHumor can be used to recommend personalized jokes for each user on their edge devices during typing or reading.

Nevertheless, the studies are limited by the availability of suitable real-world datasets. The evaluation setting is designed to cover all kinds of user preferences. However, in practice, we may only be able to collect datasets labeled according to a partial set of user preferences. In addition, the disparity in labels not only comes from people with different humor preferences, but can also result from time-delayed awareness. For example, [31] observed that the same joke could be labeled differently even by the same person after some time. This problem relates to the research field of concept shift [22], which has made a lot of progress in adapting the model to changing target

functions over time [5, 11]. Currently, there are no publicly available datasets for studying the concept shift problem in humor perception. In the future work, we will collect the appropriate datasets for studying the problem in which the objective function for the same person changes over time.

ACKNOWLEDGMENTS

This work is supported, in part, by Alibaba Group through the Alibaba Innovative Research (AIR) Program and the Alibaba-NTU Singapore Joint Research Institute (JRI), Nanyang Technological University, Singapore; the Nanyang Assistant/Associate Professorships (NAP); The RIE 2020 Advanced Manufacturing and Engineering Programmatic Fund (No. A20G8b0102), Singapore; NTU-SDU-CFAIR (NSC-2019-011); the National Natural Science Foundation of China under Grant NSFC 62106167; the National Research Foundation, Prime Minister's Office, Singapore through the AI Singapore Programme (AISG2-RP-2020-019), NRF Investigatorship Programme (NRFI Award No. NRF-NRFI05-2019-0002) and NRF Fellowship (NRF-NRFF13-2021-0006). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

REFERENCES

[1] Manoj Ghuhan Arivazhagan, Vinay Aggarwal, Aaditya Kumar Singh, and Sunav Choudhary. 2019. Federated Learning with Personalization Layers. *CoRR, arXiv:1912.00818* (2019).

[2] Muhammad Asad, Ahmed Moustafa, and Takayuki Ito. 2020. FedOpt: towards communication efficiency and privacy preservation in federated learning. *Applied Sciences* 10, 8 (2020), 2864.

[3] Simge Aykan and Erhan Nalçacı. 2018. Assessing Theory of Mind by Humor: The Humor Comprehension and Appreciation Test (ToM-HCAT). *Frontiers in psychology* 9 (2018).

[4] Kim Binsted, Anton Nijholt, Oliviero Stock, Carlo Strapparava, G Ritchie, R Manurung, H Pain, Annalu Waller, and D O'Mara. 2006. Computational humor. *IEEE Intelligent Systems* 21, 2 (2006), 59–69.

[5] Andrés Cano, Manuel Gómez-Olmedo, and Serafín Moral. 2019. A Bayesian approach to abrupt concept drift. *Knowledge-Based Systems* 185 (2019), 104909.

[6] Lei Chen and Chong MIn Lee. 2017. Convolutional neural network for humor recognition. *CoRR, arXiv:1702.02584* (2017).

[7] Peng-Yu Chen and Von-Wun Soo. 2018. Humor recognition using deep learning. In *NAACL*. 113–117.

[8] Yiqiang Chen, Xiaodong Yang, Xin Qin, Han Yu, Biao Chen, and Zhiqi Shen. 2020. FOCUS: Dealing with Label Quality Disparity in Federated Learning. *arXiv preprint arXiv:2001.11359* (2020).

[9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. https://doi.org/10.18653/v1/N19-1423

[10] Alireza Fallah, Aryan Mokhtari, and Asuman E Ozdaglar. 2020. Personalized Federated Learning with Theoretical Guarantees: A Model-Agnostic Meta-Learning Approach.. In *NeurIPS*.

[11] Isvani Frías-Blanco, Jose del Campo-Avila, Gonzalo Ramos-Jiménez, Andre CPLF Carvalho, Agustín Ortiz-Díaz, and Rafael Morales-Bueno. 2016. Online adaptive decision trees based on concentration inequalities. *Knowledge-Based Systems* 104 (2016), 179–194.

[12] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. 2020. Shortcut learning in deep neural networks. *Nature Machine Intelligence* 2, 11 (2020), 665–673.

[13] Charles R. Gruner. 1997. The Game of Humor: A Comprehensive Theory of Why We Laugh. *Transaction Publishers* (1997).

[14] Xu Guo, Boyang Li, Han Yu, and Chunyan Miao. 2021. Latent-Optimized Adversarial Neural Transfer for Sarcasm Detection. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Online, 5394–5407. https://doi.org/10.18653/v1/2021.naacl-main.425

[15] Sonja Heintz and Willibald Ruch. 2019. From four to nine styles: An update on individual differences in humor. *Personality and Individual Differences* 141 (2019), 7–12.

[16] Nabil Hossain, John Krumm, and Michael Gamon. 2019. "President Vows to Cut Hair": Dataset and Analysis of Creative Text Editing for Humorous Headlines. *CoRR, arXiv:1906.00274* (2019).

[17] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. 2019. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977* (2019).
[18] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *CoRR, arXiv:1909.11942* (2019).
[19] Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. 2020. Federated multi-task learning for competing constraints. *arXiv preprint arXiv:2012.04221* (2020).
[20] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. 2018. Federated optimization in heterogeneous networks. *arXiv preprint arXiv:1812.06127* (2018).
[21] Lizhen Liu, Donghai Zhang, and Wei Song. 2018. Exploiting syntactic structures for humor recognition. In *Proceedings of the 27th International Conference on Computational Linguistics*. 1875–1883.
[22] Jie Lu, Anjin Liu, Fan Dong, Feng Gu, Joao Gama, and Guangquan Zhang. 2018. Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering* 31, 12 (2018), 2346–2363.
[23] Lingjuan Lyu, Jiangshan Yu, Karthik Nandakumar, Yitong Li, Xingjun Ma, Jiong Jin, Han Yu, and Kee Siong Ng. 2020. Towards Fair and Privacy-Preserving Federated Deep Models. *IEEE Transactions on Parallel and Distributed Systems* 31, 11 (2020), 2524–2541.
[24] Jihang Mao and Wanli Liu. 2019. A BERT-based Approach for Automatic Humor Detection and Scoring. In *IberLEF*.
[25] Rod A Martin, Patricia Puhlik-Doris, Gwen Larsen, Jeanette Gray, and Kelly Weir. 2003. Individual differences in uses of humor and their relation to psychological well-being: Development of the Humor Styles Questionnaire. *Journal of research in personality* 37, 1 (2003), 48–75.
[26] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *AISTATS*. 1273–1282.
[27] Rada Mihalcea and Stephen Pulman. 2007. Characterizing humour: An exploration of features in humorous texts. In *International Conference on Intelligent Text Processing and Computational Linguistics*. Springer, 337–347.
[28] Rada Mihalcea and Carlo Strapparava. 2005. Making computers laugh: Investigations in automatic humor recognition. In *EMNLP*. 531–538.
[29] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NeurIPS*. 3111–3119.
[30] Tristan Miller, Christian F Hempelmann, and Iryna Gurevych. 2017. Semeval-2017 task 7: Detection and interpretation of english puns. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*. 58–68.
[31] Alejandro Mottini and Amber Roy Chowdhury. 2019. What Do You Mean I'm Funny? Personalizing the Joke Skill of a Voice-Controlled Virtual Assistant. *arXiv preprint arXiv:1912.03234* (2019).
[32] Peter Potash, Alexey Romanov, and Anna Rumshisky. 2017. SemEval-2017 Task 6: #HashtagWars: Learning a Sense of Humor. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*. Association for Computational Linguistics, Vancouver, Canada, 49–57. https://doi.org/10.18653/v1/S17-2004
[33] Swaroop Ramaswamy, Rajiv Mathews, Kanishka Rao, and Françoise Beaufays. 2019. Federated learning for emoji prediction in a mobile keyboard. *CoRR, arXiv:1906.04329* (2019).
[34] Thomas R Shultz. 1976. A cognitive-developmental analysis of humour. (1976).
[35] Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet Talwalkar. 2017. Federated multi-task learning. *arXiv preprint arXiv:1705.10467* (2017).
[36] Julia M Taylor and Lawrence J Mazlack. 2004. Computationally recognizing wordplay in jokes. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 26.
[37] Tony Veale. 2012. *Exploding the creativity myth: The computational foundations of linguistic creativity*. A&C Black.
[38] David J Wimer and Bernard C Beins. 2008. Expectations and perceived humor. *Humor* 21, 3 (2008), 347–363.
[39] Diyi Yang, Alon Lavie, Chris Dyer, and Eduard Hovy. 2015. Humor recognition and humor anchor extraction. In *EMNLP*. 2367–2376.
[40] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. 2019. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)* 10, 2 (2019), 1–19.
[41] Qiang Yang, Yang Liu, Yong Cheng, Yan Kang, Tianjian Chen, and Han Yu. 2019. *Federated Learning*. Morgan & Claypool Publishers. 189 pages.
[42] Han Yu, Zelei Liu, Yang Liu, Tianjian Chen, Mingshu Cong, Xi Weng, Dusit Niyato, and Qiang Yang. 2020. A fairness-aware incentive scheme for federated learning. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 393–399.
[43] Renxian Zhang and Naishi Liu. 2014. Recognizing humor on twitter. In *CIKM*. 889–898.