

Co-training for Demographic Classification Using Deep Learning from Label Proportions

Ehsan Mohammady Ardehaly
Department of Computer Science
Illinois Institute of Technology
Chicago, Illinois 60616
Email: emohamm1@hawk.iit.edu

Aron Culotta
Department of Computer Science
Illinois Institute of Technology
Chicago, Illinois 60616
Email: aculotta@iit.edu

Abstract—Deep learning algorithms have recently produced state-of-the-art accuracy in many classification tasks, but this success is typically dependent on access to many annotated training examples. For domains without such data, an attractive alternative is to train models with light, or distant supervision. In this paper, we introduce a deep neural network for the Learning from Label Proportion (LLP) setting, in which the training data consist of bags of unlabeled instances with associated label distributions for each bag. We introduce a new regularization layer, **Batch Averager**, that can be appended to the last layer of any deep neural network to convert it from supervised learning to LLP. This layer can be implemented readily with existing deep learning packages. To further support domains in which the data consist of two conditionally independent feature views (e.g. image and text), we propose a co-training algorithm that iteratively generates pseudo bags and refits the deep LLP model to improve classification accuracy. We demonstrate our models on demographic attribute classification (gender and race/ethnicity), which has many applications in social media analysis, public health, and marketing. We conduct experiments to predict demographics of Twitter users based on their tweets and profile image, without requiring any user-level annotations for training. We find that the deep LLP approach outperforms baselines for both text and image features separately. Additionally, we find that co-training algorithm improves image and text classification by 4% and 8% absolute F1, respectively. Finally, an ensemble of text and image classifiers further improves the absolute F1 measure by 4% on average¹.

I. INTRODUCTION

Deep learning methods have produced state-of-the-art accuracy on many different classification tasks especially for image classification [1], [2], [3], [4], [5]. Because these networks typically have millions of parameters, they rely on access to large labeled datasets such as ImageNet [6], which has over one million annotated images. While transfer learning can help adapt a network to a new domain [7], it still requires labeled data on the target domain.

An attractive alternative is Learning from Label Proportion (LLP), in which the training samples are divided into a set of bags, and only the label distribution of each bag is known. The main advantage of LLP is that it does not require annotations for individual instances. Furthermore, in many domains label proportions are readily available — for example,

by associating geolocated social media messages with county population statistics, we can fit a model of demographics without annotating individual users.

While many LLP models have been proposed based on the logistic hypothesis [8], SVM [9], and graphical models [10], there has been little work that considers deep learning. In this paper, we propose an approach which converts a deep neural network from a supervised classifier to LLP. This method can readily be implemented in popular deep learning packages by introducing a new regularization layer into the network. We propose such a layer, called the **Batch Averager**. Similar to label regularization [8], this layer computes the average of its input as the output. Like other regularization layers, this layer is only applied at training time. The **Batch Averager** is typically appended to the last layer of a network to convert it from a supervised learning to a deep LLP. We use KL-Divergence as the error function to train the network to produce predictions that match the provided label proportions.

This deep LLP framework has these key advantages:

- **Simplicity**: The **Batch Averager** is based on a very simple tensor operation (average) and, unlike the **Batch Normalizer** layer [11], it does not have any training weights and does not noticeably affect training time.
- **Compatibility**: The framework is fully compatible with almost every deep learning package, and requires only a few line of codes to implement. We have tested the approach with both Theano [12] and Tensorflow [13].
- **Availability**: The framework can convert almost every supervised classification network to LLP.
- **Accuracy**: Our empirical results indicate that the deep LLP has a comparable accuracy to supervised models subject to proper constraints to generate bags.

In addition, we also propose co-training with deep LLP to further improve accuracy. In some applications, multiple views of the data are available (e.g. text and image). For example, many social media sites contain both image and textual data. While there have been many deep learning methods that directly combine text and image features (e.g., [14]), we instead require a model that is robust to cases where one feature view is missing. For example, some users may not post images, yet we would still like to classify based on the text.

¹An expanded version of this paper is available at: <http://arxiv.org/abs/1709.04108>

An attractive alternative is co-training [15]. Because traditional co-training requires labeled data, we propose a new algorithm that is more suitable for the LLP setting. In this method, we use an image-based LLP model to create bags with estimated label proportions, which we call *pseudo-bags*. We then fit a text-based LLP model on these pseudo-bags, and use it to in turn create pseudo-bags for the image-based model.

We conduct experiments classifying Twitter users into demographic categories (gender and race/ethnicity) based on the profile image and tweets. We find that the proposed deep LLP framework outperforms both supervised and LLP baselines, and that co-training further improves accuracy of the models.

The remainder of the paper is organized as follows. In Section II, we review related work on LLP and co-training, and Section III provides our deep LLP framework and co-training algorithm. In Section IV we describe the data collected for the experiments. In Section V we present our empirical results; Section VI concludes and describes plans for future work.

II. RELATED WORK

Deep learning methods have produced state-of-the-art results on many classification tasks, but typically require many labeled training instances. For image classification, researchers usually compare their model by training on large datasets such as ImageNet [6]. For example, Xception (2016) network introduces extreme inception module by using deepwise separable convolution layers [16] and achieves 94.5% top-5 accuracy [3]. These models highly rely on the vast amount of labeled data. However, with transfer learning, the pre-trained weights (using ImageNet) can move to other image classification tasks [7]. But, still we need annotated data on the target domain, and alternative approaches such as LLP are required in the absence of labeled data.

Only a handful of researchers attempt to use deep learning for LLP settings. Kotzias et al. [17] propose a model for the particular case of LLP when the label of bags are available. For example, for text classification task, when we have the label of a bag of multiple sentences, they propose a model to infer the label of each sentence using an objective function to smooth the posterior probability of samples based on sample similarity and bag constraints. They use a convolutional neural network to infer sentence similarity for text classification.

Other researchers provide models for image segmentation. In this setting, each image is split into multiple small images (as a bag), and the classifier tries to infer labels of these regions using the label proportion of the bag. Li et al. [18] suggest a convolutional neural network with a probabilistic method to estimate labels by considering the proportion bias, using the Expectation Maximization algorithm to determine model parameters. In their approach, they use satellite images with know ice area ratio to train a classifier to predict label of small segments of images.

While these methods have promising results on a particular domain, to the best of our knowledge, no method has proposed a framework that can be readily applied to diverse classifica-

tion tasks. Inspired by label regularization [8], we fill this gap by introducing **Batch Averager** as a regularizer layer.

The traditional L2 regularization appears to not be sufficient for deep neural network because overfitting is a severe problem in these networks. Srivastava et al. [19] introduce a Dropout layer that randomly drops some units in backpropagation step and show that it can significantly reduce overfitting. Furthermore, the Batch Normalizer is introduced to normalize the output of cells and reduce internal covariate shift of weights [11]. Similar to these two normalization layers, our proposed **Batch Averager** layer applies only at training time (Batch Normalizer uses the moving average and standard deviation to normalize the output at testing time without updating the moving average and standard deviation).

Recently, demographic classification with deep learning has been proposed by researchers. Zhang et al. [20] offer a method to infer demographic attributes (gender) from the wild (unconstrained images). Similarly, Liu et al. [21] propose convolutional neural network to classify attributes such as age, gender, and race in the wild. Ranjan et al. [22] provide a fast RCNN to localize and recognize all faces in the scene with their attributes (e.g. gender, pose). Few attempt to use both textual and image features together.

Co-training is a semi-supervised method that trains on two views of features on a small set of labeled data and iteratively adds pseudo labels from a large set of unlabeled data [15]. Gupta et al. [23] demonstrate a co-training algorithm that trains on captioned image to recognize human actions with SVMs. Their model also learns from videos of athletic events with commentary.

The majority of these works use captions to leverage image classification accuracy by co-training methods, and cannot be applied to text only classification purpose. Furthermore, they require annotated labeled data. On the other hand, our proposed co-training model trains both image and text classifiers that can be used separately. Additionally, we do not need any labeled data; and if the testing sample has both image and text features, we can apply an ensemble model (by soft voting) to improve the classification accuracy.

III. MODELS

In this section, we first define our proposed regularization layer, **Batch Averager**. Then, we illustrate the deep LLP framework with this layer. Next, we provide an algorithm to create bags with appropriate label distributions for training. Finally, we offer a co-training approach for training on two views of data (e.g. text and image) in LLP settings.

A. Batch Averager Layer

In the Learning from Label Proportion setting, the training data are divided into bags, and only the label distribution for each bag is known. For bag i , let B_i be the set of samples in this bag; i.e. $B_i = \{X_{i,j}\}$ where $X_{i,j}$ is the feature vector for instance j in bag i . Let \tilde{y}_i be the provided label proportion for bag i , and T_i be the number of samples in this bag (i.e.

$T_i = |B_i|$). E.g., for binary classification, \tilde{y}_i is the fraction of positive instances in bag i .

To implement Deep LLP, we assume that each bag is assigned to exactly one batch for gradient optimization. Also, because deep neural networks are typically trained by GPU cores with limited memory, there is usually a maximum batch size determined by the number of parameters that can fit into GPU memory. As a result, if a bag is too large, we need to break it down into smaller bags.

Our work is inspired by label regularization [8], and we define a regularization layer for a neural network. The label regularization model estimates bag posterior probabilities by the average of the posterior probabilities of the instances in that bag. In a neural network, this can just be implemented by computing the average of the output of the last layer per batch. As a result, we name this layer **Batch Averager**.

Let y be the (unobserved) output of the last layer of a supervised network. Because in classification tasks, the last layer is typically **logit** function (either softmax or logistic), it returns a vector of posterior label probabilities; i.e. for bag i , instance j , and class k we have:

$$P(y = k | X_{i,j}) = y_{i,j}^{(k)} \quad (1)$$

Inspired by label regularization, we estimate the posterior probability of bag i , \bar{y}_i , as the average of posterior probability of all instances in bag i ; i.e.

$$\bar{y}_i^{(k)} = \frac{1}{T_i} \sum_{j \in B_i} P(y = k | X_{i,j}) = \frac{1}{T_i} \sum_{j \in B_i} y_{i,j}^{(k)} \quad (2)$$

We expect that the bag posterior probability (\bar{y}_i) should be close to the bag prior probability (\tilde{y}_i). Again, similar to label regularization, we use KL-divergence as the error function between posterior and prior, and the target of the neural network is to minimize this error function:

$$\Delta(\tilde{y}, \bar{y}) = \sum_k \tilde{y}^{(k)} \log \frac{\tilde{y}^{(k)}}{\bar{y}^{(k)}} \quad (3)$$

The **Batch Averager** layer is very similar to the Batch Normalizer layer [11]. The latter normalizes the output of the batch to have zero mean and unit variance; the former converts the input layer to its average. As a regularization layer, **Batch Averager** is only applied at training time, not at testing time. Typically, **Batch Averager** would be applied to the last layer of the network, immediately after the Logit layer.

Implementing the **Batch Averager** is straightforward with most current open-source deep learning packages. However, most of these packages assume that the output of the network is a tensor with the same size as batch size (T_i). As a result, we need to repeat the average for T_i times. Similarly, we need to repeat the prior (\tilde{y}_i) for T_i times. Because the number of samples per bag (batch size) is typically different per bag, we need to add sample weight $\frac{1}{T_i}$ for each sample in bag i . As a result, the entire bag has sample weight one.

More formally, for bag i , we train the network in a similar fashion as traditional supervised learning with feature

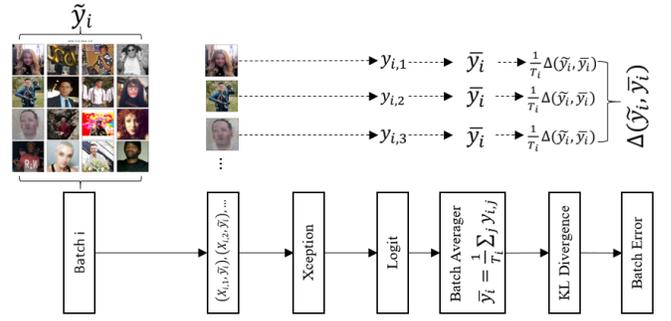


Fig. 1. The deep LLP framework applied to image classification.

vectors $[X_{i,1}, \dots, X_{i,T_i}]$, labels $[\tilde{y}_i, \dots, \tilde{y}_i]$, and sample weights $[\frac{1}{T_i}, \dots, \frac{1}{T_i}]$. We additionally use KL-divergence as the error function for back propagation.

B. Deep LLP framework

In this section, we provide a framework to implement deep LLP networks. This framework uses the **Batch Average** regularizer as the last layer, and can be applied in any classification application (e.g. text, image). We furthermore propose text classification and image classification networks.

Similar to any classification deep network, in Deep LLP framework, the input of the network is the feature vectors (i.e. textual or image feature). Then we feed the network with layers that are typically used for the classification task at hand. The next layer is the Logit (softmax) layer to compute class probabilities. Finally, we add **Batch Averager** for label regularization, which sends the average of features to the output.

For text classification, we simply use a shallow network with only one dense layer (with 16 cells), followed by a dropout layer to avoid overfitting. We also use a temperature parameter as suggested in previous work [8]. The temperature can be readily implemented in the Logit layer by a tensor operation.

For image classification, we use Xception, a state-of-the-art image classification model [3]. We use the pre-trained weights (trained on ImageNet), and freeze the first two blocks of the network to train it with the maximum batch size 32 and fit the network with color images with dimension 299×299 . Figure 1 shows a more detailed network for this model.

C. Bag creation algorithm

In some domains, bags are naturally available with soft constraints such as geolocation, allowing us to assign label proportions to groups of instances. For example, in experiments below, we attach the U.S. Census statistics of county demographics to Twitter users from the same location. However, in other domains, our constraint is a prior probability based on an attribute of an instance. For example, according to US Census, 67% of people with the last name ‘Taylor’ are white. The problem is to construct a bag that combines many of these constraints and associate an accurate label proportion for training.

Algorithm 1 : This algorithm create bags with class prior and returns them. **Parameters:** Prior probability P , class y , feature vectors X_i , threshold t , and maximum batch size N .

```

1: procedure CREATEBAGS( $P, X_i, y, t, N$ )
2:    $U \leftarrow \{X_i : P(y|X_i) > t\}$ 
3:    $U \leftarrow \text{Sorted}(U, \text{key} = P(y|X_i)) \triangleright$  Descending order.
4:    $B \leftarrow []$ 
5:    $\tilde{y} \leftarrow []$ 
6:   while  $U \neq \emptyset$  do
7:      $S \leftarrow$  first  $N$  samples in  $U$ 
8:      $B.append(S)$ 
9:      $\tilde{y}.append(\text{Average}(\{P(y|x) : \forall x \in S\}))$ 
10:     $U \leftarrow U \setminus S$ 
11:   end while
12:   return  $B, \tilde{y}$ 
13: end procedure

```

For class y and unlabeled samples $\{X_i\}$, we assume we have a prior probability for this class $P(y|X_i)$ (e.g., $P(\text{white} | \text{Taylor}) = .67$). The goal is to select instances from X_i to construct a bag, then to assign the expected proportion of that bag that have class y . Algorithm 1 shows how to create bags for this class. The algorithm takes as input a maximum bag size (N) and a threshold (t). In the results below, we set N to 64 for all experiments. The algorithm first removes samples with probability lower than the threshold t . The threshold is typically greater than .5 (in preliminary experiments, the results were not very sensitive to this parameter and we tune this threshold for each task; e.g. .8 for gender classification and .9 for race classification). Then, samples are sorted by decreasing order of their prior probability. Next, we move first N sorted samples to the first bag, and the next N samples to the next bag and continue until all remaining samples (above the threshold) are assigned to bags. Finally, we compute the bag label proportion as the average of the prior probability of samples inside the bag. These bags are then used as training data in the deep LLP model.

The advantage of this bag creation approach is that it allows us to associate label proportions with instances using pre-existing population statistics. Of course, these label proportions are likely to be inexact, and contain some selection bias because Twitter users are not representative of the population. However, several prior works have found that LLP methods are robust to this type of noise [8]; we also find this to be the case in the present work.

D. Co-training for LLP

In this section, we provide an algorithm for applying co-training in the LLP setting. This algorithm is useful when we have two conditionally independent feature views of data. An interesting case is when we have both text and image views of the data. Some views may naturally have more bags than others, or may have more accurate label proportions associated with them. Thus, our goal is to combine the advantages of each view to produce a more accurate classifier.

Algorithm 2 : The co-training algorithm for LLP settings. **Parameters:** Bags $B^{(k)}$, label proportion $\tilde{y}^{(k)}$, and unlabeled features $U^{(k)}$ for both views.

```

1: procedure COTRAININGLLP( $(B_1, \tilde{y}_1, U_1), (B_2, \tilde{y}_2, U_2)$ )
2:    $B_0 \leftarrow \emptyset$ 
3:    $\tilde{y}_0 \leftarrow \emptyset$ 
4:   while Stopping condition do
5:      $\text{dllp.train}(B_1 \cup B_0, \tilde{y}_1 \cup \tilde{y}_0) \triangleright$  Training step.
6:      $P \leftarrow \text{dllp.predict}(U_1) \triangleright$  Create posterior.
7:      $B_0, \tilde{y}_0 \leftarrow \text{CreateBags}(P, U_1) \triangleright$  Create bags.
8:      $(B_1, \tilde{y}_1, U_1) \leftrightarrow (B_2, \tilde{y}_2, U_2) \triangleright$  Switch views.
9:   end while
10: end procedure

```

To apply our algorithm, we assume we have two sets of bags, one with textual features and one with image features. Let B_k, \tilde{y}_k, U_k be bags, label proportions, and unlabeled data for view k (e.g. text or image), where U_k refer to the same set of instances. We propose Algorithm 2 for co-training in LLP settings. The algorithm proceeds by using the model trained on one view to create *pseudo bags* for the other view. The algorithm is initialized with empty pseudo bags (B_0, \tilde{y}_0). For each iteration, we first train a deep LLP model on the union of the current view (e.g. text or image) and the pseudo bags for that view. Next, it predicts the posterior probability of unlabeled samples with the current view. Then, it calls Algorithm 1 to create pseudo bags, using the current view’s posteriors as the priors $P(y|x)$. Finally, it switches views for the next iteration. In the experiments below, this algorithm tends to converge quickly (e.g., six iterations).

IV. DATA

To evaluate our deep LLP framework, we consider the task of predicting the gender and race/ethnicity of a Twitter user based on their tweets and profile image. In this section, we first describe how we collect data from Twitter and create bags, as well as the annotated data used for validation.²

A. Twitter data

It is common that image data comes with metadata. This metadata is useful to create bags. For example, images in Flickr³ have caption and description. Also, Twitter users often have a profile image. To demonstrate how the metadata can be used as a constraint to create bags, we collect data from Twitter. For the purpose of this study, we use geolocation and name in the metadata as constraints to create bags.

First, we use the Twitter Streaming API to collect roughly 120K tweets and remove users without an image profile; approximately 33K tweets with an image profile remain. However, not all of users use their own photo. To decrease noise, we want to ensure that there is only one face in the image profile. To do so, we apply the Viola-Jones object detection algorithm [24]. This algorithm detects multiple faces

²Replication code and data will be made available upon publication.

³<http://www.flickr.com>

TABLE I
THE DISTRIBUTION OF TWITTER AND CFD DATASET:

Dataset	TW	TW	CFD	CFD
Race	Male	Female	Male	Female
Hispanic	15	15	52	56
White	105	105	228	236
Black	40	40	231	295

in a scene with a low false positive rate. Finally, roughly 10.5K images with exactly one identified face remain. (Please note that this filtering only affects training data; validation data may contain multiple faces.) For each user, we download the most recent 200 tweets for use by the text-based model.

Next, we use the county and name metadata to create bags. We use the county constraint as described in recent work [25] by associating each user with a county in the U.S., based on the geolocation information provided in their tweets. This results in 85 bags with an average of 124 users per bags. We also use name attribute (where available in the user’s profile) to estimate label proportions. For gender classification, we use the first name with the data from US Social Security Administration (SSA). For example, according to SSA baby data⁴, for the first name ‘Casey’, the probability of being a man is 59%. Then we use Algorithm 1 with $N = 64$ and $t = .6$ to create bags. Additionally, for race/ethnicity classification, we use the last name as described in recent research [25] to estimate class priors for users who provide their last name, and run Algorithm 1 to create bags. For example, according to US Census, 67% of people with the last name ‘Taylor’ are white.

Finally, for evaluation purposes, we manually annotate 320 photos; the class distribution is shown in Table I. We use this dataset only for the evaluation purposes, not for training.

B. Google search

Twitter profile images are often too noisy with unbalanced class distributions (e.g. African-Americans are less frequent in both county and name constraints). This motivates a third type of constraint for image classification. We submit keywords (e.g. ‘Latino woman’ and ‘Black American man’) to Google images search to identify images that are likely to come from the desired class. Then, we apply the Viola-Jones algorithm to remove photos without exactly one face in the scene. Because Google search sort images in decreasing order of relevance, we associate decreasing label proportions as we descend the list. Specifically, we group search results into bags of size 64 for each contiguous set of results, up to a maximum of 800 results. The first bag is assigned a label proportion of 95% for the positive class, and the proportion is reduced by 5% for each subsequent bag, to a minimum of 55%.

C. CFD dataset

For additional validation, we use the Chicago Face Database (CFD), which contains high-quality frontal images (without

background) for both genders and four racial/ethnic categories (white, black, Hispanic, and Asian), and facial expressions [26]. We remove Asian photos (because only few Asian samples are in our Twitter data) and use this dataset as an additional testing set. The advantage of this database is that it has more samples for the Hispanic category than our validation set. Table I shows the class distribution of the CFD database.

However, this dataset has a different distribution of our training set. Our training set mostly has a wild condition (not frontal, with background and lower quality). As a result, we expect a different behavior of the classifier for this set. As we will discuss below in the experimental results, because these images have a very high quality, our race/ethnicity classifier is extremely accurate (with 98% F1 measure for black and white class). However, the F1 metric drops for gender classification because the lack of body features for CFD dataset.

V. RESULTS

For the purpose of this study, we train image and textual classifiers for the demographic attribute (gender and race/ethnicity) task. For race classification, we consider three categories (white, black, and Hispanic). However, according to the US Census⁵, Hispanic is a debated term that refers to Spanish culture or origin regardless of race and can be of any ethnicity. As a result, we consider another classification task restricted only to black and white classes. We refer to the latter as ‘**race2**’ and the former as ‘**race3**’ classification task. Since labeled samples for both Twitter and CFD datasets are imbalanced, we report the weighted average F1 measure to compare results.

In this section, we first provide experimental results for text classification task and compare it with the state of the art baselines. Then, we describe image classification results and compare them with third-party face APIs. Finally, we demonstrate how co-training improves both text and image classification, as well as the advantage of an ensemble approach.

A. Text classification result

Because the search constraints created by searching for Google images do not have any textual features, for text classification task we only use county and name constraints. Table II shows the weighted F1 measure for different tasks with all combination of constraints. According to this table, the county constraint results in higher accuracy for race classification, and name bags have higher accuracy for gender classification. That means that location is more informative for the race classification, and first name is more informative for the gender classification. Also, for the race task, using both constraints together increases the classification accuracy, but for the gender classification, the result of combining constraints is almost same as using only name bags.

Since using both constraints improves (or at least does not harm) the classification task, we use both name and county

⁴<https://www.ssa.gov/oact/babynames/>

⁵<http://www.census.gov/prod/cen2010/briefs/c2010br-04.pdf>

TABLE II
THE WEIGHTED AVERAGE F1 MEASURE FOR DIFFERENT TASKS AND CONSTRAINTS FOR TEXT CLASSIFICATION ON TWITTER DATASET.

Model	race2	race3	gender	Average
county	78	64	54	65
name	61	55	83	66
county-name	80	73	83	79

TABLE III
THE WEIGHTED AVERAGE F1 MEASURE WITH DIFFERENT BATCH SIZE FOR TEXT CLASSIFICATION ON TWITTER DATASET.

Max batch size	race2	race3	gender	Average
16	80	71	83	78
32	80	73	83	79
64	73	64	83	73

constraints in all subsequent text classification experiments. We use the maximum batch size of 32 to generate Table II; to do so, any larger bags are randomly split into smaller bags with the maximum size of 32 at each iteration of training.

Table III presents the weighted average F1 metric for different maximum batch sizes using county and name constraints. According to this table, while gender classification is stable for different batch sizes, the maximum batch size of 32 works better for race tasks. This result is not surprising, because a batch size of 16 is too small for a bag, and batch size of 32 has the advantage of randomly splitting bags (that is created with size 64) to avoid overfitting. An interesting case is batch size of 1, which associates a label distribution to individual instances. The problem with this approach is that it is less robust to noise — for example, if a label proportion of 90% positive is applied to a negative instance, the classifier will attempt to fit this noisy example. However, if this instance is the only negative one in a batch of size 10, then the classifier is given the flexibility of predicting it as negative while still optimizing the loss function. Indeed, we find that batch size 1 is not effective on this task.

Finally, we compare our deep LLP model with state-of-the-art shallow LLP models — we use ridge regression for LLP [25] and label regularization [8] for comparison. Table IV compares our model with maximum batch size 32 with other models for county and name constraints. According to this table, while our model has the same F1 as ridge LLP for **race2** task, it has higher F1 for all other tasks. More specifically, on average, while both ridge LLP and label regularization have the same F1 score (76), our model has the highest F1 measure (79).

B. Image classification results

In this section, we present image classification results using our proposed deep LLP model. For all experiments, we use the Xception model [3] with its training weights fit on ImageNet. The ImageNet dataset [6] that is commonly used for image classification has over a million images for 1,000 objects but does not have any class related to a human face or body. Because we have the best result with the maximum batch size 32 in the last section, we use it again for image classification

TABLE IV
COMPARISON WITH THE STATE OF THE ART LLP TEXT CLASSIFICATION BASELINES ON TWITTER DATASET.

Model Name	race2	race3	gender	Average
Deep LLP	80	73	83	79
Ridge LLP	80	69	79	76
Label Regularization	79	68	82	76

TABLE V
COMPARISON OF F1 OF IMAGE CLASSIFICATION FOR DIFFERENT CONSTRAINTS ON TWITTER AND CFD DATASETS:

Database Task	TW	CFD	TW	CFD	TW	CFD
	race2	race2	race3	race3	gender	gender
county	76	72	52	25	53	59
name	61	30	54	25	88	65
search	85	94	71	83	91	73
county-name	61	31	52	25	93	68
county-search	83	85	79	82	93	67
name-search	81	91	79	85	92	68
county-name-search	92	98	77	75	95	75

task. To reduce memory consumption, we freeze first two blocks of Xception network in the training phrase, and train the remaining layers.

To avoid overfitting and make the model robust to the wild condition of images, we apply various image distortions before each training iteration; we randomly rotate, flip (horizontally), shift (vertically and horizontally), shear, and zoom photos. Also, with probability of 80%, we randomly crop the Viola-Jones detected face to avoid overfitting the background in images. We use Adam [27], an adaptive stochastic gradient descent algorithm, as an optimization algorithm, and we train all models for up to 20 epochs and report the accuracy on the validation set (Twitter annotated image profiles).

Table V shows the weighted average F1 measure for different tasks using all constraint combinations. According to this table, the search constraint is more informative in all tasks, and using county with name constraints together has a poor result for race classification. Also, except for the **race3** task, using all constraints together has the best result. By comparing this with text classification, because of search constraints, it is apparent that image classification has higher accuracy than text classification.

This table also presents the F1 metric for CFD dataset. For **race2**, the model has a very high result of 98% F1 measure. This result reveals that race classification is much easier with high-quality frontal photos, as opposed to the noisier images from Twitter. On the other hand, while the gender classifier has very high F1 (95%) for Twitter images, it has lower F1 for CFD dataset. We believe that is in part because the CFD images omit body features, and so the classifier must rely solely on face features.

Since, on average, using all constraints has the highest average F1 measure on all tasks, in subsequent experiments we present the results of models that trains with all constraints.

Finally, we compare deep LLP with four baselines. In our first two baselines, we compare deep LLP with supervised deep learning; we use Xception with its pre-trained weights, and train it for 50 epochs using labeled data. For **Xception-**

TABLE VI
COMPARING DEEP LLP MODEL (COUNTY-NAME-SEARCH) WITH BASELINES.

Database Task	TW race2	CFD race2	TW race3	CFD race3	TW gender	CFD gender
Deep LLP	92	98	77	75	95	75
Xception-wild	—	93	—	75	—	71
Xception-HQ	85	—	71	—	82	—
Microsoft	—	—	—	—	80	90
Sightcorp	24	53	21	46	22	64

wild model, the training data is Twitter evaluation labeled data in wild conditions and we use CFD dataset with high quality images to train **Xception-HQ** model. Our next two baselines are public face APIs. These APIs can detect multiple faces in the scene with multiple attributes. The **Microsoft** face API⁶ can identify gender but does not predict race. For race classification, we use **Sightcorp** API⁷, which is the only API supporting race recognition to the best of our knowledge, but it is still in the beta phase and does not have very accurate results. Table VI compares these baselines to our deep LLP model using all constraints. Our approach outperforms Xception-wild, Xception-HQ, and Sightcorp for all tasks and outperforms the Microsoft API in the wild condition. However, the latter has the better result for CFD dataset with clean, frontal images, which our method was not trained on.

C. Co-training result

In this section, we provide experimental results for Algorithm 2. For these results, we initialize algorithm with the image view. (It is also possible to initialize it with the textual view, but we do not expect a significant difference.) In each iteration, the algorithm creates pseudo-bags, which are used in the next iteration (by switching the view). In this experiment, we use the same Twitter unlabeled data (with both image and text) as unlabeled samples. Thus, we use the same unlabeled data, but organize them into different bags with different label proportions for training. For text bags, we use county and name constraints, and for image bags, we use county, name, and search constraints.

Table VII presents the result of the co-training algorithm for image classification. In this table, the first column indicates the iteration of Algorithm 2, and the first row (iteration zero) states the F1 measure of deep LLP model without co-training. This table only shows odd iterations (image classification steps) of the co-training algorithm. According to this table, the biggest improvement comes in the first iteration, which improves absolute F1 of race2 by 2% (an error reduction of 25%), and improves the absolute F1 of race3 by 4% (an error reduction of 17%).

By applying the co-training algorithm, the most growth belongs to **race3** with an absolute improvement of 7% for Twitter and 12% for CFD dataset. We believe that is in part because that the text classifier can detect Spanish words for

⁶<https://www.microsoft.com/cognitive-services/en-us/face-api>

⁷<https://face.sightcorp.com/>

TABLE VII
THE CO-TRAINING RESULTS FOR IMAGE CLASSIFICATION. ITERATION 0 REPORTS F1 MEASURE BEFORE RUNNING THE CO-TRAINING ALGORITHM.

Database Iteration	TW race2	CFD race2	TW race3	CFD race3	TW gender	CFD gender
0	92	98	77	75	95	75
1	94	98	81	87	95	75
3	94	98	84	82	95	77
5	95	98	81	86	96	80

TABLE VIII
THE CO-TRAINING RESULTS FOR TEXT CLASSIFICATION. ITERATION 0 REPORTS F1 MEASURE BEFORE RUNNING THE CO-TRAINING ALGORITHM.

Iteration	race2	race3	gender	Average
0	80	73	83	79
2	90	81	86	86
4	91	83	86	87
6	92	83	85	87

Hispanic class and make better pseudo-bags for it. For the **race2** classification task, our co-training algorithm improves absolute F1 by 3% for the Twitter dataset. The CFD dataset has already very high F1 measure (98%), and does not have any growth by co-training steps.

Because the **gender** classification already has a very high F1 measure (95%) on Twitter, the co-training improves it by only 1%. However, it increases the F1 on the CFD dataset by 5%. We believe the lower accuracy for the **gender** classification task of the CFD database is in part because of the lack of body features.

Table VIII demonstrates the result of even iterations, text classification steps, of Algorithm 2. Again, the first row shows the result of text classification without co-training. According to this table, the second iteration of co-training has the highest improvement (7% on average across all tasks). The highest growth belongs to **race2** classes, which improves by 10% in absolute F1. This improvement is likely because the image classifier has a very high F1 for this task, and as a result, it creates accurate pseudo-bags for the text classifier. Finally, the **gender** classification grows only 3%, since it has a little improvement from image classification.

In our final experiment, we select the last iteration for text classification (6th iteration) and image classification (5th iteration) of co-training algorithm. Then, we use them on the Twitter dataset and blend their predictions by using soft voting. Table IX shows the result. According to this table, the ensemble improves image classification by average 2%, and text classification by average 6%. The ensemble produces the highest F1 for all tasks (except image classification for **gender**, which has the same accuracy).

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we introduced two enhancements to deep learning methods for image and text classification: (1) a Batch Averager layer to enable LLP deep learning, and (2) a co-training method for combining deep LLP models trained on different views. We found that a deep model that is trained on

TABLE IX

RESULTS FOR THE ENSEMBLE (SOFT VOTING) METHOD, USING THE FINAL ITERATIONS (TEXT AND IMAGE) OF THE CO-TRAINING ALGORITHM.

Method	race2	race3	gender	Average
Text	92	83	85	87
Image	95	81	96	91
Ensemble	96	86	96	93

population level data using proper constraints is comparable to traditional supervised learning. This approach decreases the burden of human annotation and can readily be implemented with almost every publicly-available deep learning software packages. We also found that for applications with two views (image and text) of features, a co-training algorithm leverages improvement of classification task for both views, and can be enhanced by an ensemble learning to achieve the highest precision.

ACKNOWLEDGMENT

This research was funded in part by the National Science Foundation under grants #IIS-1526674 and #IIS-1618244, and CCC Information Services Inc. provided a server (with multiple GPUs) to run deep learning models.

REFERENCES

- [1] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014. [Online]. Available: <http://arxiv.org/abs/1409.1556>
- [2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015. [Online]. Available: <http://arxiv.org/abs/1512.03385>
- [3] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," *CoRR*, vol. abs/1610.02357, 2016. [Online]. Available: <http://arxiv.org/abs/1610.02357>
- [4] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Computer Vision and Pattern Recognition (CVPR)*, 2015. [Online]. Available: <http://arxiv.org/abs/1409.4842>
- [5] C. Szegedy, S. Ioffe, and V. Vanhoucke, "Inception-v4, inception-resnet and the impact of residual connections on learning," *CoRR*, vol. abs/1602.07261, 2016. [Online]. Available: <http://arxiv.org/abs/1602.07261>
- [6] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015.
- [7] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" *CoRR*, vol. abs/1411.1792, 2014. [Online]. Available: <http://arxiv.org/abs/1411.1792>
- [8] G. S. Mann and A. McCallum, "Simple, robust, scalable semi-supervised learning via expectation regularization," in *Proceedings of the 24th International Conference on Machine Learning*, ser. ICML '07. New York, NY, USA: ACM, 2007, p. 593600. [Online]. Available: <http://doi.acm.org/10.1145/1273496.1273571>
- [9] K.-T. Lai, F. X. Yu, M.-S. Chen, and S.-F. Chang, "Video event detection by inferring temporal instance labels," in *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition*, ser. CVPR '14. Washington, DC, USA: IEEE Computer Society, 2014, pp. 2251–2258. [Online]. Available: <http://dx.doi.org/10.1109/CVPR.2014.288>
- [10] J. Graça, K. Ganchev, and B. Taskar, "Expectation maximization and posterior constraints," in *NIPS*, vol. 20, 2007, pp. 569–576.
- [11] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *CoRR*, vol. abs/1502.03167, 2015. [Online]. Available: <http://arxiv.org/abs/1502.03167>
- [12] R. Al-Rfou, G. Alain, A. Almahairi, C. Angermüller, and Others, "Theano: A python framework for fast computation of mathematical expressions," *CoRR*, vol. abs/1605.02688, 2016. [Online]. Available: <http://arxiv.org/abs/1605.02688>
- [13] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, and Others, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," *CoRR*, vol. abs/1603.04467, 2016. [Online]. Available: <http://arxiv.org/abs/1603.04467>
- [14] T. Sakaki, M. Okazaki, and Y. Matsuo, "Earthquake shakes twitter users: Real-time event detection by social sensors," in *Proceedings of the 19th International Conference on World Wide Web*, ser. WWW '10. New York, NY, USA: ACM, 2010, pp. 851–860. [Online]. Available: <http://doi.acm.org/10.1145/1772690.1772777>
- [15] A. Blum and T. Mitchell, "Combining labeled and unlabeled data with co-training," in *Proceedings of the Eleventh Annual Conference on Computational Learning Theory*, ser. COLT '98. New York, NY, USA: ACM, 1998, pp. 92–100. [Online]. Available: <http://doi.acm.org/10.1145/279943.279962>
- [16] L. Sifre and S. Mallat, "Rigid-motion scattering for texture classification," *CoRR*, vol. abs/1403.1687, 2014. [Online]. Available: <http://arxiv.org/abs/1403.1687>
- [17] D. Kotzias, M. Denil, N. de Freitas, and P. Smyth, "From group to individual labels using deep features," in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '15. New York, NY, USA: ACM, 2015, pp. 597–606. [Online]. Available: <http://doi.acm.org/10.1145/2783258.2783380>
- [18] F. Li and G. Taylor, "Alter-cnn: An approach to learning from label proportions with application to ice-water classification," in *Neural Information Processing Systems Workshops (NIPSW) on Learning and privacy with incomplete data and weak supervision*, 2015.
- [19] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, Jan. 2014. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2627435.2670313>
- [20] N. Zhang, M. Paluri, M. Ranzato, T. Darrell, and L. D. Bourdev, "PANDA: pose aligned networks for deep attribute modeling," *CoRR*, vol. abs/1311.5591, 2013. [Online]. Available: <http://arxiv.org/abs/1311.5591>
- [21] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," *CoRR*, vol. abs/1411.7766, 2014. [Online]. Available: <http://arxiv.org/abs/1411.7766>
- [22] R. Ranjan, V. M. Patel, and R. Chellappa, "Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition," *CoRR*, vol. abs/1603.01249, 2016. [Online]. Available: <http://arxiv.org/abs/1603.01249>
- [23] S. Gupta, J. Kim, K. Grauman, and R. Mooney, "Watch, listen & learn: Co-training on captioned images and videos," in *Proceedings of the 2008 European Conference on Machine Learning and Knowledge Discovery in Databases - Part I*, ser. ECML PKDD '08. Berlin, Heidelberg: Springer-Verlag, 2008, pp. 457–472. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-87479-9_48
- [24] P. A. Viola and M. J. Jones, "Rapid object detection using a boosted cascade of simple features," in *2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2001)*, with CD-ROM, 8-14 December 2001, Kauai, HI, USA. IEEE Computer Society, 2001, pp. 511–518. [Online]. Available: <http://dx.doi.org/10.1109/CVPR.2001.990517>
- [25] E. Mohammady and A. Culotta, "Using county demographics to infer attributes of twitter users," in *ACL Joint Workshop on Social Dynamics and Personal Attributes in Social Media*, 2014.
- [26] D. S. Ma, J. Correll, and B. Wittenbrink, "The chicago face database: A free stimulus set of faces and norming data," *Behavior Research Methods*, vol. 47, no. 4, pp. 1122–1135, 2015. [Online]. Available: <http://dx.doi.org/10.3758/s13428-014-0532-5>
- [27] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014. [Online]. Available: <http://arxiv.org/abs/1412.6980>