



Topics and Label Propagation: Best of Both Worlds for Weakly Supervised Text Classification

Sachin Pawar^{1,2}, Nitin Ramrakhiani^{1(✉)},
Swapnil Hingmire^{1,3}, and Girish K. Palshikar¹

¹ TCS Research, Tata Consultancy Services, Pune 411013, India

{sachin.pawar,nitin.ramrakhiani,swapnil.hingmire,gk.palshikar}@tcs.com

² Department of CSE, Indian Institute of Technology Bombay, Mumbai 400076, India

³ Department of CSE, Indian Institute of Technology Madras, Chennai 600036, India

Abstract. We propose a Label Propagation based algorithm for weakly supervised text classification. We construct a graph where each document is represented by a node and edge weights represent similarities among the documents. Additionally, we discover underlying topics using Latent Dirichlet Allocation (LDA) and enrich the document graph by including the topics in the form of additional nodes. The edge weights between a topic and a text document represent level of “affinity” between them. Our approach does not require document level labelling, instead it expects manual labels only for topic nodes. This significantly minimizes the level of supervision needed as only a few topics are observed to be enough for achieving sufficiently high accuracy. The Label Propagation Algorithm is employed on this enriched graph to propagate labels among the nodes. Our approach combines the advantages of Label Propagation (through document-document similarities) and Topic Modelling (for minimal but smart supervision). We demonstrate the effectiveness of our approach on various datasets and compare with state-of-the-art weakly supervised text classification approaches.

1 Introduction

Text classification is an important area of Natural Language Processing (NLP) with applications ranging from automatic request routing to text understanding. It has also been one of the most active and competitive areas of research in NLP. In this work, we propose a novel weakly supervised method to solve document classification.

We use the Label Propagation algorithm [24] which works on an undirected graph and involves iterative propagation of labels from a few labelled nodes to large number of unlabelled nodes. The algorithm stops when label distributions at all nodes have converged.

For Label Propagation, representation of documents in a graph and setting edge weights as similarity values among the documents is necessary. However, to

achieve high accuracy, providing a good number of labelled documents is necessary. Document labelling can be an expensive and time-consuming activity and would require domain expertise. To do away with the cumbersome labelling of documents we propose to add topics learned over the documents in the graph and solicit labels only for topic nodes. Hingmire et al. [12] and Razavi et al. [18] proposed labelling of topics instead of documents arguing that topic labelling invites lesser manual effort. We enrich the document similarity graph by adding labelled topics. We also introduce a topic influence parameter to control the topic enrichment process. Our algorithm LPA-TD (Label Propagation Algorithm - Topic Documents) constructs this topic enriched graph and runs Label Propagation on it to discover a classification of documents. The topic enriched graph is constructed for various configurations of the topic influence parameter and document similarities. Additionally, we experimented by constructing the topic enriched graph by dropping certain topic nodes which were incoherent and confusing to label. Closely seen, LPA-TD combines the power of topic modelling through smart manual tagging and iterative propagation of Label Propagation by harnessing document similarities.

We experiment on 4 public datasets from the 20Newsgroups (20NG) corpora and compare LPA-TD with multiple weakly supervised algorithms for text classification. We also compare LPA-TD with the performance of only Label Propagation (*OnlyLPA*) using some labelled documents. LPA-TD outperforms the *OnlyLPA* baseline on all datasets and also outperforms the other algorithms on two out of four 20NG datasets. We also perform experiments on a real-world dataset comprising of about 4000 grievances raised by employees of a large IT organization. The grievance text needs to be analysed by classifying it into four classes related to appraisals, compensation, finance and administration. Based on a manually created gold standard, LPA-TD performs at an encouraging macro-F1 of 78% on this dataset.

The paper is organized as follows. In Sect. 2 we briefly describe the background of various techniques employed in the proposed LPA-TD algorithm. In Sect. 3, we describe the construction of the topic enriched graph and the topic influence parameter. Further in Sect. 4, we present details about the datasets, experimental setup, evaluation and analysis. Relevant related work is presented in Sect. 5. We finally detail some future work and conclude the paper.

2 Background

2.1 Label Propagation Algorithm

Zhu and Ghahramani [24] proposed the Label Propagation Algorithm which is a graph based semi-supervised method. It represents labelled and unlabelled instances as nodes in a graph with edges reflecting the similarity between nodes. The label information for any node is propagated to its nearby nodes through weighted edges iteratively and finally the labels of unlabelled examples are inferred when the propagation process is converged. The detailed version of the algorithm for transductive document classification is described in Algorithm 1.

Data: 1. D_L (Set of labelled documents)
 2. D_U (Set of unlabelled documents)
 3. S ($n \times n$ Similarity matrix where $n = |D_L| + |D_U|$ and top $|D_L|$ rows correspond to labelled documents)
 4. $L = \{l_1, l_2, \dots, l_m\}$ (Set of m class labels)
Result: Label matrix $Y_{n \times m}$, where Y_{ij} represents the probability of document d_i having label l_j

/ Begin Initialization */*

- 1 Define probability transition matrix T such that $T_{ij} = \frac{s_{ij}}{\sum_k s_{kj}}$ which is the probability of jumping from l_j to l_i ;
 - 2 Define \bar{T} as the row-normalized matrix of T such that $\bar{T}_{ij} = \frac{T_{ij}}{\sum_k T_{ik}}$;
 - 3 Set iteration index $t = 0$;
 - 4 Let Y^0 be the label matrix for 0^{th} iteration and Y_L^0 be its top $|D_L|$ rows and Y_U^0 be its remaining rows;
 - 5 Set $Y_{ij}^0 = 1$ if d_i is labelled with l_j ;
 - 6 Set values of Y_U^0 arbitrarily;
- /* End Initialization */*
- 7 Propagate the labels of any node to nearby nodes by $Y^{t+1} = \bar{T}Y^t$;
 - 8 Replace the values of top $|D_L|$ rows of Y^{t+1} with Y_L^0 ;
 - 9 Set $t := t + 1$;
 - 10 Repeat steps 7 to 9 until Y converges;
 - 11 **return** Y^t ;

Algorithm 1: Label Propagation Algorithm for Transductive Document Classification

2.2 Topic Modelling

Topic modelling allows us to discover important and frequent themes or “topics” discussed in a large collection of text documents. The discovered topics provide an abstraction on the top of individual documents. Latent Dirichlet Allocation (LDA) [1] is the simplest topic model. LDA and its variants have numerous applications in natural language processing, image processing, social network analysis etc. It is widely used to browse a large corpus of documents using the most probable words of each topic and the distribution over topics for each document [3]. LDA assumes following generative process for generating documents.

1. Select word probabilities (ϕ_t) for each topic t :
 $\phi_t \sim \text{Dirichlet}(\beta)$
2. Select topic proportions (θ_d) for document d :
 $\theta_d \sim \text{Dirichlet}(\alpha)$
3. Select the topic for each word position ($z_{d,n}$):
 $z_{d,n} \sim \text{Multinomial}(\theta_d)$

4. Select the token for each word position ($w_{d,n}$):

$$w_{d,n} \sim \text{Multinomial}(z_{d,n})$$

(α and β are Dirichlet priors for document-topics and topic-words distributions respectively.)

2.3 Document Similarities

In order to assign weights to edges connecting various documents, we need to devise a way of computing similarity between any two documents. Various document similarity measures and their effect on document clustering performance are discussed in detail by Huang [13]. For all our experiments, we have employed ‘‘Cosine Similarity’’ measure.

Let $D = \{d_1, d_2, \dots, d_n\}$ be a set of n documents and $W = \{w_1, w_2, \dots, w_m\}$ be the set of m distinct words (excluding stop-words and all words occurring in only one document), constituting the vocabulary of the corpus D . We represent each document d_i by a vector V_i of length m whose j^{th} component ($V_i[j]$) corresponds to the j^{th} word in W and is computed as,

$$V_i[j] = TF(d_i, w_j) \cdot IDF(w_j)$$

where $TF(d_i, w_j)$ is number of times the word w_j occurs in the document d_i and $IDF(w_j)$ is computed using $ND(w_j)$, i.e. number of documents containing the word w_j as $IDF(w_j) = \log\left(\frac{n}{ND(w_j)}\right)$.

For any two documents (d_i, d_j) and their corresponding vector representations (V_i, V_j), Cosine similarity between them is computed as follows:

$$\text{CosSim}(d_i, d_j) = \frac{V_i \cdot V_j}{|V_i||V_j|}$$

Document Graph Construction: We construct a document graph where each node represents a document and an edge between any two documents indicates that the documents are similar. The degree of similarity between two documents is represented by assigning appropriate proportional edge weight. Higher edge weight indicates that there is a high similarity between the two documents.

It was observed that each document generally has ‘‘low’’ cosine similarity with a large number of documents. In order to prevent label propagation among the dissimilar document nodes, we need to find some threshold on the document similarities such that if the similarity is below the threshold, then no edge will be added between such documents. We determine this threshold automatically for any set of documents. The threshold is determined such that at least 90% documents within the set are connected to at least K other documents. All the similarities below this threshold are forced to 0 and hence no edge will be added in the document graph for a document pair with similarity below the threshold.

3 LPA-TD: Proposed Approach for Text Classification

In this section, we describe our novel approach for weakly supervised text classification.

3.1 Weak Supervision by Labelling Topics

The idea of manually obtaining labels for topics instead of instances was explored by Hingmire and Chakraborti [11, 12] and Razavi et al. [18]. They observed that LDA topics are easily interpretable as they can be represented by their most probable words. A human annotator can provide the most suitable class label to each topic. The level of supervision in this case is quite low, as they found that a very few topics (typically twice the number of class labels) are generally enough. LDA topics uncover underlying semantic structure of the whole set of documents. Hence, even a few labelled topics, add significant information about most of the documents. Table 1 shows the discovered topics and corresponding labels by assigned by a human annotator for the MEDICAL vs SPACE classification problem in 20 newsgroups dataset.

Table 1. Examples of topics discovered in the MEDICAL vs SPACE classification of 20 newsgroups dataset and corresponding labels assigned by a human annotator

Topic (most probable words)	Label
<u>msg</u> <u>food</u> <u>doctor</u> <u>pain</u> <u>day</u> <u>problem</u> <u>read</u> <u>evidence</u> <u>problems</u> <u>doesn</u> <u>blood</u> <u>question</u> <u>case</u> <u>body</u> <u>dye</u>	Medical
<u>space</u> <u>nasa</u> <u>science</u> <u>program</u> <u>system</u> <u>data</u> <u>research</u> <u>information</u> <u>shuttle</u> <u>technology</u> <u>station</u> <u>center</u> <u>based</u> <u>sci</u> <u>theory</u>	Space
<u>medical</u> <u>health</u> <u>water</u> <u>cancer</u> <u>disease</u> <u>number</u> <u>research</u> <u>information</u> <u>april</u> <u>care</u> <u>keyboard</u> <u>hiv</u> <u>center</u> <u>reported</u> <u>aids</u>	Medical
<u>launch</u> <u>earth</u> <u>space</u> <u>orbit</u> <u>moon</u> <u>lunar</u> <u>nasa</u> <u>high</u> <u>henry</u> <u>years</u> <u>spacecraft</u> <u>long</u> <u>cost</u> <u>mars</u> <u>pat</u>	Space

3.2 Affinity Between Topics and Documents

We use collapsed Gibbs sampling [9] to learn topics only using training documents. The collapsed Gibbs sampler for LDA gives topic-word distributions (ϕ_t) for each topic and document-topic distribution (θ_d) for each document. We use $\theta_{d,t}$ i.e. probability of generating document d by topic t as the affinity between a training document d in corpus and topic t . In other words, affinity measures the belongingness of a topic to a particular document. We use ϕ_t to infer θ_d for an unseen document using collapsed Gibbs sampling method proposed by Heinrich [10]. For a particular document, its affinities with all the topics are normalized so that they sum to 1.

Data: $D = \{d_1, d_2, \dots, d_{N_D}\} = D_{train} \cup D_{test}$ (Set of N_D documents containing training and test documents), N_T (Number of topics to be discovered and later used for enrichment), $L = \{l_1, l_2, \dots, l_m\}$ (Set of m class labels), τ (Topic influence parameter)

Result: $D_L = \{\langle d_1, l_1 \rangle, \langle d_2, l_2 \rangle, \dots, \langle d_n, l_n \rangle\}$ where $l_1, l_2, \dots, l_n \in L$

```

1  $T :=$  Discover  $N_T$  topics using LDA from documents in  $D_{train}$ ;
2  $T_L := \{\langle t_1, l_1 \rangle, \langle t_2, l_2 \rangle, \dots, \langle t_{N_T}, l_{N_T} \rangle\}$ ; /* Topic labels from human
   annotator */
3  $A_{dd} :=$  Compute document-document  $N_D \times N_D$  similarity matrix;
4  $A_{td} :=$  Compute topic-document  $N_T \times N_D$  affinity matrix;
5 for  $i = 1$  to  $N_D$  do
6    $\mu := 0$ ; /* Total influence of the neighbouring documents */
7   for  $j = 1$  to  $N_D$  do
8      $\mu := \mu + A_{dd}[i][j]$ ;
9   end
10   $c := \frac{\tau \cdot \mu}{1 - \tau}$ ; /* Multiplier for topic-document affinities so that for
   each document, fraction of influence by topic nodes is  $\tau$  */
11  for  $j = 1$  to  $N_T$  do
12     $A_{td}[j][i] := c \cdot A_{td}[j][i]$ ;
13  end
14 end
   /* Similarity Matrix for Topic-enriched Graph */
15  $A := \begin{bmatrix} A_{td} & 0 \\ A_{dd} & A_{td}^T \end{bmatrix}_{(N_T+N_D) \times (N_T+N_D)}$ ;
16  $Y_{(N_T+N_D) \times m} := LPA(T_L, D, A, L)$ ;
17  $D_L := \Phi$ ;
18 for  $i = N_T + 1$  to  $N_T + N_D$  do
19    $j :=$  Index of maximum probability in  $Y[i]$ ;
20    $D_L := D_L \cup \{\langle d_i, l_j \rangle\}$ 
21 end
22 return  $D_L$ 

```

Algorithm 2: LPA-TD: Label Propagation Algorithm on a Topic-enriched Document Graph

3.3 Topic-Enriched Graph

We propose to enrich the document graph by adding a new node corresponding to each topic. As explained earlier, all these nodes are “labelled” nodes as the human supervision is provided at the topic level. All the other nodes representing documents are the “unlabelled” nodes. Each document node is connected to all topic nodes and the edge weight between a topic and a document is proportional to the “affinity” between them.

Topic Influence Parameter (τ): During the iterations of Label Propagation Algorithm, each document receives label distributions from the topic nodes as well as document nodes it is connected with. In other words, there are two sources of label information for a document node, i.e. “labelled topics” and “similar

documents”. To control the flow of label information from the two sources, we define a Topic Influence Parameter τ . Through this parameter, the influence of topic information on a particular node can be fixed to be a specific fraction of the total edge weight incident on that node. Consider a document at which sum of incident edge weights from document nodes is μ (sum of incident edge weights from topic nodes is 1 by definition as discussed in Sect. 3.2). To achieve the desired topic influence τ , all the incident edge weights from topic nodes to the document are multiplied by a value c changing the sum of incident edge weights from topic nodes to c . The value c in turn can be expressed as a function of τ and μ as follows:

$$\tau = \frac{c}{\mu + c} \Rightarrow c = \frac{\tau * \mu}{1 - \tau} \quad (1)$$

Using a user-specified topic influence parameter, the topic-enriched document graph is constructed with appropriate edge weights. A classification of documents is then obtained by running Label Propagation Algorithm over this topic-enriched graph. Algorithm 2 describes the LPA-TD algorithm in detail.

4 Experimental Analysis

4.1 Datasets

We report the performance of our experiments on corpora from the **20News-groups (20NG)** dataset. This dataset contains messages across twenty different UseNet discussion groups, posted over a period of time. These twenty newsgroups are grouped into 6 major clusters. We use the *bydate* version of the 20Newsgroups dataset¹. This version of the 20Newsgroups dataset contains 18,846 messages and it is sorted by the date of posting of the messages. The dataset is divided into training (60%) and test (40%) sets. We employ 4 different subsets of the 20NG dataset for our experiments, namely PC vs MAC, MEDICAL vs SPACE, POLITICS vs RELIGION and POLITICS vs SCIENCE. These subsets are fairly balanced in terms of representation of individual classes.

4.2 Experimental Setup

We start by learning double the number of topics as number of classes over the training documents. Here, it is important to note that class information of training documents is not used. Training documents are used in unsupervised way only to learn the topics. Also, we do not use test documents for learning the topics to ensure fair evaluation. Test documents are only used for reporting results.

The learned topics are labelled by a single human annotator. The annotator is asked to label a topic with only one of the most appropriate class. We use the learned topics to compute a topic-document affinity matrix (A_{td}) for both the training and test documents.

¹ <http://qwone.com/~jason/20Newsgroups/>.

Next, we construct the document to document similarity graph using two configurations: (i) $K = 1$ and (ii) $K = 3$ where K is minimum number of connected nodes for 90% nodes in the graph as discussed in the Sect. 2.3. We include both the training and test documents in the document graph. To form the topic enriched graph we introduce the learned topics as additional nodes in the document similarity graph, with labels as assigned by the human annotator. We create edges from each document to all topics and experiment with multiple values of the topic influence parameter (τ) for assigning edge weights. Results on the three best values of τ each for $K = 1$ and $K = 3$ are reported.

As the process of learning topics is based on approximate inference, we carry out the topic learning, topic labelling and topic-document affinity computation processes 10 times. Hence, for a given configuration of the document similarity graph with a K and a τ value, the topic enriched graph is constructed 10 times. It is straightforward to see that the document-document similarity part remains same for all 10 runs, but topics and topic to document edges are added afresh for each run. We finally average the results over all 10 runs for a configuration and report them.

We compare the LPA-TD technique with multiple baselines presented below. In Table 2, we report the macro-F1 scores from the baselines and various configurations of LPA-TD. For the baselines requiring labelled documents, we provide them with the same number of labelled documents as number of topics to be labelled in LPA-TD.

- Expectation maximization with Naive Bayes for text classification proposed by Nigam et al. [16] with 4 randomly selected labelled documents
- GE-FL [5] with 10 labelled features, as reported by Hingmire and Chakraborti [11]
- TLC and ClassifyLDA, as proposed and reported by Hingmire and Chakraborti [11]
- Using only Label propagation: In this configuration, we only consider the document similarity graph without topic enrichment and label as many number of documents as topics in LPA-TD. We then run Label Propagation on this graph and obtain the classification results for evaluation. We ensure that the most connected documents for both classes in the graph are labelled in equal proportion to ensure fairness and getting the best from this baseline.

4.3 Incoherent Topics

As discussed earlier, to deal with approximate topic inference, we run all the experiments 10 times and report their average performance. However, for some particular runs, we observed that we get macro-F1 scores much below the average score for that configuration. Upon observing the learned topics in these runs, we found that some topics were incoherent and represent multiple classes. Hence, forcing them to one particular class label was introducing noise, in turn leading to poor performance. Table 3 shows examples of such incoherent topics.

Table 2. Experimental results

	τ	pc-mac	med-space	politics-sci	politics-rel
TLC		0.68	0.943	0.911	0.922
ClassifyLDA		0.641	0.926	0.899	0.892
NB-EM		0.429	0.99	0.474	0.466
GE-FL		0.666	0.939	0.618	0.765
Only LPA (K = 1)		0.408	0.925	0.9407	0.8355
Only LPA (K = 3)		0.37	0	0.9392	0.8325
OnlyLPA (K = 1)		0.486	0.919	0.539	0.559
OnlyLPA (K = 3)		0.374	0.953	0.601	0.657
LPA-TD (K = 1)	0.1	0.673	0.947	0.912	0.837
	0.05	0.682	0.949	0.912	0.836
	0.01	0.704	0.951	0.887	0.819
LPA-TD (K = 3)	0.2	0.661	0.948	0.918	0.840
	0.1	0.673	0.950	0.916	0.839
	0.05	0.671	0.949	0.903	0.823
LPA-TD-Coh (K = 1)	0.1	0.686	0.945	0.92	0.852
	0.05	0.696	0.950	0.918	0.854
	0.01	0.719	0.954	0.887	0.849
LPA-TD-Coh (K = 3)	0.2	0.671	0.945	0.904	0.858
	0.1	0.681	0.951	0.925	0.860
	0.05	0.674	0.953	0.918	0.853

Table 3. Examples of incoherent topics

Noisy Topic (Most probable words)	Dataset
<u>gun</u> <u>space</u> <u>nasa</u> <u>president</u> <u>launch</u> <u>weapon</u> <u>firearm</u> <u>science</u> <u>tax</u>	<u>politics-sci</u>
<u>system</u> <u>job</u> <u>earth</u> <u>orbit</u> <u>clinton</u> <u>stephanopoulo</u>	
<u>drive</u> <u>disk</u> <u>hard</u> <u>system</u> <u>controller</u> <u>floppy</u> <u>rom</u> <u>bios</u> <u>card</u> <u>port</u>	<u>pc-mac</u>
<u>sound</u> <u>board</u> <u>power</u> <u>internal</u> <u>cable</u>	
<u>msg</u> <u>food</u> <u>idea</u> <u>high</u> <u>pat</u> <u>money</u> <u>long</u> <u>read</u> <u>remember</u> <u>thought</u> <u>moon</u>	<u>med-space</u>
<u>didn</u> <u>billion</u> <u>isn</u> <u>real</u>	
<u>system</u> <u>children</u> <u>objective</u> <u>moral</u> <u>wrong</u> <u>fire</u> <u>fbi</u> <u>morality</u>	<u>politics-religion</u>
<u>koresh</u> <u>opinions</u> <u>men</u> <u>doesn</u> <u>isn</u> <u>frank</u> <u>sex</u> <u>drugs</u>	

In order to avoid adverse effect of incoherent topics on label propagation, we simply removed such topics from the topic-enriched graph while making sure that there is at least one topic mapped to each class. After removing incoherent topics from topic-enriched graph, we re-run LPA-TD algorithm. We refer this new approach as LPA-TD (Coherent) (*LPA-TD-Coh*).

4.4 Discussion

Table 2 shows experimental results of LPA-TD and LPA-TD (Coherent) along with other baselines. As we can observe from the results, LPA-TD outperforms both the configurations of the OnlyLPA baseline on all datasets. It also performs better than all other baselines on two datasets - PC vs MAC and POLITICS vs SCIENCE.

PC vs MAC is considered to be the most difficult dataset from the 20NG corpora due to significant overlap of words seen in both classes. This overlap results from high semantic similarity between the classes. On this dataset, LPA-TD outperforms all the baselines comfortably, re-iterating the merit of the new technique.

It however, doesn't perform as well as the TLC and ClassifyLDA techniques in the POLITICS vs RELIGION dataset. A look at the topics learned in the 10 runs reveals mostly complex and fuzzy topics not attributable to a single class, which brings down the overall performance. Also, on the MEDICAL vs SPACE dataset, NB-EM outperforms LPA-TD but it performs quite poorly on other datasets. On the other hand, LPA-TD demonstrates a consistent performance across all the datasets.

4.5 Case Study on Employee Grievances

We also carried out a case study to analyse a real-world industrial text dataset of grievances which were raised by employees of a major IT organization. The dataset contained about 4000 grievance descriptions related to areas like **finance**, **compensation**, **appraisals** and **administration**. However, no direct classification was available. So we got the dataset labelled from an HR executive for use as gold standard. Further, the grievances were sorted on the date of posting and we used first 70% grievances for training and the rest for testing. We tried out Naive Bayes and SVM classifiers using Weka² and obtained macro-F1 of about 80.9% and 71.2% respectively.

Here, it is important to note that both Naive Bayes and SVM are supervised classifiers and required 70% i.e. 2800 labelled grievances to achieve the above performance. Now, we employed our LPA-TD approach on this dataset. A few examples of the topics discovered and corresponding manual labels are shown in Table 4. The topic enriched graph comprised of the 4000 grievance nodes along with the 8 learned topics. From the various configurations we tried, we obtained the best macro-F1 of 77.8% for $K = 3$ and $\tau = 0.05$. This demonstrates a significant reduction in labelling effort (2800 grievances against only 8 topics) through use of the LPA-TD technique for comparable performance.

5 Related Work

Previous work in semi-supervised text classification can be broadly categorized into 4 different types based on the way supervision is provided: (i) Labelling a

² <http://www.cs.waikato.ac.nz/ml/weka/>.

Table 4. Examples of topics discovered in the Employees Grievances dataset and corresponding labels assigned by a human annotator

Topic (Most probable words)	Label
<u>basic</u> <u>salary</u> <u>grade</u> <u>compensation</u> <u>experience</u> <u>pay</u> <u>higher</u> <u>variable</u> <u>allowance</u> <u>months</u> <u>joined</u> <u>current</u> <u>designation</u> <u>letter</u> <u>mba</u>	Compensation
<u>rating</u> <u>project</u> <u>appraisal</u> <u>performance</u> <u>band</u> <u>work</u> <u>appraiser</u> <u>process</u> <u>team</u> <u>discussion</u> <u>client</u> <u>reviewer</u> <u>final</u> <u>disagreement</u> <u>worked</u>	Appraisal
<u>office</u> <u>bus</u> <u>working</u> <u>admin</u> <u>work</u> <u>food</u> <u>cab</u> <u>day</u> <u>provided</u> <u>facility</u> <u>service</u> <u>card</u> <u>transport</u> <u>canteen</u> <u>issue</u>	Admin
<u>salary</u> <u>amount</u> <u>month</u> <u>finance</u> <u>claim</u> <u>tax</u> <u>account</u> <u>months</u> <u>received</u> <u>deducted</u> <u>paid</u> <u>ticket</u> <u>allowance</u> <u>days</u> <u>payroll</u>	Finance

few documents, (ii) providing a list of features that are highly indicative of each class label, (iii) employing active learning and (iv) labelling topics.

5.1 Using Labelled and Unlabelled Documents

In this method, labelled documents and a large number of unlabelled documents are used for learning the classifier. While estimating the parameters of the classifier certain assumptions about the distribution of labelled and unlabelled documents will have to hold.

Cluster assumption: if instances are in the same cluster, they are likely to be of the same class. In other words, if the data are generated by a mixture model following a generative process and a mixture component represents one or more classes then the instances generated by a mixture component are likely to have the same class labels. Due to unlabelled data, the mixture model contains both observed and hidden variables and its parameters are estimated by Expectation-Maximization (EM) algorithm [4]. Nigam et al. [16] used EM Algorithm for semi-supervised text classification with a naive Bayes classifier.

Low-density separation assumption: the decision boundary of classification should lie in a low-density region. Using this assumption, Grandvalet and Bengio [8] proposed a maximum a posteriori (MAP) framework for learning a classifier using minimum entropy regularization. Another semi-supervised algorithm which makes this assumption for learning a text classifier using a small number of documents is *Transductive Support Vector Machines (TSVMS)* [14].

Manifold assumption: the high-dimensional data lie (roughly) on a low-dimensional manifold. Graph based semi-supervised methods make the manifold

assumption to construct a graph in which nodes are both the labelled and unlabelled instances and edge weights represent similarity between instances. The Label Propagation Algorithm [25] discussed earlier falls in this category. Other graph based text classification approaches are by Subramanya and Bilmes [21] and Wang and Zhang [22].

Multi-view assumption: each instance has two or more “different” and “independent” views and each view is sufficient for good classification individually. Co-Training [2] algorithm is based on this assumption. The Co-Training process initially constructs a weak classifier for each view using labelled instances, then each weak classifier is bootstrapped using unlabelled instances.

5.2 Incorporating Labelled Features

Sometime it is easier for human annotators to describe a class of documents using a set of features than labelling large collections of documents. Liu et al. [15] proposed a text classification algorithm by labelling the most discriminative words. Eventually, these representative words are used to create a text classifier using the combination of naive Bayes classifier and the Expectation-Maximization (EM) algorithm. Other similar approaches are Schapire et al. [19] (based on AdaBoost), Wu and Srihari [23] (generalization of SVM, Weighted Margin SVM) and Druck et al. [5] (generalized expectation criteria based maximum entropy text classifier, i.e. GE-FL).

5.3 Using Active Learning

Active learning [20] systems attempt to overcome the labelling bottleneck by asking queries in the form of unlabelled instances to be labelled by a human annotator. Some important text classification approaches using active learning are Godbole et al. [7], Raghavan et al. [17] and Druck et al. [6].

5.4 Labelling Topics

Hingmire and Chakraborti [12] proposed the idea of obtaining labels for topics instead of documents. They propose the ClassifyLDA algorithm where a topic model is learned using LDA and one class label is assigned to each topic. They use the Dirichlet distribution to aggregate all the same class label topics into a single topic and automatically classify unlabelled documents based on their similarity with the aggregated topics. Hingmire and Chakraborti [11] proposed the TLC algorithm which further improves the ClassifyLDA algorithm by allowing a topic to be labelled with multiple class labels instead of one.

6 Conclusions and Future Work

We proposed a weakly supervised text classification technique LPA-TD, based on Label Propagation and Topic Modelling. A topic enriched document graph is constructed for a set of documents where the only supervision is in the form of labelled topics. LPA-TD propagates labels over this topic enriched graph thereby exploiting benefits of both, document similarities and labelled topics. We evaluated LPA-TD on 4 datasets of the 20NG corpora and compared with multiple baselines. LPA-TD outperforms all the baselines on 2 out of these 4 datasets, including PC vs MAC which is considered to be one of the most difficult for text classification. Compared to other baselines, LPA-TD demonstrates a consistent performance across all the datasets. We also showed that the issue of incoherent topics can be handled by removing them from the topic enriched graph, without degrading LPA-TD's performance. Furthermore, removal of such incoherent topics resulted in better performance.

In future, we plan to extend LPA-TD by allowing fuzzy class labels to topics which naturally represent multiple classes. This will ease the restriction of assigning only one class per topic. Additionally, we plan to devise topic quality measures for automatic detection of incoherent topics.

References

1. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent Dirichlet allocation. *J. Mach. Learn. Res.* **3**, 993–1022 (2003)
2. Blum, A., Mitchell, T.: Combining labeled and unlabeled data with co-training. In: *Proceedings of 11th Annual Conference on Computational Learning Theory*, pp. 92–100 (1998)
3. Chaney, A.J.B., Blei, D.M.: Visualizing topic models. In: *ICWSM (2012)*
4. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Stat. Soc. B* **39**(1), 1–38 (1977)
5. Druck, G., Mann, G., McCallum, A.: Learning from labeled features using generalized expectation criteria. In: *SIGIR*, pp. 595–602 (2008)
6. Druck, G., Settles, B., McCallum, A.: Active learning by labeling features. In: *EMNLP*, pp. 81–90 (2009)
7. Godbole, S., Harpale, A., Sarawagi, S., Chakrabarti, S.: Document classification through interactive supervision of document and term labels. In: *PKDD*, pp. 185–196 (2004)
8. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. In: *NIPS (2004)*
9. Griffiths, T.L., Steyvers, M.: Finding scientific topics. *PNAS* **101**(Suppl. 1), 5228–5235 (2004)
10. Heinrich, G.: Parameter estimation for text analysis. Technical report, University of Leipzig (2008)
11. Hingmire, S., Chakraborti, S.: Topic labeled text classification: a weakly supervised approach. In: *SIGIR*, pp. 385–394. *ACM (2014)*
12. Hingmire, S., Chougule, S., Palshikar, G.K., Chakraborti, S.: Document classification by topic labeling. In: *SIGIR*, pp. 877–880. *ACM (2013)*

13. Huang, A.: Similarity measures for text document clustering. In: Proceedings of 6th New Zealand Computer Science Research Student Conference (NZCSRSC 2008), pp. 49–56 (2008)
14. Joachims, T.: Transductive inference for text classification using support vector machines. In: ICML, pp. 200–209 (1999)
15. Liu, B., Li, X., Lee, W.S., Yu, P.S.: Text classification by labeling words. In: Proceedings of 19th National Conference on Artificial Intelligence, pp. 425–430 (2004)
16. Nigam, K., McCallum, A.K., Thrun, S., Mitchell, T.: Text classification from labeled and unlabeled documents using EM. *Mach. Learn. - Special issue on Information Retrieval* **39**(2-3), 103-134 (2000)
17. Raghavan, H., Madani, O., Jones, R.: Active learning with feedback on features and instances. *JMLR* **7**, 1655–1686 (2006)
18. Razavi, A.H., Inkpen, D., Brusilovsky, D., Bogouslavski, L.: General topic annotation in social networks: a latent Dirichlet allocation approach. In: Zaïane, O.R., Zilles, S. (eds.) *AI 2013. LNCS (LNAI)*, vol. 7884, pp. 293–300. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38457-8_29
19. Schapire, R.E., Rochery, M., Rahim, M.G., Gupta, N.K.: Incorporating prior knowledge into boosting. In: ICML, pp. 538–545 (2002)
20. Settles, B.: Active learning literature survey. Computer Sciences Technical report 1648, University of Wisconsin–Madison (2009)
21. Subramanya, A., Bilmes, J.: Soft-supervised learning for text classification. In: EMNLP, pp. 1090–1099. Association for Computational Linguistics (2008)
22. Wang, F., Zhang, C.: Label propagation through linear neighborhoods. *IEEE Trans. Knowl. Data Eng.* **20**(1), 55–67 (2008)
23. Wu, X., Srihari, R.: Incorporating prior knowledge with weighted margin support vector machines. In: Proceedings of 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 326–333 (2004)
24. Zhu, X., Ghahramani, Z.: Learning from labeled and unlabeled data with label propagation. Technical report, Citeseer (2002)
25. Zhu, X., Ghahramani, Z.: Learning from labeled and unlabeled data with label propagation. Technical report, Carnegie Mellon University (2002)