

# A Robust Framework for Classifying Evolving Document Streams in an Expert-Machine-Crowd Setting

Muhammad Imran  
Qatar Computing Research Institute, HBKU  
Doha, Qatar  
mimran@qf.org.qa

Sanjay Chawla  
Qatar Computing Research Institute, HBKU  
Doha, Qatar  
schawla@qf.org.qa

Carlos Castillo  
Eurecat  
Barcelona, Spain  
chato@acm.org

**Abstract**—An emerging challenge in the online classification of social media data streams is to keep the categories used for classification up-to-date. In this paper, we propose an innovative framework based on an Expert-Machine-Crowd (EMC) triad to help categorize items by continuously identifying novel concepts in heterogeneous data streams often riddled with outliers. We unify constrained clustering and outlier detection by formulating a novel optimization problem: COD-Means. We design an algorithm to solve the COD-Means problem and show that COD-Means will not only help detect novel categories but also seamlessly discover human annotation errors and improve the overall quality of the categorization process. Experiments on diverse real data sets demonstrate that our approach is both effective and efficient.

**Keywords**-stream classification; text classification; novel concept detection; social media; outlier detection

## I. INTRODUCTION

The application that motivates our work is time-critical analysis of a social media stream. We consider a basic operation on this stream, which is to rapidly categorize messages into a series of classes of interest and also to capture novel emerging categories. This is a relevant problem during all sort of crises, such as mass convergence events and emergencies, including sudden-onset natural and man-made disasters. This problem, in practice, is addressed through automatic classification, crowdsourced classification, or a combination of both [1, 2].

We developed a system [3] that follows the latter approach. The system combines human and machine intelligence to categorize crisis-related messages on Twitter during the sudden-onset of natural or man-made disasters. The system obtains labels (for the messages) from human workers (volunteers in this case) and trains machine learning classifiers. The trained classifiers then enable the automatic classification of subsequent incoming messages into the defined categories. We refer to this unique collaboration between domain experts, crowd volunteers and machine learning classification as the Expert-Machine-Crowd (EMC) framework.

There are two important challenges that have emerged for such a system to work in an optimal fashion.

**Defining the categorization scheme:** While it is impossible to predict apriori all types of categories (e.g. previously unknown needs of affected people) that are likely to emerge during a crisis, there is increasing evidence that most disasters do have a lot in common [4]. However, the dictionary of all possible information categories people write about in social media is potentially very large. For instance, in an analysis in 2012 of social media during 4 disasters, 28 information categories were found in the messages posted on this platform [5]. Having a large set of categories is problematic from the crowdsourcing point of view as unskilled annotators cannot distinguish between very fine-grained categories, which introduce labeling errors and also drastically limits the size of the annotator pool. A large number of categories also increases burn-in (performance reduction in the pool) and drop-out (annotators leaving the pool), particularly in volunteer crowdsourcing settings [3].

**Labeling Errors:** The dynamic nature, brevity and ambiguity of messages often lead to labeling errors by annotators which can drastically reduce the accuracy and thus usefulness of the classifier. Experiments with real data have shown that even a few poorly labeled messages can often lead to dramatically divergent results.

We propose a new optimization problem COD-Means (where COD stands for Constrained Outlier Detection) to address the problem of simultaneously discovering new categories and identifying labeling errors. Initial categories and labeling are expressed in terms of constraints and an outlier detection step is used for error detection. The outlier discovery through COD-Means is carried out by extending the  $k$ -means-- algorithm [6]. The cluster refinement process is carried out by using constrained clustering which allows the generation of new and semantically distinct categories. The expert then refines the categories based on the output of COD-means. Later initializations of COD-means (i.e. once the supervised learning system is trained) can be triggered by the expert after a fixed time-interval or on observing a low classification accuracy. In that case, the COD-Means also uses machine classified items for which the machine confidence is high (e.g.  $\geq 90\%$ ).

We formally define the problem in the next Section II. A novel optimization formulation and the associated algorithm are the subject of Section III. An extensive suite of experiments were carried out to test the EMC framework and are described in Section IV. Related work is the subject of Section V and we conclude in Section VI with a summary.

## II. PROBLEM DEFINITION

We are given as input a data set  $D$  of documents which have been categorized by crowdsourcing workers (not necessarily experts) or an automatic classifier into a taxonomy  $\mathcal{T} = \{T_1, \dots, T_k, Z\}$  containing  $|\mathcal{T}| = k + 1$  categories. The taxonomy forms a partition of the documents:  $A \cap B = \emptyset \forall A, B \in \mathcal{T}, A \neq B$ .

We call categories  $T_i$  the pre-existing categories, which are defined by the expert before the data begins to arrive, based on background domain knowledge.

The category  $Z$ , instead, is the ‘‘miscellaneous’’ category, used for documents that do not fit in any of the  $T_i$  categories. In practical cases for the domain in which we focus (disasters and social media), this category contains anywhere from about 10% to 30% of the messages [7].

Our task is to produce a new taxonomy

$$\mathcal{T}' = \{T'_1, \dots, T'_k, N_1, \dots, N_n, Z'\}$$

with the following characteristics:

- There are  $n$  new categories:  $|\mathcal{T}'| = |\mathcal{T}| + n$ .
- Pre-existing categories are only slightly modified:  $T_i \approx T'_i \forall i = 1 \dots k$ .
- New categories are different from previous pre-existing categories:  $T_i \not\approx N_j \forall i = 1 \dots k, j = 1 \dots n$
- $|Z'| < |Z|$ : the size miscellaneous category is reduced

At a high level, our algorithm attempts to partition  $D$  into  $k + n + 1$  categories, guiding the clustering process in such a way that  $k$  of these categories resemble the original  $T_i$ , and none of the categories overlap with each other.

## III. PROPOSED SOLUTION

Our solution framework consists of three parts (i) the crowd effort and the supervised learning output (the populated taxonomy  $\mathcal{T}$ ) are captured as must-link (ML) and cannot-link (CL) constraints [8]; (ii) discovery of new categories is framed as a constrained clustering problem where the number of clusters specified is greater than the size of the original taxonomy  $\mathcal{T}$ ; (iii) to identify discrepancies between human and machine annotation, a new clustering problem COD-Means is defined. A  $k$ -means type algorithm is proposed to solve COD-Means.

### A. Constraints Formation

Our objective is to infer new categories and we use constrained clustering to encode the output of the crowd and the supervised learning (SL) process. We form two types

of constraints: Must-Link (ML) constraints and Cannot-Link (CL) constraints. For the taxonomy  $\mathcal{T}$  we formulate these constraints as follows. **ML constraints:** data elements belonging to a category  $T_i$  are encoded as ML constraints, i.e., if  $a$  and  $b$  both belong to  $T_i$  then a constraint  $ML(a, b)$  is created. **CL constraints:** data elements belonging to different categories  $T_i$  and  $T_j$  are encoded as CL constraints, i.e., if  $a \in T_i$  and  $b \in T_j$  and  $i \neq j$ , then  $CL(a, b)$  is created.

An important point to note is that data elements of the miscellaneous  $Z$  category are not encoded with any constraints. In the unsupervised learning process these data points have the freedom to move to any cluster.

### B. Constrained Clustering

For constrained clustering we use Davidson and Ravi’s CVEQ (Constrained Vector Quantization Error) formulation to extend the  $k$ -means objective to capture both ML and CL constraints [8].

Let  $d(x, y) : D \times D \rightarrow R$  be the distance function in the context of the application. For a clustering problem on data set  $D$  with  $k$  clusters, let  $g : D \rightarrow \{1, \dots, k\}$  be the mapping from  $D$  to cluster labels. Let  $C = \{c_1, \dots, c_k\}$  be a set of representative cluster centroids and let  $h : C \rightarrow C$  be a function such that  $h(c)$  is the nearest centroid to  $c$  if  $h(c) \neq c$ . Finally, let  $\pi$  represent a cluster set. We define the *error* objective function of assigning  $D$  to  $C$  as

$$E(C, D) = \frac{1}{2} \sum_{j=1}^k \left\{ \sum_{x_i \in \pi_j} d^2(x_i, c_j) + \right. \quad (1)$$

$$\sum_{\substack{x_i \in \pi_j \\ (x_i, x_a) \in ML \\ g(x_i) \neq g(x_a)}} d^2(c_{g(x_a)}, c_j) + \quad (2)$$

$$\left. \sum_{\substack{x_i \in \pi_j \\ (x_i, x_a) \in CL \\ g(x_i) = g(x_a)}} d^2(c_{h(g(x_a))}, c_j) \right\} \quad (3)$$

The intuition behind the design of  $E(C, D)$  is that besides the standard distortion error (the first term), if an ML constraint  $(x_i, x_a)$  is violated then the cost of the violation is equal to the distance between the two centroids that contain the instances which should have been together. Similarly if a CL constraint  $(x_i, x_a)$  is violated then the error cost is the distance between the centroid  $c$  assigned to the pair and its nearest centroid  $h(c)$ .

Based on the error objective  $E(C, D)$  a  $k$ -means style of algorithm can be designed which iterates, until convergence, between an assignment and update rule defined as follows [8]:

**Assignment Rule:**

$$\forall x_i \notin ML \cup CL : \operatorname{argmin}_j d^2(x_i, c_j) \quad (4)$$

$$\forall (x, y) \in ML : \operatorname{argmin}_{i,j} \{d^2(x, c_i) + d^2(y, c_j) + \neg\delta(x, y) * d^2(c_i, c_j)\} \quad (5)$$

$$\forall (x, y) \in CL : \operatorname{argmin}_{i,j} \{d^2(x, c_i) + d^2(y, c_j) + \neg\delta(x, y) * d^2(c_i, h(c_i))\} \quad (6)$$

Here  $\delta(x, y)$  is the Kronecker delta function, i.e.,  $\delta(x, y) = 1$  if  $x = y$  and  $\delta(x, y) = 0$  if  $x \neq y$ . Thus if  $(x, y)$  is an ML constraint and  $x \neq y$ , then  $\neg\delta(x, y) = 1$  and if  $x$  and  $y$  end up belonging to different clusters ( $c_i$  and  $c_j$ ) then an error of  $d^2(c_i, c_j)$  is incurred. Similarly, if  $(x, y)$  in a CL constraint and  $x$  and  $y$  end up in the same cluster then a cost of  $d^2(c_i, h(c_i))$  is incurred where  $h(c_i)$  is the nearest cluster centroid to  $c_i$ . Note each of the argmin operators in the above assignment rules will output the centroid or pair of centroids.

**Update Rule:**

$$c_j = \frac{\sum_{x_i \in \pi_j} [x_i + \sum_{\substack{(x_i, y) \in ML, \\ g(x_i) \neq g(y)}} c_{g(y)} + \sum_{\substack{(x_i, y) \in CL, \\ g(x_i) = g(y)}} c_{h(g(y))}]}{|\pi_j| + \sum_{\substack{(x_i, y) \in ML, \\ g(x_i) \neq g(y)}} 1 + \sum_{\substack{(x_i, y) \in CL, \\ g(x_i) = g(y)}} 1} \quad (7)$$

The update rule of  $c_j$  computes a modified average of all points that belong to  $\pi_j$ . The modification captures the number of elements in  $\pi_j$  which violated the ML and CL constraints.

*C. The COD-Means Problem*

In order to formally capture the tendency of humans (and the SL algorithm) to make errors during the labeling process, we introduce a new computational problem to capture clustering, constraints and outliers.

*Problem 1: (COD-Means)* Given a data set  $D$ , a distance function  $d : D \times D \rightarrow R$ , constraint sets  $ML$  and  $CL$ , parameters  $k$  and  $\ell$  find a set  $C = \{c_1, \dots, c_k\}$  and a set  $\mathcal{L}$  consisting of  $k \times \ell$  points ( $\ell$  points per cluster) in order to minimize the error

$$\mathcal{E}(D, C, \mathcal{L}) = E(D \setminus \mathcal{L}, C) \quad (8)$$

**Observation:** The COD-Means problem is NP-hard for  $k > 1$  and  $\ell \geq 0$ . This is clear because without the outliers the problem is standard clustering with constraints which is known to be NP-hard in the presence of CL constraints.

*D. Algorithm*

We propose a natural extension of the constrained clustering algorithm to minimize  $\mathcal{E}(D, C, L)$  shown in Algorithm 1. The algorithm is similar to a standard clustering algorithm with assignment rules given in Equation 4 - 6. The key difference is in Lines 9 - 13, where the points in each cluster  $\pi$  are first sorted based on their distance to the centroid  $c_\pi$  and the top  $\ell$  points are removed from  $\pi$  and inserted into  $L(\pi)$ . The update rule (Equation 7) is then applied to the modified  $\pi$ .

Note that because of the presence of ML and CL constraints and that they have to be processed in pairs, the running time of the algorithm is bounded by  $O(|D|^2 k^2 I)$  where  $I$  is the number of iterations. We have omitted the convergence analysis of the algorithm due to space limitations; it is an extension of [6].

---

**Algorithm 1** The COD-Means Algorithm

---

**Input:** Data  $D$ ,  $ML$  and  $CL$  constraints on  $D$ ,  $k$  number of clusters,  $\ell$  number of outliers per cluster

**Output:** Cluster sets  $\Pi$ , Outlier sets  $\mathcal{L}$

- 1: Initialize with  $c_1, \dots, c_k$  centroids
  - 2: **while** (not converged) **do**
  - 3:   **for all**  $x \notin ML \cup CL$  **do**
  - 4:     Use Assignment Rule (Eqn 4)
  - 5:   **for all**  $(x, y) \in ML$  **do**
  - 6:     Use Assignment Rule (Eqn 5)
  - 7:   **for all**  $(x, y) \in CL$  **do**
  - 8:     Use Assignment Rule (Eqn 6)
  - 9:   **for all**  $\pi \in \Pi$  **do**
  - 10:     Re-order points  $x_i$  in  $\pi$  such that
  - 11:        $d(x_1, c(\pi)) \geq d(x_2, c(\pi)) \dots \geq d(x_{|\pi|}, c(\pi))$
  - 12:        $L(\pi) = \{x_1, \dots, x_\ell\}$
  - 13:        $\pi = \pi \setminus L(\pi)$
  - 14:       Update  $c_\pi$  using Update Rule (Eqn 7)
  - 15:  $\Pi = \{\pi_1, \dots, \pi_k\}$
  - 16:  $\mathcal{L} = \{L(\pi_1), \dots, L(\pi_k)\}$
- 

IV. EXPERIMENTS

We have designed and executed an extensive set of experiments to validate the proposed Expert-Machine-Crowd (EMC) framework. In particular we would like to resolve the following questions:

- 1) Are the new clusters identified by the COD-Means algorithm genuinely different and novel compared to the topical clusters (i.e. existing categories) previously defined?
- 2) What is the nature of outliers (i.e. in our case represent labeling errors) discovered by the COD-Means algorithm. Are they genuine outliers, i.e., they do not (semantically) belong to the clusters.

- 3) What is the impact of outliers on the quality of clusters generated from the COD-Means algorithm?
- 4) Once the new clusters are added to the training process and the outliers/wrong labels removed, does the overall accuracy of the classification process improve?

#### A. Datasets

We use 8 datasets from [9], which corresponds to English messages posted on Twitter during crises in 2012 and 2013. Each crisis in this collection comprises 1,000 tweets annotated using the following categories:  $T_A$ : Affected individuals,  $T_B$ : Infrastructure and utilities,  $T_C$ : Donations and volunteering,  $T_D$ : Caution and advice,  $T_E$ : Sympathy and emotional support,  $Z$ : Miscellaneous.

Table I lists the crises and the prevalence of the topical ( $T_i$ ) and miscellaneous ( $Z$ ) categories.

Table I: Datasets details: crisis name, year, % of tweets.

Crisis name	Year	% in $T_A..T_E$	% in $Z$
Colorado wildfires (CWF)	2012	55%	45%
Alberta floods (AF)	2013	82%	18%
Australia bushfires (ABF)	2013	60%	40%
Boston bombings (BB)	2013	59%	41%
Colorado floods (CF)	2013	76%	24%
NY train crash (NYTC)	2013	54%	46%
Queensland floods (QF)	2013	70%	30%
West Texas explosion (WTE)	2013	77%	23%

#### B. Comparison with standard $k$ -means (clusters quality)

First we compare our proposed algorithm with the baseline algorithm i.e. standard  $k$ -means. For this purpose, we use standard cluster metrics: cluster cohesiveness and novelty.

**Cohesiveness.** We measure cohesiveness based on standard intra- and inter-similarity metrics. For a cluster  $C_i$ , we denote by  $\text{intra}(C_i)$  its intra-cluster distance, defined as the average distance of elements inside the cluster:

$$\text{intra}(C_i) = \frac{\sum_{a,b \in C_i} d(a,b)}{|C_i|^2}.$$

For  $C_i$ , we denote by  $\text{inter}(C_i, C_i^c)$  its average inter-cluster distance with respect to elements in the pre-existing categories:

$$\text{inter}(C_i, \cup_{i=1}^k T_i) = \frac{\sum_{a \in C_i, b \in \cup_{i=1}^k T_i} d(a,b)}{|C_i| |\cup_{i=1}^k T_i|}.$$

We ignore the  $Z$  category in this inter-cluster calculation, as it does not represent a specific topic, but corresponds to elements that do not fit in any of the existing topics. The cohesiveness of a cluster  $C_i$  is defined as a combination of the cluster intra-similarity and its inter-similarity with other clusters:  $\text{intra}(C_i) / \text{inter}(C_i, C_i^c)$ . Ideally, a cluster/category should have high cohesiveness, i.e. small intra-cluster distance and large inter-cluster distances.

**Novelty.** A candidate category must not to be similar to previously existing categories i.e. it must be novel. Similar categories confuse human annotators and reduce the effectiveness of an automatic classifier. The novelty of a cluster  $C_i$  is determined as  $\text{maxDist}(C_i) - \text{minDist}(C_i)$  where:

$$\text{maxDist}(C_i) = \max_{a \in C_i, b \in \cup_{i=1}^k T_i} d(a,b)$$

and

$$\text{minDist}(C_i) = \min_{a \in C_i, b \in \cup_{i=1}^k T_i} d(a,b).$$

A high novelty is observed when a cluster is far apart from all the other pre-existing clusters  $T_i$ .

**Comparison.** To generate clusters using both approaches ( $k$ -means and COD-means), we generate  $k_o + n + 1$  clusters, where  $k_o$  is the number of pre-existing non-miscellaneous categories (i.e. 5 in our case),  $n$  is the number of new categories we aim to generate from the  $Z$  category. In this case, we use  $n = 4$  which is set heuristically, as we observe that slightly larger values (from 5 to 10) do not yield significantly different results and can cause labeling errors (as discussed above), and smaller values tend to yield very general categories. The +1 corresponds to the new miscellaneous category  $Z'$ . We vary the number of outliers  $\ell$  from  $\ell = 0 \dots m$ .

Results generated using  $\ell = 0$  represent the standard  $k$ -means algorithm i.e. without outliers detection. We compute cohesiveness and novelty scores for the clusters generated using each value of  $\ell$ . We pick the “top” clusters from the output. We define a “top” cluster as one that is either among the  $m$  having the largest coherence, the  $m$  having the largest novelty, or both, with  $m < k$ . We use  $m = 2$  to reduce the number of annotations needed.

Figure 1 depicts the results for all the datasets. The proposed approach generates more cohesive and novel clusters by removing outliers. As the value of  $\ell$  increases, more tight and coherent clusters are observed. This also helped us determine an optimal value of the  $\ell$  parameter for each dataset (i.e. where the high coherence and novelty scores were noticed) to be used in the next experiments.

#### C. Data improvements evaluations

To achieve high classification accuracy, we aim to discover and remove incorrectly categorized items to miscellaneous and non-miscellaneous categories. We perform the following two data improvements experiments.

##### 1) Labeling errors in non-miscellaneous categories:

As described earlier COD-Means discovers local outliers for each newly generated cluster. To determine whether outliers of the non-miscellaneous clusters are semantically genuine outliers, which we also call labeling errors, we have performed a user study.

To generate clusters and discover outliers, we ran COD-Means using  $k = k_o + n + 1$  where  $k_o$  represents

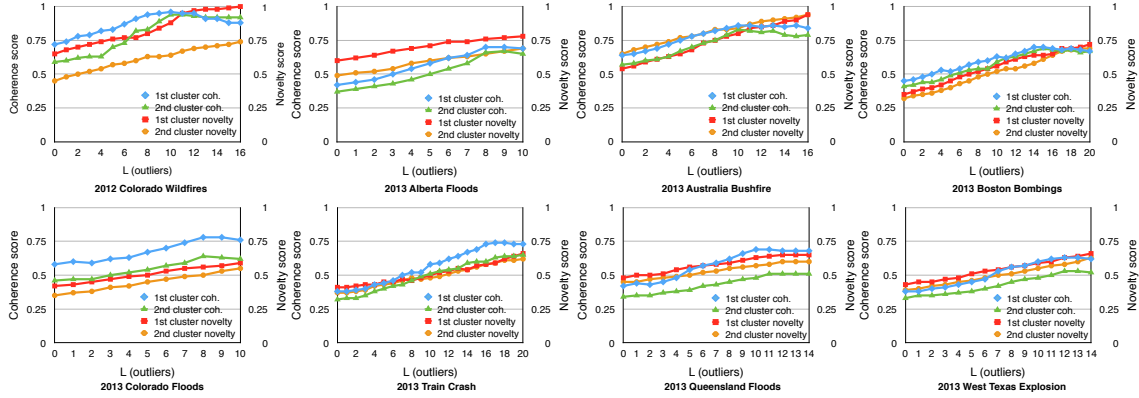


Figure 1: Coherence and novelty scores using different values of parameter  $\ell$  (outliers) of two largest clusters using simple  $k$ -means and COD-means clustering algorithms.

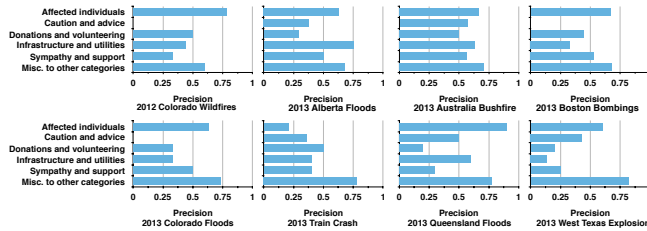


Figure 2: Precision scores of outliers identified in non-misc. categories and items moved from misc. to other categories

the number of non-miscellaneous categories i.e. 5 in all of our datasets. As in this evaluation we are not interested in generating new clusters from  $Z$ , we set  $n = 0$ , and as always we keep  $+1$ , that makes  $k = 6$ . The optimal values for  $\ell$  identified in the previous section were used for each dataset. Labeling errors from each non-miscellaneous clusters are obtained. Table II shows a few examples of wrongly labeled items identified by COD-Means.

Next, we generate crowdsourcing tasks using the identified outliers. A crowdsourcing task consists of an outlier item, its actual category (i.e. name of a  $T_i$ ), and the category description. Workers were asked to read the category description and choose whether the given item is related to the category, or not. We used the CrowdFlower<sup>1</sup> crowdsourcing platform. At least three different workers were required to finalize a task. For this task the inter-annotator agreement was 78%.

Figure 2 depicts the results obtained from the crowdsourcing task in terms of precision. The precision  $P$  is measured as:  $P = \frac{\text{number of correctly identified outliers}}{\text{number of identified outliers}}$ . From the results (Figure 2 and Table II) we can see that the proposed approach can help discover real labeling errors. To achieve a high classification accuracy, wrongly labeled items must be discarded to help classifier achieve higher

<sup>1</sup><http://crowdfower.com/>

generalization accuracy [10], which we test in section IV-D.

2) *Items incorrectly labeled as miscellaneous*: In this task, we aim to evaluate whether items moved from the miscellaneous category to one of the non-miscellaneous categories are genuinely correct. We ran the COD-means algorithm using the same settings specified in section IV-C1. We asked crowd workers to specify whether an item’s newly assigned category is correct or not. The inter-annotator agreement for this task was 76%.

Figure 2 shows the results (see precision scores against “misc. to other category”). The precision  $P$  is measured as:  $P = \frac{\text{number of correctly assigned items}}{\text{number of assigned items}}$ . The results clearly show that a large proportion of newly assigned items to the non-miscellaneous categories are indeed correct assignments. Hence, if used as a training set, they can boost classification accuracy, which we empirically prove in the next section.

#### D. Evaluation in terms of utility as training sets

To further validate the results obtained from the data improvements process, we train machine learning classifiers to compare classification accuracy before and after data improvements.

Two training sets are formed (i.e. original labels and labels after the data improvements process). We remove stop-words, URLs, and user mentions from the items. Stemming is performed using the Lovins stemmer. Uni-grams and bi-grams are used as features and we used the information gain feature selection method to select top 1k features. We use three well-known learning algorithms: SVM (support vector machines), Naive Bayes (NB), and Random forest (RF). We perform the evaluation of the learned model using 10-fold cross validation. Table III shows the classification results in terms of AUC. A substantial gain in AUC can be clearly observed in case of the improved training data.

## V. RELATED WORK

In general cluster analysis methods attempt to form disjoint groups of unlabeled data items such that items in

Table II: Outliers (labeling errors) identified by the COD-Means algorithm in the datasets (user mentions and URLs are masked)

Dataset	Actual category	Outlier
2012_QWF	Sympathy and support	What a week it's been! #highparkfire you better behave today. http://.../
2013_AF	Sympathy and support	What I didn't think about when leaving yesterday electricity shut off...goodbye food in fridge and freezer #yycflood
2013_ABF	Affected individual	RT @UserMention: Another shot of the cricket at #Faulconbridge A full pitch view Photographer- @UserMention #nswfires @UserMention http://.../
2013_CF	Affected individual	#coloradoflood : Stories of grief, generosity http://.../
2013_NYTC	Affected individual	@NTSB: Alcohol tests on #MetroNorth train crew all negative after crash
2013_QF	Infrastructure and utilities	RT @UserMention: Our partners w/@UserMention created this map of the West, TX explosion

Table III: Classification accuracies in terms of AUC for all the datasets before and after the data improvements process (Results of the Sympathy and support class are omitted due to space limitations).

Categories $T_i$		Affected individuals		Caution and advice		Donations and volunteering		Infrastructure and utilities	
Dataset	Classifier	Before	After	Before	After	Before	After	Before	After
CWF	NB	0.77	<b>0.83</b>	0.80	<b>0.92</b>	0.87	<b>0.90</b>	0.85	<b>0.89</b>
	RF	0.77	<b>0.81</b>	0.68	<b>0.90</b>	0.83	<b>0.85</b>	0.84	<b>0.86</b>
	SVM	0.76	<b>0.79</b>	0.74	<b>0.91</b>	0.81	<b>0.87</b>	0.85	<b>0.87</b>
AF	NB	0.84	<b>0.87</b>	0.72	<b>0.74</b>	0.74	<b>0.92</b>	0.57	<b>0.86</b>
	RF	0.68	<b>0.83</b>	0.58	<b>0.71</b>	0.80	<b>0.89</b>	0.70	<b>0.84</b>
	SVM	0.81	<b>0.84</b>	0.67	<b>0.72</b>	<b>0.88</b>	0.87	0.79	<b>0.82</b>
ABF	NB	0.74	<b>0.81</b>	0.76	<b>0.82</b>	0.88	<b>0.91</b>	0.78	<b>0.83</b>
	RF	0.66	<b>0.73</b>	0.77	<b>0.81</b>	0.83	<b>0.91</b>	0.72	<b>0.81</b>
	SVM	0.69	<b>0.78</b>	0.76	<b>0.79</b>	0.83	<b>0.90</b>	0.74	<b>0.81</b>
BB	NB	0.80	<b>0.84</b>	0.66	<b>0.93</b>	0.88	<b>0.91</b>	0.68	<b>0.97</b>
	RF	0.75	<b>0.87</b>	0.56	<b>0.94</b>	0.75	<b>0.87</b>	0.48	<b>0.96</b>
	SVM	0.73	<b>0.81</b>	0.55	<b>0.91</b>	0.79	<b>0.92</b>	0.59	<b>0.97</b>
CF	NB	0.91	<b>0.94</b>	0.78	<b>0.85</b>	0.91	<b>0.92</b>	0.84	<b>0.88</b>
	RF	0.89	<b>0.92</b>	0.72	<b>0.82</b>	0.90	<b>0.92</b>	0.80	<b>0.84</b>
	SVM	0.87	<b>0.90</b>	0.72	<b>0.80</b>	0.90	0.90	0.81	<b>0.85</b>
NYTC	NB	0.93	0.93	<b>0.72</b>	0.68	0.71	0.71	0.91	<b>0.92</b>
	RF	0.93	<b>0.94</b>	0.58	<b>0.60</b>	0.49	0.49	0.83	<b>0.89</b>
	SVM	<b>0.93</b>	0.92	0.53	<b>0.54</b>	0.48	0.48	0.74	<b>0.88</b>
QF	NB	0.73	<b>0.87</b>	0.77	<b>0.79</b>	0.85	<b>0.93</b>	0.76	<b>0.88</b>
	RF	0.69	<b>0.85</b>	0.69	<b>0.73</b>	0.75	<b>0.93</b>	0.77	<b>0.87</b>
	SVM	0.69	<b>0.86</b>	0.69	<b>0.72</b>	0.85	<b>0.92</b>	0.75	<b>0.87</b>
WTE	NB	0.89	<b>0.93</b>	0.80	<b>0.95</b>	0.95	<b>0.96</b>	0.82	<b>0.93</b>
	RF	0.85	<b>0.90</b>	0.61	<b>0.91</b>	0.82	<b>0.86</b>	0.69	<b>0.91</b>
	SVM	0.80	<b>0.86</b>	0.54	<b>0.91</b>	0.85	<b>0.87</b>	0.74	<b>0.92</b>

same group are similar while items in different groups are dissimilar. For instance, one such famous clustering method is  $k$ -means [11]. However, in a semi-supervised clustering approach, the algorithm uses both labeled and unlabeled data. In this particular case, the labeled data items are used as background knowledge during the clusters generation process.

Many works described the use of constraints in different ways, e.g. some use constraints at group-level (i.e. on group of items) and other use at item-level (i.e. between two items) as a way to provide background knowledge [12–14]. For instance, in [14]  $k$ -means algorithm is extended to use instance-level constraints. However, our method, in addition to the instance-level constraints, discovers and removes outliers during the cluster generation process by which more compact clusters can be obtained.

While it is well-known that outlier detection in a training set can help improve the accuracy of a classifier built using that training set (e.g. [10]), previous methods do not take into consideration the existing categories. Our method unifies constrained clustering and outlier detection by formulating a novel optimization problem and algorithm COD-Means.

## VI. CONCLUSIONS

As supervised learning systems can not be used to identify novel concepts, for this purpose, we employ unsupervised learning techniques. We presented a novel clustering algorithm COD-Means, which uses human and machine categorized items as background knowledge to form constraints and detects novel categories. The proposed algorithm, not only help detect novel categories, but also seamlessly discover outliers from each cluster by which categorization errors are fixed. Extensive experiments using real datasets demonstrate that our approach is effective and efficient.

## REFERENCES

- [1] M. Imran, I. Lykourantzou, and C. Castillo, "Engineering crowd-sourced stream processing systems," *CoRR*, vol. abs/1310.5463, 2013.
- [2] A. J. Quinn, B. B. Bederson, T. Yeh, and J. Lin, "CrowdFlow: Integrating machine learning with mechanical Turk for Speed-Cost-quality flexibility," in *CHI 2011 Workshop on Crowdsourcing and Human Computation*, 2011.
- [3] M. Imran, C. Castillo, J. Lucas, P. Meier, and S. Vieweg, "AIDR: Artificial intelligence for disaster response," in *Proc. of 23rd WWW companion*, 2014, pp. 159–162.
- [4] H. Rodriguez, E. L. Quarantelli, and R. Dynes, Eds., *Handbook of Disaster Research*, 1st ed., ser. Handbooks of Sociology and Social Research. Springer, Jul. 2006.
- [5] S. E. Vieweg, "Situational awareness in mass emergency: A behavioral and linguistic analysis of microblogged communications," 2012.
- [6] S. Chawla and A. Gionis, "k-means--: A unified approach to clustering and outlier detection," in *SDM*, 2013, pp. 189–197.
- [7] M. Imran and C. Castillo, "Towards a data-driven approach to identify crisis-related topics in social media streams," in *Proc. of 24th WWW Companion*, 2015, pp. 1205–1210.
- [8] I. Davidov and S. S. Ravi, "Clustering with constraints: Feasibility issues and the k-means algorithm," in *Proceedings of the 2005 SIAM International Conference on Data Mining*, 2005, pp. 138–149.
- [9] A. Olteanu, S. Vieweg, and C. Castillo, "What to expect when the unexpected happens: Social media communications across crises," in *In Proc. of 18th ACM CSCW*, 2015.
- [10] L. J. Cao, H. P. Lee, and W. K. Chong, "Modified support vector novelty detector using training data with outliers," *Pattern Recognition Letters*, 2003.
- [11] J. A. Hartigan and M. A. Wong, "Algorithm as 136: A k-means clustering algorithm," *Applied statistics*, pp. 100–108, 1979.
- [12] E. P. Xing, M. I. Jordan, S. Russell, and A. Y. Ng, "Distance metric learning with application to clustering with side-information," in *Advances in neural information processing systems*, 2002, pp. 505–512.
- [13] S. Basu, A. Banerjee, and R. J. Mooney, "Active semi-supervision for pairwise constrained clustering," in *SDM*, vol. 4. SIAM, 2004, pp. 333–344.
- [14] K. Wagstaff, C. Cardie, S. Rogers, and S. Schrödl, "Constrained k-means clustering with background knowledge," in *ICML*, vol. 1, 2001, pp. 577–584.