

Letter

Exponentiated Gradient Exploration for Active Learning

Djallel Bouneffouf

Received: 24 November 2015; Accepted: 5 January 2016; Published: 8 January 2016

Academic Editor: Pedro Alonso Jordá

Department of Computer Science, Télécom SudParis, UMR CNRS Samovar, 91011 Evry Cedex, France;
Djallel.Bouneffouf@it-sudparis.eu; Tel.: +33-6-7584-5063

Abstract: Active learning strategies respond to the costly labeling task in a supervised classification by selecting the most useful unlabeled examples in training a predictive model. Many conventional active learning algorithms focus on refining the decision boundary, rather than exploring new regions that can be more informative. In this setting, we propose a sequential algorithm named exponentiated gradient (EG)-active that can improve any active learning algorithm by an optimal random exploration. Experimental results show a statistically-significant and appreciable improvement in the performance of our new approach over the existing active feedback methods.

Keywords: active learning; exploration and exploitation; exponentiated gradient

1. Introduction

In active learning, the learning algorithm has access to a large set of unlabeled examples and some oracle that can query to get the label of an individual example [1]. Query oracles are assumed to be expensive, so it is only feasible to label a small subset. Thus, the goal of the algorithm is to actively select examples so that a good hypothesis is learned while using as few labeled examples as possible.

Many conventional active learning algorithms choose to label points that are near the decision boundary of the current hypothesis [2,3]. This can work well if the active learner is aware of all of the important regions of the instance space, *i.e.*, there are no large examples that the learner's hypothesis will misclassify, since it has not seen labeled examples from them. Such active learners are good at labeling examples near the boundary to refine it, but they do not conduct a random searching for large regions in the instance space that they would incorrectly classify.

Recent work done in this sense considers random exploration in active learning. The author in [4] uses two types of active learners, the first of which is dedicated to exploiting based on refining the decision boundary and the second is dedicated to exploring (random). The more the random exploration gets rewarded, the greater is the exploration.

The drawback of this approach is the fact that it takes a long time to find the optimal random exploration rate. To tackle this problem, we propose an algorithm named exponentiated gradient (EG)-active which can improve any existing active learning algorithm by using a random exploration, which is parametrized with an exponentiated gradient optimization.

We have tested EG-active with real data, where we have observed its performance regarding the state of the art.

The remainder of the paper is organized as follows. Section 2 reviews related works. Section 3 describes the model and the proposed algorithm. The experimental evaluation is illustrated in Section 4. The last section concludes the paper and points out possible directions for future works.

2. Related Works

We refer in the following to recent works that address the active learning problem.

2.1. Active Learning

A variety of active learning algorithms have been proposed in the literature employing various query strategies. One of the most popular strategies is called uncertainty sampling (US), where the active learner queries the point whose label is uncertain [5].

The uncertainty in the label is usually calculated using entropy or the variance of the label distribution [6,7]. The authors in [8] have introduced the query by committee (QBC) strategy, where a committee of potential models, which all agree with the currently labeled data, is maintained, and the point where most of the committee members disagree is considered for querying. Other strategies include the maximum expected reduction in error [9] or variance-reducing query strategies, such as [10], to query the optimal point.

Recently, different new active learning processes have been proposed. For instance, in [11], the authors propose a strategy for data-driven classifiers, which is based on an unsupervised criterion during the off-line training phase, followed by a supervised certainty-based criterion during incremental on-line training. In [12], the authors propose an effective and efficient active learning paradigm by applying an *a priori* process for the identification and organization of a small relevant subset. Furthermore, the concomitant classification and selection processes enable the classification of a very small number of samples, while selecting the informative ones. To adapt active learning to online learning, the authors in [13] propose an active learning with two concepts called “conflict and ignorance”; conflict models the extent to which a new query point lies in the conflict region between two or more classes and, therefore, reflects the level of certainty in the classifier’s prediction; ignorance represents the distance of a new query point from the training samples seen so far.

All of the above proposed approaches have just exploited the data and do not consider the random exploration that can help to find the best point to label.

2.2. Random Exploration in Active Learning

Recently, random exploration has been used in different domains, such as recommender system (RS) and information retrieval. For example, in [14,15], the authors model RS as a contextual bandit problem. The authors propose an algorithm that provides a random recommendation according to the risk of upsetting the user. However, to our knowledge, there has been only one paper addressing random exploration in active learning. The authors in [4] address this problem by randomly choosing between exploration and exploitation at each round and then receive feedback on how effective the exploration is. The impact of exploration is measured by the induced change in the learned classifier when an exploratory example is labeled and added to the training set. The active learner updates the probability of exploring in subsequent rounds based on the feedback it has received. However, none of the optimization techniques are used to compute the optimal exploration, and the work has only been done to improve the uncertainty sampling technique.

2.3. Our Contributions

As shown above, none of the mentioned works propose to improve any active learning by random exploration. This is precisely what we intend to do by exploiting the following new features:

- (1) We propose a new generic algorithm named EG-active that can improve the results of any active learning algorithm by considering the exploration at each iteration.
- (2) We propose to parametrize this exploration using an exponentiated gradient that allows EG-active to use the optimal exploration.

3. Active Learning with Random Exploration

This section focuses on the proposed model, starting by introducing the key notions used in this paper.

Pool-based active learning:

In pool-based active learning, we have provided a pool $X = \{x_1, \dots, x_n\}$ of unlabeled points and a labeling oracle O , which when queried for the label of x returns $y \sim P_Y|X = x$. Algorithms in the pool-based setting have to decide which points to query by looking at the entire pool.

Reward:

A metric is used to measure the variation of the hypotheses learned by the model between two iterations. The more the hypotheses learned by the model vary, the more is the reward. We now define the function $d(h, h')$ that we use to get the variation of the model.

Let $X = x_1, \dots, x_m = L \cup U$ be the set of labeled and unlabeled training examples that we have. Then, for each of the two real-valued hypotheses $h(\cdot), h'(\cdot)$, we define the vectors $H = (h(x_1), h(x_2), \dots, h(x_m))$ and $H' = (h'(x_1), h'(x_2), \dots, h'(x_m))$, i.e., vectors of the real-valued predictions of h and h' on X .

Now, we define $d(h, h')$,

$$d(h, h') = \frac{H \cdot H'}{\|H\| \cdot \|H'\|} \quad (1)$$

In Equation (1), we compute the cosine similarity between the two vectors H and H' . Thus, $d(h, h') \in [1, +1]$ is the cosine of the angle between H and H' , and we normalize the ratio of classes in the interval $[0, 1]$ using Equation (2).

$$r_t = \frac{2 \cdot \cos^{-1}(d(h|h'))}{\pi} \quad (2)$$

3.1. ϵ -active

In order to improve the results of any active learning algorithm, we propose to overlap any existing algorithm by an algorithm that considers at each iteration a random exploration ϵ .

In Algorithm 1, *Activelearning* can be any existing active learning, for example query by committee, uncertainty sampling or others.

Algorithm 1 ϵ -active

- 1: **Input:** X, ϵ
 - 2: **Output:** x_t, r_t
 - 3: $x_t = \begin{cases} \text{Activelearning}(X) & \text{if } (q < \epsilon) \\ \text{Random}(X) & \text{if } (q \geq \epsilon) \end{cases}$
 - 4: **if** x was not queried in the past **then** Query O for label y of x
 - 5: Observe reward r_t
-

3.2. Computing the Optimal Random Exploration

To consider the random exploration on the active learning algorithms, the proposed method updates the exploration value ϵ dynamically.

In each iteration, the algorithm runs a sampling procedure to select a new ϵ from a ϵ finite set of candidates. The probabilities associated with the candidates are uniformly initialized and updated with the exponentiated gradient (EG) [16]. This updating rule increases the probability of a candidate ϵ if it leads to a reward.

First, we assume that we have a finite number of candidate values for ϵ , denoted by $(\epsilon_1, \dots, \epsilon_T)$, and we try to learn the optimal ϵ from this set. To this end, the EG-active introduces $p = (p_1, \dots, p_T)$, where p_i stands for the probability of using ϵ_i in the ϵ -active algorithm. These probabilities are initialized to be $\frac{1}{T}$ at the beginning and then iteratively updated through iterations.

The algorithm uses a set of weights $w = (w_1, \dots, w_T)$ to keep track of the performance of each ϵ_i and updates them using the EG algorithm. The idea is to increase w_i if the algorithm receives a reward using ϵ_i . Finally, the algorithm calculates p by normalizing w with smoothing. Algorithm 2 shows EG-active.

Algorithm 2 EG-active.

Input: $(\epsilon_1, \dots, \epsilon_T)$: candidate values for ϵ
 β, τ and k : parameters for EG
 N : number of iterations
 $p_k \leftarrow \frac{1}{T}$ and $w_k \leftarrow 1, k = 1, \dots, T$
for $i=1$ **to** N **do**
 Sample d from discrete (p_1, \dots, p_T)
 Run the ϵ -active with ϵ_d
 Receive the feedback r_t
 $w_k \leftarrow w_k \exp(\frac{\tau[r_t I(k=d) + \beta]}{p_k}), k = 1, \dots, T$
 $p_k \leftarrow (1 - k)(\frac{w_k}{\sum_{j=1}^T w_j} + \frac{k}{T}), k = 1, \dots, T$
end for

In Algorithm 2, $I[z]$ is the indicator function, and τ and β are smoothing factors in the weight updating. k is a regularization factor to handle singular w_i .

4. Experimental Evaluation

To conduct our experiments, we have evaluated our algorithm in both corporate data and public datasets from the University of California, Irvine (UCI), Machine Learning Repository (<https://archive.ics.uci.edu/ml/datasets.html>).

4.1. Corporate Data

We have obtained from our company a corpus containing utterances in French of a typical communication between a customer and a call center of a telecom company.

There are 7765 utterances annotated by human experts that have been collected in four different datasets. The unannotated part consists of 3,911,695 utterances.

We use a corporate supervised algorithm (rule-based algorithm). We simulate in the experiments an expert (oracle) on the unannotated corpus by using the rule-based algorithm, which is trained by 7765 utterances. Note that the objective of this evaluation is to observe the improvement that can add the random exploration to the existing active learning.

In our experiments, we consider a version of the rule-based algorithm without training, where at each iteration, the active learning tries to select from the unannotated corpus the most interesting utterances to annotate and integrate into the training set of the rule-based algorithm.

By relating the results to the newer versions, one can verify the usefulness of the proposed approach. Moreover, we calculate the regret every 100 iterations, and we run the process during 2000 iterations, which corresponds to our budget in terms of labeling.

In addition to the randomness (baseline), we compare our methods by constructing four groups of algorithms: the first group is the state-of-the-art algorithms described in the related work (Section 2), which are the sampling by committee, request uncertainty sampling and density weight method (DW).

The second group contains modified state-of-the-art algorithms, where we have added a fixed random exploration to the existing state-of-the-art algorithms, for example 0.5-QBC means that with a probability of 0.5, the algorithm does a random exploration, and otherwise, it does request by committee.

The third group contains the proposed model EG-active tested with different existing algorithms, for example EG-active(committee) means that we have used sampling by committee in our model.

In the fourth group, we have added a dynamic random exploration to the existing state-of-the-art algorithms as is done in [4]; for example, P-US, P-QBC and P-WD are algorithms that do a random exploration with probability P , and otherwise, they do respectively P-US, P-QBC or P-WD, these algorithms use the strategy of [4] to compute the probability P of the random exploration.

In Figure 1, the horizontal axis represents the number of iterations, and the vertical axis is the performance metric.

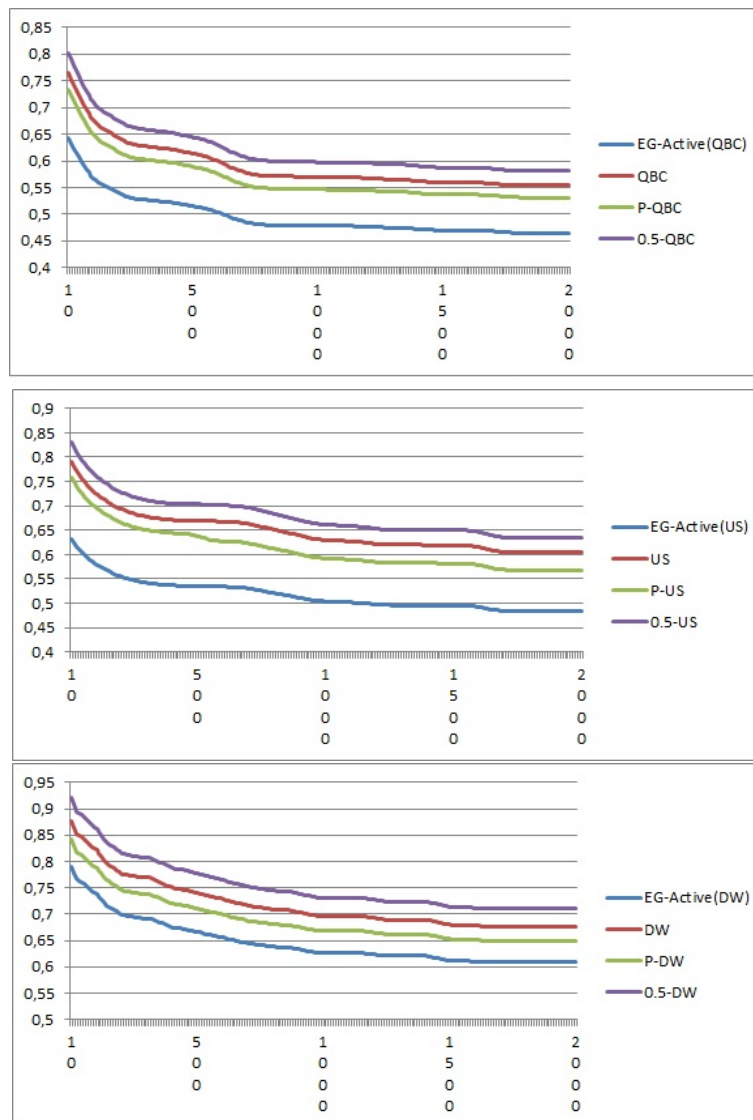


Figure 1. Average regret for active learning algorithms.

We have several observations regarding the different active learning algorithms. We observe from the plot that a fixed and non-tuned random exploration leads to a bad result. This confirms that a pure exploration is not interesting, and it justifies the need for a dynamic random exploration tuning.

A dynamic exploration leads to an improvement result of the active learning, as is shown by P-US, P-QBC and P-WD. As expected, EG-active(US), EG-active(QBC) and EG-active(WD) effectively have the best convergence rates.

EG-active(US), EG-active(QBC) and EG-active(WD) decrease the average regret respectively by a factor of 0.84, 1.3 and 0.9 over the baseline. The improvement comes from an optimization strategy for defining exploration.

These algorithms rapidly find the optimal random exploration to use, which is not the case of the P-US, P-QBC and P-WD, which take more time.

4.2. Public Benchmarks

We randomly chose nine datasets: Abalone, Breast, Ecoli, Glass, Haberman, Iris, Wine, Wdbc and Yeast. A brief summary of the datasets is listed in Table 1.

Table 1. Datasets used for benchmarking. UCI, University of California, Irvine.

UCI Datasets	Instances	Attributes	Classes
Abalone	1484	7	3
Breast	699	9	2
Ecoli	336	7	8
Glass	214	9	7
Haberman	306	3	2
Iris	150	4	3
Wine	178	13	3
Wdbc	569	32	2
Yeast	1484	6	8

We simulate in the experiments an expert (oracle) on the unannotated corpus by using support vector machine (SVM) [17], which is trained by 100% of the dataset. We consider a version of the SVM algorithm without training; where at each iteration, the active learning tries to select from the unannotated points the most interesting points to annotate and integrates them into the training set of the SVM algorithm.

We run the process until the algorithm reaches 10% of the dataset, and because active learning algorithms contain a degree of randomness, we repeat our evaluations 100 times. We measured the classification quality using the average accuracy.

We observe from Tables 2–4 that the results in public datasets confirm our expectation, where in the three algorithms tested, the proposed EG-active performs better than the other versions. The gap between the EG-active results and the original algorithms depends on the dataset. For instance, the gap is smaller in the Yeast, Ecoli and Iris datasets than in the rest of the datasets. We also observe that the performance of EG-active depends on the type of algorithm that we use.

Table 2. Accuracy of the query by committee (QBC) versions on the UCI datasets.

UCI Datasets	QBC	50-QBC	P-QBC	EG-active(QBC)
Abalone	84.82	83.74	85.69	86.09
Breast	67.85	66.86	68.83	70.86
Ecoli	69.82	67.90	69.97	70.09
Glass	71.77	60.34	72.66	74.94
Haberman	56.55	52.46	56.96	58.48
Iris	66.45	61.18	67.06	68.48
Wine	53.16	45.17	54.78	56.84
Wdbc	51.32	41.31	51.30	53.31
Yeast	67.58	60.70	68.92	69.62

Table 3. Accuracy of the uncertainty sampling (US) versions on the UCI datasets.

UCI Datasets	US	50-US	P-US	EG-active(US)
Abalone	80.69	76.09	84.44	85.91
Breast	60.83	57.86	61.86	62.25
Ecoli	60.87	59.09	62.45	61.92
Glass	55.66	59.94	73.32	76.44
Haberman	51.46	46.48	52.44	53.08
Iris	70.06	65.48	69.08	69.94
Wine	48.78	43.84	56.99	71.39
Wdbc	49.30	45.21	50.32	51.82
Yeast	66.92	57.62	68.46	67.87

Table 4. Accuracy of the density weight method (DW) versions on the UCI datasets.

UCI Datasets	DW	50-DW	P-DW	EG-active(DW)
Abalone	78.23	73.61	80.68	83.46
Breast	58.22	57.01	60.32	60.35
Ecoli	59.78	59.45	61.12	60.45
Glass	53.45	55.44	61.12	66.43
Haberman	50.09	44.28	51.34	52.18
Iris	66.44	63.08	67.98	67.98
Wine	47.53	41.99	56.84	61.36
Wdbc	47.21	43.11	48.86	49.82
Yeast	64.81	55.71	67.44	66.91

5. Conclusions

In this paper, we have proposed an improvement of active learning by considering a random exploration. We have validated our work with data from real-world applications and showed that the proposed model offers promising results. This study yields the conclusion that an optimal random exploration used with any active learning algorithm increases its result. In considering these results, we plan to try to explain why the gap between the EG-active results and the rest of algorithms depends on the dataset.

Conflicts of Interest: The author declares no conflict of interest.

References

1. Cohn, D.; Atlas, L.; Ladner, R. Improving Generalization with Active Learning. *Mach. Learn.* **1994**, *15*, 201–221.
2. Settles, B. Active Learning Literature Survey. In *Computer Sciences Technical Report 1648*; University of Wisconsin: Madison, WI, USA, 2010.
3. Koronacki, J.; Ras, Z.W.; Wierzchoń, S.T.; Kacprzyk, J. Advances in Machine Learning I. In *Studies in Computational Intelligence*; Springer: Berlin/Heidelberg, Germany, 2010; Volume 262, pp. 97–120.
4. Osugi, T.; Kim, D.; Scott, S. Balancing exploration and exploitation: A new algorithm for active machine learning. In Proceedings of the Fifth IEEE International Conference on Data Mining, Houston, TX, USA, 27–30 November 2005; p. 8.
5. Lewis, D.D.; Gale, W.A. A sequential algorithm for training text classifiers. In Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '94, Dublin, Ireland, 3–6 July 1994; Springer-Verlag: New York, NY, USA, 1994; pp. 3–12.
6. Hamed, H.; MohammadReza, K. A variance based active learning approach for named entity recognition. In *Intelligent Computing and Information Science*; Springer: Berlin/Heidelberg, Germany, 2011; pp. 347–352.

7. Jing, F.; Li, M.; Zhang, H.-J.; Zhang, B. Entropy-based active learning with support vector machines for content-based image retrieval. In Proceedings of the 2004 IEEE International Conference on Multimedia and Expo, ICME'04, Taipei, Taiwan, 27–30 June 2004; pp. 85–88.
8. Seung, H.S.; Oppier, M.; Sompolinsky, H. Query by committee. In Proceedings of the Fifth Annual Workshop on Computational Learning Theory, COLT '92, New York, NY, USA, 27–29 July 1992; pp. 287–294.
9. Zhu, X.; Lafferty, J.; Ghahramani, Z. Combining active learning and semi-supervised learning using gaussian fields and harmonic functions. In Proceedings of the ICML 2003 Workshop on The Continuum from Labeled to Unlabeled Data in Machine Learning and Data Mining, Washington, DC, USA, 21 August 2003; pp. 58–65.
10. Zhang, T.; Oles, F.J. A probability analysis on the value of unlabeled data for classification problems. In Proceedings of the 17th International Conference on Machine Learning, Stanford, CA, USA, 29 June–2 July 2000.
11. Lughofer, E. Hybrid active learning for reducing the annotation effort of operators in classification systems. *Pattern Recognit.* **2012**, *45*, 884–896.
12. Saito, P.T.M.; de Rezende, P.J.; Falcao, A.X.; Suzuki, C.T.N.; Gomes, J.F. An active learning paradigm based on a priori data reduction and organization. *Expert Syst. Appl.* **2014**, *41*, 6086–6097.
13. Lughofer, E. Single-pass active learning with conflict and ignorance. *Evol. Syst.* **2012**, *3*, 251–271.
14. Bouneffouf, D. Freshness-Aware Thompson Sampling. In Proceedings of the 21st International Conference on Neural Information, ICONIP, Kuching, Malaysia, 3–6 November 2014; pp. 373–380.
15. Bouneffouf, D. Context-Based Information Retrieval in Risky Environment. *Aust. J. Intell. Inf. Proc. Syst.* **2014**, *14*, 1–10.
16. Kivinen, J.; Warmuth, M.K. Exponentiated gradient versus gradient descent for linear predictors. *Inf. Comput.* **1995**, *132*, 1–63.
17. Mitra, V.; Wang, C.-J.; Banerjee, S. Text classification: A least square support vector machine approach. *Appl. Soft Comput.* **2007**, *7*, 908–914.



© 2016 by the author; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons by Attribution (CC-BY) license (<http://creativecommons.org/licenses/by/4.0/>).