# Pool-Based Sequential Active Learning for Regression

Dongrui Wu<sup>(D)</sup>, Senior Member, IEEE

Abstract—Active learning (AL) is a machine-learning approach for reducing the data labeling effort. Given a pool of unlabeled samples, it tries to select the most useful ones to label so that a model built from them can achieve the best possible performance. This paper focuses on pool-based sequential AL for regression (ALR). We first propose three essential criteria that an ALR approach should consider in selecting the most useful unlabeled samples: informativeness, representativeness, and diversity, and compare four existing ALR approaches against them. We then propose a new ALR approach using passive sampling, which considers both the representativeness and the diversity in both the initialization and subsequent iterations. Remarkably, this approach can also be integrated with other existing ALR approaches in the literature to further improve the performance. Extensive experiments on 11 University of California, Irvine, Carnegie Mellon University StatLib, and University of Florida Media Core data sets from various domains verified the effectiveness of our proposed ALR approaches.

*Index Terms*—Active learning (AL), inductive learning, passive sampling, ridge regression, transductive learning.

#### I. INTRODUCTION

CTIVE learning (AL) [34], a subfield of machine learning, considers the following problem: if the learning algorithm can choose the training data, then, which training samples should it choose to maximize the learning performance, under a fixed budget, e.g., the maximum number of labeled training samples? As an example, consider emotion estimation in affective computing [29]. Emotions can be represented as continuous numbers in the 2-D space of arousal and valence [31] or in the 3-D space of arousal, valence, and dominance [27]. However, emotions are very subjective, subtle, and uncertain. Therefore, usually, multiple human assessors are needed to obtain the groundtruth emotion values for each affective sample (video, audio, image, physiological signal, and so on). For example, 14-16 assessors were used to evaluate each video clip in the Database for Emotion Analysis Using Physiological Signals data set [22], 6-17 assessors for each utterance in the VAM (Vera am Mittag in German, Vera at Noon in English) spontaneous speech corpus [16], and at least 110 assessors for each sound in the International Affective Digitized Sounds 2nd Edition (IADS-2)

Manuscript received December 27, 2017; revised May 12, 2018 and July 27, 2018; accepted September 2, 2018. Date of publication September 27, 2018; date of current version April 16, 2019.

The author is with DataNova LLC, Clifton Park, NY 12065 USA (e-mail: drwu09@gmail.com).

This paper has supplementary downloadable material available at http://ieeexplore.ieee.org, provided by the author.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNNLS.2018.2868649

data set [5]. This is very time consuming and labor intensive. How should we optimally select the affective samples to label so that an accurate regression model can be built with the minimum cost (i.e., the minimum number of labeled samples)? That is the typical type of problems that AL targets at.

Many AL approaches have been proposed in the literature [1], [6], [7], [9]–[11], [15], [24], [30], [32]–[35]. According to the query scenario, they can be categorized into two groups [38]: *population based* and *pool based*. In population-based AL, the test input distribution is known, and training input samples at any desired locations can be queried. Its goal is to find the optimal training input density to generate the training input samples. In pool-based AL, a pool of unlabeled samples is given, and the goal is to optimally choose some to label, so that a model trained from them can best label the remaining samples.

Regardless of whether it is population based or pool based, typically, AL is iterative [7]. It first builds a base model from a small number of labeled training samples and then chooses a few most helpful unlabeled samples and queries for their labels. The newly labeled samples are then added to the training data set and used to update the model. This process iterates until a termination criterion is met, e.g., the maximum number of iterations, the maximum number of labeled samples, or the desired cross-validation accuracy, is reached. Depending on the number of unlabeled samples to query in each iteration, AL approaches can also be categorized into two types [7]: *sequential* AL, where one sample is queried each time, and *batch-mode* AL, where multiple samples are queried in each iteration.

This paper focuses on pool-based sequential AL for regression (ALR). Although numerous AL approaches have been proposed in the literature [34], most of them are for classification problems. Among those limited number of ALR approaches [6]–[8], [10], [15], [25], [37]–[41], only a few can be used for pool-based sequential ALR [6], [8], [40], [41]. In this paper, we will review them, point out their limitations, and propose approaches to enhance their performance.

The main contributions of this paper are as follows.

- We extend three criteria for AL—informativeness, representativeness, and diversity—from classification to regression, and propose a generic framework that can be used to enhance a baseline (BL) ALR approach.
- We instantiate several ALR approaches that consider informativeness, representativeness, and diversity simultaneously and demonstrate their promising performances in extensive application domains.

2162-237X © 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications\_standards/publications/rights/index.html for more information.



Fig. 1. Illustration of RD in pool-based sequential ALR.

The remainder of this paper is organized as follows: Section II introduces three essential criteria that should be considered in ALR and then compares several existing poolbased sequential ALR approaches against them. Section III proposes several new pool-based sequential ALR approaches. Section IV describes the data sets to evaluate the effectiveness of the proposed ALR approaches and the corresponding experimental results. Finally, Section V draws the conclusion.

# II. EXISTING POOL-BASED SEQUENTIAL ALR APPROACHES

In this section, we propose three essential criteria for selecting unlabeled samples in pool-based sequential ALR and then introduce a few existing pool-based sequential ALR approaches. We also compare these ALR approaches against the three criteria and point out their limitations.

Without loss of generality, we assume that the pool consists of *N d*-dimensional samples  $\{\mathbf{x}_n\}_{n=1}^N$ ,  $\mathbf{x}_n \in \mathbb{R}^d$ , and the first  $M_0$  samples have already been labeled with labels  $\{y_n\}_{n=1}^{M_0}$ .

#### A. Three Essential Criteria in ALR

We propose the following three criteria that should be considered in pool-based sequential ALR for selecting the most useful unlabeled sample to the label.

- Informativeness, which means that the selected sample must contain rich information, so labeling it would significantly benefit the objective function. Informativeness could be measured by uncertainty (entropy, distance to the decision boundary, confidence of the prediction, and so on), expected model change, expected error reduction, and so on [34]. For example, in query-bycommittee (QBC), a popular AL approach for both classification and regression [30], the informativeness of an unlabeled sample could be computed as the disagreement among the committee members: the more disagreement is, the more uncertain the sample is, and hence, the more informative it is.
- 2) *Representativeness*, which can be evaluated by the number of samples that are similar or close to a target sample (or its density [34]): the larger the number is, the more representative the target sample is. Clearly, the target sample should not be an outlier. For example, in Fig. 1, assume that we want to build a regression model to predict the output from  $x_1$  and  $x_2$ . The gray circle "B" is very likely to be an outlier because it is very far away

from all other samples in the input space, so labeling it could mislead the regression model and result in overall worse prediction performance. In other words, a sample like "B" should not be selected for labeling by ALR.

3) Diversity, which means that the selected samples should scatter across the full input space, instead of concentrating in a small local region. For example, in Fig. 1, the unlabeled samples from three clusters in the input space, so we should select samples from all three clusters to label, instead of focusing on only one or two of them. Assume that the two green circles have been selected and labeled. Then, selecting next a sample from the third cluster (the one contains "A") seems very reasonable.

We should point out that similar criteria have been used in AL for classification (ALC). For example, Shen et al. [36] proposed two multicriteria-based batch-mode ALC strategies, both of which considered informativeness, representativeness, and diversity simultaneously. Their Strategy 1 first chooses a few most informative samples, clusters them, and then selects the cluster centroids for labeling. Their Strategy 2 first computes a score for each sample as a linear combination of its informativeness and representativeness, selects samples with high scores, and further down selects among them the most diverse ones for labeling. Both strategies are specific to the support vector machine classifier. He et al. [19] considered uncertainty, representativeness, information content, and diversity in batch-mode ALC. Let k be the batch size. They compute the information content of an unlabeled sample as uncertainty × representativeness, select the most informative samples, cluster them into k clusters by kernel k-means clustering, and finally, select the k cluster centers for labeling.

However, to our knowledge, similar ideas have not been explored in ALR, except our recent work on enhanced batchmode ALR (EBMALR) [40]. It is not trivial to extend these concepts from classification to regression, because there could be many different strategies to integrate these three criteria, and different strategies could result in significantly different performances. EBMALR [40] is one of such strategies. However, although our previous research [40] showed that it achieved promising performance in *batch-mode* ALR in braincomputer interface, this paper (Section IV-E) shows that it does not perform well in *sequential* ALR. Therefore, how to integrate informativeness, representativeness, and diversity to design high-performance pool-based *sequential* ALR is still an open problem.

Next, we will introduce several existing pool-based sequential ALR approaches and compare their rationale against our three criteria.

#### B. Query-by-Committee

QBC is a very popular pool-based AL approach for both classification [1], [15], [26], [34], [35] and regression [6], [10], [11], [24], [30], [34]. Its basic idea is to build a committee of learners from existing labeled training data set (usually through bootstrapping and/or different learning algorithms) and then select from the pool the unlabeled samples on which the committee disagrees the most to label.

In this paper, we use the pool-based QBC for regression approach proposed by RayChaudhuri and Hamey [30]. It first bootstraps the  $M_0$  labeled samples into P copies, each containing  $M_0$  samples but with duplicates and builds a regression model from each copy, i.e., the committee consists of P regression models. Let the pth model's prediction for the nth unlabeled sample be  $y_n^p$ . Then, for each of the  $N - M_0$  unlabeled sample, it computes the variance of the P individual predictions, that is,

$$\sigma_n = \frac{1}{P} \sum_{p=1}^{P} (y_n^p - \bar{y}_n)^2, \quad n = M_0 + 1, \dots, N$$
(1)

where  $\bar{y}_n = (1/P) \sum_{p=1}^{P} y_n^p$ , and then selects the sample with the maximum variance to label.

Comparing against the three criteria for ALR, QBC only considers the informativeness, but not the representativeness and the diversity.

#### C. Expected Model Change Maximization

Expected model change maximization (EMCM) is also a very popular AL approach for classification [9], [32]–[34], regression [7], [8], and ranking [12]. Cai *et al.* [8] proposed an EMCM approach for both linear and nonlinear regression. In this section, we introduce their linear approach, as only linear regression is considered in this paper.

EMCM first uses all  $M_0$  labeled samples to build a linear regression model. Let its prediction for the *n*th sample  $\mathbf{x}_n$  be  $\hat{y}_n$ . Then, like in QBC, EMCM also uses bootstrap to construct *P* linear regression models. Let the *p*th model's prediction for the *n*th sample  $\mathbf{x}_n$  be  $y_n^p$ . Then, for each of the  $N - M_0$  unlabeled samples, it computes

$$g(\mathbf{x}_n) = \frac{1}{P} \sum_{p=1}^{P} \| (y_n^p - \hat{y}_n) \mathbf{x}_n \|, \quad n = M_0 + 1, \dots, N.$$
 (2)

EMCM selects the sample with the maximum  $g(\mathbf{x}_n)$  to label.

Comparing against the three criteria for ALR, EMCM only considers the informativeness, but not the representativeness and the diversity.

#### D. Greedy Sampling

Yu and Kim [41] proposed several very interesting passive sampling techniques for regression. Instead of finding the most informative sample based on the learned regression model, as in QBC and EMCM, they select the sample based on its geometric characteristics in the feature space. An advantage of passive sampling is that it does not require updating the regression model and evaluating the unlabeled samples in each iteration, so it is independent of the regression model.

In this paper, we use the greedy sampling (GS) for regression approach [41], which is easy to implement, and showed promising performance in [41]. Its basic idea is to select a new sample in a greedy way such that it is located far away from the previously selected and labeled samples. More specifically, for each of the  $N - M_0$  unlabeled sample { $\mathbf{x}_n$ } $_{n=M_0+1}^N$ ,

it computes its distance to each of the  $M_0$  labeled samples, that is,

$$d_{nm} = ||\mathbf{x}_n - \mathbf{x}_m||, \quad m = 1, \dots, M_0; n = M_0 + 1, \dots, N$$
(3)

then, it computes  $\underline{d}_n$  as the minimum distance from  $\mathbf{x}_n$  to the  $M_0$  labeled samples, that is,

$$\underline{d}_n = \min_{m} d_{nm}, \quad n = M_0 + 1, \dots, N \tag{4}$$

and selects the sample with the maximum  $\underline{d}_n$  to label.

Comparing against the three criteria for ALR, GS only considers the diversity, but not the informativeness and the representativeness.

#### E. Enhanced Batch-Mode ALR

We have already seen that each of the above-mentioned three ALR approaches only considers one of the three essential criteria for ALR, so there is room for improvement. In addition, all of them assume that we already have  $M_0$  initially labeled samples for training. Usually, these  $M_0$  samples are randomly selected, because the regression models cannot be constructed at the very beginning when no or very few labeled samples are available (and hence, QBC and EMCM cannot be applied). However, there can still be better initialization approaches to select more representative and diverse seedling samples, without using any label information. One such approach, EBMALR [40], was proposed recently to consider simultaneously informativeness, representativeness, and diversity to enhance QBC and EMCM. Theoretically, batch-mode ALR can also be used for sequential ALR, by setting the batch size to 1. Algorithm 1 shows the EBMALR algorithm when the batch size is 1. It first uses k-means clustering to initialize d samples that are representative and diverse and then uses a BL ALR approach, such as QBC or EMCM, to select subsequent samples sequentially.

Compared with QBC, EMCM, and GS, EBMALR identifies outliers and excludes them from being selected and considers representativeness and diversity (RD) in initializing the first *d* samples. The original EBMALR (when the batch size is larger than 1) considers both diversity and informativeness in each subsequent iteration, but when the batch size becomes 1, EBMALR is no longer able to consider the diversity among the selected samples. As a result, its performance degrades significantly, as will be demonstrated in Section IV-E.

#### F. Design of Experiments

Design of experiments (DOE) has been widely studied in statistics and used in various industries, for *exploring new processes and gaining increased knowledge of existing processes, followed by optimizing these processes for achieving world-class performance [2].* Its primary goal is usually to extract the maximum amount of information from a few observations as possible, which is very similar to ALR. There are two typical categories of DOEs [2] as follows.

1) *Screening designs*, which are smaller experiments to identify the critical few factors from the many potential trivial factors.

Algorithm 1 EBMALR Algorithm, When the Batch Size Is 1 **Input**: *N* unlabeled samples,  $\{\mathbf{x}_n\}_{n=1}^N$ , where  $\mathbf{x}_n \in \mathbb{R}^d$ ; M, the maximum number of labeled samples to query;  $\gamma$ , the threshold for outlier identification **Output**: The regression model  $f(\mathbf{x})$ . // Identify the outliers  $S = \{\mathbf{x}_n\}_{n=1}^N;$ has Outliers=True; while has Outliers do Perform k-means clustering on S to obtain d clusters,  $C_i, i = 1, \ldots, d;$ Set  $p_i = |C_i|$ ; hasOutliers=False; for i = 1, ..., k do if  $p_i \leq \max(1, \gamma N)$  then  $S = S \setminus C_i;$ hasOutliers=True; end end end // Initialize d labeled samples for i = 1, ..., d do Select the sample closest to the centroid of  $C_i$  to label; end // End initialization for m = d + 1, ..., M do Perform a BL ALR (e.g., QBC or EMCM) on S to select a sample for labeling; end Construct the regression model  $f(\mathbf{x})$  from the *M* labeled samples.

2) *Optimal designs*, which are larger experiments that investigate interactions of terms and nonlinear responses and are conducted at more than two levels for each factor.

Optimal designs are particularly relevant to ALR. They provide theoretical criteria for choosing a set of points to label, for a specific set of assumptions and objectives. Compared with optimal designs, ALR approaches are generally more heuristic. In this paper, we only consider ALR approaches.

## III. OUR PROPOSED ALR APPROACHES

In this section, we propose first a basic pool-based sequential ALR approach by considering simultaneously RD and then strategies to integrate it with QBC, EMCM, and GS to further improve the performance.

#### A. Basic RD ALR Approach

Assume that initially none of the N samples in the pool is labeled. Our proposed basic RD ALR approach consists of two parts: 1) better initialization of the first a few samples by considering RD, and 2) generating a new sample in each subsequent iteration by also considering RD.

Since the input space has *d* dimensions, it is desirable to have at least *d* initially labeled samples to construct a reasonable linear regression model.<sup>1</sup> To find the optimal locations of these *d* samples, we perform *k*-means (k = d) clustering on the *N* unlabeled samples and then select from each cluster the sample closest to the cluster centroid for labeling. This initialization ensures representativeness, because each sample is a good representation of the cluster it belongs to. It also ensures diversity, because these *d* clusters cover the full input space of **x**.

The idea of using clustering for sample selection in ALR was motivated by similar ideas in ALC. For example, Nguyen and Smeulders [28] used k-medoids clustering to select representative and diverse samples. Kang et al. [21] used k-means clustering to partition the unlabeled samples into different clusters and then selected from each cluster the sample closest to its centroid as the most representative one. Hu et al. [20] used deterministic clustering methods (farthest-first traversal, agglomerative hierarchical clustering, and affinity propagation clustering) to initialize the samples, to avoid variations introduced by nondeterministic clustering approaches such as k-medoids and k-means. Krempl et al. [23] proposed a clustering-based optimized probabilistic AL approach for online streaming ALC. However, to our knowledge, there have not yet existed any pool-based sequential ALR approaches that use clustering to initialize the samples and also perform subsequent selections.

After the first d samples are initialized by considering RD, next, we start the iterative ALR process, where a new sample is selected for labeling in each iteration. Consider the first iteration, where we already have d labeled samples and need to determine which sample from the N-d unlabeled ones should be further selected for labeling. In the basic RD algorithm, we first perform k-means clustering on all N samples, where k = d + 1. Since there are only d labeled samples, at least one cluster does not contain any labeled sample. In practice, some clusters may contain multiple labeled samples, so usually, there are more than one cluster that do not contain any labeled sample. We then identify the largest cluster that does not contain any labeled sample as the current most representative cluster and select the sample closest to its centroid for labeling. Note that this selection strategy also ensures diversity, because the identified cluster locates differently from all other clusters that already contain labeled samples. We then repeat this process to generate more labeled samples, until the maximum number of labeled samples is reached.

The pseudocode of the basic RD ALR approach is given in Algorithm 2, where *Option* 1 is used. Similar to GS, the basic RD approach also uses passive sampling, which does not require updating the regression model and evaluating the unlabeled sample in each iteration. Therefore, it is independent of the regression model.

<sup>&</sup>lt;sup>1</sup>It is also possible to initialize fewer than d samples to construct a linear regression model, by using regularized regression such as ridge regression and least absolute shrinkage and selection operator. However, here, we assume d is small and initialize d samples directly for simplicity.

| Algorithm 2 Proposed RD ALR Algorithm, and Its   |
|--|
| Variants   |
| <b>Input</b> : N unlabeled samples, $\{\mathbf{x}_n\}_{n=1}^N$ , where $\mathbf{x}_n \in \mathbb{R}^d$ ; |
| M, the maximum number of labeled samples to  |
| query  |
| <b>Output</b> : The regression model $f(\mathbf{x})$ .   |
| // Initialize $d$ labeled samples  |
| Perform k-means clustering on $\{\mathbf{x}_n\}_{n=1}^N$ , where $k = d$ ;                               |
| Select from each cluster the sample closest to its   |
| centroid, and query for its label;   |
| // End initialization  |
| for $m = d + 1,, M$ do   |
| Perform k-means clustering on $\{\mathbf{x}_n\}_{n=1}^N$ , where $k = m$ ;                               |
| Identify the largest cluster that does not already   |
| contain a labeled sample;  |
| Option 1: Select the sample closest to the cluster   |
| centroid for labeling;   |
| Option 2: Use QBC (Section II-B) to select a sample  |
| from the cluster for labeling;   |
| Option 3: Use EMCM (Section II-C) to select a  |
| sample from the cluster for labeling;  |
| Option 4: Use GS (Section II-D) to select a sample   |
| from the cluster for labeling;   |
| end  |
| Construct the regression model $f(\mathbf{x})$ from the M labeled  |
| samples.   |

# B. Integrate RD With QBC, EMCM, and GS

Interestingly, the basic RD ALR approach can be easily integrated with an existing pool-based sequential ALR approach for better performance. The pseudocode is also shown in Algorithm 2, where *Option* 2 or 3 or 4 is used. The initialization is the same as the basic RD ALR approach. In each iteration, it also selects a sample from the largest cluster that does not already contain a labeled sample for labeling. However, instead of selecting the one closest to its centroid, as in the basic RD ALR approach, now it uses QBC or EMCM or GS to select the most informative or most diverse sample to the label. We expect that when QBC or EMCM is used, the integrated RD ALR approach can achieve better performance than the basic RD ALR approach, because now informativeness, representativeness, and diversity are considered simultaneously.

## C. Differences From EBMALR

Our proposed ALR approaches have some similarity with EBMALR [40], e.g., clustering is used to ensure RD. However, there are several significant differences as follows.

 This paper considers pool-based *sequential* ALR, whereas [40] considered pool-based *batch-model* ALR. Theoretically, sequential ALR can be viewed as a special case of batch-model ALR, when the batch size equals 1. However, as pointed out in Section II-E, when the batch size becomes 1, EBMALR is no longer able to consider the diversity among the selected samples. As a result, its performance becomes significantly worse than the proposed approaches in this paper, as will be shown in Section IV-E.

- This paper explicitly defines informativeness, representativeness, and diversity as three criteria that should be considered in ALR, whereas EBMALR did not (although it implicitly used these concepts).
- 3) EBMALR considered also how to exclude outliers from being selected, but it required a user-defined parameter. Through extensive experiments, we found that this part is not critical in most applications, so this paper does not include it. As a result, our new algorithms do not require any user-defined hyperparameters, which are easier to use.
- 4) In each subsequent iteration, EBMALR (when the batch size is larger than 1) considered first the informativeness and then the diversity, but this paper considers first the diversity and then the informativeness or representativeness. Experiments showed that the latter results in better performances.
- 5) This paper introduces a GS approach (Section II-D) for ALR and also proposes a new RD approach (Section III-A), both of which were not included in [40].
- 6) This paper compares the performances of nine ALR approaches on 11 data sets from various domains, whereas [40] only compared five ALR approaches in a brain–computer interface application.

#### IV. EXPERIMENTS AND RESULTS

Extensive experiments are performed in this section to demonstrate the performance of the basic and integrated RD ALR approaches.

#### A. Data Sets

We used 10 data sets from the University of California, Irvine (UCI) Machine Learning Repository<sup>2</sup> and the Carnegie Mellon University StatLib Data Sets Archive<sup>3</sup> that have been used in previous ALR experiments [7], [8], [41]. We also used an IADS-2 data set on affective computing from the University of Florida Media Core.<sup>4</sup> It consists of 167 acoustic emotional stimuli for experimental investigations of emotion and attention. 76 acoustic features were extracted [17], and principal component analysis (PCA) was used to reduce them to 10 features. The goal was to estimate the continuous arousal value from these 10 features. The summary of these data sets is given in Table I. They cover a large variety of application domains.

Some data sets contain both numerical and categorical features. For example, the autoMPG data set contains seven raw features, among which six are numerical and one is categorical (Origin: U.S., Japan, and Germany). We used one-hot encoding to convert the categorical values into numerical values, e.g., Origin-U.S. was encoded as [1,0,0], Origin-Japan [0, 1, 0], and Origin-Germany [0, 0, 1]. In this way,

<sup>&</sup>lt;sup>2</sup>http://archive.ics.uci.edu/ml/index.php

<sup>&</sup>lt;sup>3</sup>http://lib.stat.cmu.edu/datasets/

<sup>&</sup>lt;sup>4</sup>http://csea.phhp.ufl.edu/media.html#midmedia

TABLE I SUMMARY OF THE 11 REGRESSION DATA SETS

|                           |         | No. of  | No. of   | No. of    | No. of      | No. of   |
|---------------------------|---------|---------|----------|-----------|-------------|----------|
| Dataset                   | Source  | samples | raw      | numerical | categorical | total    |
|                           |         |         | features | features  | features    | features |
| Concrete-CS1              | UCI     | 103     | 7        | 7         | 0           | 7        |
| IADS-Arousal <sup>2</sup> | UFL     | 167     | 10       | 10        | 0           | 10       |
| Yacht <sup>3</sup>        | UCI     | 308     | 6        | 6         | 0           | 6        |
| autoMPG4                  | UCI     | 392     | 7        | 6         | 1           | 9        |
| NO2 <sup>5</sup>          | StatLib | 500     | 7        | 7         | 0           | 7        |
| Housing <sup>6</sup>      | UCI     | 506     | 13       | 13        | 0           | 13       |
| CPS <sup>7</sup>          | StatLib | 534     | 11       | 8         | 3           | 19       |
| Concrete <sup>8</sup>     | UCI     | 1030    | 8        | 8         | 0           | 8        |
| Airfoil <sup>9</sup>      | UCI     | 1503    | 5        | 5         | 0           | 5        |
| Wine-red <sup>10</sup>    | UCI     | 1599    | 11       | 11        | 0           | 11       |
| Wine-white <sup>10</sup>  | UCI     | 4898    | 11       | 11        | 0           | 11       |

<sup>1</sup> https://archive.ics.uci.edu/ml/datasets/Concrete+Slump+Test

<sup>2</sup> http://csea.phhp.ufl.edu/media.html#midmedia

<sup>3</sup> https://archive.ics.uci.edu/ml/datasets/Yacht+Hydrodynamics

<sup>4</sup> https://archive.ics.uci.edu/ml/datasets/auto+mpg

5 http://lib.stat.cmu.edu/datasets/

<sup>6</sup> https://archive.ics.uci.edu/ml/machine-learning-databases/housing/

<sup>7</sup> http://lib.stat.cmu.edu/datasets/CPS\_85\_Wages

<sup>8</sup> https://archive.ics.uci.edu/ml/datasets/Concrete+Compressive+Strength

<sup>9</sup> https://archive.ics.uci.edu/ml/datasets/Airfoil+Self-Noise

<sup>10</sup> https://archive.ics.uci.edu/ml/datasets/Wine+Quality

the converted feature space has 6 + 3 = 9 dimensions. Categorical features in other data sets were converted similarly before regression. We then normalized each dimension of the feature space to mean zero and standard deviation one.

#### B. Algorithms

We compared the performances of nine different sample selection strategies as follows.

- 1) BL, which randomly selects all M samples.
- 2) RD, which is our basic RD ALR algorithm introduced in Section III-A.
- 3) QBC, which has been introduced in Section II-B. The first *d* labeled samples are randomly initialized.
- 4) RD-QBC, which is RD integrated with QBC, introduced in Section III-B.
- 5) EMCM, which has been introduced in Section II-C. The first d labeled samples are randomly initialized.
- 6) Enhanced Expected Model Change
- Maximization (EEMCM), which is the EBMALR approach introduced in Algorithm 1, when EMCM is used as the base ALR approach.
- 7) RD-EMCM, which is RD integrated with EMCM, introduced in Section III-B.
- 8) GS, which has been introduced in Section II-D. The first *d* labeled samples are randomly initialized.
- 9) RD-GS, which is RD integrated with GS, introduced in Section III-B.

All nine approaches built a ridge regression model from the labeled samples, with ridge parameter  $\sigma = 0.01$ . We used ridge regression instead of ordinary linear regression because the number of labeled samples is very small, so ridge regression, with regularization on the coefficients, generally results in better performance than the ordinary linear regression.

#### C. Evaluation Process

There could be two model evaluation strategies: 1) *inductive learning*, in which we learn a model from labeled samples,

and try to evaluate it on samples we have not seen or known about, and 2) *transductive learning*, in which we try to evaluate the model on a known (test) set of unlabeled examples. Specific to pool-based ALR, the former means labeling a small amount of samples from a fixed pool, building a regression model, and then predicting the outputs of the remaining unlabeled samples in the *same* pool, whereas the latter means labeling a small amount of samples from a fixed pool, building a regression model, and then predicting the outputs of unlabeled samples from *another* pool. This paper mainly focuses on transductive learning but will also briefly report results on inductive learning in Section IV-I (more results can be found in the Supplementary Materials). Generally, they are very similar.

The detailed evaluation process was similar to that used in our previous research on pool-based batch-mode ALR [40]. For each data set, let **P** be the pool of all samples. We first randomly selected 80% of the total samples as our training pool<sup>5</sup> (denoted as **P**<sub>80</sub>), initialized the first *d* labeled samples (since ridge regression also includes a bias term, for each data set, *d* was the number of total features in Table I plus one) either randomly or by EEMCM/RD, identified one sample to label in each subsequent iteration by different algorithms and built a ridge regression model. The maximum number of samples to be labeled, *M*, was 10% of the size of **P**<sub>80</sub>. For data sets too small or too large, we constrained  $M \in [20, 60]$ .

In the inductive learning setting, the model performance was evaluated on the 20% remaining samples that are in **P** but not in **P**<sub>80</sub>, whereas in the transductive learning setting, the model performance was evaluated on the samples in **P**<sub>80</sub>. We ran the above-mentioned evaluation process 100 times for each data set and each algorithm, to obtain statistically meaningful results.

#### D. Performance Measures

After each iteration of each algorithm, we computed the root mean squared error (RMSE) and correlation coefficient (CC) as the performance measures.

In transductive learning, because different algorithms selected different samples to label, the remaining unlabeled samples in the pool were different for each algorithm, so we cannot compare their performances based on the remaining unlabeled samples. Since the goal is to build a regression model to label all samples in the pool-based ALR as accurately as possible, we computed the RMSE and CC using all samples in the pool, where the labels for the *m* selected samples were their true labels, and the labels for the remaining N - m unlabeled samples were the predictions from the ridge regression model.

Let  $y_n$  be the true label for  $\mathbf{x}_n$ , and  $\hat{y}_n$  be the prediction from the ridge regression model. Without loss of generality,

<sup>5</sup>For a fixed pool, EEMCM, RD, RD–QBC, RD–EMCM, and RD–GS give a deterministic selection sequence, because there is no randomness involved (assume *k*-means clustering always converges to its global optimum). Therefore, we need to vary the pool in order to study the statistical properties of them. We did not use the traditional bootstrap approach, i.e., sampling with replacement to obtain the same number of samples as the original pool, because bootstrap introduces duplicate samples in the new pool, whereas in practice usually, a pool does not contain duplicates.

assume that the first m samples are selected by an algorithm, and hence, their true labels are known. Then

RMSE = 
$$\left[\frac{1}{N}\sum_{n=1}^{N} (y_n - y'_n)^2\right]^{1/2}$$
 (5)

$$CC = \frac{\sum_{n=1}^{N} (y_n - \bar{y}) (y'_n - \bar{y}')}{\sqrt{\sum_{n=1}^{N} (y_n - \bar{y})^2} \sqrt{\sum_{n=1}^{N} (y'_n - \bar{y}')^2}}$$
(6)

where

$$y'_{n} = \begin{cases} y_{n}, & n = 1, \dots, m \\ \hat{y}_{n}, & n = m + 1, \dots, N \end{cases}$$
(7)

$$\bar{y} = \frac{1}{N} \sum_{n=1}^{N} y_n; \quad \bar{y}' = \frac{1}{N} \sum_{n=1}^{N} y'_n.$$
 (8)

Note that we should consider the RMSE as the primary performance measure, because it is directly optimized in the objective function of ridge regression (CC is not). Generally, as the RMSE decreases, the CC should increase, but there is no guarantee. In other words, we expect that an ALR approach performing well on the RMSE should also perform well on the CC, but this is not always true. Therefore, the CC can only be viewed as a secondary performance measure.

In inductive learning, the RMSE and CC can be computed directly on the 20% samples that are in **P** but not in  $P_{80}$ .

#### E. Experimental Results

The RMSEs and CCs for the nine algorithms on the 11 data sets in transductive learning, averaged over 100 runs, are shown in Fig. 2. We observe the following.

- Generally, as *m* increased, all nine algorithms achieved better performance (smaller RMSE and larger CC), which is intuitive, because more labeled training samples generally result in a more reliable ridge regression model.
- RD, QBC, EMCM, EEMCM, and GS achieved better performances than BL on almost all data sets, suggesting that all these ALR approaches were effective.
- 3) Generally, RD-QBC achieved better performance than both RD and QBC, RD-EMCM achieved better performance than both RD and EMCM, and RD-GS achieved better performance than both RD and GS. These results suggest that our proposed RD ALR approach is complementary to QBC, EMCM, and GS, and hence, integrating them can outperform each individual ALR approach.

To see the forest for the trees, we also define an aggregated performance measure called the area under the curve (AUC) for the average RMSE and the average CC on each of the 11 data sets in Fig. 2. The AUCs for the RMSEs are shown in Fig. 3(a), where for each data set, we used the AUC of BL to normalize the AUCs of the other eight algorithms, so the AUC of BL was always 1. For the RMSE, a smaller AUC indicates a better performance. Similarly, we also show the AUCs of the CCs in Fig. 3(b), where a larger AUC indicates a better performance. We observe that the following.

1) RD achieved smaller AUCs for the RMSE than BL on 10 of the 11 data sets, and larger AUCs for the CC than

TABLE II Ranks of the Nine Approaches on the 11 Data Sets in Transductive Learning

|      |              |             |     |      |       |    |    | RD- | RD-  | RD- |
|------|--------------|-------------|-----|------|-------|----|----|-----|------|-----|
|      | Dataset      | $_{\rm BL}$ | QBC | EMCM | EEMCM | GS | RD | QBC | EMCM | GS  |
|      | Concrete-CS  | 9           | 7   | 8    | 6     | 5  | 4  | 1   | 2    | 3   |
|      | IADS-Arousal | 9           | 8   | 7    | 6     | 1  | 5  | 4   | 2    | 3   |
|      | Yacht        | 6           | 8   | 7    | 2     | 9  | 1  | 3   | 5    | 4   |
|      | autoMPG      | 9           | 7   | 8    | 5     | 1  | 6  | 4   | 3    | 2   |
| RMSE | NO2          | 8           | 5   | 6    | 7     | 2  | 9  | 4   | 1    | 3   |
|      | Housing      | 9           | 7   | 6    | 5     | 8  | 4  | 2   | 1    | 3   |
|      | CPS          | 8           | 7   | 6    | 9     | 1  | 5  | 4   | 2    | 3   |
|      | Concrete     | 9           | 8   | 7    | 6     | 4  | 5  | 2   | 1    | 3   |
|      | Airfoil      | 9           | 7   | 8    | 6     | 5  | 2  | 4   | 3    | 1   |
|      | Wine-red     | 9           | 5   | 6    | 7     | 3  | 8  | 2   | 1    | 4   |
|      | Wine-white   | 8           | 5   | 6    | 7     | 2  | 9  | 4   | 1    | 3   |
|      | Average      | 9           | 7   | 8    | 6     | 4  | 5  | 3   | 1    | 2   |
|      | Concrete-CS  | 9           | 7   | 8    | 6     | 5  | 1  | 2   | 3    | 4   |
|      | IAPS-Arousal | 9           | 6   | 5    | 8     | 7  | 3  | 2   | 4    | 1   |
|      | Yacht        | 9           | 8   | 7    | 2     | 6  | 1  | 5   | 4    | 3   |
|      | autoMPG      | 9           | 6   | 7    | 5     | 1  | 8  | 4   | 3    | 2   |
| CC   | NO2          | 8           | 4   | 7    | 3     | 5  | 9  | 6   | 2    | 1   |
|      | Housing      | 9           | 6   | 8    | 7     | 5  | 4  | 3   | 1    | 2   |
|      | CPS          | 8           | 7   | 5    | 9     | 1  | 6  | 3   | 2    | 4   |
|      | Concrete     | 8           | 7   | 6    | 9     | 5  | 4  | 2   | 1    | 3   |
|      | Airfoil      | 9           | 8   | 7    | 6     | 5  | 1  | 4   | 3    | 2   |
|      | Wine-red     | 4           | 7   | 9    | 6     | 8  | 5  | 2   | 1    | 3   |
|      | Wine-white   | 3           | 1   | 2    | 9     | 7  | 8  | 5   | 4    | 6   |
|      | Average      | 9           | 6   | 8    | 7     | 5  | 4  | 3   | 1    | 2   |

BL on 9 of the 11 data sets, suggesting that RD is indeed effective.

- 2) Among the four existing ALR approaches, GS achieved the best average performance on both RMSE and CC. The reason can be explained as follows. In pool-based ALR, we compute the RMSE and CC on all remaining unlabeled samples, and a large error on a single sample may significantly deteriorate the overall performance, i.e., the samples make unequal contributions to the RMSE and CC. A diverse sample, which is far away from currently selected samples, is more likely to give such a large error (its neighborhood has not been sufficiently modeled). GS considers only the diversity and makes sure the selected samples are somewhat uniformly distributed in the entire input space, i.e., all neighborhoods in the input space are considered, and hence, large errors are less likely to occur. This is different from ALC, in which all misclassified samples make equal contributions to the classification error, no matter how far away they are from the decision boundary.
- 3) Generally, RD-QBC, RD-EMCM, and RD-GS achieved the best performances among the nine algorithms.

The ranks of the nine approaches on the 11 data sets, according to the AUCs, are shown in Table II. Observe that on average, RD-EMCM, RD-GS, and RD-QBC ranked among the top three on both RMSE and CC, RD and GS ranked the next, EEMCM slightly outperformed EMCM, and BL was the last. This again confirms the superiority of our proposed approaches.

## F. Statistical Analysis

To determine if the differences between different pairs of algorithms were statistically significant, we also performed nonparametric multiple comparison tests on the AUCs using



Fig. 2. Performances of the nine algorithms on the 11 data sets in transductive learning, averaged over 100 runs. (a) Concrete-CS. (b) IADS-Arousal. (c) Yacht. (d) autoMPG. (e) NO2. (f) Housing. (g) CPS. (h) Concrete. (i) Airfoil. (j) Wine-red. (k) Wine-white.

Dunn's procedure [13], [14], with a *p*-value correction using the false discovery rate method [4]. The *p*-values for the AUCs of RMSEs and CCs are shown in Table III, where the statistically significant ones are marked in bold. We observe the following.

- 1) All ALR approaches had statistically significantly better RMSEs and CCs than BL.
- 2) Among the four existing ALR approaches, GS had statistically significantly better RMSEs and CCs than EMCM and EEMCM and also statistically significantly better RMSE than QBC.
- RD had statistically significantly better RMSE and CC than QBC, EMCM, and EEMCM and also statistically significantly better CC than GS.
- 4) RD-QBC, RD-EMCM, and RD-GS all had statistically significantly better RMSE and CC than the other six approaches, suggesting again that RD is complementary to QBC, EMCM, and GS, and hence, integrating RD with any of the latter three can further improve the performance.
- 5) There were no statistically significant differences among RD-QBC, RD-EMCM, and RD-GS.

# G. Visualization

It is also interesting to visualize the sample selection results of different algorithms to confirm the superiority of the RD-based ALR approaches. However, because the feature spaces had at least seven dimensions, it is difficult to visualize



Fig. 3. AUCs of the nine algorithms on the 11 data sets in transductive learning. (a) RMSE. (b) CC.

TABLE III *p*-Values of Nonparametric Multiple Comparisons on the AUCS of RMSEs and CCS in Transductive Learning

|      |         |       |       |       |       |       |       | RD-   | RD-   |
|------|---------|-------|-------|-------|-------|-------|-------|-------|-------|
|      |         | BT.   | OBC   | FMCM  | FFMCM | GS    | ВD    | OBC   | FMCM  |
|      | ODC     | 0000  | QDC   | шиси  | ппсы  | 05    | цр    | брс   | шиси  |
|      | QBC<br> | .0000 | 1000  |       |       |       |       |       |       |
|      | EMCM    | .0000 | .4220 |       |       |       |       |       |       |
|      | EEMCM   | .0000 | .0738 | .1036 |       |       |       |       |       |
|      | GS      | .0000 | .0000 | .0000 | .0000 |       |       |       |       |
| RMSE | RD      | .0000 | .0000 | .0000 | .0000 | .2780 |       |       |       |
|      | RD-QBC  | .0000 | .0000 | .0000 | .0000 | .0000 | .0000 |       |       |
|      | RD-EMCM | .0000 | .0000 | .0000 | .0000 | .0000 | .0000 | .1449 |       |
|      | RD-GS   | .0000 | .0000 | .0000 | .0000 | .0000 | .0000 | .3607 | .2478 |
|      | QBC     | .0000 |       |       |       |       |       |       |       |
|      | EMCM    | .0000 | .0924 |       |       |       |       |       |       |
|      | EEMCM   | .0000 | .0000 | .0017 |       |       |       |       |       |
|      | GS      | .0000 | .0353 | .0008 | .0000 |       |       |       |       |
| CC   | RD      | .0000 | .0000 | .0000 | .0000 | .0023 |       |       |       |
|      | RD-QBC  | .0000 | .0000 | .0000 | .0000 | .0000 | .0057 |       |       |
|      | RD-EMCM | .0000 | .0000 | .0000 | .0000 | .0000 | .0001 | .1167 |       |
|      | RD-GS   | .0000 | .0000 | .0000 | .0000 | .0000 | .0000 | .0508 | .3187 |

them directly. Therefore, we performed PCA on each data set and represented all samples as their projections on the first two principal components. Due to the page limit, we only show the results for the Concrete-CS data set (more results can be found in the Supplementary Materials). The red asterisks in Fig. 4(a) indicate the initial seven selected samples. Observe that random initialization, which was used in BL, QBC, EMCM, and GS, may leave a large portion of the feature space unsampled. However, RD, RD–QBC, RD–EMCM, and RD–GS, which used our proposed initialization approach, initialized the samples more uniformly in the entire feature space. The red asterisks in Fig. 4(b) indicate the final 20 samples selected by different algorithms. We observe the following.

1) BL still left a large region of the feature space unsampled, even after 20 samples.



Fig. 4. PCA visualization of the selected samples (red asterisks) by different algorithms on the Concrete-CS data set. (a) Initial seven selected samples. (b) Final 20 selected samples.

- 2) QBC, EMCM, and GS selected one or more samples near the boundary of the feature space, which may be outliers. On the contrary, EEMCM and RD selected samples uniformly distributed in the whole feature space, and no selected samples were obvious outliers.
- 3) The samples selected by RD-QBC distributed in the feature space more uniformly than those selected by QBC. Similar patterns can also be observed between RD-EMCM and EMCM, and between RD-GS and GS.

In summary, the PCA visualization results confirm that the RD-based ALR approaches selected more reasonable samples, which resulted in better regression performances.

#### H. Individual Improvements

Table II shows that RD-EMCM achieved the best average RMSE among the nine algorithms. Recall that RD-EMCM has three enhancements over the random sampling approach (BL) as follows.

Enhancement 1: RD-EMCM considers both the representativeness and the diversity in initializing the first d samples, but BL does not consider either of them.

*Enhancement 2:* RD-EMCM considers both the representativeness and the diversity in selecting the new sample in each iteration, but BL does not consider either of them.

*Enhancement 3:* RD-EMCM considers also the informativeness in selecting the new sample in each iteration, but BL does not consider it.



Fig. 5. Effect of the individual enhancements on the concrete data set.



Fig. 6. AUCs of the five algorithms over the 11 data sets, averaged over 100 runs. (a) RMSE. (b) CC.

It is interesting to study if each of the three enhancements is necessary, and if so, what their individual effect is.

For this purpose, we constructed three modified versions of RD-EMCM, by considering each enhancement individually: E1, which employs only the first enhancement on more representative and diverse initialization; E2, which employs only the second enhancement on more representative and diverse sampling in each iteration; and E3, which employs only the third enhancement on more informative sampling in each iteration. We then compared their performances with BL and RD-EMCM. Due to the page limit, we only show the results on the concrete data set (averaged over 100 runs) in Fig. 5 (more results can be found in the Supplementary Materials). The AUCs for RMSEs and CCs for all 11 data sets are shown in Fig. 6.

Observe that all three enhancements outperformed BL on most data sets, especially for the RMSE, which was directly optimized in the objective function of ridge regression. More specifically, Fig. 5 shows that the first enhancement on more



Fig. 7. AUCs of the nine algorithms on the 11 data sets in inductive learning, averaged over 100 runs. (a) RMSE. (b) CC.

representative and diverse initialization helped when m was very small; the second and third enhancements helped when m became larger. By combining the three enhancements, RD-EMCM achieved the best performance at both small and large m. This suggests that the three enhancements are complementary, and they are all essential to the improved performance of RD-EMCM.

# I. Inductive Learning Results

The results presented in this section so far focused only on transductive learning. This section presents the inductive learning results, i.e., testing results on the 20% samples in **P** but not in **P**<sub>80</sub>. The AUCs of the nine algorithms on the 11 data sets are shown in Fig. 7 (more results can be found in the Supplementary Materials). Observe that Fig. 7 is very similar to Fig. 3. Our conclusions are drawn in transductive learning still hold in inductive learning.

# J. Computational Cost

Finally, we study the computational cost of the proposed algorithms.

First, consider RD. Observe from Algorithm 2 that the most time-consuming operation in RD is *k*-means clustering, whose computational complexity is O(kndi) [18], where *k* is the number of clusters, *n* is the number of data samples, *d* is the number of features, and *i* is the number of iterations before convergence. *i* is mainly determined by *n*, and in the worst case,  $i = 2^{\Omega(\sqrt{n})}$  [3]. However, for a data set that does have a clustering structure, *i* is usually small, and in practice, we can set an upper bound for it.

TABLE IV COMPUTATION TIME (SECONDS) OF DIFFERENT ALR APPROACHES

| No. of  | No. of   |      |      |       |      |      | RD-  | RD-  | RD-  |
|---------|----------|------|------|-------|------|------|------|------|------|
| Samples | Features | QBC  | EMCM | EEMCM | GS   | RD   | QBC  | EMCM | GS   |
|         | 10       | 0.01 | 0.01 | 0.02  | 0.00 | 0.12 | 0.12 | 0.13 | 0.11 |
| 100     | 30       | 0.01 | 0.01 | 0.03  | 0.00 | 0.12 | 0.13 | 0.13 | 0.10 |
|         | 50       | 0.02 | 0.02 | 0.03  | 0.00 | 0.12 | 0.14 | 0.14 | 0.11 |
|         | 10       | 0.06 | 0.06 | 0.08  | 0.03 | 0.21 | 0.23 | 0.23 | 0.19 |
| 1,000   | 30       | 0.07 | 0.07 | 0.09  | 0.03 | 0.25 | 0.26 | 0.26 | 0.23 |
|         | 50       | 0.07 | 0.08 | 0.11  | 0.04 | 0.29 | 0.30 | 0.30 | 0.26 |
|         | 10       | 0.54 | 0.59 | 0.76  | 2.03 | 2.43 | 2.43 | 2.43 | 2.41 |
| 10,000  | 30       | 0.64 | 0.70 | 1.01  | 2.21 | 3.98 | 4.00 | 3.97 | 4.01 |
|         | 50       | 0.74 | 0.82 | 1.29  | 2.52 | 5.57 | 5.50 | 5.70 | 5.44 |

In summary, the computational complexity of k-means clustering, and hence RD, is more significantly influenced by n, and then d and k as follows.

- 1) The number of data samples, *n*, is fixed, given a specific data set, and we can do nothing about it.
- 2) The number of features, *d*, could be large in practice. In such cases, we could use preprocessing, such as PCA, to reduce the dimensionality, as we have done for data set IADS-Arousal in our experiments. We expect that usually PCA needs at most tens of scores to capture enough variance (e.g., 95%), and hence, *d* could be kept small.
- The number of clusters, k, equals the number of currently available labeled samples plus one, so it increases as the query goes on, i.e., RD will become slower as it iterates.

Next, consider the variants of RD. RD-QBC, RD-EMCM, and RD-GS should have similar computational complexities as RD, as the computational complexities of QBC, EMCM, and GS are low compared with RD.

To compare the actual computation time, we performed an experiment using the Online News Popularity Data Set<sup>6</sup> in the UCI Machine Learning Repository, which includes 39797 samples (n = 39, 797) and 58 numerical features (d = 58). We picked  $n = \{100, 1000, 10000\}$  (using the first *n* samples in the data set) and  $d = \{10, 30, 50\}$  (randomly selected from the 58 features) to study the effects of n and don the computation time. The platform was a Dell XPS15 laptop with Intel Core i7-6700HQ CPU @ 2.60 Hz, 16 GB memory, and 512 GB SSD, running Windows 10 x64 and MATLAB 2017a. For each combination of n and d, we increased k from 60 to 70, recorded the total running time of different algorithms and repeated this process 10 times to compute the average. The results are shown in Table IV. They are consistent with our qualitative analyses mentioned above: the computation time of RD increases rapidly with n but is about linear to d when  $d \in [10, 50]$ ; and the computation time of RD-QBC, RD-EMCM, and RD-GS is close to RD.

In summary, although the computational costs of our proposed algorithms are much higher than QBC, EMCM, and GS, they are intended for offline applications, in which computation time may not be an obstacle.

#### V. CONCLUSION

AL has been frequently used to reduce the data labeling effort in machine learning. However, most existing AL approaches are for classification. This paper studied ALR. We proposed three essential criteria that should be considered in selecting a new sample in pool-based sequential ALR, which are informativeness, representativeness, and diversity. An ALR approach called RD was proposed, which considers RD in both initialization and subsequent iterations. The RD approach can also be integrated with existing pool-based sequential ALR approaches, such as QBC, EMCM, and GS, to further improve the performance. Extensive experiments on 11 public data sets from various domains confirmed the effectiveness of our proposed approaches.

#### REFERENCES

- N. Abe and H. Mamitsuka, "Query learning strategies using boosting and bagging," in *Proc. 15th Int. Conf. Mach. Learn. (ICML)*, Madison, WI, USA, Jul. 1998, pp. 1–9.
- [2] J. Antony, *Design of Experiments for Engineers and Scientists*, 2nd ed. London, U.K.: Elsevier, 2014.
- [3] D. Arthur and S. Vassilvitskii, "How slow is the k-means method?" in *Proc. 22nd Annu. Symp. Comput. Geometry*, Sedona, AZ, USA, Jun. 2006, pp. 144–153.
- [4] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: A practical and powerful approach to multiple testing," J. Roy. Statist. Soc., Ser. B, vol. 57, pp. 289–300, Jan. 1995.
- [5] M. M. Bradley and P. J. Lang, "The international affective digitized sounds (IADS-2): Affective ratings of sounds and instruction manual," 2nd ed., Center Study Emotion Attention, Univ. Florida, Gainesville, FL, USA, Tech. Rep. B-3, 2007.
- [6] R. Burbidge, J. J. Rowland, and R. D. King, "Active learning for regression based on query by committee," in *Intelligent Data Engineering and Automated Learning* (Lecture Notes in Computer Science), vol. 4881. 2007, pp. 209–218.
- [7] W. Cai, M. Zhang, and Y. Zhang, "Batch mode active learning for regression with expected model change," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 28, no. 7, pp. 1668–1681, Jul. 2017.
- [8] W. Cai, Y. Zhang, and J. Zhou, "Maximizing expected model change for active learning in regression," in *Proc. IEEE 13th Int. Conf. Data Mining*, Dallas, TX, USA, Dec. 2013, pp. 51–60.
- [9] W. Cai, Y. Zhang, S. Zhou, W. Wang, C. Ding, and X. Gu, "Active learning for support vector machines with maximum model change," in *Machine Learning and Knowledge Discovery in Databases* (Lecture Notes in Computer Science). Berlin, Germany: Springer-Verlag, 2014, pp. 211–216.
- [10] D. A. Cohn, Z. Ghahramani, and M. I. Jordan, "Active learning with statistical models," J. Artif. Intell. Res., vol. 4, pp. 129–145, Mar. 1996.
- [11] B. Demir and L. Bruzzone, "A multiple criteria active learning method for support vector regression," *Pattern Recognit.*, vol. 47, pp. 2558–2567, Jul. 2014.
- [12] P. Donmez and J. G. Carbonell, "Optimizing estimated loss reduction for active sampling in rank learning," in *Proc. 25th Int. Conf. Mach. Learn. (ICML)*, Helsinki, Finland, Jul. 2008, pp. 248–255.
- [13] O. J. Dunn, "Multiple comparisons among means," J. Amer. Statist. Assoc., vol. 56, no. 293, pp. 62–64, 1961.
- [14] O. J. Dunn, "Multiple comparisons using rank sums," *Technometrics*, vol. 6, no. 3, pp. 214–252, 1964.
- [15] Y. Freund, H. S. Seung, E. Shamir, and N. Tishby, "Selective sampling using the query by committee algorithm," *Mach. Learn.*, vol. 28, no. 2, pp. 133–168, 1997.
- [16] M. Grimm, K. Kroschel, and S. Narayanan, "The Vera am Mittag German audio-visual emotional speech database," in *Proc. Int. Conf. Multimedia Expo (ICME)*, Hannover, Germany, Jun. 2008, pp. 865–868.
- [17] C. Guo and D. Wu, "Feature dimensionality reduction for video affect classification: A comparative study," in *Proc. 1st Asian Conf. Affect. Comput. Intell. Interact.*, Beijing, China, May 2018.
- [18] J. A. Hartigan and M. A. Wong, "Algorithm AS 136: A k-means clustering algorithm," *Appl. Statist.*, vol. 28, no. 1, pp. 100–108, 1979.

<sup>&</sup>lt;sup>6</sup>https://archive.ics.uci.edu/ml/datasets/online+news+popularity

- [19] T. He *et al.*, "An active learning approach with uncertainty, representativeness, and diversity," *Sci. World J.*, vol. 2014, Aug. 2014, Art. no. 827586.
- [20] R. Hu, B. M. Namee, and S. J. Delany, "Off to a good start: Using clustering to select the initial training set in active learning," in *Proc.* 23rd Int. Florida Artif. Intell. Res. Soc. Conf., Daytona Beach, FL, May 2010.
- [21] J. Kang, K. R. Ryu, and H.-C. Kwon, "Using cluster-based sampling to select initial training set for active learning in text classification," in *Proc. Pacific–Asia Conf. Knowl. Discovery Data Mining*, Sydney, NSW, Australia, May 2004, pp. 384–388.
- [22] S. Koelstra *et al.*, "DEAP: A database for emotion analysis; Using physiological signals," *IEEE Trans. Affective Comput.*, vol. 3, no. 1, pp. 18–31, Oct./Mar. 2012.
- [23] G. Krempl, T. C. Ha, and M. Spiliopoulou, "Clustering-based optimised probabilistic active learning (COPAL)," in *Proc. 18th Int. Conf. Discov*ery Sci. Banff, AB, Canada: Springer, Oct. 2015, pp. 101–115.
- [24] A. Krogh and J. Vedelsby, "Neural network ensembles, cross validation and active learning," in *Proc. Neural Inf. Process. Syst. (NIPS)*, Denver, CO, USA, Nov. 1995, pp. 231–238.
- [25] D. J. C. MacKay, "Information-based objective functions for active data selection," *Neural Comput.*, vol. 4, no. 4, pp. 590–604, Jul. 1992.
- [26] A. D. Marathe, V. J. Lawhern, D. Wu, D. Slayback, and B. J. Lance, "Improved neural signal classification in a rapid serial visual presentation task using active learning," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 24, no. 3, pp. 333–343, Mar. 2016.
- [27] A. Mehrabian, Basic Dimensions for a General Psychological Theory: Implications for Personality, Social, Environmental, and Developmental Studies. Cambridge, MA, USA: Oelgeschlager, Gunn & Hain, 1980.
- [28] H. T. Nguyen and A. Smeulders, "Active learning using pre-clustering," Proc. 25th Int. Conf. Mach. Learning, Jul. 2004, p. 79.
- [29] R. Picard, Affective Computing. Cambridge, MA, USA: MIT Press, 1997.
- [30] T. RayChaudhuri and L. G. C. Hamey, "Minimisation of data collection by active learning," in *Proc. IEEE Int. Conf. Neural Netw.*, Perth, WA, Australia, vol. 3, Nov./Dec. 1995, pp. 1338–1341.
- [31] J. A. Russell, "A circumplex model of affect," J. Personality Social Psychol., vol. 39, no. 6, pp. 1161–1178, Dec. 1980.
- [32] B. Settles and M. Craven, "An analysis of active learning strategies for sequence labeling tasks," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, Honolulu, HI, USA, Oct. 2008, pp. 1070–1079.
- [33] B. Settles, M. Craven, and S. Ray, "Multiple-instance active learning," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, Vancouver, BC, Canada, Dec. 2008, pp. 1289–1296.
- [34] B. Settles, "Active learning literature survey," Comput. Sci., Univ. Wisconsin–Madison, Madison, WI, USA, Tech. Rep. 1648, 2009.
- [35] H. S. Seung, M. Opper, and H. Sompolinsky, "Query by committee," in *Proc. 5th Annu. Workshop Comput. Learn. Theory*, Pittsburgh, PA, USA, Jul. 1992, pp. 287–294.

- [36] D. Shen, J. Zhang, J. Su, G. Zhou, and C.-L. Tan, "Multi-criteriabased active learning for named entity recognition," in *Proc. 42nd Annu. Meeting Assoc. Comput. Linguistics*, Barcelona, Spain, Jul. 2004, Art. no. 589.
- [37] M. Sugiyama, "Active learning in approximately linear regression based on conditional expectation of generalization error," *J. Mach. Learn. Res.*, vol. 7, pp. 141–166, Jan. 2006.
- [38] M. Sugiyama and S. Nakajima, "Pool-based active learning in approximate linear regression," *Mach. Learn.*, vol. 75, no. 3, pp. 249–274, 2009.
- [39] R. Willett, R. Nowak, and R. M. Castro, "Faster rates in regression via active learning," in *Proc. Adv. Neural Inf. Process. Syst.*, Vancouver, BC, Canada, Dec. 2006, pp. 179–186.
- [40] D. Wu, V. J. Lawhern, S. Gordon, B. J. Lance, and C.-T. Lin, "Offline EEG-based driver drowsiness estimation using enhanced batch-mode active learning (EBMAL) for regression," in *Proc. IEEE Int. Conf. Syst.*, *Man, Cybern.*, Budapest, Hungary, Oct. 2016, pp. 730–736.
- [41] H. Yu and S. Kim, "Passive sampling for regression," in *Proc. IEEE Int. Conf. Syst., Man, Cybern.*, Sydney, NSW, Australia, Dec. 2010, pp. 1151–1156.



**Dongrui Wu** (S'05–M'09–SM'14) received the Ph.D. degree in electrical engineering from the University of Southern California, Los Angeles, CA, USA, in 2009.

He was a Lead Researcher with the GE Global Research, Niskayuna, NY, USA. He is currently a Chief Scientist with DataNova LLC, Clifton Park, NY, USA. He has authored or co-authored more than 100 publications including a book *Perceptual Computing* (Wiley-IEEE, 2010). His current research interests include affective computing,

brain-computer interface, computational intelligence, and machine learning.

Dr. Wu was a recipient of the IEEE CIS Outstanding Ph.D. Dissertation Award in 2012, the IEEE Transactions on Fuzzy Systems Outstanding Paper Award in 2014, the NAFIPS Early Career Award in 2014, and the IEEE SMC Society Early Career Award in 2017. He is an Associate Editor of the IEEE TRANSACTIONS ON FUZZY SYSTEMS, the IEEE TRANSACTIONS ON HUMAN–MACHINE SYSTEMS, and *IEEE Computational Intelligence Magazine*.