

© 2014 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Title: Cost-Sensitive Active Learning With Lookahead: Optimizing Field Surveys for Remote Sensing Data Classification

This paper appears in: IEEE Transactions on Geoscience and Remote Sensing

Date of Publication: 2014

Author(s): Claudio Persello, Abdeslam Boularias, Michele Dalponte, Terje Gobakken, Erik Næsset, Bernhard Schölkopf.

Volume: -, Issue: -

Page(s): -

DOI: 10.1109/TGRS.2014.2300189

Cost-Sensitive Active Learning with Lookahead: Optimizing Field Surveys for Remote Sensing Data Classification

Claudio Persello, *Member, IEEE*, Abdeslam Boularias, Michele Dalponte, Terje Gobakken, Erik Næsset, Bernhard Schölkopf

Abstract—Active learning (AL) typically aims at minimizing the number of labeled samples to be included in the training set to reach a certain level of classification accuracy. Standard methods do not usually take into account the real annotation procedures and implicitly assume that all samples require the same effort to be labeled. Here, we consider the case where the cost associated with the annotation of a given sample depends on the previously labeled samples. In general, this is the case when annotating a queried sample is an action that changes the state of a dynamic system, and the cost is a function of the state of the system. In order to minimize the total annotation cost, the active sample selection problem is addressed in the framework of a Markov Decision Process (MDP), which allows one to plan the next labeling action on the basis of an expected long-term cumulative reward. This framework allows us to address the problem of optimizing the collection of labeled samples by field surveys for the classification of remote sensing data. The proposed method is applied to the ground sample collection for tree species classification using airborne hyperspectral images. Experiments carried out in the context of a real case study on forest inventory show the effectiveness of the proposed method.

Index Terms—Active Learning, Markov Decision Process, Support Vector Machine, Image Classification, Hyperspectral Data, Field Surveys, Forest Inventories.

I. INTRODUCTION

In supervised classification, the amount and quality of training samples are crucial for obtaining accurate results. Furthermore, considering that sample labeling is usually expensive and time consuming, tools for selecting the most informative samples can significantly reduce the effort associated with the labeling of redundant or unnecessary samples. AL methods provide a way to iteratively select the samples that are expected to lead to the highest gain in predictive performances, once they are labeled and added to the training set. An expert is guided in the collection of an effective training set reducing the annotation cost compared to a passive approach. The aim is typically to minimize the number of samples to be labeled and added to the training set in order to reach a certain level of accuracy. AL has been applied to a variety of real-world problem domains, including text classification, information extraction, video classification and retrieval, speech recognition [1] and recently also to remote sensing (RS) classification problems [2], [3]. However, most of the AL methods have been developed for general purposes and do not take into account the real annotation procedures and costs, which usually depend on the application.

In this study, we consider problems where the annotation cost for a given sample is a function of previously collected samples. This type of scenario occurs for instance when a human expert has to visit different locations to retrieve the labels of the queried samples, and the cost is thus a function of the travelled distance from the previous sample location. In this situation, our goal is to optimize the sample selection not just with respect to the number of labelings, but considering the total cost of the annotation procedure. The considered scenario models the problem of optimizing the collection of labeled samples through field surveys for RS data classification. In this context, several classification problems require the collection of a training set through ground sample collection. This is typically the case of hyperspectral images classification, where sample labels cannot usually be obtained by visual inspection of false color image compositions. In contrast to photo-interpretation, *in situ* surveys are typically very expensive and need to be performed by human experts. In such applications, the advantage of using an effective AL strategy that can guide the human expert in an optimized sequence of site visits is particularly important, given the significant savings in terms of time and money.

In view of this, we propose a novel Cost-Sensitive Active Learning (CSAL) method, i.e., an AL strategy that explicitly takes into account the cost for obtaining sample labels in the selection process. This strategy is used for guiding the user in the annotation procedure for selecting the most informative samples while minimizing the cost. It is worth noting that in standard AL methods the query function usually selects the sample that maximizes an immediate *utility*, i.e., the expected gain in predictive performances at the very next iteration. As opposed to this, the considered cost-sensitive setting requires us to plan the sample labeling procedure by taking into account several steps ahead. For this reason, we modeled the CSAL problem in the framework of a Markov Decision Process (MDP), which allows us to optimize the next labeling action on the basis of an expected long-term cumulative *reward*. The considered problem, that we call here the *traveling annotator problem*, is depicted in Figure 1: the human expert has to decide on the next sample to collect starting from the current position. Collecting the sample that maximizes the immediate expected utility would be a myopic strategy. A better decision can be taken by “looking ahead” and optimizing a long-term cumulative *reward* that considers both the prediction utility and the traveling cost.

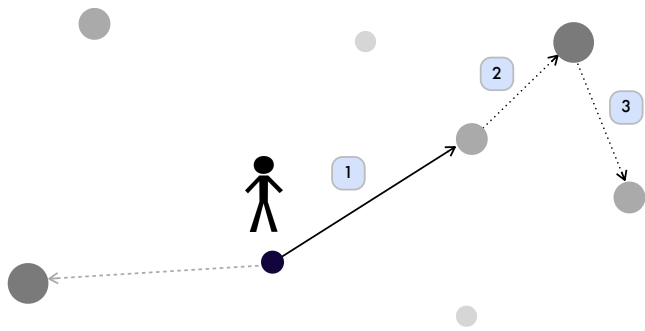


Fig. 1. Graphical representation of the traveling annotator problem. The human expert has to decide on the next sample to collect starting from the current position. Samples are depicted as grey balls, whose expected prediction utility is associated with their size (and grey level). Collecting the sample that maximizes the immediate expected utility (grey dashed arrow) would be a myopic and sub-optimal strategy in view of a long-term reward, which considers both the prediction utility and the traveling cost. A better decision can be taken by looking ahead (for three steps in this example) and performing the action that optimizes the long-term cumulative *reward* (black solid arrow).

As an application of particular interest, we consider here the use of hyperspectral data for forest inventory purposes. The introduction of very high geometrical resolution RS images allows one to improve the spatial resolution of RS-based forest inventories and much research effort is currently devoted to Individual Tree Crown (ITC) level inventories [4], [5]. ITC-inventories provide information about stem volume, height, species, etc. for trees present in the dominant layer of the canopy. Regarding the species classification, airborne hyperspectral sensors, thanks to their very high spectral and spatial resolution, can be powerful instruments for these kinds of inventories. Several studies have showed their effectiveness in distinguishing very similar tree species, in various environments (e.g., [6]–[8]). However, an expensive part of a forest inventory at ITC level is represented by the ground sample collection. Usually, the sample collection is carried out inside field plots (circular areas of a given radius) distributed over the area of interest according to a predefined sampling strategy (e.g., systematic sampling strategy, random sampling strategy, stratified sampling strategy, etc.) [9]. Inside each field plot, trees are measured according to a predefined rule. Usually, trees for which the Diameter at Breast Height (DBH) is higher than a certain threshold are measured. Obviously, this procedure is largely time consuming and it represents the main cost of an inventory. Thus, effective tools that can reduce the cost of the ground samples collection, without losing accuracy in the inventory, are needed.

The main contributions of this study consist in: 1) modeling the problem of optimizing the collection of labeled samples for the classification of RS data as the described *traveling annotator problem*, 2) proposing a novel CSAL method with lookahead to effectively solve the problem, 3) applying the proposed method to a real study on forest inventory using airborne hyperspectral images. In particular, we propose two different *query functions with lookahead* to address the considered problem. The remainder of the paper is organized as follows. The next section reviews basic concepts about AL and about its use for RS image classification. Section

III introduces Markov Decision Process. Section IV presents the proposed CSAL method and two query functions with lookahead. Section V and VI reports the considered data set and the experimental analysis, respectively. Finally, Section VII draws the conclusion of the paper.

II. ACTIVE LEARNING

Active learning is an iterative procedure where the user is repeatedly asked to annotate new samples that are selected by a *query function*. In the *pool-based* setting, the query function is used to select the samples from a pool $\mathcal{U} = (x^i)_{i=1}^n$ of n candidate unlabeled samples that are expected to lead to the highest gain in predictive performances once they are labeled and added to the training set. The classification algorithm is retrained using the training set that contains the new samples and this procedure is repeated until a stopping criterion is met. Most of existing works have focused on the selection of one sample to be labeled in each iteration. To this end, different criteria have been adopted for selecting the (expected) most informative sample. One of the first strategies introduced in the literature is based on uncertainty sampling [10], which aims at selecting the closest sample to the decision boundary. The same principle has also been used in the context of support vector machine (SVM) classification [11]–[13]. Other strategies are query by committee [14] and expected error reduction [15]. A survey of several existing methods is available in [1].

Other studies have focused on the selection of *batches of samples* at each iteration, which allow one to speed up the learning process. In this latter setting, the overlap of information among the selected samples has to be considered in order to evaluate their expected information content. Brinker introduced an SVM-based batch approach, which selects a batch of samples that minimizes the margin while maximizing their *diversity* [16]. The diversity is assessed considering the kernel cosine-angular distance between points. Another approach to consider the diversity in the query function is the use of clustering [17], [18]. In [17], an AL heuristic is presented, which explores the clustering structure of samples and identifies uncertain samples avoiding redundancy. In [19], the authors chose the batch of instances that maximizes the Fisher information of the classification model, which leads to a tradeoff between uncertainty and diversity. In [20], batch active learning is formulated as an optimization problem that maximizes the discriminative classification performance while taking into consideration the unlabeled examples. Azimi *et. al* [21] used Monte-Carlo simulation to estimate the distribution of unlabeled examples selected by a sequential policy and query the samples that best matched such a distribution.

In recent years, active learning has attracted the interest of the remote sensing community, and it has mainly been applied to the classification of multispectral and hyperspectral images [2], [3], [22]. In [2], an AL technique is presented, which selects the unlabeled sample that maximizes the information gain between the posterior probability distribution estimated from the current training set and the training set obtained by including that sample into it. The information gain is

measured by the Kullback-Leibler (KL) divergence. In [23], two AL techniques for multi-class RS classification problems are proposed. The first technique is margin sampling by closest support vector, which selects the most uncertain unlabeled samples that do not share the closest support vector. In the second technique, samples are selected according to the maximum disagreement between a committee of classifiers, which is obtained by bagging: different training sets are drawn with replacement from the original training data and used for training different supervised classifiers. In [24], a supervised Bayesian approach to hyperspectral image segmentation with active learning is presented. The adopted AL method is based on a multinomial logistic regression model, which is used to learn the class posterior probability distributions. In [3], different batch-mode AL techniques for the multi-class classification of RS images with SVM are investigated. The investigated techniques exploit different query functions, which are based on both the uncertainty and diversity criteria. One of the investigated techniques is called Multiclass-Level Uncertainty - Angle Based Diversity (MCLU-ABD), which effectively extend the method present in [16] to deal with multi-class classification problems. Moreover, a query function that is based on a kernel-clustering technique for assessing the diversity of samples and a strategy for selecting the most informative representative sample from each cluster is proposed. Such a technique is called Multiclass-Level Uncertainty with Enhanced Clustering Based Diversity (MCLU-ECBD). Di *et al.* [25] investigate AL methods based on the idea of query by committee, where the committee of classifiers is derived by using multiple views, i.e., different disjoint subsets of features. The paper investigates different approaches for view generation from hyperspectral images, including clustering, random selection and uniform slicing methods. Recent studies adopt AL to address domain adaptation problems [26]–[28], i.e., for adapting the supervised classifier trained on a given RS image to classify another similar image acquired on a different geographical area. The method presented in [28] iteratively selects the most informative samples of the target image to be included in the training set, while the source-image samples are re-weighted or possibly removed from the training set on the basis of their disagreement with the target-image classification problem. In this way, the consistent information of the source image can be effectively exploited for the classification of the target image and for guiding the selection of new samples to be labeled, whereas the inconsistent information is automatically detected and removed.

Very few studies addressed the active learning problem in a cost-sensitive setting for optimizing the ground sample collection. Moreover, the use of AL in the context of forest inventory applications was not yet investigated. In [29], the CSAL problem was modeled as a traveling salesman problem with profits, where the profit for visiting a given sample is its uncertainty score. Nevertheless, the proposed heuristics are suboptimal and the overlap of information between subsequent samples is not taken into account. In [30], a batch-mode active learning method that considers uncertainty, diversity, and cost in the definition of the query function is proposed. The proposed heuristic is however suboptimal, given that the

selection in the first step is based only on the uncertainty. Moreover, the optimal parameter values of the method are difficult to be set in advance in the real application.

To the best of our knowledge, state-of-the-art methods focus on selecting the sample(s) that maximize(s) an immediate reward, i.e., an estimate of the accuracy gain at the very next iteration. But considering that active learning is actually a sequential decision-making problem, it appears natural to consider the reward of a decision after a finite number of steps in view of a long-term optimization. In the standard setting, where the labeling cost is not taken into account or does not depend on previously labeled samples, the advantage of looking ahead, beyond the next iteration, may not be evident. Nevertheless, its advantage becomes particularly important in the considered traveling annotator problem, where a reward function can be defined on the basis of both the annotation cost and the expected prediction utility of the selected sample.

III. MARKOV DECISION PROCESS

Markov Decision Processes (MDPs) [31] are a powerful tool that provides a natural mathematical formalization of sequential decision-making problems. MDPs are used in a variety of areas, including robotics, automated control, manufacturing, games, and economics. A survey of applications of MDPs can be found in [32]. In this paper, we formalize the problem of sample selection in active learning as a Markov decision process. We start by briefly recalling MDPs in the remainder of this section.

Formally, a bounded-horizon MDP is defined by a tuple $(\mathcal{S}, \mathcal{A}, T, R, H, \gamma)$, where \mathcal{S} is a set of *states* and \mathcal{A} is a set of *actions*. T is a *transition function* that returns a next state s' when action a is executed in state s . In general, T is a stochastic function and $T(s, a, s')$ denotes the probability of going to state s' after applying action a in state s , i.e., $T(s, a, s') = P(s_{t+1} = s' | s_t = s, a_t = a)$. T can also be deterministic, in which case $s' = T(s, a)$ denotes the next state after choosing action a in state s . R is a *reward function* where $R(s, a)$ is the numerical reward given for choosing action a in state s . $H \in \mathbb{N}^+$ is a planning horizon, i.e., the number of future time-steps considered in choosing an action, and $\gamma \in [0, 1]$ is a discount factor used to weigh less rewards received further in the future. A *policy* π is a function that returns an action $a = \pi(s)$ for each state s . The expected *value* of policy π in state s is the expected sum of rewards that will be received when starting from s and following policy π for H time-steps. The value function is denoted by V_H^π and defined as

$$V_H^\pi(s) = \sum_{t=0}^{H-1} \gamma^t \mathbb{E}_{s_t} [R(s_t, a_t) | s_0 = s, a_t = \pi(s_t)]. \quad (1)$$

An optimal policy π^* is one satisfying $\pi^* \in \arg \max_{\pi} V_H^\pi(s), \forall s \in \mathcal{S}$. The expected value of executing action a in state s and then following policy π for the next H steps is called a *Q-value*, and defined as

$$Q_H^\pi(s, a) = R(s, a) + \gamma \sum_{s' \in \mathcal{S}} T(s, a, s') V_{H-1}^\pi(s'). \quad (2)$$

Therefore, $\pi^*(s) = \arg \max_{a \in \mathcal{A}} Q_H^{\pi^*}(s, a)$. An optimal policy π^* can be found by using dynamic programming techniques, such as *policy iteration* or *value iteration* [32]. The computational complexity of these methods is exponential in the number of states [33], which is not suited for applications with large state spaces. However, given a fixed initial state s_0 , an optimal policy for a bounded horizon can be found in polynomial time using a lookahead tree search.

IV. COST-SENSITIVE ACTIVE LEARNING WITH LOOKAHEAD

The goal of the proposed CSAL method is to optimize the sample collection process in order to maximize the long-term classification accuracy while minimizing the total annotation cost. The annotation cost takes into account the expected traveling cost (e.g., time) to reach the location of the next sample and a cost for obtaining its label. This defines a multi-objective optimization problem. In order to address this problem, we propose two possible strategies to define the cost-sensitive *query function*. The first strategy consists in maximizing a long-term reward that is defined as a tradeoff between the expected prediction utility of the samples and the associated annotation cost. The tradeoff between the two terms is regulated by pre-defined parameters. This strategy results in modeling the query function as a bounded-horizon MDP. The second strategy addresses the problem by maximizing a long-term reward, which is defined as the expected prediction utility of the samples, while constraining the total annotation cost to be smaller than a pre-defined budget B . This second strategy results in modeling the query function as an MDP where the number of future steps is not bounded by a fixed horizon H , but the looking-ahead procedure is stopped when the accumulated cost exceeds the available budget B . We refer to the two query functions as: 1) *query with bounded horizon*, and 2) *query with constrained budget*. Both strategies share the same formalization based on MPD, which is described in the next subsection.

A. Query functions as Markov Decision Process

We model the proposed *query functions* on the basis of an MDP $(\mathcal{S}, \mathcal{A}, T, R)$, where the state $s_t \in \mathcal{S}$ is the sequence of labeled samples $x_0 y_0 x_1 y_1 \dots x_t y_t$ that contains the last collected sample $x_0 y_0$ when the query was called and those that are planned to be selected and labeled in the future up to time-step t . The samples $x_0 \dots x_t \in \mathcal{X}$ are vectors that contain both the features used for the classification task and for computing the annotation cost of the sample, e.g., the spectral signatures and the geographical positions of the samples, which are necessary to compute the traveling cost. Both sets of features can be extracted from geo-referenced RS data. The label $y_0 \in \mathcal{Y}$ is the last actually annotated label, whereas $y_1 \dots y_t \in \mathcal{Y}$ can just be predicted during the planning procedure by a transition function, which can be considered a prior probability on the distribution of the labels. An action $a \in \mathcal{A}$ corresponds to enquiring the label of a sample x , therefore $\mathcal{A} = \mathcal{U}$. Any action is associated with: 1) a prediction utility $u(s_t, x_{t+1})$ that evaluates the expected gain in predictive

performance of the classifier after labeling the selected sample x_{t+1} in state s_t and including it in the training set, and 2) an annotation cost $\Theta(s_t, x_{t+1})$ that evaluates the actual cost for obtaining the label of sample x_{t+1} in state s_t . The transition function $T(s_t, x_{t+1}, s_{t+1})$ is the probability of labeling sample x_{t+1} with y_{t+1} given the sequence of labeled examples s_t , i.e., $T(s_t, x_{t+1}, s_{t+1}) = P(y_{t+1} | x_0 y_0 \dots x_t y_t, x_{t+1})$. The action of annotating the sample x_{t+1} in state s_t is associated with a reward $R(s_t, x_{t+1})$. Different definitions of the reward function will be considered in the next subsections for the two proposed queries.

The aforementioned general model, theoretically casts the query problem into a stochastic MDP. However, the probability of future labels $P(y_{t+1} | x_0 y_0 \dots x_t y_t, x_{t+1})$ represented by the transition function T can hardly be estimated and used in the design of the query function. For this reason, we consider here an agnostic approach, where a uniform distribution is used for the transition function. We also do not use the predicted labels in the reward function. This results in an equivalent representation of the MDP, where the transition function is a deterministic function and the state $s_t \in \mathcal{S}$ is the sequence of only unlabeled samples $x_0 x_1 \dots x_t$. The prediction utility $u(s_t, x_{t+1})$ and the annotation cost $\Theta(s_t, x_{t+1})$ are functions of both the selected sample x_{t+1} and the current sequence of samples $s_t = x_0 x_1 \dots x_t$. In order to estimate the prediction utility $u(s_t, x_{t+1})$, the overlap of information between x_{t+1} and the samples in the sequence $x_1 \dots x_t$ is taken into account by considering a tradeoff between uncertainty and diversity. Different strategies usually adopted for batch-mode active learning can be employed for this purpose in the proposed architecture.

B. Query with Bounded Horizon

In the query with bounded horizon, the reward is defined on the basis of both the expected prediction utility $u(s_t, x_{t+1})$ and the annotation cost $\Theta(s_t, x_{t+1})$. Using a linear model, the reward function is defined as:

$$R(s_t, x_{t+1}) = \alpha u(s_t, x_{t+1}) - \beta \Theta(s_t, x_{t+1}), \quad (3)$$

where α and β are user-defined parameters that tune the trade-off between prediction utility and annotation cost.

The proposed query function requires the exploration of a decision tree in order to select the action that leads to the maximum expected long-term cumulative reward. Figure 2 shows an example of such a tree. The tree exploration starts with the last annotated sample, which is always denoted by x_0 and used in calculating the reward function. The value of labeling each sample $x^i \in \mathcal{U}$ is calculated by summing the reward $R(x_0, x^i)$, given in Equation 3, and the subsequent discounted rewards for $H - 1$ steps, which are found by exploring the sub-tree starting from state $x_0 x^i$. Finally, the query function selects the sample $\hat{x} \in \mathcal{U}$ that maximizes the cumulative reward, i.e.,

$$\hat{x} = \operatorname{argmax}_{x^i \in \mathcal{U}} [R(x_0, x^i) + \gamma V_{H-1}^*(x_0 x^i)], \quad (4)$$

where $V_{H-1}^*(x_0 x^i)$ is the value of the maximum cumulative reward that can be obtained starting from state $x_0 x^i$ and

summing the discounted rewards for the subsequent $H - 1$ steps.




Fig. 2. Decision tree for sample selection with the proposed query function with bounded horizon.

Algorithm 1 illustrates the recursive planning procedure used for computing the query function with bounded horizon.

Algorithm 1 Query(s_t, \mathcal{U}, H) with bounded Horizon

Input: State $s_t = x_0x_1x_2 \dots x_t$
Pool $\mathcal{U} = (x^i)_{i=1}^n$ of n unlabeled samples
Horizon H

```

foreach  $x \in \mathcal{U}$  do
   $R(s_t, x_{t+1}) = \alpha u(s_t, x_{t+1}) - \beta \Theta(s_t, x_{t+1})$ 
   $s_{t+1} = x_0x_1x_2 \dots x_t x$ 
  if  $H > 1$  then
     $V_{H-1}^*(s_{t+1}) = \text{Query}(s_{t+1}, \mathcal{U} - \{x\}, H - 1)$ 
  else
     $V_{H-1}^*(s_{t+1}) = 0$ 
  end
   $Q_H^*(s_t, x) = R(s_t, x) + \gamma V_{H-1}^*(s_{t+1})$ 

```

end
 $x_{t+1} = \arg \max_{x \in \mathcal{U}} Q_H^*(s_t, x),$
 $V_H^*(s_t) = Q_H^*(s_t, x_{t+1})$

Output: A sample x_{t+1} with value $V_H^*(s_t)$

C. Query with Constrained Budget

In the second proposed strategy, the reward is defined as only the expected prediction utility of the selected sample $u(s_t, x_{t+1})$. In this case, we do not use a bounded-horizon MDP with a fixed number of future time-steps to explore. Instead, we stop the planning when the accumulated cost exceeds the available budget B . In this way, we maximize the long-term classification accuracy with respect to the available budget. This strategy is more appropriate in real applications, where fixing a priori a good tradeoff between the prediction utility and the annotation cost can be difficult. Also in this case, the query requires the exploration of a decision tree. The value of labeling each sample $x \in \mathcal{U}$ in state $s_t = x_0x_1 \dots x_t$ is calculated by summing the reward $R(s_t, x)$ and the subsequent rewards that are found by the exploration of the sub-tree until the cumulative cost exceeds the available budget. The depth of each sub-tree depends on B but is not the same for all sub-trees (as in the previous case). Finally, the

Algorithm 2 Query(s_t, \mathcal{U}, B) with constrained Budget

Input: State $s_t = x_0x_1x_2 \dots x_t$
Pool $\mathcal{U} = (x^i)_{i=1}^n$ of n unlabeled samples
Budget B

```

foreach  $x \in \mathcal{U}$  do
   $R(s_t, x) = u(s_t, x)$ 
   $s_{t+1} = x_0x_1x_2 \dots x_t x$ 
  if  $B - \Theta(s_t, x) \geq 0$  then
     $V^*(s_{t+1}) = \text{Query}(s_{t+1}, \mathcal{U} - \{x\}, B - \Theta(s_t, x))$ 
  else
     $V^*(s_{t+1}) = 0$ 
  end
   $Q^*(s_t, x) = R(s_t, x) + V^*(s_{t+1})$ 

```

end

$x_{t+1} = \arg \max_{x \in \mathcal{U}} Q^*(s_t, x)$
 $V^*(s_t) = Q^*(s_t, x_{t+1})$

Output: A sample x_{t+1} with value $V^*(s_t)$

Algorithm 3 CSAL with Lookahead

- 1: Train the classifier using the initial training set \mathcal{D}
 - 2: Initialize the state s to the last collected sample
 - 3: **repeat**
 - 4: $x = \text{Query}(s, \mathcal{U}, H)$ or $x = \text{Query}(s, \mathcal{U}, B)$
 - 5: The user labels the selected sample x with y
 - 6: $\mathcal{D} \leftarrow \mathcal{D} \cup \{(x, y)\}, \mathcal{U} \leftarrow \mathcal{U} - \{x\}, s \leftarrow x$
 - 7: Re-train the supervised classifier with \mathcal{D}
 - 8: **until** a stopping criterion is satisfied
-

query function selects the sample $\hat{x} = \arg \max_{x \in \mathcal{U}} Q(s_0, x)$. Algorithm 2 illustrates the recursive planning procedure used by the query with constrained budget.

D. Cost-Sensitive Active Learning with Lookahead

The workflow of the proposed CSAL is described in algorithm 3 for both the proposed query functions. At every iteration of the AL loop, the human expert is requested to travel to the location of the queried sample (e.g., tree) and provide its label. The sample x is then removed from the set of unlabeled samples \mathcal{U} and added to the training set \mathcal{D} . At the next iteration, the classifier is re-trained using the updated training set and the query function is called again. In this way, the CSAL method effectively guides the human expert in the field survey. Note that the looking-ahead procedure is used for selecting only the next sample (not for planning the whole future path in advance) and it is repeated at each iteration of the AL loop.

V. DATA SET DESCRIPTION

We applied the proposed CSAL methods to the optimization of the ground sample collection for tree species classification in the context of a real case study. In this section we describe the data set that we used for our experiments. The same area and data set were previously used in [8].

A. Study area

The study area is located in the municipality of Aurskog-Høland, southeastern Norway, 120-390 m above sea level (Figure 3). Approximately three quarters of the total land area (890 km²) is characterized by managed, productive forests dominated by *Pinus sylvestris* L. (Scots pine; 50%), *Picea abies* (L.) Karst (Norway spruce; 35%), and deciduous tree species (15%), such as *Betula spp.* L. (birch) and *Populus tremula* L. (aspen).

B. Remote sensing data

Hyperspectral images were acquired simultaneously with two different sensors: the HySpex VNIR-1600 and HySpex SWIR 320i sensor. The VNIR (visible and near infrared) sensor acquired data for 160 bands between 400 and 990 nm, with a spatial resolution of 0.4 m. The SWIR (short-wavelength infrared) sensor acquired data for 147 bands between 930 and 1700 nm, with a spatial resolution of 1.5 m. The sensors were mounted on a Piper Chieftain PA-31-350 fixed-wing aircraft flying at an altitude of 1500 m above ground level at a speed of 70 m/s. The two hyperspectral sensors are both line (pushbroom) scanners. Three hyperspectral images were acquired with each sensor.

C. Field data

Tree-species labels were collected through ground surveys on the basis of 23 circular sample plots. Among them 11 were located in spruce dominated forest and the remaining 12 in pine dominated forest. The size of each plot was 1000 m², except for one located in young forest, where the plot size was reduced to 500 m² due to a very high stem density. Tree species, DBH, and tree coordinates were recorded for all trees with DBH \geq 5 cm. For the purpose of this study, only the annotated tree species and coordinates are used. Tree positions were determined by measuring the azimuth and distance of the trees from the plot center with a total station (Topcon SokkiaSET5F). Plot center coordinates were determined using differential Global Navigation Satellite Systems with two Topcon Legacy E+ receivers as base and rover receivers, respectively. A total of 2407 trees were measured inside the 23 plots. The field surveys took approximately one day per plot involving the work of two experts.

D. Data preprocessing

The hyperspectral images were orthorectified using a digital terrain model, atmospherically corrected [34] and normalized [35] in order to reduce the spectral differences among the images. The SWIR images were resampled at 0.4 m in order to have the same spatial resolution of the VNIR ones.

Starting from the tree positions measured on the ground, individual tree crowns (ITCs) were manually identified and delineated on the hyperspectral data (see Figure 4). A total of 1001 out of the 2407 field measured trees were identified in the hyperspectral images. It is worth noting that during the field measurements also trees located in the undercanopy were measured, and thus these trees were not visible in the

TABLE I
NUMBER OF LABELED ITCs

| Class | Training set | Test set |
|---------------|--------------|----------|
| Pine | 201 | 198 |
| Spruce | 258 | 243 |
| Birch | 42 | 31 |
| Other species | 14 | 14 |

hyperspectral data. The ITCs were grouped into the following classes according to tree species: 1) Pine, 2) Spruce, 3) Birch, and 4) Other species. The samples were divided into a training set and a test set (see Table I). The two sets were defined in order to have a similar composition in terms of tree species and spatial distribution within the images. The values of the pixels inside each ITC segment were averaged in order to obtain a single spectral signature for each ITC. All the available bands were used in the classification process.

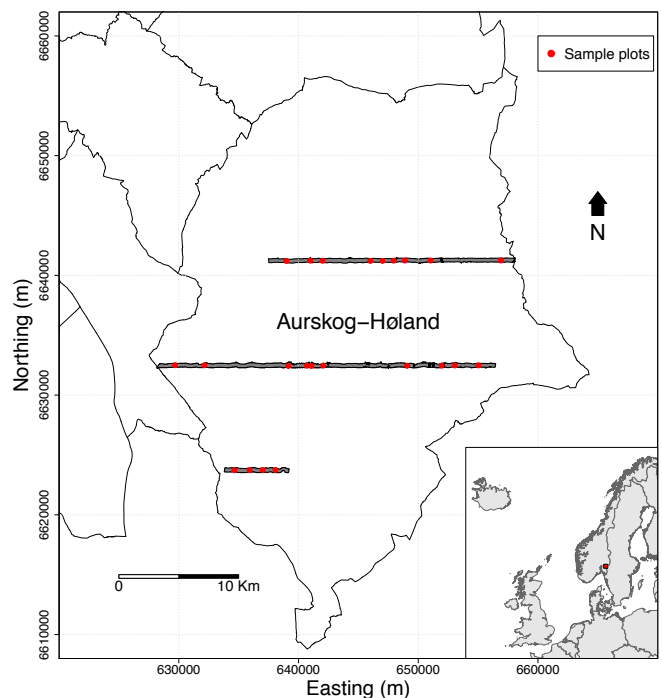


Fig. 3. Map of the study area showing the location of the 23 sample plots and hyperspectral data coverage. Inset map shows the location of the study area (red rectangle) in the map of northwestern Europe.

VI. EXPERIMENTAL ANALYSIS

In this section we present the experimental analysis that we carried out with the proposed CSAL method using both the proposed strategies: 1) query with bounded horizon, and 2) query with constrained budget.

A. Experimental Setting

The annotation cost $\Theta(s_t, x)$ was defined as the sum of two terms: 1) the traveling cost \mathcal{T} for reaching the location of the tree x (starting from the location of the last annotated tree x_t),




Fig. 4. False color representation of HySpex VNIR 1600 data acquired over a plot. The manually delineated ITCs appear in yellow.

and 2) the labeling cost \mathcal{L} for assigning the correct label to it:

$$\Theta(s_t, x) = \mathcal{T}(x_t, x) + \mathcal{L}(x). \quad (5)$$

The traveling cost \mathcal{T} considers the traveling time required by the human expert to reach the location of x , depending on the distance from x_t and the speed of the considered transportation mode. Two transportation modes are considered: by car or foot. A precise evaluation of the traveling cost may also take into account fuel expenses as well as any other costs. The labeling cost \mathcal{L} defines the cost of acquiring the right label for x once the human expert is already on the site. This cost does not depend on previously collected samples and here we assume that it is the same for all the samples, i.e., $\mathcal{L}(x) = \mathcal{L}$. In the most general framework, the different costs may be expressed in monetary terms. For simplicity, in our experiments we expressed them in terms of elapsed time. We considered a labeling time \mathcal{L} of 10 minutes and the traveling time \mathcal{T} was calculated according to the distance between samples and the speed of the transportation mode. The human expert walks from a tree to the next one if they are both in the same plot, whereas he travels by car for moving to another plot.

It is worth noting that the considered annotation cost $\Theta(s_t, x)$ represents a prediction of the expected real cost for labeling sample x . Several unexpected factors related to the accessibility of certain areas may affect the real annotation cost. Nevertheless, reasonable estimation of the annotation cost can be obtained by considering the distance between samples, the mean traveling speed and some prior information about the considered geographical area. One can also consider the altitude of the samples for obtaining a better estimate of the traveling cost by considering a digital elevation model (DEM)

of the area. Since the area considered in our experiments is relatively flat, the altitude information was not taken into account.

The classification was performed using an SVM with a *one-against-all* (OAA) multi-class architecture. Feature vectors associated with ITCs contain the VNIR and SWIR spectral channels as well as the position coordinates. All the experiments were performed in ten trials with initial training sets made up of 54 labeled samples (ITCs) randomly selected from three adjacent plots. 456 and 439 samples taken from the other 20 plots were used as pool \mathcal{U} and test set, respectively. The model selection of the SVM was carried out on the basis of the accuracy obtained on a validation set made up of 52 samples taken from the same plots of the initial training sets.

The prediction utility was calculated considering a tradeoff between uncertainty and diversity. The uncertainty criterion was evaluated considering a confidence measure defined for the multi-class case as:

$$c(x) = f_1(x) - f_2(x) \quad (6)$$

where $f_1(x)$ and $f_2(x)$ are the first and second highest output score of the binary SVMs in the OAA architecture. Querying the sample that minimizes $c(x)$ results in the selection of the closest sample to the boundary between the two most probable classes. The diversity was computed considering the kernel cosine-angular similarity between points [16]:

$$k^*(x, x_i) = \frac{k(x, x_i)}{\sqrt{k(x, x)k(x_i, x_i)}}, \quad (7)$$

where $k(\cdot, \cdot)$ is a positive semidefinite kernel function. In our experiments we adopted an RBF kernel function for both the diversity assessment and the SVM classification.

B. Query with Bounded Horizon: Results

In the experiments with the bounded-horizon query, the immediate prediction utility was computed as:

$$u(s_t, x) = - \left((1 - \rho)c(x) + \rho \sum_{i=1}^t k^*(x, x_i) \right), \quad (8)$$

where $\rho \in [0, 1]$ tunes the tradeoff between uncertainty and diversity, and $s_t = x_0 \dots x_t$. The reward function was finally calculated as:

$$R(s_t, x_{t+1}) = (1 - \lambda)u(s_t, x_{t+1}) - \lambda\Theta(s_t, x_{t+1}) \quad (9)$$

with $\lambda \in [0, 1]$. The values of $c(x)$, $\sum_{i=1}^t k^*(x, x_i)$ and $\Theta(s_t, x_{t+1})$ for all samples in the pool were normalized between zero and one to make them comparable.

It is worth noting that a full exploration of the general decision tree as described in algorithm 1 can be computationally prohibitive (for a large pool and horizon), and moreover, we should consider that the query function is run in real-time, i.e., while the human expert is on the field. For this reason, one should use a very fast heuristic to explore the decision tree. Here, we considered a heuristic for pruning samples from the decision tree in order to speed up the query function. The heuristic is based on restricting the search to the set of samples with the m highest immediate rewards.

We compared the results of the proposed CSAL method using the query with bounded horizon (CSAL Horizon Lookahead) with the following baseline methods: 1) random selection (Random), 2) selection by minimum confidence (Uncertainty), 3) selection of the closest geographical sample (Minimum Distance), 4) myopic CSAL (Myopic CSAL Horizon). One sample is selected for labeling at each iteration. The myopic CSAL selects the next sample according to just the immediate reward, as defined in (9), without looking ahead. This is equivalent to the proposed query with bounded horizon imposing $H = 1$. Note that baselines 2 and 3 are special cases of the myopic CSAL by setting $\lambda = 1$ and $\lambda = 0$, respectively. Selection of the closest geographical sample represents a simulation of traditional (passive) sampling procedures, which are not making use the learner's feedback. For the CSAL methods (with and without lookahead), the value of the parameter λ , tuning the tradeoff between prediction utility and annotation cost, was set to 0.2. For the proposed CSAL with lookahead, we set the planning horizon $H = 3$, the discount factor $\gamma = 0.9$, the parameter tuning the tradeoff between uncertainty and diversity $\rho = 0.8$, and the parameter for heuristic pruning of the tree search $m = 100$.

Figure 5 shows the obtained curves of the overall accuracy (OA), averaged over the ten trials starting from the initial training sets, versus the number of labeled samples added to the original training set. Not surprisingly, the standard active learning method based on uncertainty (i.e., not considering the annotation cost) resulted in the best performances in this case. However, it is more important to consider Figure 6, which reports the average OA with respect to the time spent by the human annotator to travel and collect sample labels in the field. The learning curves are reported up to 80 hours of field work in order to highlight the difference between the proposed method and the other baselines. These results show that the proposed method substantially improves the classification accuracy of the considered baselines with respect to the annotation cost. The myopic CSAL improves the other baselines; CSAL with lookahead further improves the classification accuracy with respect to the myopic method. The advantage is particularly evident in the first 30 hours of the simulations. It is worth noting that the proposed CSAL method leads to the top accuracy of 92.6% in 50 simulated hours, whereas a traditional approach like minimum distance requires approximately 80 hours to get the same accuracy. Since the minimum distance approach reasonably represents how the real field campaigns are carried out with a standard sampling approach, the difference between the learning curve obtained with that approach and of the proposed CSAL method gives a quantitative evaluation of the potential improvement that the proposed method can offer in real applications. Uncertainty-based method and random selection take much longer to converge to the top accuracy (150 and 340 hours, respectively). Figure 7 reports the graphs of the time spend by the human expert to travel and collect samples versus the number of collected samples. As one can observe from these graphs, using the uncertainty-based method and random selection, the human expert takes in average much longer to reach the location of the next sample, leading to very high traveling costs.

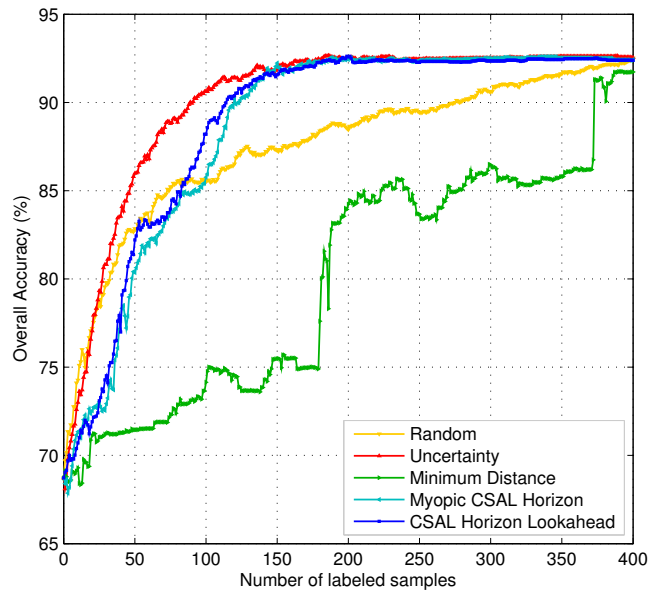


Fig. 5. Average overall accuracy versus the number of labeled samples added to the initial training set.

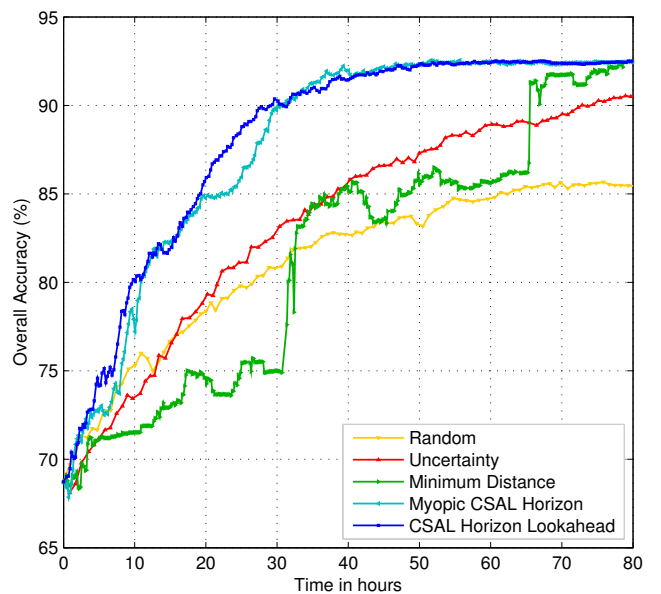


Fig. 6. Average overall accuracy versus the time spent by the human annotator to travel and collect sample labels in the field.

C. Query with Constrained Budget: Results

In the experiments with the constrained-budget query, the immediate prediction utility was computed as:

$$u(s_t, x) = \max \left(0, 1 - c(x) - \rho \sum_{i=1}^t k^*(x, x_i) \right), \quad (10)$$

where $\rho \in [0, 1]$ tunes the trade-off between uncertainty and diversity and $s_t = x_0 \dots x_t$. The reward was defined as the prediction utility, i.e., $R(s_t, x_{t+1}) = u(s_t, x_{t+1})$. This particular definition of the prediction utility was considered in order to have only positive reward values and for stopping the

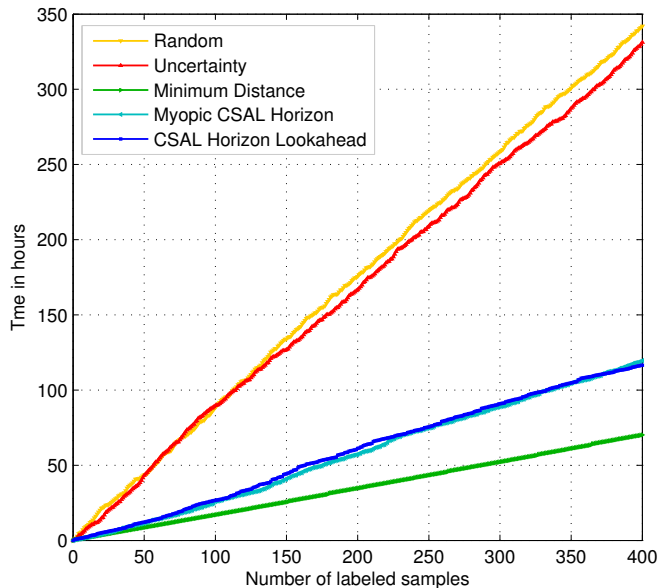


Fig. 7. Time spend by the human expert in the annotation process versus the number of collected samples.

recursive tree exploration when $R(s_t, x_{t+1}) = 0$. This results in an additional stopping criterion for the lookahead planning that speeds up the computation of algorithm 2.

We compared the results of the proposed CSAL method using the query with constrained budget (CSAL Budget Lookahead) with the following baseline methods: 1) random selection (Random), 2) selection by maximum uncertainty (Uncertainty), 3) selection of the closest geographical sample (Minimum Distance), 4) myopic CSAL (Myopic CSAL Budget). One sample is selected for labeling at each iteration. The myopic CSAL method selects the most uncertain sample among the ones whose immediate cost does not exceed the budget B . We set $B = 30$ minutes for both the myopic and the nonmyopic CSAL method. We set the tuning parameter $\rho = 0.3$.

Figure 8 shows the obtained average OA versus the number of labeled samples added to the original training set. Figure 9 reports the average OA as a function of the time spent by the human annotator to collect sample labels in the field. The obtained learning curves confirm the effectiveness of the proposed method, which results in significantly higher classification accuracies with respect to the traveling cost compared with the considered baselines. We observe that also in this case the proposed CSAL method leads to an OA of 92% in 50 hours of field work, whereas a traditional approach like minimum distance requires approximately 80 hours to reach the same accuracy. Figure 10 shows the averaged producer's accuracies (PA) of the classes obtained with the proposed CSAL method with lookahead versus the time spent by the human annotator for the sample collection. Table II reports the averaged accuracies of the classes obtained at different times with the proposed method CSAL with lookahead and the myopic strategy. Accuracies are reported in terms of producer's accuracies for the single classes [36]

and OA (mean values and standard deviations are reported). Results show that the proposed method lead in general to better accuracies with a lower standard deviation on the ten trials. The practical usefulness of the proposed method in the context of tree species classification is remarkable. As underlined in the introduction, a reduction of the time for the field measurements can noticeably reduce the inventory costs. From Table II it is clear that if only a detailed classification of Pine and Spruce species is needed, 30 hours of field work are enough for obtaining accurate classification results, i.e., producer's accuracies higher than 90%. These two species represent in Norway more than 80% of the total forest volume [37], and 99% of the total harvested volume [38] and thus a proper distinction of them is very important. The classification of classes "Birch" and "Other species" is problematic because very few samples are available for those classes and the classifier cannot therefore accurately discriminate them from the other classes (species different from Pine, Spruce, and Birch are very rare in boreal forests).

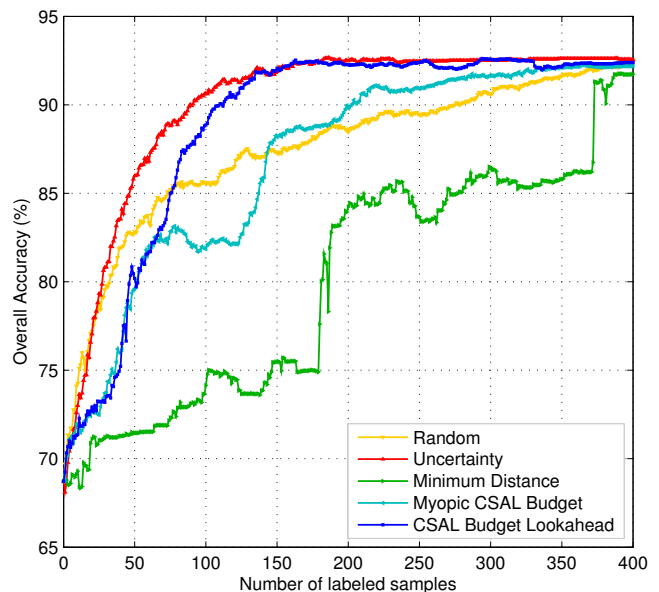


Fig. 8. Average overall accuracy versus the number of labeled samples added to the initial training set.

D. Operational Considerations

The choice between the two proposed query functions may depend on the available prior information that can be used for setting the free parameters of the methods. In case of using the query with bounded horizon, the λ parameter can be set by taking into account prior information about the expected average traveling cost (time) $\mathbb{E}\{\mathcal{T}\}$ and the fixed labeling cost \mathcal{L} . If $\mathcal{L} \gg \mathbb{E}\{\mathcal{T}\}$, the effect of using the cost Θ in the query function becomes negligible. If $\mathcal{L} \ll \mathbb{E}\{\mathcal{T}\}$, we expect that the optimal λ should be closer to one. On the basis of this observation, a possible heuristic to choose λ is the following:

$$\hat{\lambda} = \frac{\mathbb{E}\{\mathcal{T}\}}{\mathbb{E}\{\mathcal{T}\} + \mathcal{L}}. \quad (11)$$

TABLE II
AVERAGE PRODUCER'S ACCURACIES, AND OVERALL ACCURACIES OBTAINED AT DIFFERENT TIMES WITH THE PROPOSED METHOD CSAL BUDGET WITH LOOKAHEAD (CSLA) AND THE MYOPIC CASE (CSMY), AVERAGED OVER THE TEN TRIALS (MEAN VALUES AND STANDARD DEVIATIONS ARE REPORTED).

| Class | 20 hours | | 40 hours | | 60 hours | |
|---------------|-------------------------|------------------------|------------------------|------------------------|------------------------|------------------------|
| | CSLA | CSMY | CSLA | CSMY | CSLA | CSMY |
| Pine | 81.1% ± 8.6 | 78.1% ± 7.1 | 95.0% ± 1.0 | 90.7% ± 1.5 | 96.5% ± 0.9 | 95.9% ± 0.5 |
| Spruce | 92.0% ± 2.2 | 92.0% ± 2.7 | 94.5% ± 1.6 | 94.1% ± 0.9 | 95.0% ± 1.3 | 93.8% ± 1.4 |
| Birch | 40.0% ± 15.7 | 58.5% ± 7.6 | 62.3% ± 8.8 | 51.9% ± 5.5 | 66.2% ± 4.7 | 59.6% ± 5.2 |
| Other Species | 12.5% ± 16.7 | 2.5% ± 5.3 | 11.3% ± 4.0 | 12.5% ± 0.0 | 12.5% ± 0 | 17.5% ± 6.5 |
| OA | 83.0% ± 3.1 | 82.6% ± 2.9 | 91.3% ± 1.0 | 88.7% ± 1.2 | 92.4% ± 0.6 | 91.2% ± 0.9 |




Fig. 9. Average overall accuracy versus the time spent by the human annotator to travel and collect sample labels in the field.

Figure 11 reports the learning curves obtained using the query with bounded horizon setting different values for the parameter λ . These graphs show that the results are not particularly sensitive to the parameter value. However, in several real applications, the query with constrained budget might be more appropriate. In such a case, the tradeoff between immediate utility and annotation cost should not be fixed in advance. The value of the budget B can be set not critically by considering mainly the minimum annotation cost (time) and computational issues. Figure 11 reports the learning curves obtained by the query with constrained budget using different values for B . Better accuracies are obtained with small values of B , because the estimation of the expected prediction utility becomes obviously less reliable after several time-steps planned in the future. Typical values are in the order of few (e.g., two or three) times the minimum annotation time. The same consideration applies to the choice of H in the query with bounded horizon (typical values are two or three in that case). Typical values of the parameter ρ are between 0.3 and 0.7. However, also the choice of this parameter does not significantly affect the performance of the proposed method. This behavior is confirmed by the graphs reported in Figure 13, which reports

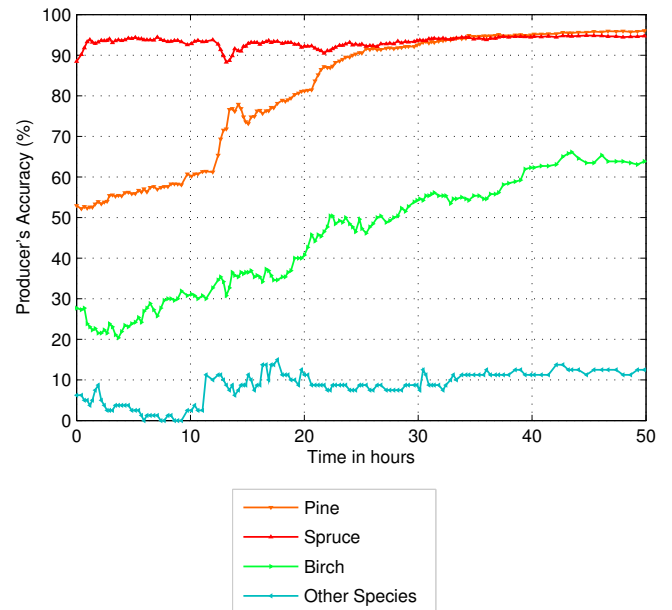


Fig. 10. Average producer's accuracies of the classes obtained with the proposed query with constrained budget versus the time spent by the human annotator to travel and collect sample labels in the field.

the learning curves of the proposed CSAL method with budget constraint for different values of such parameter.

Regarding the selection of the initial training set, several possible sampling procedures can be adopted, which depends on the application field. In the specific case of forest inventories, the field data collection is usually carried out inside field plots. The initial training samples can therefore be collected by annotating tree labels from few (e.g., 2 or 3) sample plots. This is the procedure adopted in our experimental analysis. The proposed interactive CSAL method is independent from the sampling strategy adopted for selecting the initial plots. The authors in [39] introduces an unsupervised method for plot or tree selection that can be adopted as the initial step of the proposed interactive CSAL method. It is worth noting that the size of the initial training set is usually substantially smaller than the size of the training set at convergence. In our experiments, the initial training set is made up of 52 samples, while the convergence is reached with approximately 200 training samples. Moreover, the initial training set is collected from three adjacent plots, for a cost of 9.2 hours of field work (calculated according to the collection with a

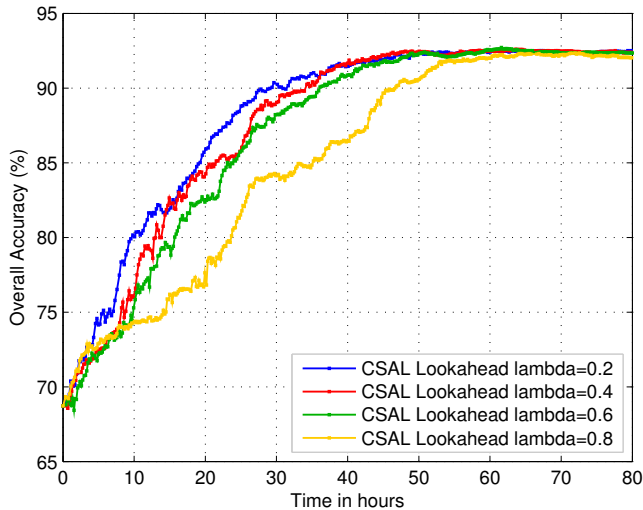


Fig. 11. Average overall accuracy versus the annotation time obtained by the query with bounded horizon using different values for the tradeoff parameter λ .

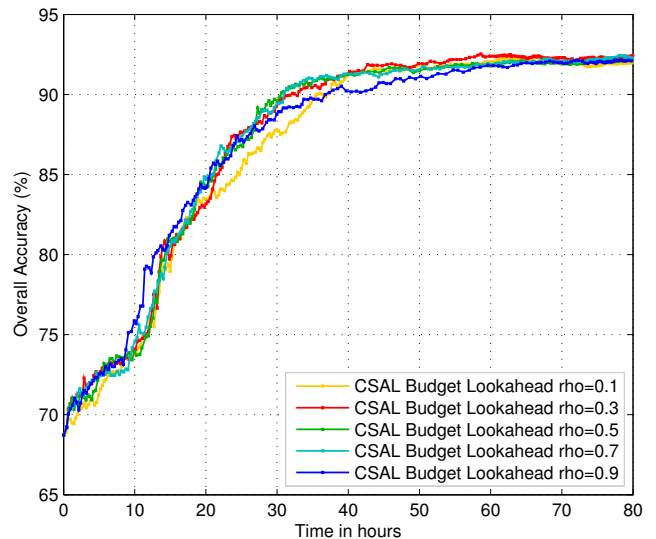


Fig. 13. Average overall accuracy versus the annotation time obtained by the query with constrained budget using different values of ρ .

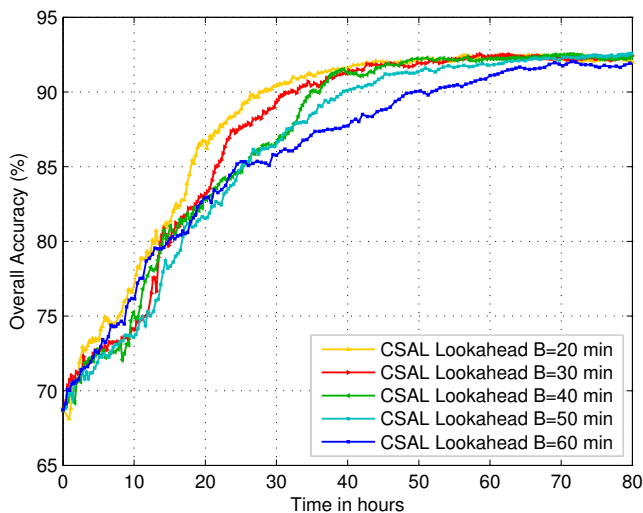


Fig. 12. Average overall accuracy versus the annotation time obtained by the query with constrained budget using different values of budget B .

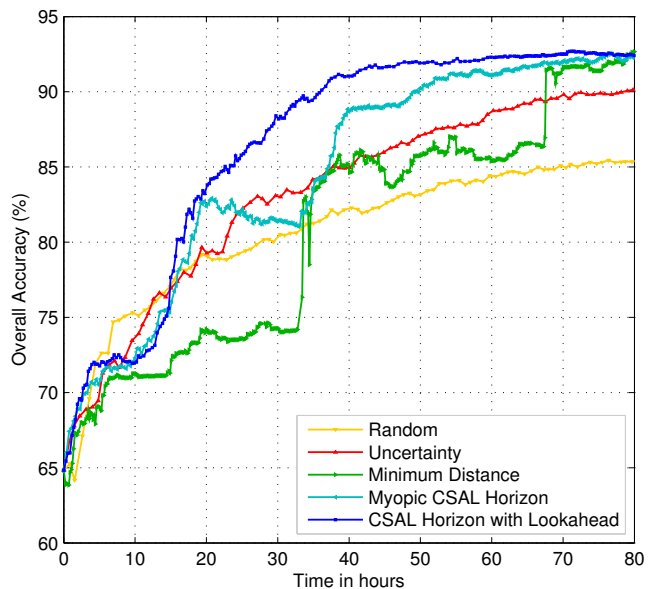


Fig. 14. Average overall accuracy versus the time spent by the human annotator to travel and collect sample labels in the field. Initial training sets are collected on 2 adjacent plots (41 initial training samples) corresponding to 7 hours of field work.

minimum distance approach). The size and the quality of the initial training set is not particularly critical for the subsequent selection process. We carried out additional experiments, considering initial trainings sets collected in only 2 adjacent plots (41 labeled samples) corresponding to 7 hours of field work (calculated as in the previous case). Figure 14 shows the learning curves obtained by the CSAL with constrained budget and the baselines, starting from these smaller initial training sets. From this results, we can observe that the size of the initial training set does not affect the convergence capability of the proposed CSAL method. This result is in agreement with other studies about AL reported in the literature (e.g., [3], [25], [26]).

Table III reports the mean query and annotation times for the considered selection strategies, respectively. The query time considers the mean computational time for retraining the SVM with the last labeled sample and for selecting the next

one. The experiments have been carried out using a laptop computer with a 2.2 GHz CPU and 8 GB of RAM. The annotation time refers to the mean time taken by the human expert to reach the location of the selected sample and to label it. The computational time required by query with bounded horizon is obviously higher than the one of the myopic method and the other competing strategies, but it is still negligible with respect to the annotation time. The implementation of the query with constrained budget resulted in a very fast computation. Both the proposed methods are therefore suitable (in terms of computational requirements) for the considered application.

TABLE III

MEAN QUERY AND ANNOTATION TIMES FOR THE CONSIDERED SELECTION STRATEGIES. QUERY TIMES CONSIDER THE MEAN COMPUTATIONAL TIME FOR RETRAINING THE SVM WITH THE LAST LABELED SAMPLE AND FOR SELECTING THE NEXT ONE. THE ANNOTATION TIMES REFER TO THE MEAN TIME TAKEN BY THE HUMAN EXPERT TO REACH THE LOCATION OF THE SELECTED SAMPLE AND TO LABEL IT.

| Selection Strategy | Query | Annotation |
|------------------------|---------|------------|
| Random | 0.073 s | 50.2 min |
| Uncertainty | 0.120 s | 48.9 min |
| Minimum Distance | 0.085 s | 10.5 min |
| Myopic CSAL Horizon | 0.123 s | 16.9 min |
| Myopic CSAL Budget | 0.090 s | 13.2 min |
| CSAL Horizon Lookahead | 1.248 s | 17.0 min |
| CSAL Budget Lookahead | 0.092 s | 16.5 min |

VII. CONCLUSION

In this paper, we proposed a novel cost-sensitive active learning method that allows one to plan the sequential annotation process in order to maximize the long-term classification accuracy while minimizing the total annotation cost. To this end, the active learning problem was modeled in the framework of a Markov Decision Process. Two possible strategies to address the optimization problem have been proposed, giving rise to the definition of two different query functions: 1) query with bounded horizon, and 2) query with limited budget. The proposed method was applied to the problem of optimizing the ground sample collection by a human expert for the classification of forest tree species using hyperspectral images. The experimental results show the effectiveness of the proposed query functions with lookahead and their significant improvements over myopic strategies and other baseline methods. We observed that the proposed queries converge to a high classification accuracy with a significant lower cost compared to state-of-the-art methods.

It is worth noting that the implementation of the proposed method in the considered scenario of forest inventories requires a change in the standard protocols used for ground surveys. New and more effective protocols can be defined by leveraging the information of geo-referenced hyperspectral images to extract tree features and coordinates. The information of the tree coordinates extracted from RS data is generally not considered in standard protocols, but can be effectively considered by the CSAL method. Field surveys can then be guided by the proposed strategy to interactively classify the species of the trees identified in the hyperspectral image with the aid of a portable computer or a tablet. This may reduce significantly the costs of the inventory.

In general, the proposed method can be applied to any learning problem where the cost associated with the annotation of a given sample depends on the previously labeled samples. This is the case when annotating a queried sample is an action that changes the state of a dynamic system, and the cost is a function of the state of the system. We are currently investigating the use of the proposed method for different applications.

ACKNOWLEDGMENT

The work of Dr. Claudio Persello and of Dr. Michele Dalponte has been supported by the Autonomous Province of Trento and the European Community in the framework of the research programme Marie Curie actions - COFUND "Trentino - The Trentino programme of research, training and mobility of post-doctoral researchers - PCOFUND-GA-2008-226070 (call 3 - post-doc 2010 Outgoing)".

REFERENCES

- [1] B. Settles, "Active learning literature survey," University of Wisconsin-Madison, Computer Sciences Technical Report 1648, 2009.
- [2] S. Rajan, J. Ghosh, and M. M. Crawford, "An active learning approach to hyperspectral data classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 46, no. 4, pp. 1231-1242, April 2008.
- [3] B. Demir, C. Persello, and L. Bruzzone, "Batch-mode active-learning methods for the interactive classification of remote sensing images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 3, pp. 1014-1031, 2011.
- [4] T. Brandtberg, T. A. Warner, R. E. Landenberger, and J. B. McGraw, "Detection and analysis of individual leaf-off tree crowns in small footprint, high sampling density lidar data from the eastern deciduous forest in North America," *Remote Sensing of Environment*, vol. 85, no. 3, pp. 290-303, 2003.
- [5] J. Holmgren and A. Persson, "Identifying species of individual trees using airborne laser scanner," *Remote Sensing of Environment*, vol. 90, no. 4, pp. 415-423, 2004.
- [6] M. Clark, D. Roberts, and D. Clark, "Hyperspectral discrimination of tropical rain forest tree species at leaf to crown scales," *Remote Sensing of Environment*, vol. 96, no. 3-4, pp. 375-398, 2005.
- [7] M. Dalponte, L. Bruzzone, and D. Gianelle, "Tree species classification in the southern alps based on the fusion of very high geometrical resolution multispectral/hyperspectral images and lidar data," *Remote Sensing of Environment*, vol. 123, pp. 258-270, 2012.
- [8] M. Dalponte, H. Ørka, T. Gobakken, D. Gianelle, and E. Næsset, "Tree species classification in boreal forests with hyperspectral data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 51, no. 5, pp. 2632-2645, 2013.
- [9] T. G. Gregoire and H. Valentine, *Sampling Strategies for Natural Resources and the Environment*. Chapman and Hall/CRC, 2008.
- [10] D. D. Lewis and W. A. Gale, "A sequential algorithm for training text classifiers," in *Proceedings of the 17th annual international ACM SIGIR conference on Research and development in information retrieval*, 1994.
- [11] G. Schohn and D. Cohn, "Less is more: Active learning with support vector machines," *International Conference on Machine Learning*, pp. 839-846, 2000.
- [12] C. Campbell, N. Cristianini, and A. Smola, "Query learning with large margin classifiers," *Proceedings of the Seventeenth International Conference on Machine Learning*, pp. 111-118, 2000.
- [13] S. Tong and D. Koller, "Support vector machine active learning with applications to text classification," *Journal of Machine Learning Research*, vol. 2, pp. 45-66, March 2002.
- [14] H. S. Seung, M. Opper, and H. Sompolinsky, "Query by committee," in *Proceedings of the fifth annual workshop on Computational learning theory*, ser. COLT '92, 1992, pp. 287-294.
- [15] N. Roy and A. McCallum, "Toward optimal active learning through sampling estimation of error reduction," in *International Conference on Machine Learning*, 2001, pp. 441-448.
- [16] K. Brinker, "Incorporating diversity in active learning with support vector machines," *International Conference on Machine Learning*, pp. 59-66, 2003.
- [17] Z. Xu, K. Yu, V. Tresp, X. Xu, and J. Wang, "Representative sampling for text classification using support vector machines," in *Proceedings of the 25th European Conference on Information Retrieval Research*, ser. ECIR '03, 2003.
- [18] H. T. Nguyen and A. Smeulders, "Active learning using pre-clustering," in *Proceedings of the twenty-first international conference on Machine learning*, ser. ICML '04. New York, NY, USA: ACM, 2004, pp. 79-.
- [19] S. C. H. Hoi, R. Jin, J. Zhu, and M. R. Lyu, "Batch mode active learning and its application to medical image classification," in *International Conference on Machine Learning*, 2006.
- [20] Y. Guo and D. Schuurmans, "Discriminative batch mode active learning," in *Neural Information Processing Systems*, 2007.

- [21] J. Azimi, A. Fern, X. Z. Fern, G. Borradaile, and B. Heeringa, "Batch active learning via coordinated matching," in *International Conference on Machine Learning*, 2012.
- [22] W. Di and M. M. Crawford, "Active learning via multi-view and local proximity co-regularization for hyperspectral image classification," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 3, pp. 618–628, 2011.
- [23] D. Tuia, F. Ratle, F. Pacifici, M. Kanevski, and W. Emery, "Active learning methods for remote sensing image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 7, pp. 2218–2232, 2009.
- [24] J. Li, J. M. Bioucas-Dias, and A. Plaza, "Hyperspectral image segmentation using a new bayesian approach with active learning," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 10, pp. 3947–3960, 2011.
- [25] W. Di and M. M. Crawford, "View generation for multiview maximum disagreement based active learning for hyperspectral image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 50, no. 5, pp. 1942–1954, 2012.
- [26] C. Persello and L. Bruzzone, "Active learning for domain adaptation in the supervised classification of remote sensing images," *IEEE Transactions on Geoscience and Remote Sensing*, in press.
- [27] G. Matasci, D. Tuia, and M. Kanevski, "Svm-based boosting of active learning strategies for efficient domain adaptation," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 5, no. 5, pp. 1335–1343, 2012.
- [28] C. Persello, "Interactive domain adaptation for the classification of remote sensing images using active learning," *IEEE Geoscience and Remote Sensing Letters*, vol. 10, no. 4, pp. 736 – 740, 2013.
- [29] A. Liu, G. Jun, and J. Ghosh, "Spatially cost-sensitive active learning," in *SIAM International Conference on Data Mining*, 2009, pp. 814–825.
- [30] B. Demir, L. Minello, and L. Bruzzone, "A cost-sensitive active learning technique for the definition of effective training sets for supervised classifiers," in *IEEE International Conference on Geoscience and Remote Sensing*, 2012, pp. 1781–1784.
- [31] M. L. Puterman, *Markov Decision Processes Discrete Stochastic Dynamic Programming*. Wiley & Sons, 2005.
- [32] W. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. John Wiley and Sons, Inc., 2007.
- [33] Y. Mansour and S. Singh, "On the complexity of policy iteration," *proceedings of the 15th international conference on uncertainty in AI*, pp. 401–408, 1999.
- [34] L. Bernstein, R. Sundberg, R. L. T. Perkins, and A. Berk, "A new method for atmospheric correction and aerosol optical property retrieval for VIS-SWIR multi- and hyperspectral imaging sensors: Quac (quick atmospheric correction)," in *IEEE International Conference on Geoscience and Remote Sensing*, 2005, pp. 3549–3552.
- [35] B. Yu, I. Ostland, P. Gong, and R. Pu, "Penalized discriminant analysis of in situ hyperspectral data for conifer species recognition," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 37, pp. 2569–2577, 1999.
- [36] R. G. Congalton and K. Green, *Assessing the Accuracy of Remotely Sensed Data: Principles and Practices*. CRC Press, 1999.
- [37] J. Y. Larsson and G. Hysten, "Statistics of forest conditions and forest resources in norway," Viten fra Skog og landskap, url: <http://www.skogoglandskap.no/filearchive/viten-1-07.pdf>, 2007.
- [38] Statistics Norway, "Skogstatistikk 2008," Oslo–Kongsvinger, 2008.
- [39] M. Dalponte, L. T. Ene, H. Ørka, T. Gobakken, and E. Næsset, "Unsupervised selection of training plots and trees for tree species classification with hyperspectral data," in *IEEE International Conference on Geoscience and Remote Sensing*, 2013.