Chemometrics and Intelligent Laboratory Systems 132 (2014) 91–102

Contents lists available at ScienceDirect

Chemometrics and Intelligent Laboratory Systems

journal homepage: www.elsevier.com/locate/chemolab

Active classification with arrays of tunable chemical sensors

Rakesh Gosangi, Ricardo Gutierrez-Osuna

Department of Computer Science and Engineering, Texas A&M University, United States

ARTICLE INFO

Article history:
Received 20 October 2013
Received in revised form 5 January 2014
Accepted 6 January 2014
Available online 15 January 2014

Keywords:
Active sensing
Sensor arrays
Chemical classification
Metal-oxide sensors
Fabry–Perot interferometers

ABSTRACT

This paper presents Posterior-Weighted Active Search (PWAS), an active-sensing algorithm for classification of volatile compounds with arrays of tunable chemical sensors. The algorithm combines concepts from feature subset selection and sequential Bayesian filtering to optimize the sensor array tunings on-the-fly based on information from previous measurements. Namely, the algorithm maintains an estimate of the posterior probability associated with each chemical class, and updates it sequentially upon arrival of each new sensor observations. The updated posteriors are then used to bias the selection of the next sensor tunings towards the most likely classes, in this way reducing the number of measurements required for discrimination. We characterized PWAS on a database of infrared absorption spectra with 250 analytes, and then validated it experimentally on an array of metal-oxide sensors. Our results show that PWAS outperforms passive-sensing approaches based on sequential forward selection, both in terms of classification performance and robustness to noise in sensor measurements.

1. Introduction

Chemical sensors are generally used as first-order devices, where one measures the sensor's response at a fixed setting, e.g., absorption of an optical sensor at a specific wavelength, or conductivity of a solid-state sensor at a specific operating temperature [1]. In many cases, additional information can be extracted by modulating some internal property of the sensor. As an example, measuring the conductivity of a metal-oxide chemical sensor at different temperatures can provide a wealth of discriminatory information [2]. However, this additional information comes at a cost, such as sensing times or power consumption. For this reason, feature subset selection (FSS) techniques are commonly used to identify a subset of the most informative sensor configurations.

Over the past decade, a handful of investigators in the chemical sensor community have explored active sensing as an alternative to FSS [3–7]. In contrast with FSS, where the sensor configurations are optimized off-line, active sensing adapts the sensor configurations in real-time based on information obtained from previous measurements. In previous work [6,7], we showed that active sensing can achieve higher classification performance than FSS with fewer measurements and provides a trade-off between sensing costs and classification performance. Unfortunately, these active-sensing methods were developed for individual sensors, and do not scale up to sensor arrays. First, the number of operating configurations for a sensor array grows exponentially with the size of the array; given an N-sensor array with D configurations per sensor, there exist D^N unique configurations. Second, chemical sensor arrays are notoriously collinear (i.e., their response across multiple chemicals is correlated), so additional strategies are needed to account for correlation among sensors.

This article proposes Posterior-Weighted Active Search (PWAS), an active-sensing algorithm for sensor arrays that addresses both issues (combinatorial explosion and sensor collinearity). PWAS performs active sensing by optimizing the sensors' tunings towards the most likely classes at each sensing step; for this purpose, PWAS uses the sequential Bayesian filter of our prior work [6,7] to update the posterior probability of each class upon arrival of each new measurement. To cope with the combinatorial explosion in sensor array tunings, PWAS uses local search to build the sensor array configurations incrementally (one sensor at a time). Finally, to cope with sensor collinearity PWAS uses one of the two objective functions we have developed for this purpose. The first objective function is a parametric filter derived from the multivariate Fisher score [8], and weighs the within-class and between-class scatter matrices according to the estimated class posteriors. The second objective function is a non-parametric information-theoretic filter that measures feature relevance and feature redundancy with respect to the class posteriors. PWAS operates following a 'search-sense-update' sequence. During the 'search' step, the algorithm uses the local search and objective functions to build the next sensor array configuration. During the 'sense' step, the algorithm takes sensor measurements using the selected configuration. During the final 'update' step, the algorithm re-estimates the class posteriors by feeding the measured sensor responses to a sequential Bayesian update equation. This search-sense-update process is continued until a predefined stopping

E-mail addresses: rakesh@cse.tamu.edu (R. Gosangi), rgutier@cse.tamu.edu (R. Gutierrez-Osuna).

0169-7439/$ – see front matter © 2014 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.chemolab.2014.01.001

criterion is met, at which point the final class label is declared based on maximum a posteriori (MAP) criterion.

The rest of the paper is organized as follows. Section 2 provides background material on active sensing, and its applications to chemical sensors. Section 3 describes the proposed PWAS algorithm with a focus on the two objective functions, which are novel contributions of this paper. Section 4 provides a thorough evaluation of PWAS against FSS and random feature selection on a database of low-resolution absorption spectra containing 250 chemicals. Section 5 describes the experimental setup and results from validating the approach on an array of commercial MOX sensors. The article concludes with a discussion of results and directions for future work.

2. Background

The idea of active sensing originates from the theory of 'active perception' [9,10], which states that an organism actively probes the environment to enhance its ability to extract behaviorally relevant information. The concept caught on during the 1980s in the robotics and vision community [11], where it was used to denote control strategies that dynamically adapted sensing configurations as the sensor interacted with its environment. Since then, active-sensing principles have been used widely in vision, robotics and target tracking to address various computational problems such as classification, detection, estimation, sampling, and tracking. This prior work has shown that active sensing can manage sensing resources more efficiently than passive sensing, and can also provide a balance between sensing costs and sensing accuracy [12].

In a classic paper on active vision, Aloimonos et al. [13] showed that several computer vision problems that are ill-posed and non-linear with passive observers become well-posed and linear by use of an active observer (i.e., one that can control the parameters of its apparatus, such as focal length or orientation). Over the last two decades, active sensing has also been used for motion tracking [14], scene exploration and reconstruction [15], face recognition [16], vision-based localization and mapping [17], and scene segmentation [18].

Active-sensing strategies have also been broadly used in robotic navigation [19], localization [20,21], simultaneous localization and mapping [17], and robotic exploration [22]. A classical active-sensing problem in robotics is to decide where to move the robot (location decisions) and how to reconfigure its sensors (sensing decisions) [23]. These problems arise from the exploration-exploitation dilemma, which involves a trade-off between immediate rewards (exploitation) such as bringing the robot closer to its goal, and long-term effects (exploration) such as gathering information through landmarks, surrounding obstacles, or reading signs.

Along these lines, active sensing has also received attention for use in military scenarios, specifically for tracking dynamic targets with stationary [24] and mobile sensors [25–27]. The target-tracking problem involves estimating locations and velocities of multiple moving targets (e.g., ground vehicles) using surveillance sensors such as radars, sonars, or electro-optical sensors. One of the central challenges in target tracking is selecting the next sensing action; this involves choosing sensors, setting their configurations (such as pointing angles, dwell lengths, etc.), or possibly moving them to another location.

2.1. Prior work in active chemical sensing

Though not as broadly as in vision, robotics and target tracking, active-sensing principles have been applied to various chemical sensing problems as well, including odor generation, chemical discrimination, and data collection. To our knowledge, the earliest use of active sensing in the chemical/olfaction domain is the work of Nakamoto et al. [28,29] on odor generation. The objective of this work was to reproduce an odor blend by creating a mixture from its individual components. The authors developed an active-control algorithm that adjusted the mixture ratio so that the response of a gas sensor array to the mixture matched the response of the array to the odor blend.

Active sensing has also been used for chemical discrimination problems. As an example, Priebe et al. [3] developed a statistical pattern classification method termed Integrated Sensing and Processing Decision Trees (ISPDT). This method builds a decision tree to partition feature space hierarchically; nodes close to the root provide good clustering of examples regardless of class labels, whereas nodes at the leaves seek to discriminate examples from different classes. Each internal node defines a sensor configuration (a feature) and its children the possible observations. The decision tree is used to guide the sensing process as follows. First, the sensor is operated according to the feature at the root node. The resulting observation falls into one of the child nodes, which determines the next step: either acquire new measurements (if it is an internal node), or to classify the sample and terminate sensing (in case of a leaf node). The authors evaluated ISPDT on a dataset containing the response of an optical sensor array to trichloroethylene (a carcinogenic industrial solvent) in complex backgrounds; ISPDT reduced misclassification rates by 50%, while requiring only 20% of all the sensors to make any individual classification.

More recently, Lomasky et al. [30] developed an "active class selection" method to optimize the generation of training datasets for e-nose applications. Their approach was based on principles from active learning, a machine-learning technique where the learning algorithm chooses the training samples from which it learns. Active learning assumes that many training instances are readily available and that the cost lies in labeling them (e.g. through human annotation). However, in e-nose applications the costs are not associated with labeling existing samples but with the more laborious process of collecting new ones. Therefore, the active class selection problem involves choosing the class of the next training instance, whereas the active learning problem deals with choosing the next training instance to be labeled. Lomasky's approach consists of generating the next set of n training instances in proportion to the instability of class boundaries, measured in terms of the number of test instances whose classification labels change upon inclusion of the previous set of n training instances. The authors validated the approach on an experimental dataset from an array of fluorescent micro-bead sensors exposed to six organic chemicals and their mixtures. The results show that active class selection can minimize the number of new training instances needed to obtain the maximal classification performance.

An optical implementation of active-sensing principles was proposed by Dinakarababu et al. [4] for rapid identification of chemicals. In this work, the authors developed an Adaptive Feature Specific Spectrometer (AFSS), a digital micro-mirror device capable of multiplexing certain spectral bands and directing them onto a photo-detector. In this fashion, the system is able to measure the projection of the incoming spectral density onto a set of basis vectors, rather than measure the spectral density directly. The basis vectors are the eigenvectors of a probabilistically-weighted covariance matrix, with the probabilities corresponding to the likelihoods of different classes based on previous measurements.

Our early investigations of active sensing focused on the problem of discriminating M chemicals at fixed concentration with a single temperature-modulated metal-oxide sensor. In [6], we presented a partially observable Markov decision process (POMDP) solution to this problem, and proposed a myopic policy that selected sensing actions based on the expected reduction in Bayesian risk. In subsequent work [5], we reformulated the problem to not only identify chemical samples but also estimate their concentrations using Fabry–Perot interferometers. This new approach used nonnegative matrix factorization [31] to create concentration-independent absorption profiles of different chemicals, and linear least squares to fit sensor observations to the response profiles. In latter work, we extended the active-sensing method to estimate the concentration of mixtures with known components [32], and the more challenging problem of estimating concentrations of mixtures with unknown components [33].