Energy-Aware Active Chemical Sensing

Rakesh Gosangi and Ricardo Gutierrez-Osuna
Department of Computer Science and Engineering
Texas A&M University, College Station, Texas 77840
{rakesh,rgutier}@cse.tamu.edu

Abstract— **We propose an adaptive sensing framework for metal-oxide (MOX) sensors that seeks to minimize energy consumption through temperature modulation. Our approach generates temperature programs by means of an active-sensing strategy combined with an objective function that penalizes power consumption. The problem is modeled as a partially observable Markov decision process (POMDP) and solved with a myopic policy that operates in real time. The policy selects sensing actions (i.e., temperature pulses) that balance misclassification costs (e.g., chemicals identified as the wrong target) and sensing costs (i.e., power consumption). We experimentally validate the method on a ternary chemical discrimination problem, and compare it against a "passive classifier." Our results show that, for a given energy budget, the active-sensing strategy selects temperatures with more discriminatory information than those of the passive classifier by penalizing pulses of higher temperature and longer durations.**

I. INTRODUCTION

Tunable chemical sensors are usually cross-selective and are equipped with operating parameters that regulate their selectivity. Modulating these parameters can produce informative gas-specific response patterns. For example, the responses obtained by modulating the working temperature of MOX sensors during gas exposure contain more discriminatory information than their isothermal responses. However, this form of sensing (involving parameter modulation) is usually power-intensive, time consuming, and requires careful tuning of the parameters. This creates a need for management techniques that can adaptively select the operating parameters of chemical sensors such that (*i*) the discriminatory information obtained from the sensor responses is maximized while (*ii*) the sensing costs are minimized, where sensing costs may correspond to power consumption, response time, or any other form of measurable expenditure.

This paper presents a systematic approach to actively modulate the operating temperature of MOX sensors according to the two criteria mentioned above. MOX chemical sensors are robust, inexpensive, and highly sensitive. However, their lack of selectivity has been a significant factor in limiting their use in large-scale chemical sensor networks. To counter this problem of low selectivity, researchers have often used temperature modulation [1]. Traditional research on temperature modulation of MOX sensors has generally been empirical: researchers evaluate a number of temperature programs (e.g. sine wave, ramp) and select the programs that produce maximum discrimination between the target gases. However, these empirically-derived temperature programs are not necessarily power-efficient or time-aware since in many cases one does not need to capture the sensor's response to an entire program in order to identify a chemical. This creates a need for methods that generate optimized temperature programs for MOX sensors.

To our knowledge, there are two studies that have presented systematic approaches to optimizing temperature profiles [2, 3]. The objective of these two papers was to find optimal temperature programs (for micro-hotplate metal-oxide gas sensors) that maximized distance between the temperature-modulated sensor responses to target gases, but these papers did not consider the time or energy consumption of the sensors. Recently, we presented an active-sensing strategy to optimize the temperature program of MOX sensors in real time [4]. The main focus of this prior study was to develop a systematic approach that can optimize the temperature programs on-line, that is, as the sensor collects data from its environment. In this paper, we extend this active-sensing approach by considering a cost function that incorporates power awareness. We operate the MOX sensor with a variety of temperature pulses (sensing action), where each pulse involves operating the sensor at a fixed temperature for a fixed duration. Every pulse has a sensing cost that captures the power consumed by the sensor. We use a decision-theoretic formulation that identifies the best temperature pulse to be applied to the MOX sensor at a given point of time. This decision is based on three factors: (*i*) the information acquired from previous measurements (*ii*) the cost of misclassifying a chemical and (*iii*) the cost of applying new temperature pulses. At a high level, our method consists of two steps. First, we model the dynamic response of the sensor to a sequence of temperature pulses by means of an Input-Output Hidden Markov Model (IOHMM) [5]. Second, we cast the IOHMM into a Partially Observable Markov Decision Process (POMDP) [6] and use a greedy strategy to find a policy (i.e., a mapping from states to actions.)

The paper is organized as follows. In section II, we present a formal definition of the problem and describe other potential approaches to solve the problem. Section III summarizes the

algorithmic methods used in the paper. Namely, we describe how the IOHMMs are used to model the dynamic responses of the sensor and explain how the parameters of the IOHMMs can be used to formulate a POMDP. Section IV presents experimental results on a commercial gas sensor for a ternary classification problem and a comparison with a "passive" maximum likelihood classifier. We conclude the paper with a discussion and potential directions for future work.

II. Problem Definition

Consider the problem of classifying an unknown gas sample into one out of M known categories $\{\omega^{(1)}, \omega^{(2)}, \dots, \omega^{(M)}\}$ using a MOX sensor. The MOX sensor can be pulsed at T different operating temperatures $\{\rho_1, \rho_2, \dots, \rho_T\}$ with D different pulse durations $\{\xi_1, \xi_2, \dots, \xi_D\}$. Each sensing action consists of operating the sensor at a fixed temperature ρ_j for a fixed duration ξ_k. Therefore, there are $T*D$ different sensing actions $\{a_1 = (\rho_1, \xi_1), a_2 = (\rho_1, \xi_2), \dots, a_{T*D} = (\rho_T, \xi_D)\}$ associated with the MOX sensor. Each sensing action has a predefined sensing cost $C(a_i = \rho_j, \xi_k) = c_{jk}$ that reflects the power consumed by the MOX sensor. In addition to sensing actions, we have classification actions $\{\hat{\rho}_1, \hat{\rho}_2, \dots, \hat{\rho}_M\}$, where the action $\hat{\rho}_u$ corresponds to declaring the chemical to be of class u. The costs associated with each classification action is defined as $C(a = \hat{\rho}_u) = c_{uv}$, where c_{uv} is the cost of declaring a chemical to be of class u, if in reality it is of class v. The classification actions and the sensing actions have the same units of measurement. The goal is to identify a sequence of actions $\boldsymbol{a} = [a_1, a_2, \dots, a_S]$ such that we accumulate sufficient evidence to identify the gas sample while minimizing the cumulative cost of sensing and classification $\mathrm{E}(\boldsymbol{a})$:

$$\mathrm{E}(\boldsymbol{a}) = \sum_{i=1}^{S} C(a_i) \tag{1}$$

The sensing costs of each action are assumed to be proportional to the product of the operating temperature and the pulse duration:

$$C(a_i = \rho_j, \xi_k) = c_{jk} = E_c * \rho_j * \xi_k \tag{2}$$

and E_c is a normalization constant.

We can solve this problem using standard techniques like feature subset selection (FSS). FSS involves evaluating the performance of a large number of feature subsets in terms of a scoring metric (such as correlation, mutual information, or inter-class distance) *and* the total sensing cost (1). However, there are two drawbacks to using this form of "passive" feature selection: (*i*) it is computationally intensive and (*ii*) it cannot handle sequential interaction between features. Our aim is to determine an optimal sequence of actions $\boldsymbol{a} = [a_1, a_2, \dots, a_S]$ in an adaptive fashion based on sensing cost constraints and evidence accumulated from previous measurements.

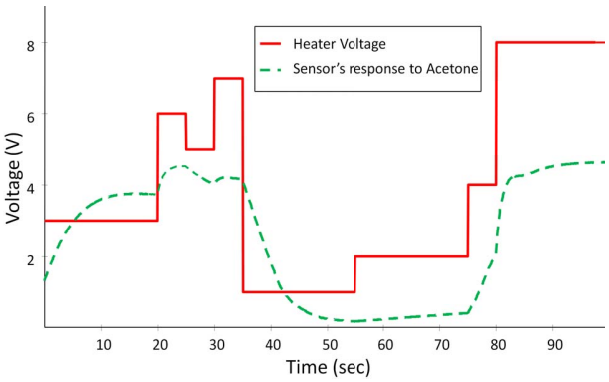

Figure 1. Response of the MOX sensor to temperature modulation. The sensor was driven for 100 seconds with pulses of varying amplitude and durations in the presence of acetone.

III. Methods

A. Sensor Modelling

We model the steady state response of the MOX sensor to a pulse a_i (of temperature ρ_j and duration ξ_k) exposed to a chemical $\omega^{(c)}$ with a Gaussian mixture:

$$p(x_i|\omega^{(c)}) = \sum_{m_i=1}^{M_i} \alpha_{i,m_i}^{(c)} N\left(x_i \middle| \mu_{i,m_i}^{(c)}, \Sigma_{i,m_i}^{(c)}\right) \tag{3}$$

where x_i is the sensor's response, M_i is the number of Gaussians, and $\alpha_{i,m_i}^{(c)}, \mu_{i,m_i}^{(c)}, \Sigma_{i,m_i}^{(c)}$ are the mixing coefficient, mean vector and covariance matrix of each Gaussian for class $\omega^{(c)}$, respectively. Given a sequence of actions $\boldsymbol{a} = [a_1, a_2, \dots, a_S]$, we assume that the sensor progresses through a series of states $\boldsymbol{s} = [s_1, s_2, \dots, s_S]$ to produce an observation sequence $\boldsymbol{o} = [o_1, o_2, \dots, o_S]$. Each state s_i represents a Gaussian in (3) and is therefore hidden. We model the sensor dynamics with an input-output hidden Markov model (IOHMM) [7], a generalization of the traditional HMM. In an HMM the future state of the system is assumed independent of its history, and dependent only on its current state. In an IOHMM the future state not only depends on the current state but also on an external input to the system. In our case, the additional input is a sensing action (i.e. temperature pulse), and the output is the sensor conductance.

Formally, an IOHMM can be defined as a 6-tuple $\{S, A, O, \pi, \tau, \phi\}$ where S is a finite set of states, each state corresponding to a Gaussian in (3); A is a finite set of discrete actions, each action corresponding to selecting a combination of one of T sensor temperatures and one of the D pulse durations; O is a set of observations, each corresponding to the sensor's response for a given sensing action; $\pi(s)$ is the initial state distribution, which captures any prior knowledge; $\tau(s'|s, a)$ is the state transition function, which describes the probability of a transition from state s to state s' given a sensing action a (i.e., the sensor dynamics); and $\phi(o|s)$ is the observation function, which describes the probability of making observation o at state s (i.e., the sensor response at a given state).

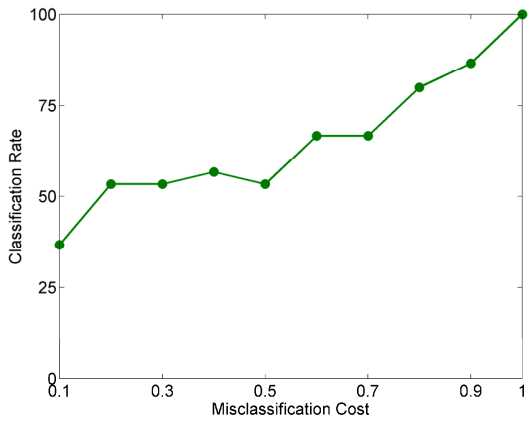

Figure 2. Classification rate as a function of the misclassification cost.

B. Classification using IOHMMs

We train a separate IOHMM for each individual chemical class using the Baum-Welch method. The training data is obtained by driving the sensor with a random sequence of pulses in the presence of the chemical and recording the corresponding responses. This is repeated multiple times to capture the variance in the sensor response due to history effects. Once trained, we can use the IOHMMs as classifiers. Namely, given an action sequence $[a_1 \dots a_S]$ and the corresponding observation sequence $[o_1 \dots o_S]$, we compute the log-likelihood that each trained IOHMM gives to the action-observation sequence. This log-likelihood is calculated using the Forward-Backward algorithm [8]. If the i^{th} IOHMM is the most likely, then we declare the sequence belongs to class *i*.

C. Eneryg aware sensing as a POMDP

Though we can use IOHMMs to classify chemicals, our goal is to determine a sequence of actions $\boldsymbol{a} = [a_1 \dots a_S]$ that produces an observation sequence $[o_1 \dots o_S]$ that provides sufficient information to identify the chemical, while simultaneously minimizing the total cost $\mathrm{E}(\boldsymbol{a})$ (1). Therefore, we formulate the problem (section II) as a POMDP. POMPDs have been proven very successful in a variety of sequential planning applications such as robot navigation, machine maintenance, machine vision, or moving target search. A POMDP provides a framework for an agent decision process. The agent can access the observations and observational probabilities but cannot observe the underlying state, and therefore must maintain a probability distribution over all possible states known as the *belief state*. A POMDP is formally defined as a 7-tuple $\{S, A, O, b_0, T, \Omega, C\}$, where *S*, *A*, and *O* are the finite set of states, sensing actions and observations respectively. $b_0(s)$ is an initial belief state, $T(s'|s,a)$ is the probability of a transition from state *s* to state *s'* given action *a*, $\Omega(o|s)$ is the probability of making observation *o* at state *s*, and $C(s,a)$ is the cost of executing action *a* at state *s*.

Most of the POMDP parameters can be obtained directly from the IOHMMs as follows. First, define *S* and *O* as the union over all states and observations corresponding to the IOHMMs, and *A* as the union of two sets of actions: sensing actions and classification actions. Sensing actions are obtained from the IOHMMs. The classification actions ($a = \hat{\rho}_c$) are terminal actions which assign a sample to a particular chemical class (section II). Except for the cost function, the remaining POMDP parameters are obtained from the IOHMMs as follows:

- Initial belief: $b_0(s) = p(\omega^{(c)})\pi^{(c)}(s)$, for $s \in S^{(c)}$
- State transition: $T(s'|s,a) = \tau^{(u)}(s'|s,a)$ for $s, s' \in S^{(u)}$; zero otherwise.
- Observation function: $\Omega(o|s) = \phi^{(c)}(o|s)$, for $s \in S^{(c)}$

The cost function *C* assigns a specific cost to each action. The cost of sensing actions is defined in (2). The cost of classification actions $C(a = \hat{\rho}_u) = c_{uv}$ represents a misclassification penalty whenever $u \neq v$. We set $c_{uv} = 0$ for $u = v$, since this represents a correct classification.

After performing a sequence of actions $[a_1 \dots a_T]$ and making a sequence of observations $[o_1 \dots o_T]$, the POMPD must update its belief state $b_T(s)$. Given the initial belief $b_0(s)$ and the history of actions and observations the belief state can be computed as:

$$b_T(s) = p(s|o_1 \dots o_T, a_1 \dots a_T, b_0) = p(s|o_T, a_T, b_{T-1}) \quad (4)$$

Since $b_T(s)$ is a *sufficient statistic*, it can be updated incrementally from its previous estimate $b_{T-1}(s)$ by incorporating the latest action a_T and observation o_T:

$$b_T(s') = \frac{p(o_T|s', a_T)\sum_s p(s'|a_T, s) b_{T-1}(s)}{p(o_T|a_T, b_{T-1})} \quad (5)$$

where the denominator $p(o_T|a_T, b_{T-1})$ can be treated as a normalization term to ensure that $b_T(s')$ sums up to 1. Using this POMDP formulation, the energy-aware sensing problem becomes one of finding a function that maps belief states into actions; this function is referred to as a *policy*.

D. Finding a Sensing Policy

We employ a myopic policy that only takes sensing action if the cost of sensing is lower than the expected future reduction in Bayes risk. Given belief state $b_T(s)$, the expected risk of a classification action is:

$$R_C(b_T(s)) = \min_u \sum_v c_{uv} \sum_{s \in S^{(v)}} b_T(s) \quad (6)$$

where *u* represents the class with minimum Bayes risk, and $\sum_{s \in S^{(v)}} b_T(s)$ is the belief across all states that belong to chemical *v*. In turn, the expected risk of a sensing action is:

$$R_S(b_T(s), a) = \sum_{\forall o} \min_u \left(\sum_v c_{uv} \sum_{s' \in S^{(v)}} \sum_s p(o|s', a) p(s'|s, a)\, b_T(s) \right) \quad (7)$$

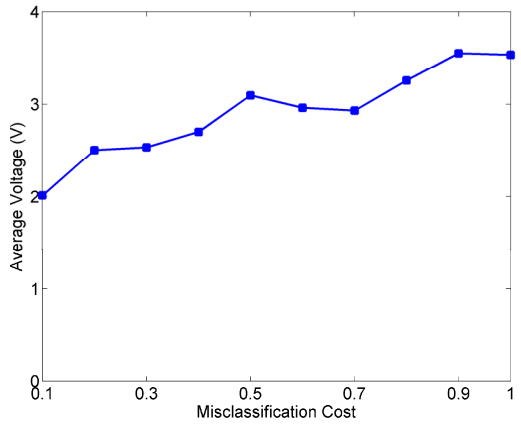

Figure 3. Average heater voltage as a function of the misclassification cost.

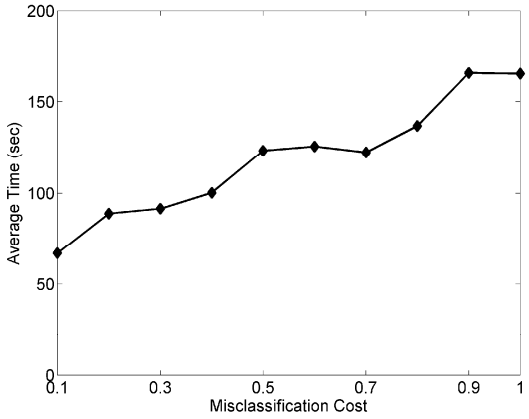

Figure 4. Average classification times as a function of the misclassification cost.

which averages the minimum Bayes risk over all observations that may result from the sensing action. Hence, the utility of sensing action a can be computed as:

$$U(b_T(s), a) = \left[R_C\left(b_T(s)\right) - R_S(b_T(s), a)\right] - C(a) \qquad (8)$$

If $U(b_T(s), a)$ is negative for all sensing actions, then the cost of sensing exceeds the expected reduction in risk $[R_C(\cdot) - \mathcal{RS}\cdot$, and a classification action is taken. Otherwise, the action with maximum utility is taken. The classification decision is based on (6).

IV. Experiments and Results

The energy-aware method was validated on two types of experiments. The first set of experiments characterized the approach in terms of its classification performance, energy consumption, and duration (time) as a function of misclassification costs. The second set of experiments provides a comparison with a passive classification approach. We used a commercial MOX sensor (Figaro TGS 2600) exposed to three different analytes (acetone, ammonia, and isopropyl alcohol). As an approximation, we used the heater voltage V_H as a proxy for the sensor's operating temperature. We operated the sensor at eight different heater voltages ranging from $V_H = 1V$ to $V_H = 8V$. We used three pulse durations $\xi = 5{,}10{,}20\ sec$, where $\xi = 5$ is the sensor's time constant[1]. Therefore, there were 24(8*3) different sensing actions, where each action corresponds to applying a fixed heater voltage for a fixed duration. For the experiments reported here, the information from each transient response of the sensor was reduced to a single observation by computing its integral response. Typically, the sensor's response to a pulse of longer duration contains more information than that of a shorter duration because when the sensor is driven with a longer pulse it has sufficient time to reach its steady state. On the other hand, when driven with shorter pulses the information mostly lies in transient response (transient response is also obtained with a longer pulse). Thus, longer pulses are informative but costlier, whereas shorter pulses are inexpensive but generally less informative.

A short description of our experimental setup follows. The sensor was introduced in 30ml glass vials through a tight aperture on the cap and allowed to equilibrate with the static headspace of 10ml of the analyte. The sensor was interfaced through a National Instruments data acquisition card (USB 6009) controlled using Matlab. Each analyte was serially diluted (1/3 dilution ratio) to determine a set of dilutions at which the sensor provided similar isothermal responses for all the analytes; this ensured that the analytes could not be discriminated without temperature modulation. To obtain the isothermal response, the sensor was exposed to each analyte for 100 seconds under a constant heater voltage of 5V. The final dilution factors were 1/9 for ammonia (two serial dilution), 1/729 for acetone (six serial dilutions), and 1/2187 for isopropyl alcohol (seven serial dilutions).

We trained a separate IOHMM for each analyte. We collected sensor responses to 45 random pulse sequences, 15 per analyte. Each pulse sequence was a random permutation of the 24 sensing actions. Fig. 1 shows the sensor's response to a random sequence of pulses in the presence of acetone.

A. Experiment 1: Classification performance

To test the POMDP, we varied misclassification costs from $c_{uv} = 0.1$ to $c_{uv} = 1.0$ (if $u \neq v$) in steps of 0.1; correct classification costs were assumed zero($c_{uv} = 0; u = v$), and sensing costs $C(a_i)$ were assigned as described in (2). The normalization term $E_c = 0.001$ was used to ensure that sensing costs and risk values (as estimated in (6) and (7)) had similar range of values. If the value of E_c is set very high, the utility of the actions would be negative (8), which would prevent the algorithm from taking any additional measurements. On the other hand, if the value of E_c is set too low, the myopic approach would tend to collect the sensor's response to all the available pulses. The value of E_c used in these experiments was determined empirically. We initialized each belief state $b_0(s)$ with a constant value. We tested the

[1] We obtained an estimate of the sensor's time constant $\xi = 5sec$ as the average time required by the sensor to reach 90% value of the steady state. Therefore, the three pulse durations correspond to one, two and four time constants.

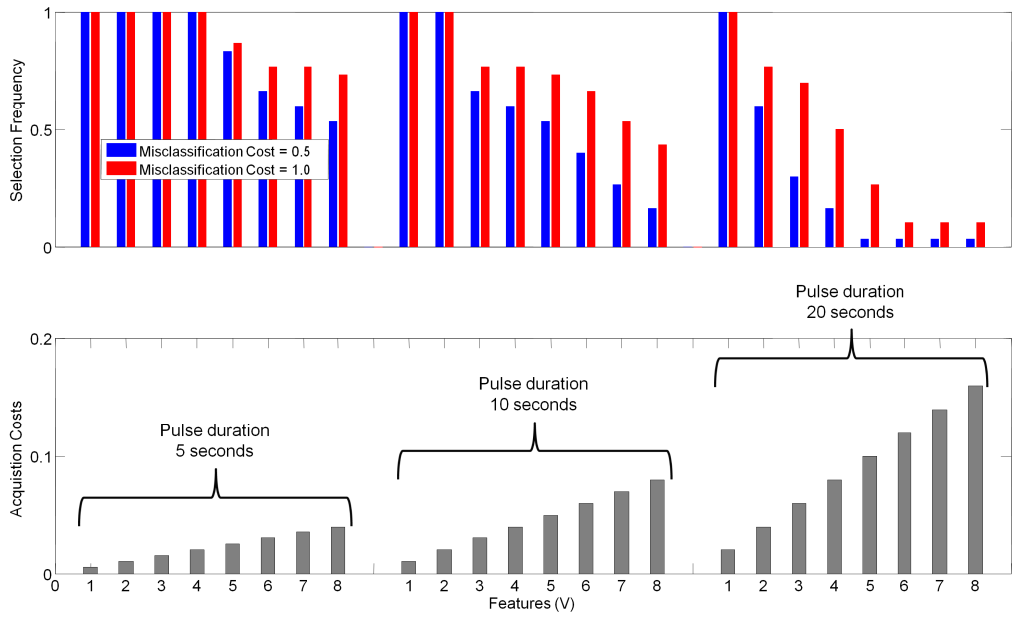

Figure 5. (top) Frequency of selecting each feature (across classes) for two misclassification cost $\boldsymbol{c_{uv}} = \{\mathbf{0.5}, \mathbf{1.0}\}$. (bottom) Sensing costs associated with each action as a function of heater voltage and pulse duration.

POMDP 30 times (10 times per chemical) at each misclassification cost setting, resulting in 300 test cases.

Results are summarized in Fig. 2, Fig. 3, and Fig. 4 in terms of classification performance, average voltage, and average response time as a function of misclassification costs, respectively. Shown in Fig. 2, as misclassification costs increase, the POMPD tends to select more pulses in order to improve classification performance. Likewise, the average heater voltage also increases with increasing misclassification costs (Fig. 3). The average heater voltage is estimated as the weighted average of all the pulses selected by the POMDP, the weights being the duration of the pulses. This suggests that the POMPD tends to use pulses of higher amplitude with increasing misclassification costs. Similarly, the average time (estimated as the average duration of all pulses used) increases with misclassification costs (Fig. 4), suggesting that the POMDP selects pulses of longer duration to ensure high classification performance. A classification rate of 100% is obtained at $c_{uv} = 1.0$.

Fig. 5 shows the frequency at which the POMDP selects each feature (across all the chemicals) for misclassification costs $c_{uv} = 0.5$ and 1.0. These results show how the POMDP balances the cost of applying additional pulses against the risk of misclassifying the chemical analyte. As an example, the POMDP tends to select the costlier pulses more often at $c_{uv} = 1.0$ than at $c_{uv} = 0.5$.

B. Experiment 2: Active vs. Random feature selection

We also compared our approach with a passive classification strategy. The idea behind this experiment was to fix a sensing cost budget c_b and compare the performance of active and random feature selection techniques. Both the POMDP and the passive classifier are based on a maximum likelihood criterion; their differences lay in their respective feature selection strategies: the POMDP actively selects the features based on costs and previous measurements, whereas the passive classifier selects features randomly. As described in section III.B we train a separate IOHMM per each class. For the passive classifier, we select random sequences of sensing actions (pulses) that sum up to the fixed cost threshold c_b. We drive the sensor with this random sequence, record the corresponding observation sequence, and estimate the class label based on the log-likelihood calculation described in section III.B. To ensure a fair comparison between POMDP and the passive classifier, we changed the stopping criterion of the myopic policy: it takes a classification (terminal) action whenever (*i*) the utility of all the sensing actions is negative *or* (*ii*) the total sensing cost has exceeded the predefined threshold c_b.

We varied the cost threshold c_b between 0 and 0.8 in steps of 0.2. At each threshold setting, we tested both methods 30 times (10 times per each chemical). Fig. 6 shows the classification performance of the two methods as a function of the threshold c_b. As expected, both classifiers perform better with increasing thresholds because they can obtain more observations without exceeding the threshold. However, the POMDP consistently outperforms its passive counterpart. This could be attributed to the active feature selection strategy used by the POMDP. For a given sensing budget, the POMDP selects pulses that tend to have more discriminatory information than a set of pulses randomly selected by the passive classifier.

V. CONCLUSION

We have presented an active sensing framework for metal-oxide (MOX) sensors that minimizes energy consumption through temperature modulation. Our approach generates temperature programs (sequences of temperature pulses) by incorporating a cost function that penalizes power

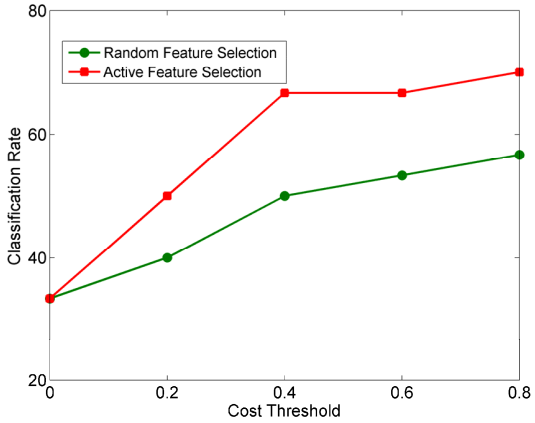

Figure 6. Classification performance of the active classifier (POMDP) compared with a classifier that uses a random feature selection policy.

consumption. The method was validated experimentally on a single tunable MOX sensor, but the approach could be adapted to other tunable chemical sensors. Also, we can extend this approach to manage multiple tunable chemical sensors, where a centralized algorithm would sequentially select sensors and their operating parameters in an adaptive fashion.

There is a great potential for deploying tunable chemical sensors in large-scale sensor networks. However, a number of chemical detectors (e.g., MOX sensors, infrared sensors, preconcentrators) are power intensive, which limits their use in distributed wireless networks. Current algorithms for wireless sensor networks assume that data transmission is the main source of energy expenditures. Although this may be true for the measurement of physical parameters (e.g., pressure, temperature, light, and vibration), a number of chemical detectors are power intensive and may require different assumptions about the relative energy budgets of sensing versus data transmission [9]. This creates a need for sensor management techniques that can help manage sensing costs by selecting operating parameters judiciously.

REFERENCES

[1] A.P. Lee and B.J. Reedy, "Temperature modulation in semiconductor gas sensing," Sens. Actuators, B, vol. 60, no. 1, 1999, pp. 35-42.

[2] A. Vergara, et al., "Optimized temperature modulation of micro-hotplate gas sensors through pseudorandom binary sequences," IEEE Sens. J., vol. 5, no. 6, 2005, pp. 1369-1378.

[3] T.A. Kunt, T.J. McAvoy, R.E. Cavicchi and S. Semancik, "Optimization of temperature programmed sensing for gas identification using micro-hotplate sensors," Sensors and Actuators B: Chemical, vol. 53, no. 1-2, 1998, pp. 24-43.

[4] R. Gosangi and R. Gutierrez-Osuna, "Active Temperature Programming for Metal-Oxide Chemoresistors," Sensors Journal, IEEE, vol. 10, no. 6, 2010, pp. 1075-1082.

[5] Y. Bengio and P. Frasconi, "An input output HMM architecture," Advances in Neural Information Processing Systems, 1995, pp. 427-434.

[6] L.P. Kaelbling, M.L. Littman and A.R. Cassandra, "Planning and acting in partially observable stochastic domains," Artificial Intelligence, vol. 101, no. 1-2, 1998, pp. 99-134.

[7] S. Ji and L. Carin, "Cost-sensitive feature acquisition and classification," Pattern Recognit., vol. 40, no. 5, 2007, pp. 1474-1485.

[8] L. Rabiner, "A tutorial on hidden Markov models and selected applications inspeech recognition," Proceedings of the IEEE, vol. 77, no. 2, 1989, pp. 257-286.

[9] R. Gutierrez-Osuna and A. Hierlemann, "Adaptive Microsensor Systems," Annual Review of Analytical Chemistry, vol. 3, no. 1, 2010, pp. 255-276.