

Odor recognition in robotics applications by discriminative time-series modeling

Frank-Michael Schleif · Barbara Hammer · Javier Gonzalez Monroy ·
Javier Gonzalez Jimenez · Jose-Luis Blanco-Claraco ·
Michael Biehl · Nicolai Petkov

Received: 15 March 2014 / Accepted: 21 December 2014 / Published online: 8 January 2015
© Springer-Verlag London 2015

Abstract Odor classification by a robot equipped with an electronic nose (e-nose) is a challenging task for pattern recognition since volatiles have to be classified quickly and reliably even in the case of short measurement sequences, gathered under operation in the field. Signals obtained in these circumstances are characterized by a high-dimensionality, which limits the use of classical classification techniques based on unsupervised and semi-supervised settings, and where predictive variables can be only identified using wrapper or post-processing techniques. In this paper, we consider generative topographic mapping through time (GTM-TT) as an unsupervised model for time-series inspection, based on hidden Markov models regularized by topographic constraints. We further extend the model such that supervised classification and relevance learning can be integrated, resulting in supervised GTM-TT. Then, we evaluate the suitability of this new technique for the odor classification problem in robotics applications. The performance is compared with classical techniques as nearest neighbor, as an absolute baseline, support vector

machine and a recent time-series kernel approach, demonstrating the eligibility of our approach for high-dimensional data. Additionally, we exploit the learning system introduced in this work, providing a measure of the relevance of each sensor and individual time points in the classification process, from which important information can be extracted.

Keywords Electronic nose · Volatile classification · Odor recognition · time-series · Prototype learning · Relevance learning

1 Introduction

Olfaction plays an important role in the development of many applications, such as quality control in food processing chains, detection and diagnosis in medicine, finding drugs and explosives, and the more common estimation of blood alcohol content (BAC) for drivers. Among them,

F.-M. Schleif (✉)
School of Computer Science, University of Birmingham,
Birmingham B15 2TT, UK
e-mail: fschleif@techfak.uni-bielefeld.de;
schleify@cs.bham.ac.uk

B. Hammer
Center of Excellence, University of Bielefeld, Universitätsstrasse
21-23, 33615 Bielefeld, Germany
e-mail: bhammer@techfak.uni-bielefeld.de

J. G. Monroy · J. G. Jimenez
Dpto. Ingenieria de Sistemas y Automatica E.T.S.I. Informatica,
Telecomunicacion, Universidad de Malaga, Campus
Universitario de Teatinos, 29071 Malaga, Spain
e-mail: jgmonroy@uma.es

J. G. Jimenez
e-mail: javiergonzalez@uma.es

J.-L. Blanco-Claraco
Universidad de Almeria, Edificio CITE II-A, 04120 La Caada,
Almeria, Spain
e-mail: joseluisblancoc@gmail.com

M. Biehl · N. Petkov
Johann Bernoulli Institute for Mathematics and Computer
Science, University of Groningen, P.O. Box 407,
9700 AK Groningen, The Netherlands
e-mail: m.biehl@rug.nl

N. Petkov
e-mail: N.petkov@rug.nl

there are some applications like pollution monitoring or leak detection that require to measure the environment continuously and at different locations. For such scenarios, the use of a mobile robot with the capability of identifying and measuring the volatiles' concentration is of great help as already reported in [25, 26]. Furthermore, olfaction also plays a key role in the development of more intelligent and useful robots at home, for example, by recognizing activities and environmental conditions, or improving social interaction [16].

Three are the main fields within robotics olfaction: gas distribution mapping (GDM) [4, 22], where the objective is to obtain a truthful representation of how volatiles are dispersed in the inspected area and their respective concentrations, gas source localization (GSL) where the robot is commanded to localize the emission sources [13], and odor recognition which deals with the problem of identifying which set of categories a new volatile sample belongs to [45].

The discrimination of gases performed with a robot equipped with an array of gas sensors presents a number of additional challenges when compared to standard analyte identification applications, mostly due to the differences in the measurement conditions. While standard classification tasks usually host gas sensors inside a chamber with controlled humidity, temperature and airflow conditions, in robotics olfaction there is no control over the sensing conditions. This entails that the sensor signals to be processed are noisy and dominated by the signal transient behavior.

Only few modeling methods are available to obtain interpretable, compact and precise predictive models for such type of data like [7, 23]. This is mainly due to the following reasons: (1) the number of time points is often low, while the dimensionality of the data is rather high, (2) the number of time sequences is often low, leading to a sparsely populated data space and (3) the sequences may have missing values, and may be of different length.

In this paper, we demonstrate the suitability of a novel approach based on generative topographic mapping through time (GTM-TT) to the problem of volatile identification in robotics. The model extends classical GTM-TT by integrating supervised classification and relevance learning, resulting in supervised GTM-TT (SGTM-TT). More precisely, we have tested the SGTM-TT method with an e-nose comprising an array of MOX (metal oxide sensors) to classify samples of seven different volatiles under uncontrolled conditions. The performance is compared with techniques as nearest neighbor (NN), support vector machine (SVM) and a reservoir computing time-series kernel (RTK). We illustrate one of the main advantages of the proposed method when classifying odors based on short data sequences, providing the predictive classification

accuracy for sequences of reduced lengths (1, 10 and 20 s). Furthermore, we highlight the introduced relevance learning system for temporal high-dimensional data, by studying the relevance of sensors and time points on the classification performance.

2 Related works

Odor discrimination with electronic noses has received growing attention and many studies have been done on how to classify odors using an array of gas sensors and a pattern recognition algorithm. In [8, 11, 18, 40] the principal methods for chemical classification with an array of gas sensors are reviewed, including NN, mahalanobis linear discriminant analysis, neural networks (ANN), cluster analysis with self-organizing maps (SOM) and SVM.

More recently, approaches based on ensembles of classifiers have been reported to improve the classification accuracy [44], to improve the earliness on the classification [20], or to deal with the common problems of sensor drift and sensor replacement. In [47], a SVM based ensemble of classifiers is used to solve the gas discrimination problem over a period of 3 years by training different classifiers at different points of time. Similarly, in [27], a flexible classification strategy based on cooperative classifiers is proposed to increase the robustness of chemo-sensory systems against failures in their constituent sensing elements, postponing the necessity of replacing a sensor in the array, as well as facilitating the insertion of newly sensing elements.

Nevertheless, little attention has been given to the problem of classification in uncontrolled conditions, as revealed by the few works found in the literature that perform classification focusing only on the transient phase of the sensor signals. An evaluation for the suitability of different feature extraction techniques for such scenarios is provided in [45], where Trincavelli et al. propose a preprocessing stage to isolate the relevant parts of the sensor signals that can then be passed to the pattern recognition algorithm. More recently, in [12] a SVM is applied to a set of features obtained from changes of the spectral sensor signal characteristics (frequency components, phase shift and energy sums), reporting a substantial increase of the classification performance.

Gas sensor data have been analyzed by many different machine learning techniques with typically substantial preprocessing steps, limiting an out of sample extension, as discussed in more detail later on. Recent work [5] regarding the classification of gas sensor data is based on density estimates or models of the time-series using decision trees [10].

Time series processing constitutes an advanced field of research with many existing powerful statistical analysis tools (see for example [41]). However, their methods usually require a sufficient length of the time-series as compared to their dimensionality or consider only one-dimensional time-series. Further the focus is often on modeling a timeseries, by means of a longer sequence to explore trends and predict future measurement values. In this work we are interested on discriminative models between different groups of time-series and we would like to predict the class of the timeseries.

A few machine learning techniques exist to investigate high-dimensional time-series: topographic mappings such as the self-organizing map (SOM) (see [1] for a recent review) were extended by a recursive context which accounts for the temporal dynamics [43]. A probabilistic counterpart is provided by the GTM-TT which combines hidden Markov models with a constraint mixture model induced by a low-dimensional latent space. This approach is extended to better take the relevance of the feature components into account in [31], but relying on an unsupervised model. The identification of relevant dimensions is very important as outlined, e.g. in [23, 31] to obtain a better understanding of the data, to reduce the processing complexity, and to improve the overall prediction accuracy. A supervised relevance weighting scheme which singles out relevant features in a wrapper approach based on hidden markov models has been proposed in [23]. In [7], a similar approach introducing class-wise constraints in the hidden Markov model is presented. In [23], applications to life science data are presented resulting in 85 % prediction accuracy on a multiple sclerosis (MS) data set, but the approach makes multiple, restrictive assumptions regarding the used hidden Markov model (HMM). The approach [7] is evaluated in the same scenario with improved performance for the sclerosis data set. Ongoing work in the field reflects the high demand for effective methods for short but high-dimensional time-series data [33]. This is not limited to the bio-medical domain [7, 23] but covers a broader field of applications in industry and geo-science [31, 43]. In this work, we employ a supervised variant of GTM-TT (SGTM-TT) as introduced in [36] and extended in [37].

3 Method

3.1 Generative topographic mapping

As outlined before the complexity of the considered data requests for a strong regularizing and interpretable model. Topographic maps appear to be a good choice and especially the generative topographic mapping (GTM)

combines multiple necessary features. GTM was first introduced in [2] and models a given set of data vectors $\mathbf{x} \in \mathbb{R}^D$ in form of a mapping based on a constrained mixture of Gaussians. The mixture is induced by a lattice of points \mathbf{w} in a low-dimensional, so called, latent space which can also be used for visualization. The low-dimensional lattice points are mapped by a projection $\mathbf{w} \mapsto \mathbf{t} = y(\mathbf{w}, \mathbf{W})$ into the high-dimensional data space. The corresponding mapping function is parametrized by the parameters \mathbf{W} , which usually are chosen in form of a generalized linear regression

$$y : \mathbf{w} \mapsto \Phi(\mathbf{w}) \cdot \mathbf{W} \tag{1}$$

with basis functions Φ as equally spaced Gaussians. The high-dimensional points $y(\mathbf{w}, \mathbf{W})$ are called prototypes and are determined in the original data space. The prototypes define a quantization of the original data space, representing the data with minimum possible error and can be inspected directly. For more recent work on prototype-based learning and topographic maps see [1].

Every grid point of the GTM induces a Gaussian

$$p(\mathbf{x}|\mathbf{w}, \mathbf{W}, \beta) = \left(\frac{\beta}{2\pi}\right)^{D/2} \exp\left(-\frac{\beta}{2}\|\mathbf{x} - y(\mathbf{w}, \mathbf{W})\|^2\right) \tag{2}$$

with variance β^{-1} . Assuming a Dirac distribution of the prototypes, the data are modeled by a mixture of K modes

$$p(\mathbf{x}|\mathbf{W}, \beta) = \sum_{k=1}^K p(\mathbf{w}^k) p(\mathbf{x}|\mathbf{w}^k, \mathbf{W}, \beta) \tag{3}$$

with $p(\mathbf{w}^k) = 1/K$, assuming equal probabilities of the modes. We optimize the data log-likelihood

$$\ln\left(\prod_{n=1}^N \left(\sum_{k=1}^K p(\mathbf{w}^k) p(\mathbf{x}^n|\mathbf{w}^k, \mathbf{W}, \beta)\right)\right) \tag{4}$$

by means of an expectation maximization (EM) strategy with respect to the model parameters \mathbf{W} and β with data dimensionality D and number of points N as detailed in [2]. Finally an unsupervised restricted Gaussian mixture model (GMM), induced by a low-dimensional latent space, is defined.

3.2 GTM through-time

For temporal data the original GTM formulation is limited because it does not account for the dependency between different time points leading to quite complex and redundant GTM models (see [31]). An extension was provided by the GTM through time (GTM-TT) [2] where the entries over time are no longer independent. It basically provides an advanced time-series clustering using a constrained hidden Markov model, which is useful under our given

constraints. It is assumed that the data are time-series in the D -dimensional metric space, i.e. $\mathbf{x} = \mathbf{x}(1) \dots \mathbf{x}(T) \in (\mathbb{R}^D)^*$ where, $T \geq 1$ is the length of the time-series. A data point of the training data will be referred to as \mathbf{x}^i . We assume that entries, consecutive in time, $\mathbf{x}(t)$ and $\mathbf{x}(t + 1)$ are strongly correlated. In the GTM-TT the observation space (over time) is represented by a topographic mapping as described before but, its time dependence is modeled in the form of a hidden Markov model (HMM). In the GTM-TT model the hidden states are given by the lattice points \mathbf{w}^j . The concept of the GTM-TT is depicted in Fig. 1. Let us assume a given sequence \mathbf{x} of observations and an underlying sequence of hidden states of the same length $\mathbf{z} = \mathbf{z}(1) \dots \mathbf{z}(T)$ where, $\mathbf{z}(i)$ is equivalent to a point \mathbf{w}^j . Then, the probability of the observations and a corresponding path of hidden states \mathbf{z} can be described by $p(\mathbf{x}, \mathbf{z}|\Theta) =$

$$p(\mathbf{z}(1)) \prod_{t=2}^T p(\mathbf{z}(t)|\mathbf{z}(t-1), \mathbf{W}, \beta) \prod_{t=1}^T p(\mathbf{x}(t)|\mathbf{z}(t)) \quad (5)$$

with the conditional probability $p(\mathbf{x}(t)|\mathbf{z}(t)) := p(\mathbf{x}(t)|\mathbf{z}(t), \mathbf{W}, \beta)$ is as before (2) [2]. This results in the overall probability of \mathbf{x} : $p(\mathbf{x}|\Theta) = \sum_{\mathbf{z} \in \{\mathbf{w}^1, \dots, \mathbf{w}^k\}^T} p(\mathbf{x}, \mathbf{z}|\Theta)$

For the parametrization of the GTM-TT ($\Theta = (\mathbf{W}, \beta, \pi, \mathbf{P})$) we rely on the assumption of the standard Markov property and stationarity of the dynamics. With additional parameters for the initial state probabilities $\pi = (\pi_j)_{j=1}^K$ where $\pi_j = p(\mathbf{z}(1) = \mathbf{w}^j)$ and transition probabilities $\mathbf{P} = (p_{ij})_{i,j=1}^K$ where $p_{ij} = p(\mathbf{z}(t) = \mathbf{w}^j | \mathbf{z}(t-1) = \mathbf{w}^i)$

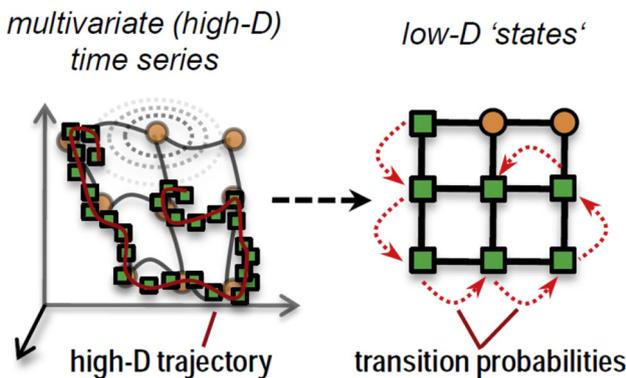


Fig. 1 GTM-TT consisting of a hidden Markov model, which hidden states are constrained to be organized on a grid topology (the latent points of the GTM model). The emission probabilities are governed by the GTM mixture distribution [2]. In the left figure a data distribution is given in a 3D space with an intrinsic low-dimensional support. Additionally, these data are not i.i.d. but dependent over time leading to some trajectory. GTM is used to project the data to a low-dimensional grid (here 2D, right plot). The prototypes (circles left) are generated by the latent points (in 2D, right) as HMM constrained Gaussians (left, dotted circles). Here we consider nine hidden states organized on a 3×3 grid. The data distribution may change over time and hence also the mapping of the GTM is effected over time, assuring smooth transitions within the HMM

\mathbf{w}^i), the latter one characterizing the temporal correlations of subsequent states. The data log likelihood is optimized by: $\ln(\prod_{n=1}^N p(\mathbf{x}^n|\Theta))$, using an EM approach. Like standard HMMs the hidden parameters (responsibilities) are defined by a forward–backward procedure [48]. Based on these parameters \mathbf{W} and β can be determined as specified before. The probability of being in state \mathbf{w}^k at time t , given the observation sequence \mathbf{x}^n (responsibilities) is given as:

$$r^{kn}(t) = p(\mathbf{z}(t) = \mathbf{w}^k | \mathbf{x}^n, \Theta) = \frac{A_{kt} B_{kt}}{p(\mathbf{x}^n | \Theta)} \quad (6)$$

Using the joint probability $p(\mathbf{x}^n(1) \dots \mathbf{x}^n(t), \mathbf{z}(t) = \mathbf{w}^k | \Theta)$ and the subsequent equation:

$$A_{kt} = \sum_{i=1}^K A_{it-1} p_{ik} p(\mathbf{x}^n(t) | \mathbf{w}^k, \Theta) \quad (7)$$

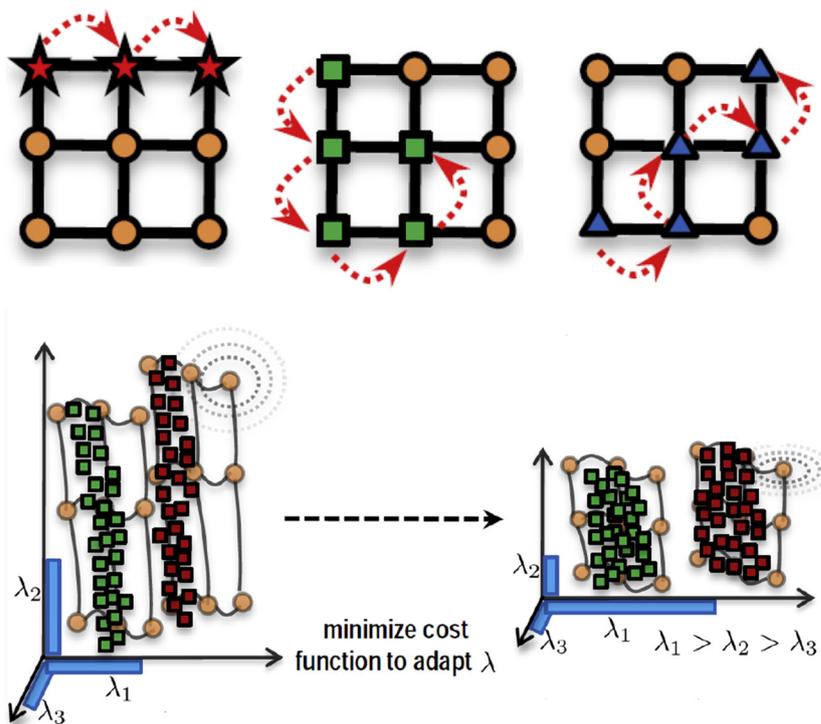
we get the forward variable A_{kt} with the start condition $A_{k1} = \pi_k p(\mathbf{x}^n(1) | \mathbf{w}^k, \Theta)$. The variable B_{kt} is the joint probability $p(\mathbf{x}^n(t+1) \dots \mathbf{x}^n(t_n), \mathbf{z}(t) = \mathbf{w}^k | \Theta)$ and is calculated using $B_{kt} = \sum_{i=1}^K p_{ik} p(\mathbf{x}^n(t+1) | \mathbf{w}^i, \Theta) B_{it+1}$ where $B_{kT} = 1$, B_{kt} defines the backward variable. The transition parameters are trained using the standard Baum–Welch training. As usual the underlying HMM also permits to deal with missing values and sequences of arbitrary length [3]. A more detailed description of the GTM-TT is given in [42].

For an input time-series $\mathbf{x}^n(1) \dots \mathbf{x}^n(T)$, GTM-TT specifies a time-series of responsibilities $r^{kn}(1) \dots r^{kn}(T)$ of neuron k . This can be used to define a winner for every time step t : $\text{argmax}_k r^{kn}(t)$.

3.3 Supervised GTM-TT

In the considered problem scenario our time-series data provide additional label information, such that \mathbf{x} is equipped with a label l , element of a finite label set $\{1, \dots, L\}$. We also assume that the given label is constant over time. Now, we would like to incorporate the label information in the optimization process of the GTM-TT leading to an extended supervised classification scheme. Given a labeled training set, we learn a separate GTM-TT for every class, whereby the models are linked by the same bandwidth β and the same underlying topological grid. We also use the same basis functions Φ and the Dirac distribution on the latent space. However, the prototype parameters \mathbf{W}_l , the initial state probability π_l and the transition probabilities \mathbf{P}_l are learned individually for every model representing label l . We refer to this model as the Supervised GTM-TT (SGTM-TT) as depicted schematically in Fig. 2. Accordingly, we will have a quantitative model for every class l after training.

Fig. 2 Illustration of the SGTM-TT. It consists of multiple GTM-TT models. It behaves similar to the regular GTM-TT but the training is classwise and the β parameter is common over the different models. The different classwise models are used to represent the data distribution over time (here for three classes). The SGTM-TT with relevance learning is shown at the bottom. The relevance of the input-dimensions is weighted over time during training. And only relevant dimensions with large λ -values are kept. In the figure, the λ_1 dimension discriminates the two groups and is pronounced by metric adaptation



In the recall or test phase we have to analyze a novel time-series \mathbf{x} and obtain L time-series of predicted responsibilities according to every model which will be denoted by $r_l^k(\mathbf{x}(t))$ (responsibilities of model l for input \mathbf{x} at time point t). We can summarize the responsibilities in an aggregated form as:

$$r_l(\mathbf{x}) := \sum_{k=1}^K \sum_{t=1}^T r_l^k(\mathbf{x}(t)) / (KT) \tag{8}$$

and one can select the label l as predicted output for which this value is largest.

3.4 Relevance learning for SGTM-TT

Metric adaptation for discriminative prototype-based learning has been introduced in [19], it is often also referred to as relevance learning. The basic idea is to parametrize the distance measure to incorporate auxiliary information. For the squared Euclidean metric one can define a parametrized, weighted, variant:

$$d_\lambda(\mathbf{x}, \mathbf{t}) = \sum_{d=1}^D \lambda_d^2 (x_d - t_d)^2. \tag{9}$$

For the GTM such a parametrization was already discussed in [15] for i.i.d. data resulting in relevance GTM (R-GTM). However, having temporal data some adaptations are necessary and also the supervision has to be handled in an alternative way. To keep the approach simple and to limit

the number of free parameters we will restrict our approach to a global diagonal weighted distance, in which case a weight λ_i directly corresponds to the relevance of dimension i . Here we assume normalized data with mean 0 and a standard deviation of 1 for each dimension. For GTM-(TT), the distance used to compute local probabilities is replaced by the previously discussed weighted Euclidean distance:

$$p_\lambda(\mathbf{x}|\mathbf{w}, \mathbf{W}, \beta) = \left(\frac{\beta}{2\pi}\right)^{D/2} \exp\left(-\frac{\beta}{2} d_\lambda(\mathbf{x}, y(\mathbf{w}, \mathbf{W}))\right). \tag{10}$$

Accordingly the data log likelihood considers the relevance of the data dimensions and, hence we obtain a corresponding topographic mapping.

A main difference of this approach to a standard integration of a data correlation matrix into the Gaussians consists in the fact that we prefer to adapt the relevance parameters in a supervised way according to the given label information, resulting in a discriminative approach.

The relevance parameters λ are optimized as suggested in [15] using the class information in an additional update step, interleaved with the standard adaptation of the SGTM-TT using the parametrized distance.

The discriminative learning of the metric parameters is controlled by the cost function of the generalized learning vector quantization (GRLVQ) which is a large margin technique [39]. We assume a classification based on a finite set of prototypes \mathbf{t}^j which are equipped with class labels

and represent the given data. A classification is done by means of a winner takes all scheme: the predicted label corresponds to the prototype with smallest distance $d_\lambda(\mathbf{x}, \mathbf{t}^j)$. For standard GTM, our prototypes are given by latent points $\mathbf{t}^j = y(\mathbf{w}^j, \mathbf{W})$, and the distances determine the responsibilities of the data points. The relevance terms λ are adapted such that the costs

$$E(\lambda) = \sum_n \text{sgd} \left(\frac{d_\lambda(\mathbf{x}^n, \mathbf{t}^+) - d_\lambda(\mathbf{x}^n, \mathbf{t}^-)}{d_\lambda(\mathbf{x}^n, \mathbf{t}^+) + d_\lambda(\mathbf{x}^n, \mathbf{t}^-)} \right) \quad (11)$$

are minimized. The closest prototype with the correct labeling is denoted by \mathbf{t}^+ and the one with the incorrect label by \mathbf{t}^- , for a given input \mathbf{x}^n . The sigmoid function (sgd) is defined as: $\text{sgd}(x) = \frac{1}{1+\exp(-\sigma \cdot x)} \in [0, 1]$. This optimization scheme can be integrated into the vectorial GTM, simultaneously adapting the GTM parameters, optimizing the data log-likelihood, and the metric parameters optimizing the classification margin. The update equations for the parameters λ can be derived from (11), taking the derivatives. To keep a quadratic form in the distance measure, the metric parameters are normalized after each adaptation step.

Given an input sequence \mathbf{x} we get a prototype representation of this time-series by evaluating the SGTM-TT in the following way. For every class label we consider the time-series of prototypes of the corresponding GTM-TT model according to the winner prototypes over time:

$$\mathbf{t}_l = (\mathbf{t}_l(1) \dots \mathbf{t}_l(T)). \quad (12)$$

where

$$\mathbf{t}_l(t) = y(\mathbf{w}^k, \mathbf{W}_l) \quad \text{with } k = \text{argmax}_k r_l^k(\mathbf{x}(t)). \quad (13)$$

Now the time-series \mathbf{x} and the corresponding time-series of prototypes representing a correct or a wrong class label can be used in (11) to adapt the underlying metric. If we assume an appropriate metric for the comparison of two time-series, a well defined cost function results.

Several reasonable distance measures for time-series can be considered, whereby the only property which we will use is differentiability. For simplicity we will also assume, that the time-series have equal length, although the model can be generalized to time-series of different length.

A very simple distance for such time-series would be to average over the Euclidean distances in each time point. This however is inappropriate, because it will completely neglect the functional form of the data. An appropriate measure, designed for the comparison of timeseries, was proposed in [21] and will be used instead. Further alternative time-series metrics are possible see, e.g. [10], but the chosen one has been found to be effective in prior work [38] and can be calculated at low costs.

The considered distance measure integrates the functional form of three subsequent time steps in comparing $\mathbf{x}(t)$ and $\mathbf{t}(t)$. Let us assume we have a real valued time-series $\mathbf{v} = v(1) \dots v(T)$, then the functional L_p norm can be defined as [21]:

$$\mathcal{L}_p^f(\mathbf{v}) = \left(\sum_{t=1}^T (\Delta A_t(\mathbf{v}) + \Delta B_t(\mathbf{v}))^p \right)^{\frac{1}{p}} \quad (14)$$

with

$$\Delta A_k(\mathbf{v}) = \begin{cases} \frac{\tau}{2} |t| & \text{if } 0 \leq v(t)v(t-1) \\ \frac{\tau}{2} \frac{(t)^2}{|v(t)| + |v(t-1)|} & \text{if } 0 > v(t)v(t-1) \end{cases} \quad (15)$$

$$\Delta B_k(\mathbf{v}) = \begin{cases} \frac{\tau}{2} |v(t)| & \text{if } 0 \leq v(t)v(t+1) \\ \frac{\tau}{2} \frac{v(t)^2}{|v(t)| + |v(t+1)|} & \text{if } 0 > v(t)v(t+1) \end{cases} \quad (16)$$

representing the triangles on the right and the left sides of $v(t)$ and boundary points are set to 0. This norm accounts for entries which change the sign in subsequent time steps. We obtain a weighted distance, for vectorial data \mathbf{x} and \mathbf{t} over time with equal dimensionality D at each time point:

$$d_\lambda(\mathbf{x}, \mathbf{t}) = \sum_{i=1}^D \lambda_i \mathcal{L}_p^f(\mathbf{x}_i - \mathbf{t}_i) \quad (17)$$

where $\mathbf{x}_i - \mathbf{t}_i$ refers to the time-series of real numbers given by the distance of the entries in dimension i . As a special property of this distance measure the similarity of the curvature of the sequences is taken into account. Again, each dimension is weighted by the normalized relevance parameters λ_i .

This weighted metric (17) is used in the cost function (11). If we take the derivatives (see [38] for \mathcal{L}_p -norm) with respect to the relevance terms an adaptive weighting for the input dimensions is obtained taking the functional form of the data into account. Again the λ are normalized after every adaptation to obtain non-negative values, summing up to 1.

3.4.1 Relevant time points:

Since SGTM-TT relies on HMMs, every time point depends on its predecessor only. Thus, it is not reasonable to adapt the relevance of time points to obtain a better representation of data in the GTM-TT models. However, it is reasonable to judge the relevance of time points resulting from the GTM-TT models for the final classification, in

particular if time-series are of the same or a similar length. This method offers insights into the model to identify time points which are particularly discriminative for the given task at hand.

We obtain a relevance profile in the following way: denote by $r_l(\mathbf{x}(t)) := \sum_{k=1}^K (r_l^k(\mathbf{x}(t)))/K$ the accumulated responsibility of the GTM-TT model l for data point \mathbf{x}^n at time point t . Based on this value, a classification can be based on the maximum responsibility $r_l(\mathbf{x}(t))$ in time point t . For every time point t , we simply count the number of data points which are classified correctly as belonging to class l based on the classification for time point t only, averaged over all data. A global relevance profile results thereof as a sum over all labels.

4 Odor measurement system

The analyte measurement system employed to gather the data presented in this article is shown in Fig. 3. It consists of an array of metal oxide semiconductor (MOX) gas sensors hosted inside a measurement chamber, a pneumatic circuit to control the exposition of the sensors to the volatile molecules dispersed in the environment, and the electronics necessary to power up the sensors and respective measurement circuits.

The election of MOX as the gas sensing technology has been made attending at its high sensitivity, commercial availability and low price. However, they present some shortcomings including: poor selectivity, influence by environmental factors such as humidity and temperature [30] and major limitations in their response speed [29]. Among these drawbacks, their poor selectivity is of the largest concern for odor classification. To overcome this, it is a common practice to build the e-nose upon

an array of MOX sensors with different and partially overlapping sensitivities. The output of the array is then processed with a pattern recognition algorithm to find out which substance the e-nose is exposed to. Based on this concept, we choose five different MOX gas sensors to compose the sensor array: TGS-2600, TGS-2602, TGS-2611 and TGS-2620 from Figaro Sensors,¹ and MiCS-5135 from e2V Sensors.²

In order to enable sensors to interact with the volatile molecules dispersed in the environment, the e-nose employs a pump to enforce a constant airflow through the sensors array. The aspiration and release of the air samples are accomplished through tubes, conveniently separated one from another to avoid cross contamination. Additionally, the aspiration through flexible tubes allows the displacement of the aspiration entry without the need to move the complete system. This advantage is particularly useful in robotics to easily sample the space, for example, by attaching the e-nose aspiration to the hand of an arm robot as shown in the “experimental section”.

4.1 Signal conditioning and data preprocessing

The data, as provided by the e-nose, present a measurement intrinsic baseline, which can be seen as a signal offset. Here, we estimate the baseline value as the median signal intensity within the first 5–20 s, and then, remove it from each measurement truncating values to zero when necessary.

More sophisticated preprocessing, by means of advanced baseline correction algorithms, smoothing strategies or normalization techniques [35] are possible but out of focus of this paper. We also do not further explore specific feature extraction techniques for spectral data but focus on the obtained normalized intensities.

5 Experimental results

This section describes the setups and classification results for three different experiments designed with increasing classification challenge. Furthermore, a comparison of results with SVM, NN and a very recent reservoir computing based time-series classifier (RTK) as proposed in [6] is provided. For RTK the core idea is to transform the time-series into a higher-dimensional dynamical feature space via reservoir computation models. Subsequently varying aspects of the signal are represented through variation in the linear readout models trained in such dynamical feature spaces, for details see [6].

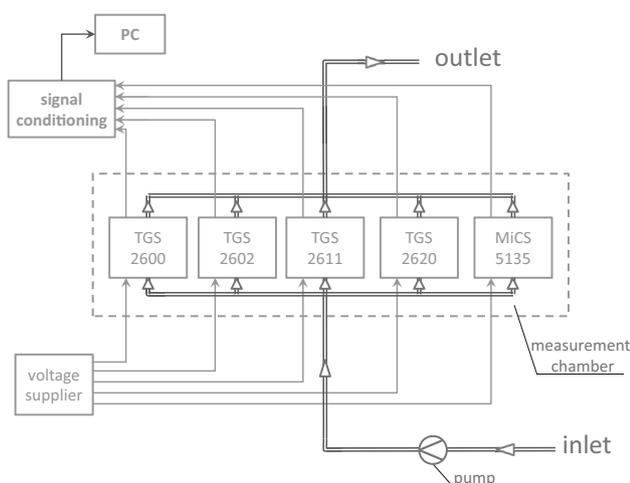


Fig. 3 Measurement system

¹ Figaro engineering inc. <http://www.figaro.co.jp>.

² e2v. <http://www.e2v.com/>.

Table 1 Average test set accuracy for the first and second experiments in a fivefold cross-validation

CV-accuracy	SGTM-TT (%)	SVM (%)	NN (%)	RTK (%)
SIM	<i>94.00 ± 4.18</i>	90.00 ± 5.00	55.00 ± 13.54	66.30 ± 8.54
DS1	88.03 ± 9.72	86.36 ± 9.66	80.49 ± 11.90	<i>96.67 ± 4.56</i>
DS-UCI-1	87.78 ± 5.76	<i>93.89 ± 4.97</i>	86.81 ± 7.98	64.44 ± 4.12
DS-UCI-2	79.55 ± 9.15	83.03 ± 18.47	76.33 ± 18.15	<i>94.70 ± 8.05</i>

Significant better results are italicized

In general, we are interested on simple methods or at least methods which provide direct interpretation of the model parameters and results. For example it is very desirable to have direct links to the input features to find channels which are most discriminative for a specific substance, relevant over all classes but also the other way, being not very relevant. The later is an important characteristic for systems with limited resources, like mobile robotics, where it would be desirable to power on only the relevant sensors. Accordingly (local) linear methods are interesting in contrast to black box non-linear kernel mappings. We are also interested on approaches which permit an easy and quick out of sample extension to, in our case, substantially shorter sequences in the test phase. This rules out multiple complicated time-series models.

For SVM we used a linear kernel with optimal C determined over the training data on a grid search. Since SVM can not directly be applied to temporal data, nor can it be used for sequences of different length in a direct way, for the comparison we simply concatenate the measurements of the different channels to remove the time dimension. More complex strategies of applying SVM, e.g. by using a dynamic time warping (DTW) kernel could be done but are not in the focus of this paper and out of sample extensions are often not immediate which is an issue for online robotic sensor systems. For more recent work around DTW or kernel related time-series analysis see, e.g. [6, 28, 32]. Additionally we would like to avoid more complex preprocessing steps to permit an easy out of sample extension in practical settings. Although maximum classification performance is not our main objective, we also provide a comparison with a very recent reservoir computing kernel [6]. This approach is known to be very effective for timeseries but on the other hand is less interpretable nor is the out of sample extension for very short sequences immediate. For RTK there are three parameters optimized on the training data within a grid search³ as detailed in [6].

³ Grids: $\lambda, \gamma = [0, 10^{-6}, \dots, 10^{-1}, 0.5, 1, \dots, 5, 10, 30, 50, 100]$
costs = $[0.1, 10, 10^2, 5 \times 10^2, 10^3, 5 \times 10^3, 10^4, 5 \times 10^4]$.

5.1 Experiment 1: simulated data

The first experiment is based on the simulated data proposed in [23] with the only intention of validating the proposed algorithm under known conditions.

The simulated data (SIM) consist of 100 samples separated into two classes of 50 samples each. Each point is located in a 100 dimensional feature space with 8 time points. From the given features, only 10 are expected to differentiate between the classes. Details about the data and the generation procedure are given in [23].

We applied SGTM-TT with relevance learning using nine hidden states and four basis functions. We observe an overall prediction accuracy of $94 \pm 4\%$. The relevance profile identified all known ten features and effectively pruned out the remaining irrelevant data dimensions. Our results are slightly better than those reported in [23] (90%) and in [7] (92%).

The dataset is a particular short time-series with a rather large number of input dimensions. Especially the small number of time points can be quite challenging for other time-series models but may actually occur in the context of electronic nose experiments, where short sensing cycles would be very desirable. The prediction results of the different methods are summarized in Table 1. With the exception of NN most methods perform reliably well but SGTM-TT was significantly better.

5.2 Experiment 2: controlled gas exposure

The second experiment aims to test the proposed method with real odor data under restrained environmental conditions. To this end, a dataset of real odor samples is gathered in a scenario as controlled as possible. The dataset is comprised of 39 samples generated by exposing the e-nose to gas pulses of four different analytes: a commercial spirit (*Larios Gin*), a polish remover based on acetone, standard ethanol and lighter gas (butane mixed with propane). Acetone was given by nine samples and the other classes by ten samples each.

Each sample is collected according to the following three-phase procedure: (1) for the initial 30 s, baseline value is estimated by measuring the sensor response in the

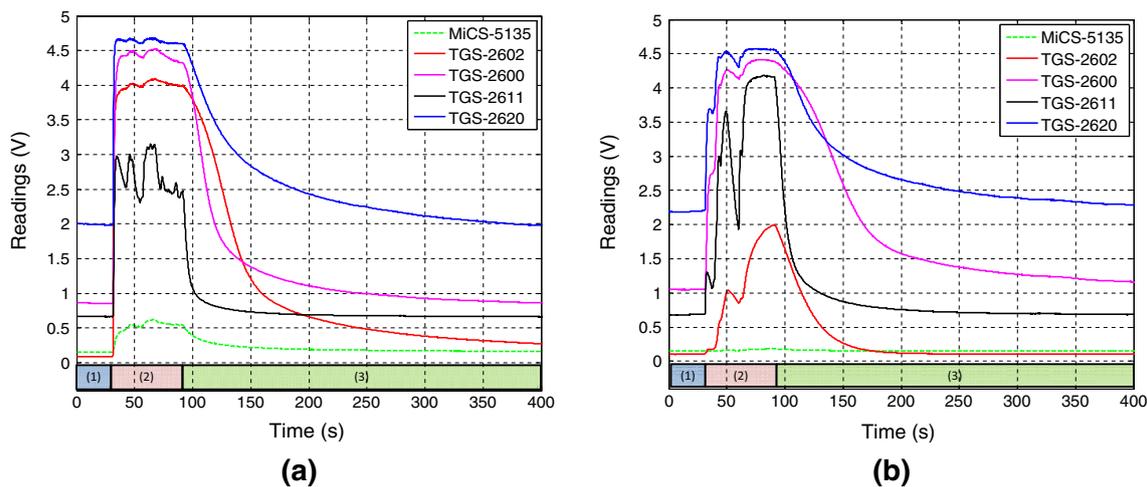


Fig. 4 Two different samples of the olfaction dataset gathered in the second experiment. The three phases in which the samples can be decomposed are marked at the *bottom* of each figure as (1), (2) and (3)

absence of the target gas, (2) and for a duration of 60 s the e-nose is placed next to the gas source (about 10 cm) exposing the sensor array to the volatile. (3) the gas source is removed allowing the sensor array to recover to its initial state (baseline).

Figure 4 shows two different samples of such dataset. Notice that although the gas exposure was “controlled” by time exposure and distance to the source, strong fluctuations in the sensor readings occur due to the chaotic nature of the gas dispersion.

The SGTM-TT is inherently capable of dealing with measurement sequences of different length in time, using the HMM mapping functionality. However, to permit fair comparison with other approaches like vector embeddings, we consider only the first 100 s of the data. That is, we built a first dataset (DS1) using the initial 100 s of each sample, which corresponds to 487 sampling points.

For comparison, we also use two public domain data sets of similar type (electronic nose data) from the UCI database. The DS-UCI-1 data set is given by the two sources gas data [14]. The data are measured using a chemical detection platform composed of eight chemoresistive gas sensors which were exposed to turbulent gas mixtures generated naturally in a wind tunnel. It consists of 180 time-series of ethylene (Eth), carbon monoxide (CO) and methane (Me) mixtures at different concentrations. We use the data as a two class prediction problem to predict the whether Eth was mixed with CO or Me. Available features are temperature, humidity and the eight sensor channel outputs. Each time-series is given with 2,970 sampling points.

The DS-UCI-2 data set is given by the pulmon data [49]. The data are measured using a chemical sensing system based on an array of 16 metal-oxide gas sensors and an

external mechanical ventilator to simulate the biological respiration cycle. The tested gas classes are mixtures of acetone and ethanol. Data have been normalized to zero-mean, and intensity and considered again as a prediction problem to identify whether the mixture contains Me or CO.

The classification accuracy for DS1, DS-UCI-1, DS-UCI-2 is given in Table 1 in comparison to some standard approaches. We observe that the SGTM-TT performs reliably well although the best prediction accuracy for DS1 and DS-UCI-2 is obtained by the RTK approach. For the DS-UCI-1 dataset RTK is significantly worse than the other approaches and the SVM obtained the best performance. Hence there is no clear winner regarding the classification accuracy but SGTM-TT represents a good approach with a reliable and consistent performance. Furthermore, as previously commented, the classification performance is not the only point that matters but also the simplicity of the model and the interpretability of the results. Neither RTK nor SVM provide additional insight into the relevance of the sensor channels.⁴ Here we are mainly interested in interpretable models [24] which also simplify a later transfer of the approach to an embedded system or the sensor platform. In Fig. 5 we show the averaged (global) sensor relevance profile of DS-UCI-1 and DS-UCI-2.

Subsequently we give a detailed analysis for our own dataset—DS1, where we have more background information to provide a specific in-depth discussion of the results. For the analysis of the sensor relevance and time points relevance, the whole measurement sequence of each

⁴ Approaches for feature ranking by SVM are available but not for this type of data and not directly for multi-class problems as studied for DS1.

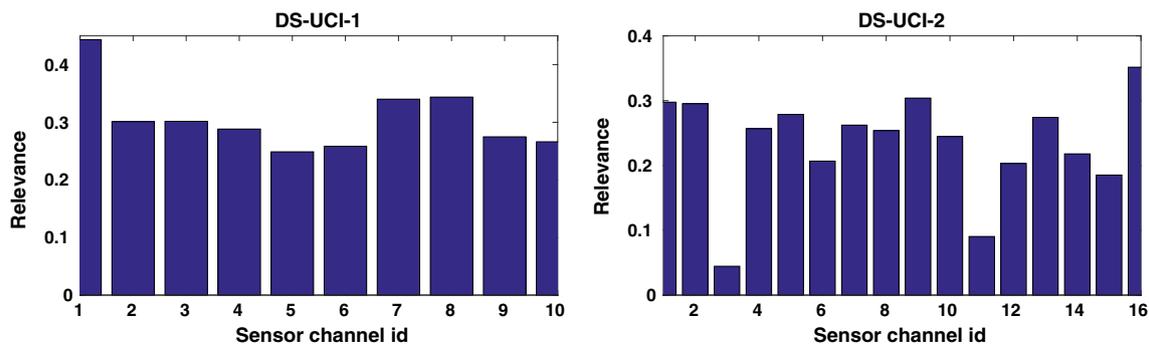


Fig. 5 Relevance profile of the sensor input for DS-UCI-1 (*left*) and DS-UCI-2 (*right*). For both profiles the information is distributed over the various sensors but some sensors are more important, e.g. sensor 1 for DS-UCI-1 and sensor 3 for DS-UCI-2

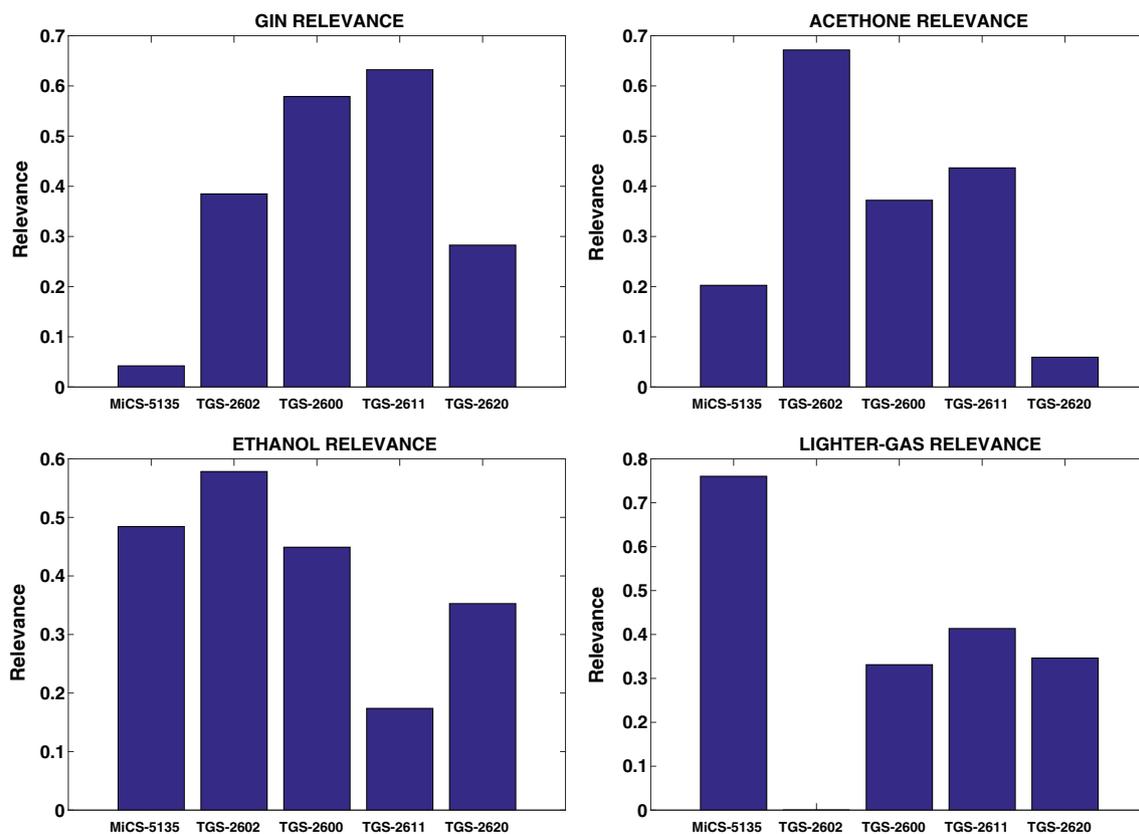


Fig. 6 Sensor relevance indexes for the four odor classes used on dataset DS2

sample was down-sampled to 800 time points each (DS2). The SGTM-TT was then trained in a fivefold cross validation with four hidden states and four basis functions. In Fig. 6 we show the relevance indexes of the five gas sensors of the e-nose for the different target volatiles of DS2 as obtained by SGTM-TT. Different conclusions can be drawn from the study of such relevance plot:

- In general, the five MOX sensors are relevant for the classification of the different volatiles, being sensor

TGS-2620 the less relevant one, and so the most expendable.

- Sensor TGS-2602 is the most relevant one when classifying acetone and ethanol samples, with a notable difference with respect to the other sensors in the case of acetone. This characteristic is already reported in the manufacturer's datasheet, indicating the high sensitivity to volatile organic compounds (VOCS) of this sensor model.

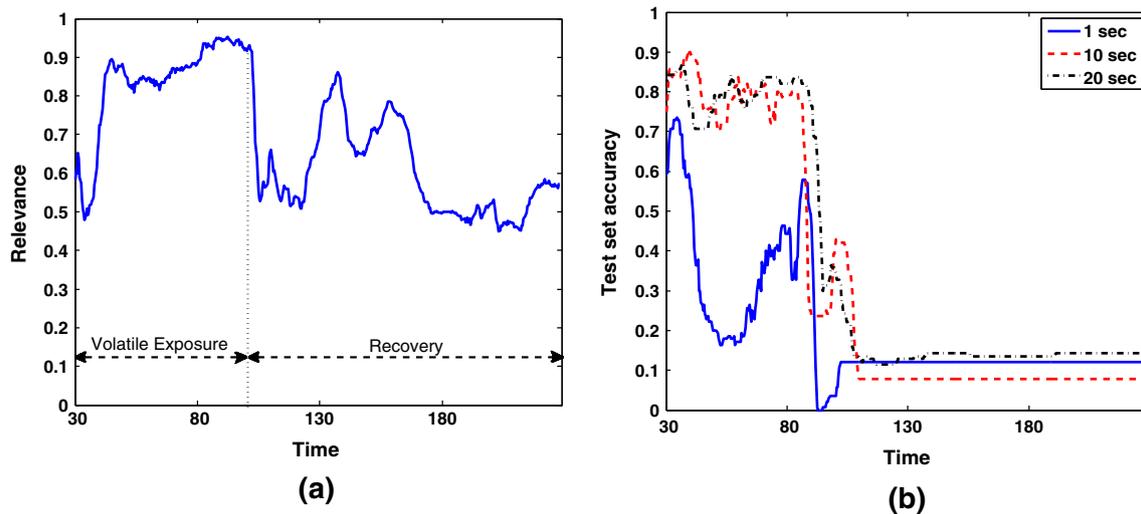


Fig. 7 Time points relevance profile (a), averaged over all classes and mean prediction accuracy over time with window lengths of $\approx 1, 10, 20$ s (b)

Fig. 8 a The robotic arm used in the third experiment mounted over a mobile platform, and a detailed view of the attached e-nose aspiration. **b** Picture of the proposed setup for the third experiment. Each of the *black plastic vessels* contains a different substance



- As expected from the low selectivity characteristic of MOX sensors, each sensor presents a high relevance index for more than one odor class.

We also explore the relevance of individual time points of the dataset DS2, depicted in Fig. 7a. As expected, the time interval under volatile exposition, the first 100 seconds, is the most discriminating. Furthermore, and as already reported in [9], it is a noticeable fact that the relevant information for classification purposes can be found in the recovery phase, after the volatile has been removed.

Since in real robotics conditions the classifier is expected to work on small data sequences, a second configuration for the dataset DS2 was tested. Here, the test data consist only of short sensor readings over time. Figure 7a depicts the accuracy in the classification for three different window lengths (1, 10 and 20 s). We observe that given the highly dynamic response of MOX sensors in addition to the inherent signal noise, very small windows (1 s) do not carry

enough information for a reasonable classification, but for data sequences of ten seconds the accuracy in the prediction achieves very good results (values near 0.8). Furthermore, window lengths over ten seconds seem to not improve the accuracy, which indicates that long sequences encode a lot of noise contributions, hampering the model in the prediction. Finally, it must be noticed that the classification accuracy is usually higher when using data from the transient parts of the signal (rise and decay) than when steady-state data are employed, as denoted by the accuracy peaks found around 30 and 90 s.

5.3 Experiment 3: robotics experiment—uncontrolled gas pulses

Finally, and with the aim to validate the classification performance in a more challenging robotic scenario, a third experiment is presented. In this case, the e-nose aspiration

Fig. 9 MCE-nose gathered signals of the classification experiment with a robotic arm, and the “ground-truth” sequence of the employed analytes. The active chamber [0, 1, 2, 3] is switched every 20 s. Signals are shown for the four different sensor channels as described before

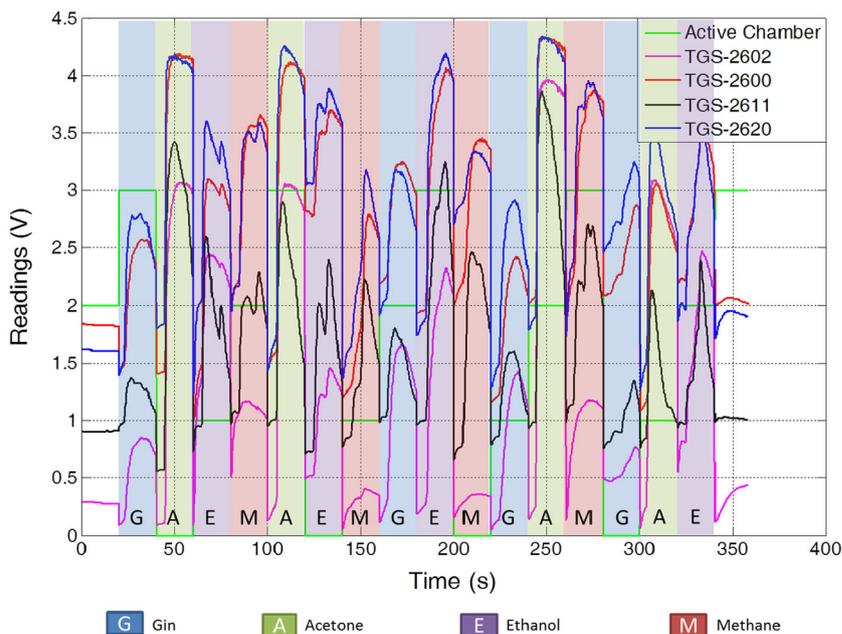


Table 2 Predictions for the external evaluation data using the first respective crossvalidation model

Time	40	60	80	100	120	140	160	180	200	220	240	260	280	300	320	340
True	G	A	E	M	A	E	M	G	E	M	G	A	M	G	A	E
SGTM-TT																
Pred.	G	M	E	M	A	M	M	G	E	M	G	A	M	G	A	M
Error		o														o
NN																
Pred.	G	E	G	M	E	G	G	G	G	M	G	A	G	M	E	G
Error		o	o		o	o	o		o				o	o	o	o

The ‘o’ in the line labeled with error indicates mismatches

(see Fig. 8) is attached to the hand of a robotic arm [34] which is commanded to approximate the e-nose aspiration to each of four recipients containing different substances (acetone, ethanol, butane⁵ and gin).

To avoid waiting for the sensors to recover their baseline levels after each exposure (which would take more than a minute), we have employed a specially designed e-nose, called MCE-nose [17], that allows the measurement of fast changing gas concentrations.

The robotic arm is commanded to approximate to the containers following a predefined sequence. The exposition to each of the substances takes 20 s, after which the arm moves to another container. The volatile sequence and the gathered signals during the experiment are depicted in Fig. 9. A video of a similar experiment is additionally available at <http://mapir.isa.uma.es/mapirwebsite/index.php/2008-tep-4016-media>.

⁵ Since butane is found at gas state at ambient temperature, the content of a lighter was released when the e-nose aspiration moved over the container.

Each of the short sequences was pre-processed such that the baseline is removed. Then the sequences have been matched with the SGMT-TT or NN model as obtained from DS1.⁶ This can be considered to be a test of the model on an independently measured hold out dataset.

The ground-truth and predicted labels of the sequences are given in Table 2 with only 3 errors out of the 16 test samples. In the experiment the SGMT-TT classifier was continuously online and fed by new data every 20 s according to the measurement protocol. This experiment is interesting because the input data processed by the SGMT-TT method are substantially shorter than the training dataset, with around 30 sampling points for the core measurement. The SVM model cannot be applied here due to the varying length of the input data and for the RTK model the sequence is also too short to get reliable predictions as the method is not designed for this type of test inputs. For NN we applied a local DTW alignment between each training and test sample using the best local fit.

⁶ Here we simply used the model from the first crossvalidation run.

6 Conclusion

A novel approach for the analysis of high-dimensional and rather short temporal sequences was presented. It is based on the idea to introduce available meta information into the modeling process of a GTM through time, given in form of supervised information and relevance learning. We have analyzed the suitability of such model for the odor classification problem in robotics applications, providing comparative results with SVM, NN and the reservoir time-series kernel (RTK) for three different scenarios (with increasing classification challenge), and demonstrating that the proposed method is effective for solving such high-dimensional data problem.

Other remarkable advantages of the method in the context of odor classification in robotics are on the one hand, the possibility for the robot to perform rapid classification of chemical substances using a short data sequence. On the other hand, the SGTM-TT method outputs relevance values for both the sensors being used as well as the time-points of the signal, which provide very valuable information to configure the e-nose and to carry out the robot smelling.

In future work, it will be of interest to analyze the SGTM-TT in the context of drift problems as recently discussed in [46, 47] and how the method can be further improved by early decision strategies [20].

Acknowledgments The first author likes to thank Peter Tino, University of Birmingham, for interesting discussions about probabilistic modeling and Tien-ho Lin, Carnegie Mellon University, USA for support with the simulation data. Further, Ivan Olier, University of Manchester, UK; Iain Strachan, AEA Technology, Harwell, UK and Markus Svensen, Microsoft Research, Cambridge, UK for providing code for GTM and GTM-TT. Further we would like to thank Fengzhen Tang, University of Birmingham for providing invaluable support with the RTK method. This work was supported by the DFG project HA2719/4-1 to BH, by the DFG-NSF project TO 409/8-1, and by the Cluster of Excellence 277 CITEC funded in the framework of the German Excellence Initiative. Further, a Marie Curie Intra-European Fellowship (IEF): FP7-PEOPLE-2012-IEF (FP7-327791-ProMoS) is gratefully acknowledged. Additional support was provided by funds from the Andalucía Regional Government and the European Union (FEDER) under research project: TEP08-4016.

References

1. Astudillo CA, Oommen BJ (2014) Topology-oriented self-organizing maps: a survey. *Pattern Anal Appl* 17(2):1–26
2. Bishop CM, Svensén M, Williams CKI (1998) Gtm: the generative topographic mapping. *Neural Comput* 10(1):215–234
3. Bishop CM (2006) *Pattern recognition and machine learning*. Information science and statistics. Springer, New York
4. Blanco J-L, Monroy JG, González-Jiménez, J, Lilienthal A (2013) A Kalman filter based approach to probabilistic gas distribution mapping. In: 28th symposium on applied computing (SAC), March 2013.
5. Brahim-Belhouari S, Bermak A (2005) Gas identification using density models. *Pattern Recognit Lett* 26(6):699–706
6. Chen H, Tang F, Tino P, Yao X (2013) Model-based kernel for efficient time-series analysis. In: *Proceedings of the 19th ACM SIGKDD conference on knowledge discovery and data mining (KDD) (KDD'13)*, Chicago, USA
7. Costa IG, Schönhuth A, Hafemeister C, Schliep A (2009) Constrained mixture estimation for analysis and robust classification of clinical time-series. *Bioinformatics* 25(12):i6–i14
8. Distanto C, Siciliano P, Persaud KC (2002) Dynamic cluster recognition with multiple self-organising maps. *Pattern Anal Appl* 5(3):306–315
9. Distanto C, Leo M, Siciliano P, Persaud KC (2002) On the study of feature extraction methods for an electronic nose. *Sens Actuators B Chem* 87(2):274–288
10. Douzal-Chouakria A, Amblard C (2012) Classification trees for time-series. *Pattern Recognit* 45(3):1076–1091
11. Dumitrescu D, Lazzarini B, Marcelloni F (2000) A fuzzy hierarchical classification system for olfactory signals. *Pattern Anal Appl* 3(4):325–334
12. Ehret B, Safenreiter K, Lorenz F, Biermann J (2011) A new feature extraction method for odour classification. *Sens Actuators B Chem* 158(1):75–88
13. Ferri G, Caselli E, Mattoli V, Mondini A, Mazzolai B, Dario P (2009) Spiral: a novel biologically-inspired algorithm for gas/odor source localization in an indoor environment with no strong airflow. *Robot Auton Sys* 57(4):393–402
14. Fonollosa J, Rodríguez-Luján I, Trincavelli M, Vergara A, Huerta R (2014) Chemical discrimination in turbulent gas mixtures with mox sensors validated by gas chromatography-mass spectrometry. *Sensors* 14(10):19336–19353
15. Gisbrecht A, Hammer B (2011) Relevance learning in generative topographic mapping. *Neurocomputing* 74(9):1359–1371
16. González-Jiménez J, Monroy JG, Blanco J-L (2013) Robots that can smell: motivation and problems. In: *Submitted to 15th international symposium on olfaction and electronic nose (ISOEN)*
17. Gonzalez-Jimenez J, Monroy JG, Blanco J-L (2011) The multi-chamber electronic nose. An improved olfaction sensor for mobile robotics. *Sensors* 11(6):6145–6164
18. Gutierrez-Osuna R (2002) Pattern analysis for machine olfaction: a review. *Sens J IEEE* 2(3):189–202
19. Hammer B, Villmann TH (2002) Generalized relevance learning vector quantization. *Neural Netw* 15(8–9):1059–1068
20. Hatami N, Chira C (2013) Classifiers with a reject option for early time-series classification. In: *Proceedings of the IEEE symposium on computational intelligence and ensemble learning, CIEL 2013, IEEE symposium series on computational intelligence (SSCI)*, 16–19 April 2013. IEEE, Singapore, pp 9–16
21. Lee J, Verleysen M (2005) Generalizations of the lp norm for time-series and its application to self-organizing maps. In: Cottrell M (ed) *5th workshop on self-organizing maps*, vol 1, pp 733–740
22. Lilienthal AJ, Reggente M, Trincavelli M, Blanco J-L, Gonzalez J (2009) A statistical approach to gas distribution modelling with mobile robots—the kernel dmV algorithm. In: *Intelligent robots and systems, 2009. IROS 2009. IEEE/RSJ international conference*, pp 570–576
23. Lin T-H, Kaminski N, Bar-Joseph Z (2008) Alignment and classification of time-series gene expression in clinical studies. In: *ISMB*, pp 147–155
24. Lisboa PJG (2013) Interpretability in machine learning—principles and practice. In: Masulli F, Pasi G, Yager RR (eds) *Fuzzy logic and applications—10th international workshop, WILF 2013, Genoa, Italy, 19–22 Nov 2013. Proceedings, Lecture Notes in Computer Science*, vol 8256. Springer, pp 15–21

25. Loutfi A, Coradeschi S, Lilienthal AJ, Gonzalez J (2009) Gas distribution mapping of multiple odour sources using a mobile robot. *Robotica* 27(2):311–319
26. Marques L, Nunes, U, de Almeida AT (2002) Olfaction-based mobile robot navigation. *Thin Solid Films*, 418(1):51–58. Proceedings from the international school on gas sensors in conjunction with the 3rd European school of the NOSE network
27. Martinelli E, Magna G, Vergara A, Di Natale C (2014) Cooperative classifiers for reconfigurable sensor arrays. *Sens Actuators B Chem* 199:83–92
28. Michalak M (2011) Adaptive kernel approach to the time-series prediction. *Pattern Anal Appl* 14(3):283–293
29. Monroy JG, Gonzalez-Jimenez J, Blanco J-L (2012) Overcoming the slow recovery of mox gas sensors through a system modeling approach. *Sensors* 12(10):13664–13680
30. Moseley P, Tofield B (eds) (1987) *Solid state gas sensors*. Adam Hilger, Bristol
31. Olier I, Vellido A (2008) Advances in clustering and visualization of time-series using gtm through time. *Neural Netw* 21(7):904–913
32. Petitjean F, Ketterlin A, Ganarski P (2011) A global averaging method for dynamic time warping, with applications to clustering. *Pattern Recognit* 44(3):678–693
33. Prieto OJ, Alonso-González CJ, Rodríguez JJ (2013) Stacking for multivariate time-series classification. *Pattern Anal Appl* 1–16. doi:10.1007/s10044-013-0351-9
34. Robotics K. The jaco research edition robotic arm. <http://www.kinovarobotics.com/>
35. Sauve AC, Speed TP (2004) Normalization, baseline correction and alignment of high-throughput mass spectrometry data. In: *Proceedings of Gensips*. <http://www.stat.berkeley.edu/~terry/Group/publications.html>
36. Schleif F-M, Gisbrecht A, Hammer B (2012) Relevance learning for short high-dimensional time-series in the life sciences. *Proceedings of IJCNN*, pp 2069–2076
37. Schleif F-M, Mokbel B, Gisbrecht A, Theunissen L, Dürr V, Hammer B (2012) Learning relevant time points for time-series data in the life sciences. In: *Proceedings of ICANN*, pp 531–539
38. Schleif F-M, Ongyerth F-M, Villmann T (2009) Supervised data analysis and reliability estimation for spectral data. *Neurocomputing* 72(16–18):3590–3601
39. Schneider P, Biehl M, Hammer B (2009) Distance learning in discriminative vector quantization. *Neural Comput* 21:2942–2969
40. Shaffer RE, Rose-Pehrsson SL, Andrew McGill R (1999) A comparison study of chemical sensor array pattern recognition algorithms. *Anal Chim Acta* 384(3):305–317
41. Shumway RH, Stoffer DS (2000) *Time-series analysis and Its applications*. Springer Texts In Statistics. Springer, New York
42. Strachan IGD (2002) *Latent variable methods for visualization through time*. Ph.D. thesis, University of Edinburgh, Edinburgh
43. Strickert M, Hammer B (2005) Merge SOM for temporal data. *Neurocomputing* 64:39–72
44. Szczurek A, Krawczyk B, Maciejewska M (2013) VOCs classification based on the committee of classifiers coupled with single sensor signals. *Chemometr Intell Lab Syst* 125:1–10
45. Trincavelli M, Coradeschi S, Loutfi A (2009) Odour classification system for continuous monitoring applications. *Sens Actuators B Chem* 58:265–273
46. Vergara A, Fonollosa J, Mahiques J, Trincavelli M, Rulkov N, Huerta R (2013) On the performance of gas sensor arrays in open sampling systems using inhibitory support vector machines. *Sens Actuators B Chem* 185:462–477
47. Vergara A, Vembu S, Ayhan T, Ryan MA, Homer ML, Huerta R (2012) Chemical gas sensor drift compensation using classifier ensembles. *Sens Actuators B Chem* 166–167:320–329
48. Welch LR (2003) Hidden Markov models and the Baum–Welch algorithm. *IEEE Inf Theory Soc Newsl* 53(4). http://www.itsoc.org/publications/nltr/it_dec_03final.pdf
49. Ziyatdinov A, Fonollosa J, Fernandez L, Gutierrez-Glvez A, Marco S, Perera A (2015) Bioinspired early detection through gas flow modulation in chemo-sensory systems. *Sens Actuators B Chem* 206:538–547