

1 **Odor recognition in robotics applications**  
2 **by discriminative time series modeling**

3 **F.-M. Schleif, B. Hammer,**  
**J. G. Monroy, J. Gonzalez-Jimenez, J. L. Blanco,**  
**M. Biehl, N. Petkov**

4  
5

6 **Abstract** Odor classification by a robot equipped with an electronic nose (e-nose) is  
7 a challenging task for pattern recognition since volatiles have to be classified quickly  
8 and reliably even in the case of short measurement sequences, gathered under op-  
9 eration in the field. Signals obtained in these circumstances are characterized by a  
10 high dimensionality, which limits the use of classical classification techniques based  
11 on unsupervised and semi-supervised settings, and where predictive variables can be  
12 only identified using wrapper or post-processing techniques. In this paper, we con-  
13 sider generative topographic mapping through time (GTM-TT) as an unsupervised  
14 model for time series inspection, based on hidden Markov models regularized by  
15 topographic constraints. We further extend the model such that supervised classifica-  
16 tion and relevance learning can be integrated, resulting in supervised GTM-TT. Then,  
17 we evaluate the suitability of this new technique for the odor classification problem  
18 in robotics applications. The performance is compared with classical techniques as  
19 nearest neighbor (NN), as an absolute baseline, support vector machine (SVM) and a  
20 recent time series kernel approach, demonstrating the eligibility of our approach for  
21 high dimensional data. Additionally, we exploit the learning system introduced in this  
22 work, providing a measure of the relevance of each sensor and individual time points  
23 in the classification process, from which important information can be extracted.

24 **Keywords** electronic nose, volatile classification, odor recognition, time series,  
25 prototype learning, relevance learning

26 **1 Introduction**

27 Olfaction plays an important role in the development of many applications, such as  
28 quality control in food processing chains, detection and diagnosis in medicine, find-  
29 ing drugs and explosives, and the more common estimation of blood alcohol content

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

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

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

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

56 In this paper we demonstrate the suitability of a novel approach based on genera-  
57 tive topographic mapping through time (GTM-TT) to the problem of volatile identi-  
58 fication in robotics. The model extends classical GTM-TT by integrating supervised  
59 classification and relevance learning, resulting in supervised GTM-TT (SGTM-TT).  
60 More precisely, we have tested the SGTM-TT method with an e-nose comprising an  
61 array of MOX (metal oxide sensors) to classify samples of seven different volatiles  
62 under uncontrolled conditions. The performance is compared with techniques as near-  
63 est neighbor (NN), support vector machine (SVM) and a reservoir computing time  
64 series kernel (RTK) We illustrate one of the main advantages of the proposed method  
65 when classifying odors based on short data sequences, providing the predictive clas-  
66 sification accuracy for sequences of reduced lengths (1s, 10s and 20s). Furthermore,  
67 we highlight the introduced relevance learning system for temporal high dimensional  
68 data, by studying the relevance of sensors and time points on the classification per-  
69 formance.

## 70 **2 Related works**

71 Odor discrimination with electronic noses has received growing attention and many  
72 studies have been done on how to classify odors using an array of gas sensors and

73 a pattern recognition algorithm. In [8, 11, 40, 18] the principal methods for chemical  
74 classification with an array of gas sensors are reviewed, including nearest neighbor  
75 (NN), mahalanobis linear discriminant analysis, neural networks (ANN), cluster analysis  
76 with self-organizing Maps (SOM) and Support Vector Machines (SVM).

77 More recently, approaches based on ensembles of classifiers have been reported  
78 to improve the classification accuracy [44], improve the *earliness* on the classifica-  
79 tion [20], or to deal with the common problems of sensor drift and sensor replace-  
80 ment. In [47], a SVM based ensemble of classifiers is used to solve the gas dis-  
81 crimination problem over a period of three years by training different classifiers at  
82 different points of time. Similarly, in [27], a flexible classification strategy based on  
83 cooperative classifiers is proposed to increase the robustness of chemo-sensory sys-  
84 tems against failures in their constituent sensing elements, postponing the necessity  
85 of replacing a sensor in the array, as well as facilitating the insertion of newly sensing  
86 elements.

87 Nevertheless, little attention has been given to the problem of classification in un-  
88 controlled conditions, as revealed by the few works found in literature that perform  
89 classification focusing only on the transient phase of the sensor signals. An evalua-  
90 tion for the suitability of different feature extraction techniques for such scenarios is  
91 provided in [45], where Trincavelli et al. propose a preprocessing stage to isolate the  
92 relevant parts of the sensor signals that can then be passed to the pattern recognition  
93 algorithm. More recently, in [12] a Support Vector Machine is applied to a set of fea-  
94 tures obtained from changes of the spectral sensor signal characteristics (frequency  
95 components, phase shift and energy sums), reporting a substantially increase of the  
96 classification performance.

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

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

110 A few machine learning techniques exist to investigate high dimensional time  
111 series: Topographic mappings such as the self-organizing map (SOM) (see [1] for a  
112 recent review) were extended by a recursive context which accounts for the temporal  
113 dynamics [43]. A probabilistic counterpart is provided by the Generative Topographic  
114 Mapping Through Time (GTM-TT) which combines hidden Markov models with a  
115 constraint mixture model induced by a low dimensional latent space. This approach is  
116 extended to better take the relevance of the feature components into account in [31],  
117 but relying on an unsupervised model. The identification of relevant dimensions is  
118 very important as outlined e.g. in [31, 23] to obtain a better understanding of the data,

119 to reduce the processing complexity, and to improve the overall prediction accuracy.  
120 A supervised relevance weighting scheme which singles out relevant features in a  
121 wrapper approach based on hidden Markov models has been proposed in [23]. In [7]  
122 a similar approach introducing class-wise constraints in the hidden Markov model  
123 is presented. In [23], applications to life science data are presented resulting in 85%  
124 prediction accuracy on a multiple sclerosis (MS) data set, but the approach makes  
125 multiple, restrictive assumptions regarding the used Hidden Markov Model (HMM).  
126 The approach [7] is evaluated in the same scenario with improved performance for  
127 the sclerosis data set. Ongoing work in the field reflects the high demand for effective  
128 methods for short but high dimensional time series data [33]. This is not limited to  
129 the bio-medical domain [23, 7] but covers a broader field of applications in industry  
130 and geo-science [31, 43]. In this work, we employ a supervised variant of GTM-TT  
131 (SGTM-TT) as introduced in [36] and extended in [37].

### 132 3 Method

#### 133 3.1 Generative Topographic Mapping

134 As outlined before the complexity of the considered data requests for a strong regu-  
135 larizing and interpretable model. Topographic maps appear to be a good choice and  
136 especially the Generative Topographic Mapping (GTM) combines multiple neces-  
137 sary features. GTM was first introduced in [2] and models a given set of data vectors  
138  $\mathbf{x} \in \mathbb{R}^D$  in form of a mapping based on a constrained mixture of Gaussians. The  
139 mixture is induced by a lattice of points  $\mathbf{w}$  in a low dimensional, so called, latent  
140 space which can also be used for visualization. The low dimensional lattice points  
141 are mapped by a projection  $\mathbf{w} \mapsto \mathbf{t} = y(\mathbf{w}, \mathbf{W})$  into the high-dimensional data  
142 space. The corresponding mapping function is parametrized by the parameters  $\mathbf{W}$ ;  
143 which usually are chosen in form of a generalized linear regression

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

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

149 Every grid point of the GTM induces a Gaussian

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

150 with variance  $\beta^{-1}$ . Assuming a Dirac distribution of the prototypes, the data are  
151 modeled by a mixture of  $K$  modes

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

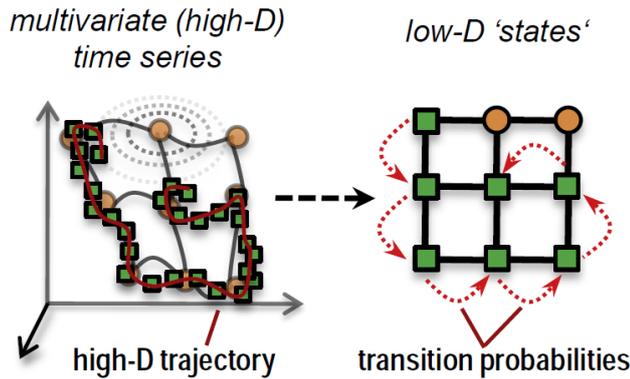


Fig. 1: GTM-TT consisting of a hidden Markov model, which hidden states are constrained to be organized on a grid topology (the latent points of the GTM model). The emission probabilities are governed by the GTM mixture distribution [2]. In the left figure a data distribution is given in a 3D space with 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 9 hidden states organized on a  $3 \times 3$  grid. The data distribution may change over time and hence also the mapping of the GTM is effected over time, assuming smooth transitions within the HMM.

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

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

154 by means of an expectation maximization (EM) strategy with respect to the model  
 155 parameters  $\mathbf{W}$  and  $\beta$  with data dimensionality  $D$  and number of points  $N$  as detailed  
 156 in [2].

157 Finally an unsupervised restricted Gaussian mixture model (GMM), induced by  
 158 a low dimensional latent space, is defined.

### 159 3.2 GTM Through-Time

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

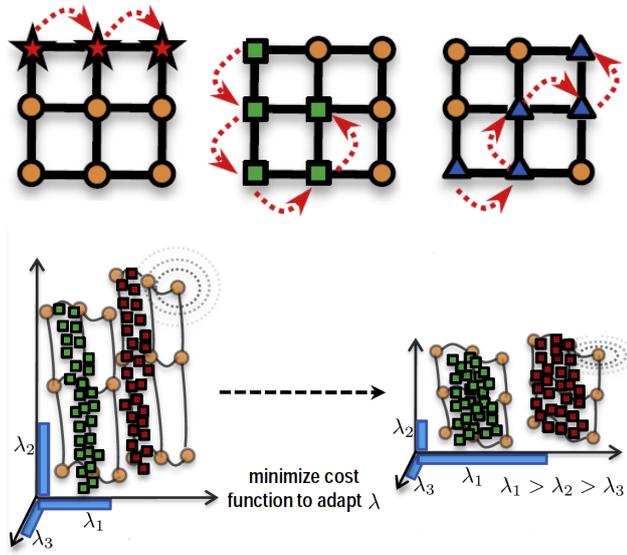


Fig. 2: Illustration of the SGTM-TT. It consists of multiple GTM-TT models. It behaves similar to the regular GTM-TT but the training is classwise and the  $\beta$  parameter is common over the different models. The different classwise models are used to represent the data distribution over time (here for three classes). In the bottom the SGTM-TT with relevance learning is shown. The relevance of the input-dimensions is weighted over time during training. And only relevant dimensions with large  $\lambda$ -values are kept. In the figure the  $\lambda_1$  dimension discriminates the two groups and is pronounced by metric adaptation.

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

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

177 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)  
 178 [2]. This results in the overall probability of  $\mathbf{x}$ :  $p(\mathbf{x}|\Theta) = \sum_{\mathbf{z} \in \{\mathbf{w}^1, \dots, \mathbf{w}^K\}^T} p(\mathbf{x}, \mathbf{z}|\Theta)$

179 For the parametrization of the GTM-TT ( $\Theta = (\mathbf{W}, \beta, \pi, \mathbf{P})$ ) we rely on the as-  
 180 sumption of the standard Markov property and stationarity of the dynamics. With  
 181 additional parameters for the initial state probabilities  $\pi = (\pi_j)_{j=1}^K$  where  $\pi_j =$   
 182  $p(\mathbf{z}(1) = \mathbf{w}^j)$  and transition probabilities  $\mathbf{P} = (p_{ij})_{i,j=1}^K$  where  $p_{ij} = p(\mathbf{z}(t) =$   
 183  $\mathbf{w}^j | \mathbf{z}(t-1) = \mathbf{w}^i)$ , the latter one characterizing the temporal correlations of subse-  
 184 quent states. The data log likelihood is optimized by:  $\ln \left( \prod_{n=1}^N p(\mathbf{x}^n | \Theta) \right)$ . using an  
 185 EM-approach. Like for standard HMMs the hidden parameters (responsibilities) are  
 186 defined by a forward-backward procedure [48]. Based on these parameters  $\mathbf{W}$  and  $\beta$   
 187 can be determined as specified before. The probability of being in state  $\mathbf{w}^k$  at time  $t$ ,  
 188 given the observation sequence  $\mathbf{x}^n$  (responsibilities) is given as:

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

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

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

191 we get the forward variable  $A_{kt}$  with the start condition  $A_{k1} = \pi_k p(\mathbf{x}^n(1) | \mathbf{w}^k, \Theta)$ .  
 192 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  
 193 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  $B_{kT} = 1$ ,  $B_{kt}$   
 194 defines the backward variable. The transition parameters are trained using the stan-  
 195 dard Baum-Welch training. As usual the underlying HMM also permits to deal with  
 196 missing values and sequences of arbitrary length [3]). A more detailed description of  
 197 the GTM-TT is given in [42].

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

### 201 3.3 Supervised GTM-TT

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

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

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

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

### 219 3.4 Relevance learning for SGTM -TT

220 Metric adaptation for discriminative prototype based learning has been introduced in  
 221 [19], it is often also referred to as relevance learning. The basic idea is to parametrize  
 222 the distance measure to incorporate auxiliary information. For the squared Euclidean  
 223 metric one can define a parametrized, weighted, variant:

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

224 For the GTM such a parametrization was already discussed in [15] for i.i.d. data  
 225 resulting in relevance GTM (R-GTM). However, having temporal data some adapta-  
 226 tions are necessary and also the supervision has to be handled in an alternative way.  
 227 To keep the approach simple and to limit the number of free parameters we will re-  
 228 strict our approach to a global diagonal weighted distance, in which case a weight  $\lambda_i$   
 229 directly corresponds to the relevance of dimension  $i$ . Here we assume normalized data  
 230 with mean 0 and a standard deviation of 1 for each dimension. For GTM-(TT), the  
 231 distance used to compute local probabilities is replaced by the previously discussed  
 232 weighted Euclidean distance:

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

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

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

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

242 The discriminative learning of the metric parameters is controlled by the cost  
 243 function of the generalized learning vector quantization (GRLVQ) which is a large  
 244 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  $\lambda$  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  $\lambda$  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

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

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 alternatives 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]:

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

281 with

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

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

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

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

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

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

295 *Relevant time points:*

296 Since SGTM-TT relies on HMMs, every time point depends on its predecessor only.  
 297 Thus, it is not reasonable to adapt the relevance of time points to obtain a better rep-  
 298 resentation of data in the GTM-TT models. However, it is reasonable to judge the  
 299 relevance of time points resulting from the GTM-TT models for the final classifica-  
 300 tion, in particular if time series are of the same or a similar length. This method offers  
 301 insights into the model to identify time points which are particularly discriminative  
 302 for the given task at hand.

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

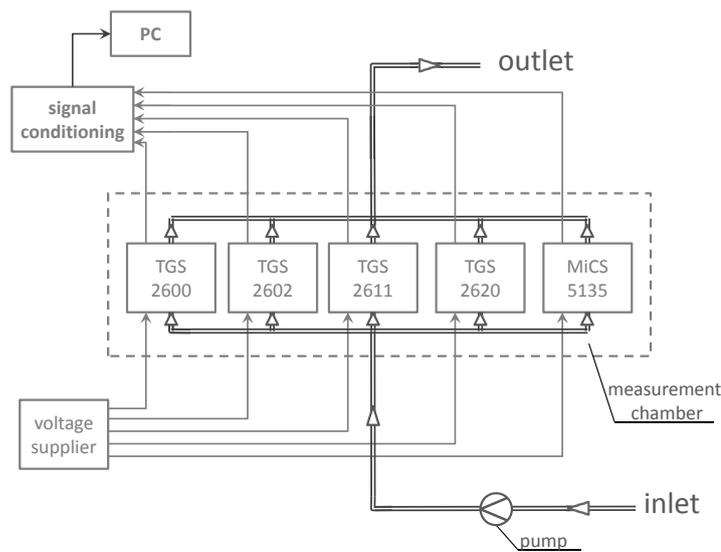


Fig. 3: Measurement system.

#### 310 4 Odor Measurement System

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

317 The election of MOX as the gas sensing technology has been made attending at its  
318 high sensitivity, commercial availability and low price. However, they present some  
319 shortcomings including: poor selectivity, influence by environmental factors such as  
320 humidity and temperature [30] and major limitations in their response speed [29].  
321 Among these drawbacks, their poor selectivity is of the largest concern for odor clas-  
322 sification. To overcome this, it is a common practice to build the e-nose upon an array  
323 of MOX sensors with different and partially overlapping sensitivities. The output of  
324 the array is then processed with a pattern recognition algorithm to find out which  
325 substance the e-nose is exposed to. Based in this concept, we choose five different  
326 MOX gas sensors to compose the sensor array: TGS-2600, TGS-2602, TGS-2611 and  
327 TGS-2620 from Figaro Sensors<sup>1</sup>, and MiCS-5135 from e2V Sensors<sup>2</sup>.

328 In order to enable sensors to interact with the volatile molecules dispersed in the  
329 environment, the e-nose employs a pump to enforce a constant airflow through the  
330 sensors array. The aspiration and release of the air samples are accomplished through  
331 tubes, conveniently separated one to another to avoid cross contamination. Addition-

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

<sup>2</sup> e2v. <http://www.e2v.com/>.

332 ally, the aspiration through flexible tubes allows the displacement of the aspiration  
333 entry without the need to move the complete system. This advantage is particularly  
334 useful in robotics to easily sample the space, for example, by attaching the e-nose  
335 aspiration to the hand of an arm robot as shown in the experimental section.

#### 336 4.1 Signal conditioning and data preprocessing

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

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

### 345 5 Experimental Results

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

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

364 For SVM we used a linear kernel with optimal  $C$  determined over the training  
365 data on a grid search. Since SVM can not directly be applied to temporal data, nor  
366 can it be used for sequences of different length in a direct way, for the comparison we  
367 simply concatenate the measurements of the different channels to remove the time  
368 dimension. More complex strategies of applying SVM, e.g. by using a dynamic time  
369 warping (DTW) kernel could be done but are not in the focus of this paper and out of  
370 sample extensions are often not immediate which is an issue for online robotic sensor  
371 systems. For more recent work around DTW or kernel related time series analysis

372 see e.g. [32,28,6]. Additionally we would like to avoid more complex preprocessing  
373 steps to permit an easy out of sample extension in practical settings. Although max-  
374 imum classification performance is not our main objective, we also provide a compar-  
375 ison with a very recent reservoir computing kernel [6]. This approach is known to  
376 be very effective for timeseries but on the other hand is less interpretable nor is the  
377 out of sample extension for very short sequences immediate. For RTK there are three  
378 parameters optimized on the training data within a grid search<sup>3</sup> as detailed in [6].

## 379 5.1 Experiment 1: Simulated data

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

382 The simulated data (SIM) consist of 100 samples separated into two classes of  
383 50 samples each. Each point is located in a 100 dimensional feature space with 8  
384 time points. From the given features, only 10 are expected to differentiate between  
385 the classes. Details about the data and the generation procedure are given in [23].

386 We applied SGTm-TT with relevance learning using 9 hidden states and 4 ba-  
387 sis functions. We observe an overall prediction accuracy of  $94 \pm 4\%$ . The relevance  
388 profile identified all known 10 features and effectively pruned out the remaining ir-  
389 relevant data dimensions. Our results are slightly better than those reported in [23]  
390 (90%) and in [7] (92%).

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

## 397 5.2 Experiment 2: Controlled gas exposure

398 The second experiment aims to test the proposed method with real odor data un-  
399 der restrained environmental conditions. To this end, a dataset of real odor samples  
400 is gathered in a scenario as controlled as possible. The dataset is comprised by 39  
401 samples generated by exposing the e-nose to gas pulses of four different analytes: a  
402 commercial spirit (*Larios Gin*), a polish remover based on Acetone, standard ethanol  
403 and lighter gas (butane mixed with propane). Acetone was given by 9 samples and  
404 the other classes by 10 samples each.

405 Each sample is collected according to the following three-phases procedure: (1)  
406 for the initial 30s, baseline value is estimated by measuring the sensor response in  
407 absence of the target gas, (2) then, for a duration of 60s the e-nose is placed next  
408 to the gas source (about 10cm) exposing the sensor array to the volatile. Finally (3),

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

409 the gas source is removed allowing the sensor array to recover to its initial state  
410 (baseline).

411 Figure 4 shows two different samples of such dataset. Notice that although the  
412 gas exposure was "controlled" by time exposure and distance to the source, strong  
413 fluctuations in the sensor readings occur due to the chaotic nature of the gas disper-  
414 sion.

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

420 For comparison we also use two public domain data sets of similar type (elec-  
421 tronic nose data) from the UCI database. The DS-UCI-1 data set is given by the  
422 *two sources gas* data [14]. The data are measured using a chemical detection plat-  
423 form composed of 8 chemo-resistive gas sensors which were exposed to turbulent  
424 gas mixtures generated naturally in a wind tunnel. It consists of 180 time series of  
425 Ethylene (Eth), Carbon Monoxide (CO) and Methane (Me) mixtures at different con-  
426 centrations. We use the data as a two class prediction problem to predict the whether  
427 Eth was mixed with CO or Me. Available features are temperature, humidity and the  
428 8 sensor channel outputs. Each time series is given with 2970 sampling points.

429 The DS-UCI-2 data set is given by the *pulmon* data [49]. The data are measured  
430 using a chemical sensing system based on an array of 16 metal-oxide gas sensors and  
431 an external mechanical ventilator to simulate the biological respiration cycle. The  
432 tested gas classes are mixtures of acetone and ethanol. Data have been normalized  
433 to zero-mean and intensity and considered again as a prediction problem to identify  
434 whether the mixture contains Me or CO.

435 The classification accuracy for DS1, DS-UCI-1, DS-UCI-2 is given in Table 1 in  
436 comparison to some standard approaches. We observe that the SGTM-TT performs

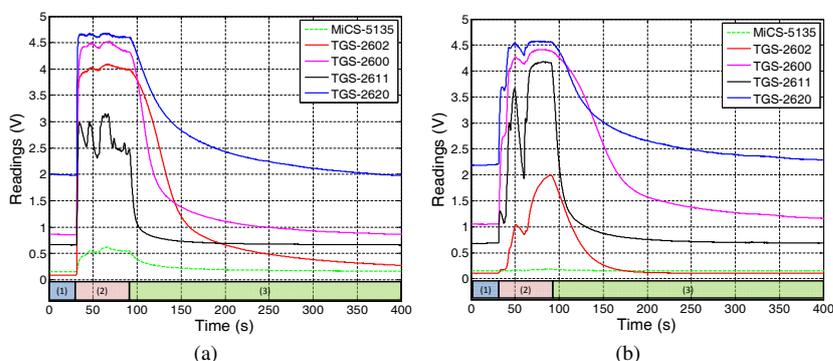


Fig. 4: Two different samples of the olfaction dataset gathered in the second experi-  
ment. The three phases in which the samples can be decomposed are marked at the  
bottom of each figure as (1),(2) and (3).

CV-Accuracy	SGTM-TT	SVM	NN	RTK
SIM	<u>94.00</u> ± 4.18%	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%	<u>96.67</u> ± 4.56%
DS-UCI-1	87.78 ± 5.76%	<u>93.89</u> ± 4.97%	86.81 ± 7.98%	64.44 ± 4.12%
DS-UCI-2	79.55 ± 9.15%	83.03 ± 18.47%	76.33 ± 18.15%	<u>94.70</u> ± 8.05%

Table 1: Average test set accuracy for the first and second experiment in a 5 fold cross-validation. Significant better results are underlined.

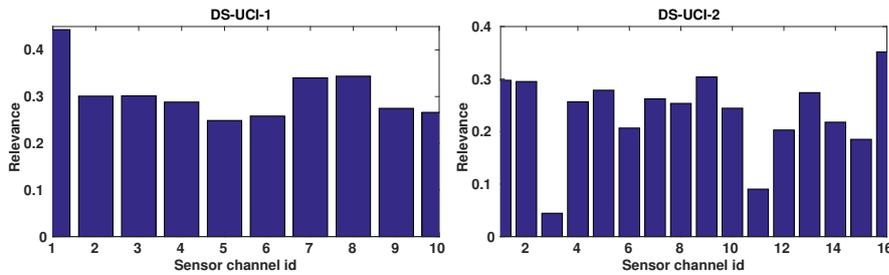


Fig. 5: Relevance profile of the sensor input for DS-UCI-1 (left) and DS-UCI-2 (right). For both profiles the information is distributed over the various sensors but some sensors are more important e.g. sensor 1 for DS-UCI-1 and sensor 3 for DS-UCI-2.

437 reliable well although the best prediction accuracy for DS1 and DS-UCI-2 is ob-  
 438 tained by the RTK approach. For the DS-UCI-1 dataset RTK is significantly worse  
 439 than the other approaches and the SVM obtained the best performance. Hence there  
 440 is not a clear winner regarding the classification accuracy but SGTM-TT represents  
 441 a good approach with a reliable and consistent performance. Furthermore, as previ-  
 442 ously commented, the classification performance is not the only point that matters but  
 443 also the simplicity of the model and the interpretability of the results. Neither RTK  
 444 nor SVM provide additional insight in the relevance of the sensor channels<sup>4</sup>. Here we  
 445 are mainly interested on interpretable models [24] which also simplify a later transfer  
 446 of the approach to an embedded system or the sensor platform. In Figure 5 we show  
 447 the averaged (global) sensor relevance profile of DS-UCI-1 and DS-UCI-2.

448 Subsequently we give a detailed analysis for our own dataset - DS1, where we have  
 449 more background information to provide a specific in depth discussion of the results.  
 450 For the analysis of the sensor relevance and time points relevance, the whole measure-  
 451 ment sequence of each sample was down-sampled to 800 time points each (DS2). The  
 452 SGTM-TT was then trained in a 5-fold crossvalidation with 4 hidden states and 4 ba-  
 453 sis functions. In Figure 6 we show the relevance indexes of the five gas sensors of the  
 454 e-nose for the different target volatiles of DS2 as obtained by SGTM-TT. Different  
 455 conclusions can be drawn from the study of such relevance plot:

<sup>4</sup> 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.

- 456 – In general, the five MOX sensors are relevant for the classification of the differ-  
457 ent volatiles, being sensor *TGS – 2620* the less relevant one, and so the most  
458 expendable.  
459 – Sensor *TGS – 2602* is the most relevant one when classifying Acetone and  
460 Ethanol samples, with a notable difference with respect the other sensors in the  
461 case of Acetone. This characteristic is already reported in the manufacturer's  
462 datasheet, indicating the high sensitivity to volatile organic compounds (VOCS)  
463 of this sensor model.  
464 – As expected from the low selectivity characteristic of MOX sensors, each sensor  
465 presents a high relevance index for more than one odor class.

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

471 Since in real robotics conditions the classifier is expected to work on small data  
472 sequences, a second configuration for the dataset DS2 was tested. Here, the test data  
473 consist only of short sensor readings over time. Figure 7(b) depicts the accuracy in the  
474 classification for three different window lengths (1s, 10s and 20s). We observe that  
475 given the highly dynamic response of MOX sensors in addition to the inherent sig-  
476 nal noise, very small windows (1s) do not carry enough information for a reasonable  
477 classification, but for data sequences of ten seconds the accuracy in the prediction

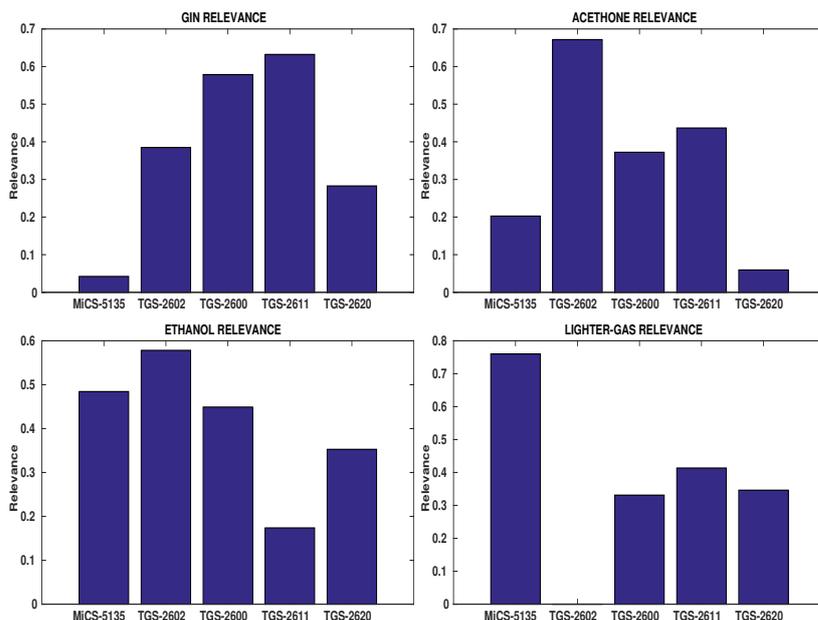


Fig. 6: Sensor relevance indexes for the four odor classes used on dataset DS2.

478 achieves very good results (values near 0.8). Furthermore, window lengths over ten  
479 seconds seems to not improve the accuracy, which indicates that long sequences en-  
480 code a lot of noise contributions, hampering the model in the prediction. Finally, it  
481 must be noticed that the classification accuracy is usually higher when using data  
482 from the transient parts of the signal (rise and decay) than when steady state data is  
483 employed, as denoted by the accuracy peaks found around seconds 30 and 90.

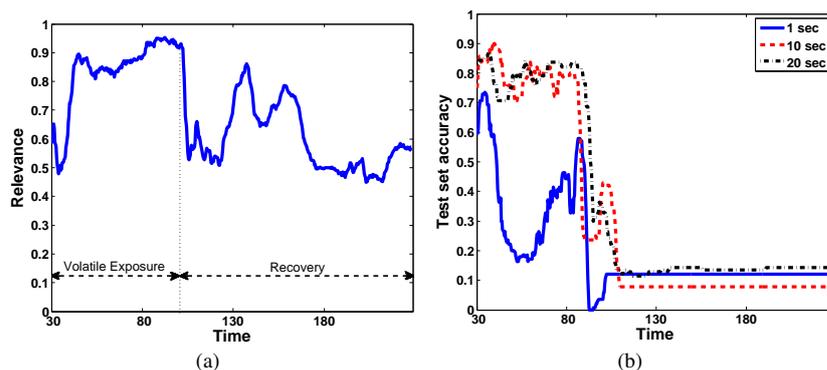


Fig. 7: Time points relevance profile 7(a) averaged over all classes and mean prediction accuracy over time with window length's of  $\approx 1, 10, 20sec$  7(b).



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.

### 484 5.3 Experiment 3: Robotics experiment - uncontrolled gas pulses

485 Finally, and with the aim to validate the classification performance in a more chal-  
486 lenging robotic scenario, a third experiment is presented. In this case, the e-nose  
487 aspiration (see Figure 8) is attached to the hand of a robotic arm [34] which is com-  
488 manded to approximate the e-nose aspiration to each of four recipients containing  
489 different substances (Acetone, Ethanol, Butane<sup>5</sup> and Gin).

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

494 The robotic arm is commanded to approximate to the containers following a pre-  
495 defined sequence. The exposition to each of the substances takes 20sec, after which  
496 the arm moves to another container. The volatile sequence and the gathered signals  
497 during the experiment are depicted in Figure 9. A video of a similar experiment  
498 is additionally available at [http://mapir.isa.uma.es/mapirwebsite/](http://mapir.isa.uma.es/mapirwebsite/index.php/2008-tep-4016-media)  
499 [index.php/2008-tep-4016-media](http://mapir.isa.uma.es/mapirwebsite/index.php/2008-tep-4016-media)

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

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

## 514 6 Conclusion

515 A novel approach for the analysis of high dimensional and rather short temporal se-  
516 quences was presented. It is based on the idea to introduce available meta information  
517 into the modeling process of a Generative Topographic Mapping through time, given  
518 in form of supervised information and relevance learning. We have analyzed the suit-  
519 ability of such model for the odor classification problem in robotics applications,  
520 providing comparative results with support vector machine (SVM), nearest neighbor  
521 (NN) and the reservoir time series kernel (RTK) for three different scenarios (with

<sup>5</sup> 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.

<sup>6</sup> Here we simply used the model from the first crossvalidation run.

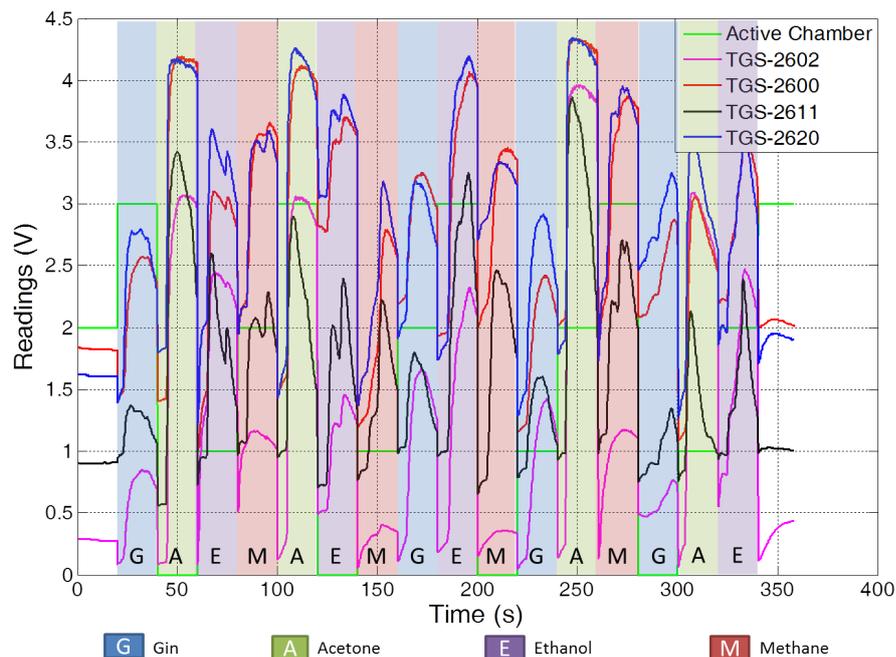


Fig. 9: MCE-nose gathered signals of the classification experiment with a robotic arm, and the "ground-truth" sequence of the employed analytes. The active chamber [0,1,2,3] is switched every 20 sec. Signals are shown for the 4 different sensor channels as described before.

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										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	o

Table 2: Predictions for the external evaluation data using the first respective cross-validation model. The 'o' in the line labeled with *Error* indicates mismatches.

522 increasing classification challenge), and demonstrating that the proposed method is  
 523 effective for solving such highly dimensional data problem.

524 Other remarkable advantages of the method in the context of odor classification  
 525 in robotics are: on the one hand, the possibility for the robot to perform rapid classifi-  
 526 cation of chemical substances by using a short data sequence. On the other hand, the  
 527 SGTM-TT method outputs relevance values for both the sensors being used as well

528 as the time-points of the signal, which provide very valuable information to configure  
529 the e-nose and to carry out the robot smelling.

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

### 533 Acknowledgment

534 The first author likes to thank: Tien-ho Lin, Carnegie Mellon University, USA for  
535 support with the simulation data. Further, Ivan Olier, University of Manchester, UK;  
536 Iain Strachan, AEA Technology, Harwell, UK and Markus Svensen, Microsoft Re-  
537 search, Cambridge, UK for providing code for GTM and GTM-TT. Further we would  
538 like to thank Fengzhen Tang, University of Birmingham for providing invaluable sup-  
539 port with the RTK method.

540 *Funding:* This work was supported by the DFG project HA2719/4-1 to BH, by  
541 the DFG-NSF project TO 409/8-1, and by the Cluster of Excellence 277 CITEC  
542 funded in the framework of the German Excellence Initiative. Further, a Marie Curie  
543 Intra-European Fellowship (IEF): FP7-PEOPLE-2012-IEF (FP7-327791-ProMoS) is  
544 gratefully acknowledged. Additional support was provided by funds from the An-  
545 dalucía Regional Government and the European Union (FEDER) under research  
546 project: TEP08-4016.

### 547 References

- 548 1. C.A. Astudillo and B.J. Oommen. Topology-oriented self-organizing maps: a survey. *Pattern Analysis*  
549 *and Applications*, pages 1–26, 2014.
- 550 2. Christopher M. Bishop, Markus Svensén, and Christopher K. I. Williams. Gtm: The generative topo-  
551 graphic mapping. *Neural Computation*, 10(1):215–234, 1998.
- 552 3. C.M. Bishop. *Pattern recognition and machine learning*. Information science and statistics. Springer,  
553 2006.
- 554 4. José-Luis Blanco, Javier G. Monroy, Javier González-Jiménez, and Achim Lilienthal. A kalman filter  
555 based approach to probabilistic gas distribution mapping. In *28th Symposium On Applied Computing*  
556 *(SAC)*, mar 2013.
- 557 5. Sofiane Brahim-Belhouari and Amine Bermak. Gas identification using density models. *Pattern*  
558 *Recognition Letters*, 26(6):699 – 706, 2005.
- 559 6. H. Chen, F. Tang, P. Tino, and X. Yao. Model-based kernel for efficient time series analysis. In  
560 *Proceedings of the 19th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*  
561 *(KDD'13)*, Chicago, USA, 2013.
- 562 7. Ivan G. Costa, Alexander Schönhuth, Christoph Hafemeister, and Alexander Schliep. Constrained  
563 mixture estimation for analysis and robust classification of clinical time series. *Bioinformatics*, 25(12),  
564 2009.
- 565 8. C. Distanto, P. Siciliano, and K.C. Persaud. Dynamic cluster recognition with multiple self-organising  
566 maps. *Pattern Analysis and Applications*, 5(3):306–315, 2002.
- 567 9. Cosimo Distanto, Marco Leo, Pietro Siciliano, and Krishna C. Persaud. On the study of feature  
568 extraction methods for an electronic nose. *Sensors and Actuators B: Chemical*, 87(2):274 – 288,  
569 2002.
- 570 10. Ahlame Douzal-Chouakria and C. Amblard. Classification trees for time series. *Pattern Recognition*,  
571 45(3):1076 – 1091, 2012.

- 572 11. D. Dumitrescu, B. Lazzarini, and F. Marcelloni. A fuzzy hierarchical classification system for olfac-  
573 tory signals. *Pattern Analysis and Applications*, 3(4):325–334, 2000.
- 574 12. Bernd Ehret, Konstantin Safenreiter, Frank Lorenz, and Joachim Biermann. A new feature extraction  
575 method for odour classification. *Sensors and Actuators B: Chemical*, 158(1):75 – 88, 2011.
- 576 13. Gabriele Ferri, Emanuele Caselli, Virgilio Mattoli, Alessio Mondini, Barbara Mazzolai, and Paolo  
577 Dario. Spiral: A novel biologically-inspired algorithm for gas/odor source localization in an indoor  
578 environment with no strong airflow. *Robotics and Autonomous Systems*, 57(4):393 – 402, 2009.
- 579 14. Jordi Fonollosa, Irene Rodriguez-Lujn, Marco Trincavelli, Alexander Vergara, and Ramn  
580 Huerta. Chemical discrimination in turbulent gas mixtures with mox sensors validated by gas  
581 chromatography-mass spectrometry. *Sensors*, 14(10):19336–19353, 2014.
- 582 15. A. Gisbrecht and B. Hammer. Relevance learning in generative topographic mapping. *Neurocomput-*  
583 *ing*, 74(9):1359–1371, 2011.
- 584 16. Javier González-Jiménez, Javier G. Monroy, and José-Luis Blanco. Robots that can smell: motivation  
585 and problems. In *Submitted to 15th International Symposium On Olfaction and Electronic Nose*  
586 *(ISOEN)*, 2013.
- 587 17. Javier Gonzalez-Jimenez, Javier G. Monroy, and Jose Luis Blanco. The multi-chamber electronic  
588 nose. an improved olfaction sensor for mobile robotics. *Sensors*, 11(6):6145–6164, 2011.
- 589 18. R. Gutierrez-Osuna. Pattern analysis for machine olfaction: a review. *Sensors Journal, IEEE*, 2(3):189  
590 – 202, jun 2002.
- 591 19. B. Hammer and Th. Villmann. Generalized relevance learning vector quantization. *Neural Networks*,  
592 15(8-9):1059–1068, 2002.
- 593 20. Nima Hatami and Camelia Chira. Classifiers with a reject option for early time-series classification.  
594 In *Proceedings of the IEEE Symposium on Computational Intelligence and Ensemble Learning, CIEL*  
595 *2013, IEEE Symposium Series on Computational Intelligence (SSCI), 16-19 April 2013, Singapore*,  
596 pages 9–16. IEEE, 2013.
- 597 21. J. Lee and M. Verleysen. Generalizations of the lp norm for time series and its application to self-  
598 organizing maps. In Marie Cottrell, editor, *5th Workshop on Self-Organizing Maps*, volume 1, pages  
599 733–740, 2005.
- 600 22. A.J. Lilienthal, M. Reggente, M. Trincavelli, J.-L. Blanco, and J. Gonzalez. A statistical approach to  
601 gas distribution modelling with mobile robots - the kernel dm+v algorithm. In *Intelligent Robots and*  
602 *Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 570 –576, 2009.
- 603 23. Tien-ho Lin, Naftali Kaminski, and Ziv Bar-Joseph. Alignment and classification of time series gene  
604 expression in clinical studies. In *ISMB*, pages 147–155, 2008.
- 605 24. Paulo J. G. Lisboa. Interpretability in machine learning - principles and practice. In Francesco Ma-  
606 sulli, Gabriella Pasi, and Ronald R. Yager, editors, *Fuzzy Logic and Applications - 10th International*  
607 *Workshop, WILF 2013, Genoa, Italy, November 19-22, 2013. Proceedings*, volume 8256 of *Lecture*  
608 *Notes in Computer Science*, pages 15–21. Springer, 2013.
- 609 25. Amy Loutfi, Silvia Coradeschi, Achim J. Lilienthal, and Javier Gonzalez. Gas distribution mapping  
610 of multiple odour sources using a mobile robot. *Robotica*, 27(2):311–319, June 4 2009.
- 611 26. L. Marques, U. Nunes, and A. T. de Almeida. Olfaction-based mobile robot navigation. *Thin Solid*  
612 *Films*, 418(1):51 – 58, 2002. Proceedings from the International School on Gas Sensors in conjunction  
613 with the 3rd European School of the NOSE Network.
- 614 27. Eugenio Martinelli, Gabriele Magna, Alexander Vergara, and Corrado Di Natale. Cooperative classi-  
615 fiers for reconfigurable sensor arrays. *Sensors and Actuators B: Chemical*, 199(0):83 – 92, 2014.
- 616 28. M. Michalak. Adaptive kernel approach to the time series prediction. *Pattern Analysis and Applica-*  
617 *tions*, 14(3):283–293, 2011.
- 618 29. Javier G. Monroy, Javier Gonzlez-Jimnez, and Jose Luis Blanco. Overcoming the slow recovery of  
619 mox gas sensors through a system modeling approach. *Sensors*, 12(10):13664–13680, 2012.
- 620 30. P. Moseley and B. Tofield, editors. *Solid state gas sensors*. Bristol, UK: Adam Hilger, 1987.
- 621 31. Iván Olier and Alfredo Vellido. Advances in clustering and visualization of time series using gtm  
622 through time. *Neural Networks*, 21(7):904–913, 2008.
- 623 32. F. Petitjean, A. Ketterlin, and P. Ganarski. A global averaging method for dynamic time warping, with  
624 applications to clustering. *Pattern Recognition*, 44(3):678–693, 2011.
- 625 33. O.J. Prieto, C.J. Alonso-Gonzlez, and J.J. Rodriguez. Stacking for multivariate time series classifica-  
626 tion. *Pattern Analysis and Applications*, pages 1–16, 2013.
- 627 34. Kinova Robotics. The jaco research edition robotic arm. <http://www.kinovarobotics.com/>.
- 628 35. A.C. Sauve and T.P. Speed. Normalization, baseline correction and alignment of high-throughput  
629 mass spectrometry data. In *Proc. Gensips*, 2004.

- 630 36. F.-M. Schleif, A. Gisbrecht, and B. Hammer. Relevance learning for short high-dimensional time  
631 series in the life sciences. In *Proceedings of IJCNN 2012*, pages 2069–2076, 2012.
- 632 37. F.-M. Schleif, B. Mokbel, A. Gisbrecht, L. Theunissen, V. Dürr, and B. Hammer. Learning relevant  
633 time points for time-series data in the life sciences. In *Proceedings of ICANN 2012*, pages 531–539,  
634 2012.
- 635 38. F.-M. Schleif, F.-M. Ongyerth, and T. Villmann. Supervised data analysis and reliability estimation  
636 for spectral data. *NeuroComputing*, 72(16-18):3590–3601, 2009.
- 637 39. P. Schneider, M. Biehl, and B. Hammer. Distance learning in discriminative vector quantization.  
638 *Neural Computation*, 21:2942–2969, 2009.
- 639 40. Ronald E. Shaffer, Susan L. Rose-Pehrsson, and R. Andrew McGill. A comparison study of chemical  
640 sensor array pattern recognition algorithms. *Analytica Chimica Acta*, 384(3):305 – 317, 1999.
- 641 41. R.H. Shumway and D.S. Stoffer. *Time Series Analysis and Its Applications*. Springer Texts In Statis-  
642 tics. Springer, 2000.
- 643 42. I. G. D. Strachan. *Latent Variable Methods for Visualization Through Time*. PhD thesis, University of  
644 Edinburgh, Edinburgh, UK, 2002.
- 645 43. M. Strickert and B. Hammer. Merge SOM for temporal data. *Neurocomputing*, 64:39–72, 2005.
- 646 44. Andrzej Szczurek, Bartosz Krawczyk, and Monika Maciejewska. {VOCs} classification based on the  
647 committee of classifiers coupled with single sensor signals. *Chemometrics and Intelligent Laboratory  
648 Systems*, 125(0):1 – 10, 2013.
- 649 45. Marco Trincavelli, Silvia Coradeschi, and Amy Loutfi. Odour classification system for continuous  
650 monitoring applications. *Sensors and Actuators B: Chemical*, 58:265 – 273, 2009.
- 651 46. Alexander Vergara, Jordi Fonollosa, Jonas Mahiques, Marco Trincavelli, Nikolai Rulkov, and Ramn  
652 Huerta. On the performance of gas sensor arrays in open sampling systems using inhibitory support  
653 vector machines. *Sensors and Actuators B: Chemical*, 185(0):462 – 477, 2013.
- 654 47. Alexander Vergara, Shankar Vembu, Tuba Ayhan, Margaret A. Ryan, Margie L. Homer, and Ramon  
655 Huerta. Chemical gas sensor drift compensation using classifier ensembles. *Sensors and Actuators  
656 B: Chemical*, 166-167(0):320 – 329, 2012.
- 657 48. Lloyd R. Welch. Hidden Markov Models and the Baum-Welch Algorithm. *IEEE Information Theory  
658 Society Newsletter*, 53(4), December 2003.
- 659 49. Andrey Ziyatdinov, Jordi Fonollosa, Luis Fernndez, Agustn Gutierrez-Glvez, Santiago Marco, and  
660 Alexandre Perera. Bioinspired early detection through gas flow modulation in chemo-sensory sys-  
661 tems. *Sensors and Actuators B: Chemical*, 206(0):538 – 547, 2015.