

On Mappings Between Electronic Noses

Oded Shaham Liran Carmel David Harel*
Dept. of Computer Science and Applied Mathematics
The Weizmann Institute of Science, Rehovot, Israel
{odedsh,liran,dharel}@wisdom.weizmann.ac.il

June 24, 2003

Abstract

We consider the task of finding a mapping between two eNoses that employ two different sensor technologies, quartz microbalance and conducting polymers. Such a mapping is a model that predicts the response of one eNose based on the response of the other. eNose mappings are important for odor communication and synthesis, as well as for eNose data integration. We investigated a number of methods for performing this task, including principal components regression, partial least squares, neural networks and tessellation-based linear interpolation. Our measure of success is the percentage of predictions that are correctly-classifiable. Using two different techniques for splitting our data set, we achieved success rates of 67% and 100%.

1 Introduction

Our group recently introduced a scheme that will enable *odor communication* [1]. Electronic noses (eNoses) play a fundamental role in the scheme, at the basis of which lies the concept of an *odor space*. An odor space is the collection of all possible response patterns of a nose, be it the human nose or an eNose.

The odor communication system described in [1] requires the realization of an algorithm — the *mix-to-mimic (MTM)* algorithm — which is as yet not fully realizable. One of its core components is an algorithm capable of mapping from the space of an eNose into the space of the human perception, termed the *psychophysical space*.

While eNose responses can be directly measured by applying chemical samples to them, the way to obtain measurements from the psychophysical space is to conduct large-scale human panel experiments. Due to the difficulty in carrying out these experiments and directly developing MTM, the suggestion in [1] is to construct three sub-algorithms, each adding a further complication. Upon the completion of the third, the full MTM will be available. The second of these sub-algorithms — the MTM2 algorithm — involves the construction of a mapping between two different eNoses.

In this paper, we discuss our work on constructing a mapping between two eNoses that employ two very different sensor technologies, quartz microbalance and conducting polymers. The technology differences are important, as, in a way, they represent the differences between eNoses and the human nose.

Mappings between eNoses are also important for purposes other than odor communication. Using such mappings we would be able to integrate response patterns of different eNoses into a unified

*corresponding author

database [4]. This would allow one to combine data from different types of eNoses, to maintain continuity when replacing sensor modules, and to overcome drift effects due to sensor aging.

It appears that there is almost no work done on eNose mappings. One example we found is [4], which describes a linear mapping between two eNoses that share the same sensor technology. We have not been able to find any research on mappings between eNoses with different sensor technologies.

2 Experimental

At the center of our laboratory setup are two eNoses: MOSESII ([6, 7]) with an eight-sensors quartz microbalance (*QMB*) module, and Cyranose 320 ([8, 9]) which is a palm-top device that contains 32 conducting polymer (*CP*) sensors. *CP* and *QMB* are described in depth in [10].

The sample loop in our experiments consists of 3 components connected in a row: the HP7694 headspace sampler, followed by MOSESII, followed by Cyranose 320. The duration of sample exposure that the headspace sampler dictated, approximately 20 seconds, is not long enough to reach saturation at the sensors, thus we measure the transient response. Between sample injections, the sample loop is purged for 20 minutes with synthetic air.

The MOSESII eNose is designed for connecting to the HP7694 headspace sampler. In contrast, the connection of MOSESII to Cyranose 320 had to be improvised, and our solution is shown in Figure 1. The pressure in the headspace sampler was set to 1.52 bars, which is higher than the regular pressure of 0.96-0.97 bars that we have regularly used when measuring with MOSESII alone. However, increasing the pressure was necessary in order to achieve sufficiently strong responses of the *CP* sensors of Cyranose 320.



Figure 1: The connection between MOSESII and Cyranose 320. The out-pipe of the *QMB* module is connected to the 2" needle at the sample inlet of Cyranose 320. This connection is mediated by a tailor made non-sealing adapter, which directs part of the sample into Cyranose 320, and allows the rest to evaporate in the room.

Recording the sensor readings was done differently in MOSESII and in Cyranose 320. The ability of MOSESII to synchronize with the headspace sampler allowed us to start the recording in MOSESII directly at the signal start. In Cyranose 320, which could not be synchronized with the headspace sampler, we had to continuously record the entire batch of measurements. The number of vials in a batch was up to 21, making a batch recording last as long as seven hours. We later detected the *CP* signals within the continuous recordings using a sliding window algorithm. The time interval between consecutive readings was 1.2 seconds in *QMB*, and one second in *CP*.

The list of samples for our measurements (Table 1) includes organic chemicals with varied chemical properties, with the intention to obtain measurements from a wide portion of the eNose spaces. Each chemical was measured in at least seven consecutive measurements (one vial per measurement).

List of Chemicals	
(1s)-(-)-alpha pinene	(1s)-(-)-beta pinene
(r)-(+)-limonene	(s)-(-)-limonene
1-propanol	2,6-dimethylpyridine
2-acetylpyridine	4-methylanisole
allyl hexanoate	butyl butyrate
cis-3-hexen-1-ol	cis-3-hexenyl acetate
decyl acetate	dipentene
ethyl valerate	heptyl alcohol
hexyl alcohol	isoamyl propionate
n-hexane	n-octane
pentyl acetate	piperidine
terpiolene	

Table 1: The 23 chemicals in our data set.

3 Methods

3.1 Feature Extraction

Prior to feature extraction, we gathered the raw data of both QMB and CP into a uniform database, containing an entry per measurement. In addition to the sensors’ raw data, each entry holds the sensors’ baseline, which was calculated in advance in order to be used during the feature extraction step.

When choosing features to extract, the main objective is to represent the raw data as well as possible, using as few features as possible. A central aspect of representation is the ability to discriminate between different odorants. We considered two features: *height*, which is the difference between the signal maximum and its baseline, and *relative height*, which is the height divided by the baseline. Both features are widely used in eNoses for a variety of applications, and are simple to extract. For example, Our group had successfully used the height feature of QMB for tasks that include classification and concentration prediction [3]. For Cyranose 320, It is suggested by the designers of its CP sensors to use the relative height feature [9].

We visualized the scatter of the data with height and relative height, using principal component analysis (PCA), as shown in Figure 2. In order to compare the discrimination power of the two features, we referred to repeated measurements of the same odorant as a cluster, and for each feature we calculated a *cluster discrimination measure (CDM)*, defined as:

$$CDM = tr(S_b)/tr(S_w), \quad (1)$$

where S_b is the weighted covariance matrix of the cluster centroids (the *between variance*), and S_w is the weighted sum of cluster covariance matrices (the *within variance*). The cluster weight in both S_b and S_w is the fraction of the cluster size out of the total data. The higher the CDM value, the better the discrimination.

	height	relative height
QMB	199.3	189.0
CP	31.7	30.2

Table 2: Cluster discrimination measure (CDM) for height and relative height.

Prior to CDM calculation, the variance of all sensors was unit-scaled, thus enabling the comparison of CDM values between features and between sensor sets. The CDM values for height and relative

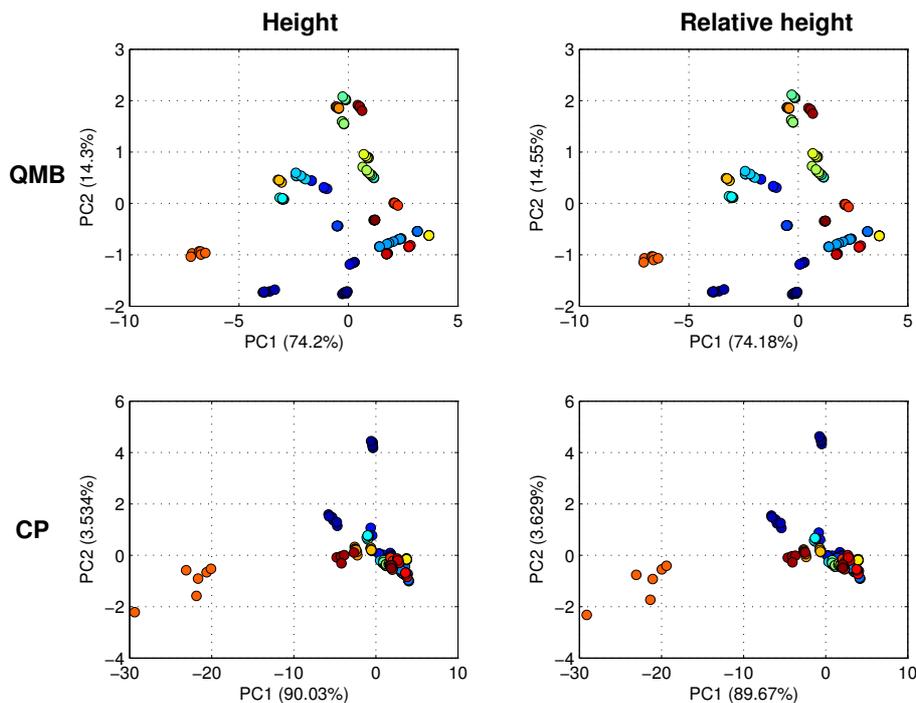


Figure 2: The scatter of the data set with the height and relative height features, in CP and QMB. The colors represent odorants. The data was unit-scaled prior to applying the principal component analysis (PCA).

height are shown in Table 2. Between the two features there is only a small difference; therefore we do not expect the choice between them to have much influence on our results. There is, however, a large CDM difference between QMB and CP; QMB has much higher CDM, in accordance with the visual impression conveyed by Figure 2.

Since the CDM of the two features is roughly the same, we chose to use the 'traditional' feature for each of the sensor sets, namely the height for QMB and the relative height for CP.

3.2 Mapping Methods

Our task is to learn a mapping between two eNose spaces. Mappings are directional. Thus, in every mapping one of the spaces is the *source space* and the other is the *destination space*. Denote by X and Y the two matrices containing the readings in the source space and the destination space, respectively. X is an n by d_S matrix and Y is an n by d_D matrix, where n is the number of measurements and d_S , d_D are the dimension (number of sensors) of the source space and the destination space, respectively. In X and Y , each measurement is represented by a row that contains a single feature per sensor. We seek a mapping model that will explain Y using X .

3.2.1 Linear regression methods

We tried three linear regression methods: multiple linear regression (*MLR*), principal component regression (*PCR*) and partial least squares (*PLS*). These methods are explained in detail in [11, 12], and a wide statistical background for them can be found in [13].

3.2.2 Nonlinear neural networks

The actual nature of the mapping between eNose spaces is unknown to us, and we would like to be able to account for nonlinear effects as well as for linear ones. Feed-forward back-propagation neural networks (*NN*) are a powerful learning tool, which can be used for learning nonlinear functions. We used a network with a sigmoid hidden layer and a linear output layer. This network architecture can approximate arbitrarily well any function with a finite number of discontinuities, including nonlinear ones, given sufficient neurons in the hidden layer [14]. In practice, we used networks whose hidden layer size does not exceed 32.

Since we use *NN* for prediction, it is necessary to provide means for generalization; that is, avoiding over-fitting the model to the calibration data. For this purpose, we used the Bayesian regularization training method. In addition, we applied the early stopping technique, using a validation set consisting of 40% of the calibration data. The description of these methods can be found in [14].

3.2.3 Tessellation-based linear interpolation

The Tessellation-based linear interpolation method (*TLI*) consists of two stages: Tessellation and prediction. First, we tessellate the convex hull of the X points for which we know the corresponding Y points (the calibration set) with (d_S+1) -vertex shapes (namely, *simplices*; single: *simplex*). For example, with $d_S=2$, this means tessellating an area with triangles. In this tessellation, the vertices of all simplices are calibration set X points. Figure 3 shows tessellations of the QMB space and the CP space.

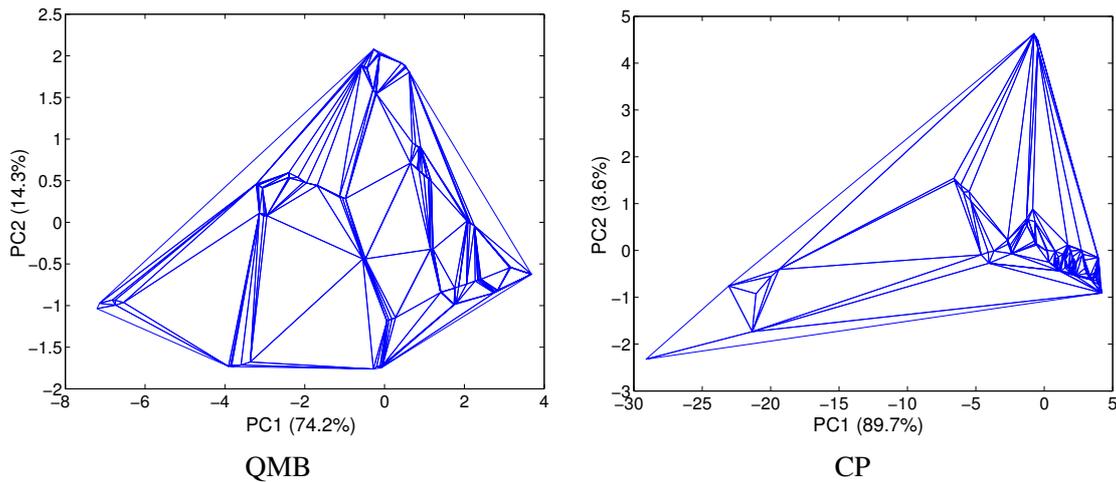


Figure 3: The two-dimensional Delaunay tessellation of QMB and CP. The axes are the first two principal components.

We predict the Y point for a new X point in the following manner: we first locate the simplex enclosing the new X point, which is assured to exist provided that the point is within the convex hull of the calibration set. We then calculate the *barycentric coordinates* of the point relative to the vertices of the simplex enclosing it. The barycentric coordinates of a point p within a simplex can be intuitively described as weights, which, if placed at the simplex's vertices, their center of mass will coincide with p . We normalize the sum of coordinates to 1, achieving unique coordinates which can be used as interpolation weights. The Y point we predict for the new X point is then the linear interpolation from the simplex's vertices, which are calibration set points with known Y values. In other words,

the prediction is the average Y value of the simplex’s vertices, weighted according to the barycentric coordinates.

The TLI method is local; this is the main property differentiating it from all the other methods we use. By local, we mean that the prediction is based only on calibration set points in the vicinity of the predicted point. This property protects TLI from global topological distortions in the mapping.

A method similar to TLI was recently presented in a different context in [5], termed there the Law of Mixtures (LM). The difference between TLI and LM is in the tessellation method. TLI uses Delaunay tessellation, which is the unique tessellation such that the circumsphere of each simplex has no vertices in its interior. This property works to prevent, when possible, the use of elongated simplices, and thus contributes to the locality of TLI.

We find that it is better to apply TLI in a low dimension, because the number of outliers (test set points which lie outside the convex hull) rises quickly with the dimension. For that reason we applied TLI in the space defined by only the first few principal components of the source space. As an alternative, we used PLS weights instead of principal components (we termed this variant TLI_{PLS}). Furthermore, when applying TLI in two dimensions, we also predict the Y values of outliers, using a method which was proposed in [5]. This method is based on mirroring the outlier into the convex hull, predicting the Y value for the mirrored point using interpolation, as previously described, and then predicting the Y value of the outlier by extrapolating from the mirrored point.

3.3 Assessment of the Prediction Quality

Our task is to map from the source space into the destination space. A way to evaluate our success is to check whether the prediction of a sample from odorant o can be classified, in the destination space, as odorant o . Specifically, we find the center of mass (centroid) in the destination space for all odorants. Then, we calculate the distance from the prediction to all the odorant centroids. A prediction of a sample from odorant o for which the closest centroid is the centroid of o is counted as a “hit”, since it can be correctly classified. When assessing the prediction quality over the entire test set, our indicator is the percent of predictions that gave a hit, which we term $\%hits$. In all calculations, we use Euclidean distance in the mean-centered and unit-scaled destination space.

3.4 Splitting the Data Set into Calibration and Test Parts

Our data set contains multiple samples from each odorant. Naturally, samples of the same odorant are generally closer to each other than samples of different odorants. When applying a mapping model to a new sample (a sample from the test set), the *proximity* of the calibration set samples to the new sample affects the prediction’s accuracy. It is not desirable to blur this effect by disregarding the heterogenic nature of the data set. Therefore, we separately examined the two extremes of the proximity range, by using two different techniques for splitting the data set into calibration and test parts: the *representative test set (REP)* and *leave-one-out (LO)*. In the REP technique, the test set consists of a random representative from each odorant in the data set. The LO technique, on the other hand, does not take a random odorant, but iterates through all odorants, using - in each iteration - all samples of that odorant as the test set. In this way, LO actually “leaves out” one odorant each time around. A crude way of viewing the difference is this: REP tests the mapping on samples that are “close” to the calibration set, while LO tests the mapping on more “distant” samples.

4 Results

Following the feature extraction step described in Section 3.1, and after excluding from the data set odorants with responses that are too weak and other types of noisy samples, the data set consisted of 141 different samples from 23 odorants, with some six repeats per odorant. To this data set we applied the methods described in Section 3.2 for learning mappings in both directions, that is, both from QMB into CP and from CP into QMB. In PCR, PLS, NN and TLI we tested a range of model dimensionalities. The size of the test set was 23 with the REP technique (one sample from each odorant) and 141 with the LO technique (the agglomeration of test sets from all iterations). The results are summarized in Table 3.

Mapping Direction		<i>MLR</i>	<i>PCR</i>	<i>PLS</i>	<i>NN</i>	<i>TLI</i>	<i>TLI_{PLS}</i>
From QMB into CP	<i>LO</i>	21%	23%	23%	27%	21%	26%
	<i>REP</i>	35%	43%	39%	65%	96%	91%
From CP into QMB	<i>LO</i>	57%	67%	64%	65%	15%	16%
	<i>REP</i>	83%	91%	91%	100%	78%	78%

Table 3: Prediction results of the mapping methods. The entries are the %*hits* values of each mapping method, in each of the two data splitting techniques. In PCR, PLS and NN the results are those of the most accurate model in the dimensionality range tested. In TLI, The results are those of the two-dimensional model, and they include the prediction for outliers.

In mapping from QMB to CP, the best prediction performance is 27% hits in LO and 96% hits in REP. In mapping from CP to QMB, the best prediction performance is 67% hits in LO and 100% hits in REP. Based on these results, we estimate that mapping from QMB into CP is harder than in the opposite direction. Illustrations for the predictions from CP into QMB are shown in Figure 4.

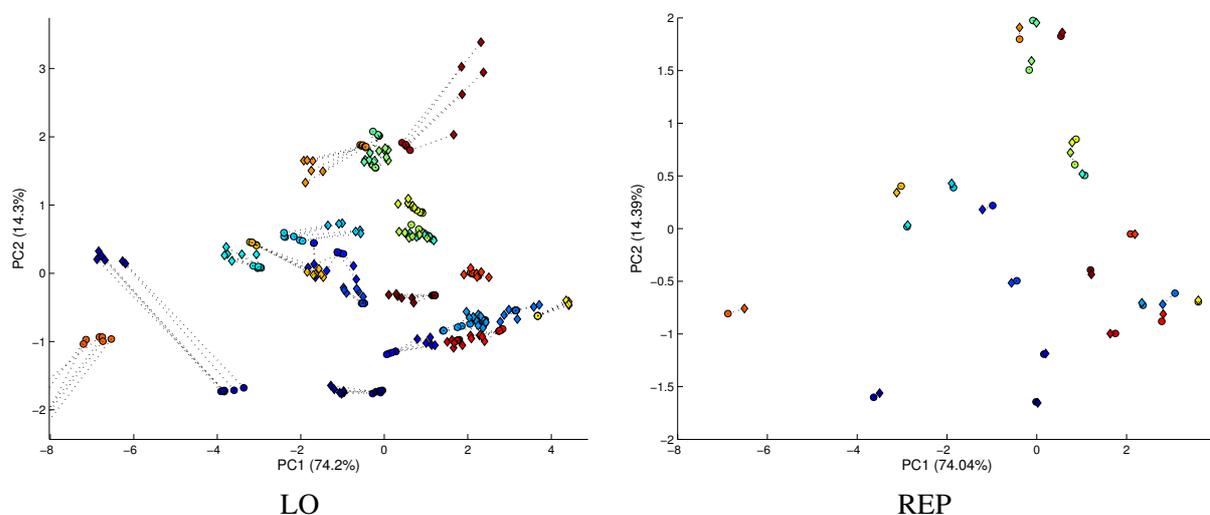


Figure 4: Prediction from CP into QMB. Colors represent odorants. Actual measurements are marked with circles, predictions are marked with diamonds, and the dashed lines link measurements to their corresponding predictions. In LO, the best method was PCR, achieving 67% hits. In REP, the best method was NN, achieving 100% hits.

The comparison of the results of PCR and PLS in Table 3 shows that the performance of the two methods is equivalent. A closer look into the differences can be obtained from inspecting the %*hits* performance of PCR and PLS simultaneously as a function of the model dimensionality. This analysis

for mappings from CP into QMB is shown in Figure 5. The best model dimensionality of PCR and PLS was similar in most of our experiments. Furthermore, the performance of PCR and PLS does not significantly differ in any place along the dimensionality range. Based on this analysis, we can conclude that despite the theoretical advantage of PLS over PCR, the methods are practically equivalent in the context of the current challenge.

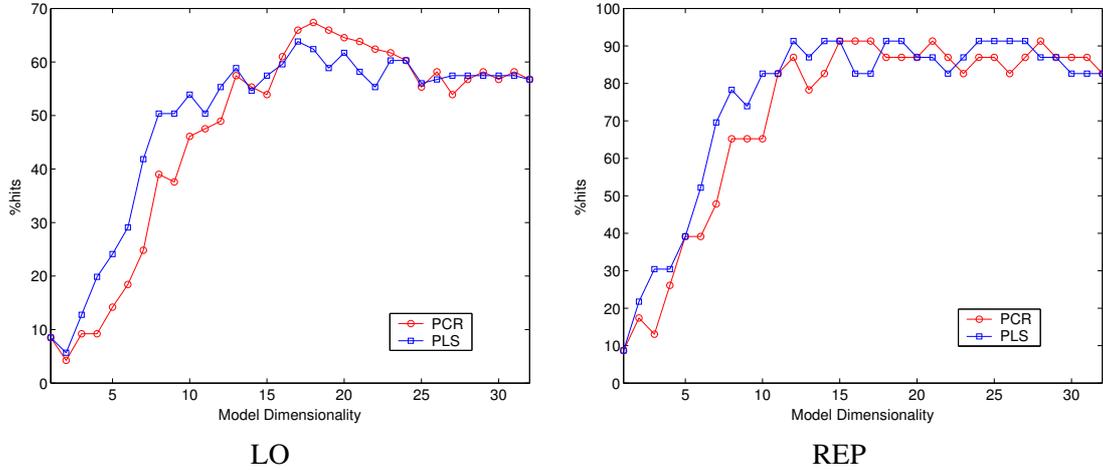


Figure 5: The dependency of the $\%hits$ performance of PCR and PLS on the model dimensionality, in mappings from CP into QMB. The best dimensionality with LO was 18/17 and with REP 15/12, in PCR/PLS respectively.

Another view into the performance difference between PCR and PLS can be obtained from comparing the results of TLI and TLI_{PLS} . The results do not exhibit any consistent advantage of either method over the other.

The performance of the NN method is consistently good, relative to the other methods, in all experiments. In one of the experiments, the mapping from CP into QMB with the REP splitting technique, it turned out to be the best method, achieving 100 $\%hits$ (Table 3). The NN performance monotonically increases as additional nodes are added to the network’s hidden layer, up to a certain level (Figure 6 (a)). This is expected, since generally a larger network possesses a better expressive power. However, we observe that a hidden layer size of eight or 16 is sufficient, and beyond this size the performance does not improve. It can be estimated that this network size reflects the complexity of the underlying mapping from CP into QMB.

The performance of the TLI method (including the TLI_{PLS} variant) is outstanding in the mappings from QMB to CP with the REP technique, achieving 96 $\%hits$, while the next best method in this experiment is NN with only 65 $\%hits$ (Table 3). However, in other situations the performance of TLI is relatively low, such as in mappings from CP into QMB with the LO technique, where it achieves only 16 $\%hits$, while all the other methods achieve 57-67 $\%hits$. We attribute both the high performance in REP and the low performance in LO to the local nature of TLI, as discussed in Section 3.2.3.

The TLI performance does not change significantly, nor consistently, in the model dimensionality range that we tried (between two and five). The number of outliers rises quickly with the dimensionality, and with five dimensions over 50% of the test set were outliers. We can thus conclude that complicating the TLI model beyond two dimensions does not seem necessary.

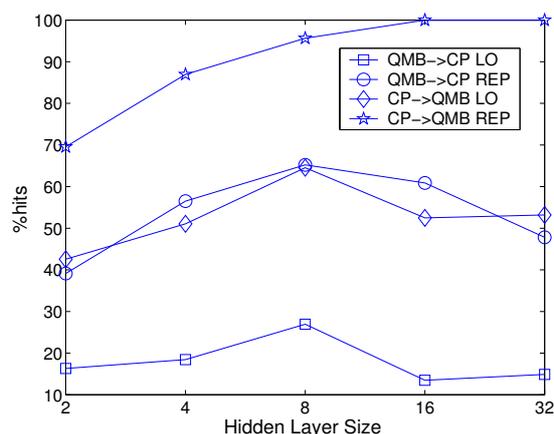


Figure 6: $\%hits$ performance of the NN method as a function of the network’s hidden layer size. Note that the hidden layer sizes tested are powers of two.

5 Conclusions and Future Work

In this paper we investigated methods for learning mappings between Cyranose 320, a conducting polymers eNose, and the quartz microbalance module of the MOSESII eNose. In mappings from CP into QMB, the percent of predictions that could have been classified correctly ($\%hits$) was 67% for samples that are “distant” from the calibration set (LO) and 100% for samples that are “close” to the calibration set (REP). This mapping direction gave higher prediction accuracy than the opposite direction. A possible explanation for this is that QMB possesses a high power of odor discrimination, thus with QMB as a target space more predictions can be correctly classified. The prediction results provide information about the appropriateness of the different mapping methods to this task, as well as information about the nature of the underlying mapping.

Some of the mapping methods we evaluated did not have consistent performance in all experiments. The most prominent example is the TLI method, which in the REP experiments was far better than in the LO experiments. We concluded that TLI exploits the proximity of new samples to calibration set samples better than the other methods. From a practical point of view, it would be advisable to dynamically choose a mapping method for the prediction of each new sample, based on the proximity of that sample to the calibration set.

The advantage of the nonlinear method we used, NN, over the linear regression methods, is apparent from the results, even though the gaps are not always large. This does not suffice for determining whether the underlying mapping is a nonlinear one. However, the results support the preference of NN over the linear regression methods as a general strategy.

Another indication as to the nature of the underlying mapping is obtained from inspecting the model dimensionality. In PCR and PLS, the optimal prediction accuracy was achieved using a dimensionality that was lower than the total number of sensors. Similarly, in NN we identified the hidden layer size beyond which the accuracy did not improve. This type of information is helpful in choosing an appropriate mapping model.

We have some ideas for improved mappings between eNoses. Feature extraction approaches which capture more information, such as the one presented in [2], could be applied. In addition, changes in the current mapping methods or application of new ones may yield an improvement. In the context of NN, different network architectures may prove better.

Mapping between eNoses is important as a preliminary step to mapping an eNose into the psy-

chophysical (human panel) space, which is a key task to realizing odor communication. The encouraging results we have achieved confirm the viability of mapping between odor spaces. The representation of the psychophysical odor space that eventually will be obtained from human panel experiments will no doubt have different characteristics than an eNose space, and will thus probably require different mapping methods. Nevertheless, we expect the conclusions we have drawn from the current work to be helpful in the next stage.

References

- [1] Harel D., Carmel L. and Lancet D., "Towards an Odor Communication System", *Computational Biology and Chemistry (formerly Computers & Chemistry)* 27 (2003), 121-133. (Early version titled "An Algorithmic Approach to Odor Communication and Reproduction", Technical Report MCS01-16, The Weizmann Institute of Science, 2001.)
- [2] Carmel L., Levy S., Lancet D. and Harel D., "A Feature Extraction Method for Chemical Sensors in Electronic Noses", *Sensors and Actuators B: Chemical* 93 (2003), 66-75.
- [3] Carmel L., Sever N. and Harel D., "On Predicting Response to Mixtures in QMB Sensors", submitted to *Sensors and Actuators B: Chemical* (2003).
- [4] Balaban M. O., Korel F., Odabasi A. Z. and Folkes G., "Transportability of data between electronic noses: mathematical methods", *Sensors and Actuators B: Chemical* 71 3 (2000), 203-211.
- [5] Jin L., Pierna J. A. F., Wahl F., Dardenne P., Massart D. L., "The Law of Mixtures method for multivariate calibration", *Analytica Chimica Acta* 476 (2003), 73-84.
- [6] <http://www.lennartz-electronic.de/>
- [7] Mitrovics J., Ulmer H., Weimar U. and Gopel W., "Modular Sensor Systems for Gas Sensing and Odor Monitoring: The MOSES Concept", *Acc. Chem. Res.* 31 (1998), 307.
- [8] <http://cyranosciences.com/>
- [9] Lonergan M. C., Severin E. J., Doleman B. J., Beaber S. A., Grubbs R. H. and Lewis N. S., "Array-Based Vapor Sensing Using Chemically Sensitive, Carbon Black-Polymer Resistors", *Chem. Mater.* 8 (1996), 2298-2312.
- [10] Gardner J. W. and Bartlett P. N., *Electronic Noses, Principles and Applications*, Oxford University Press (1999), Oxford.
- [11] Geladi P. and Kowalski B. R., "Partial Least-Squares Regression: A Tutorial", *Analytica Chimica Acta* 185 (1986), 1-17.
- [12] Martens, H. and Næs, T., *Multivariate Calibration*, John Wiley & Sons (1989), New York, NY.
- [13] Everitt B. S. and Dunn G., *Applied Multivariate Data Analysis*, Arnold (1991), London.
- [14] Matlab Neural Network Toolbox 4.0.1 User's Guide.