On the Decomposition of the Self-Organizing Map Distortion Measure

Juha Vesanto, Mika Sulkava, Jaakko Hollmén

Helsinki University of Technology, Laboratory of Computer and Information Science, P.O. Box 5400, FIN-02015 HUT, Finland
Tel. +358-(0)9-451 5290, Fax. +358-(0)9-451 3277
e-mail: Juha.Vesanto@hut.fi,Mika.Sulkava@hut.fi,Jaakko.Hollmen@hut.fi

Keywords: Self-Organizing Map, quantization error, topological ordering, distortion measure

Abstract— Assessment of model properties with respect to data is important for reliable analysis of data. After training, Self-Organizing Map (SOM) can be assessed, for instance, with respect to its quantization or its topology preservation properties with one-number summaries. In this paper, we present a decomposition of the SOM distortion measure for measuring different aspects of the SOM for map units locally. The terms measure quantization quality, the goodness of topological preservation, and the stress between these two aspects. Experiments are used to illustrate the behavior of the distortion measure terms in different error scenarios.

1 Introduction

The Self-Organizing Map (SOM) [11] is a neural network model based on unsupervised, competitive learning. The model combines aspects of vector quantization with a topology-preserving ordering of the quantization vectors. For reliable analysis of data and for drawing meaningful inferences about the data analysis problem, model assessment with regard to both of these properties should be carefully performed. In the remainder of the section, we describe the basic algorithm and motivate our view on the assessment problem of the SOM.

A SOM consists of m units located on a regular, low-dimensional grid of map units. The map unit positions $\mathbf{r}_j$ on the regular grid are fixed; each map unit is connected to a number of neighboring map units with a neighborhood relation. The closest set of map units are called its neighbors. Each (map) unit j has an associated d-dimensional prototype vector $\mathbf{m}_j = [m_{j1}, \ldots, m_{jd}]$. During training, the map adjusts to the data by adapting the prototype vectors according to the following training rule:

$$\mathbf{m}_j := \mathbf{m}_j + \alpha h_{b_i j}(\mathbf{x}_i - \mathbf{m}_j) \tag{1}$$

where $\mathbf{x}_i$ is a sample vector, α is a learning rate, and $h_{b_i j}$ is the value of neighborhood function between map units j and b_i, the best-matching prototype to the sample vector $\mathbf{x}_i$: $b_i = \arg\min_j\{\|\mathbf{x}_i - \mathbf{m}_j\|^2\}$. The most usual neighborhood function is the Gaussian:

$$h_{ij} = e^{-\|\mathbf{r}_i - \mathbf{r}_j\|^2/2r^2}, \tag{2}$$

where r is a neighborhood radius [1].

The prototype vectors together with the search for the best-matching unit define a tessellation of the input space into a set of Voronoi regions or sets

$$V_j = \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{m}_j\| < \|\mathbf{x} - \mathbf{m}_k\| \, \forall k \neq j\}$$

and by construction, each map unit is associated with one such Voronoi set. In the experiments in this paper, an alternative training algorithm called batch map is used to train the maps. In this algorithm, the best-matching map unit (BMU) of each data vector $\mathbf{x}_i$ is found first. Secondly, new prototype vectors are calculated as weighted averages of all data samples:

$$\mathbf{m}_j = \frac{\sum_{i=1}^{n} h_{b_i j}\mathbf{x}_i}{\sum_{i=1}^{n} h_{b_i j}} = \frac{\sum_{i=1}^{m} h_{ij} N_i \mathbf{n}_i}{\sum_{i=1}^{m} h_{ij} N_i} \tag{3}$$

where n is the number of data samples, N_i is the number of data samples in Voronoi set V_i, and $\mathbf{n}_i = \sum_{\mathbf{x} \in V_i} \mathbf{x}/N_i$ is their centroid. Eq. 1 is a stochastic approximation of the batch map algorithm.

A mathematical difficulty in the analysis of the SOM algorithm (Eq. 1) is that it has been shown not to be the gradient of any cost function in the general case [2]. The absence of a cost function has lead researchers to find other ways to measure the overall quality of SOMs, for example based on classification error [12], maximum likelihood [16], and topological path lengths [8]. Also the distortion measure

$$E_d = \sum_{i=1}^{n}\sum_{j=1}^{m} h_{b_i j}\|\mathbf{x}_i - \mathbf{m}_j\|^2 \tag{4}$$

[1] A generalization of the SOM architecture is a network of map units with no output space locations at all. Examples of such networks include growing grid [4] and GCS [3]. By defining the neighborhood function in terms of path lengths along the network, the measures described in this paper are applicable to such variants, too.

investigated in this paper can be used for measuring the overall quality of a given SOM. Two results speak for its usage: (1) in case of a discrete data set and fixed neighborhood kernel, it can be shown to be a local energy function of the SOM [10], and (2) the SOM training rule (Eq. 1) has been shown to be a computationally efficient approximation for the gradient of the distortion measure [6].

However, there are two sides to this question of SOM evaluation: the quality of quantization and the quality of topology representation. Since these are competing, but not directly opposed objectives of the SOM, we must not restrict ourselves to an arbitrary combination of their quality measures. Rather, to fully understand a given SOM, both of these must be measured.

We would like to know what kind of errors there are and where they are located. Note that there are two kinds of errors in topological representation: (1) errors in preservation of the original topological relationships between items, and (2) errors in trustworthiness of the representation of relationships [17, 5]. Errors in the first occur when data items originally close to each other are mapped far from each other in the representation. Errors in the latter occur when data items close to each other on the representation are actually far from each other in the original space.

2 Decomposition of the SOM distortion measure

The SOM distortion measure E_d can be divided to two component terms [14]:

$$E_d = \sum_{i=1}^{n} H_{b_i} \|\mathbf{x}_i - \mathbf{n}_{b_i}\|^2 + \sum_{i=1}^{m} N_i \sum_{j=1}^{m} h_{ij} \|\mathbf{n}_i - \mathbf{m}_j\|^2. \tag{5}$$

where $H_{b_i} = \sum_j h_{b_i j}$. The second term can be further divided such that E_d can be expressed as a sum of three component terms $E_d = E_{qx} + E_{nb} + E_{nv}$:

$$\begin{aligned} E_d &= \sum_{j=1}^{m} N_j H_j \Big(\mathrm{Var}\{\mathbf{x}|j\} + \|\mathbf{n}_j - \bar{\mathbf{m}}_j\|^2 + \mathrm{Var}_h\{\mathbf{m}|j\} \Big) \\ &= \underbrace{\sum_{j=1}^{m} N_j H_j \mathrm{Var}\{\mathbf{x}|j\}}_{E_{qx}} + \underbrace{\sum_{j=1}^{m} N_j H_j \|\mathbf{n}_j - \bar{\mathbf{m}}_j\|^2}_{E_{nb}} \\ &\quad + \underbrace{\sum_{j=1}^{m} N_j H_j \mathrm{Var}_h\{\mathbf{m}|j\}}_{E_{nv}} \end{aligned} \tag{6}$$

where $\mathrm{Var}\{\mathbf{x}|j\}$ is the (biased) local variance of the data $\mathrm{Var}\{\mathbf{x}|j\} = \sum_{\mathbf{x} \in V_j} \|\mathbf{x} - \mathbf{n}_j\|^2 / N_j$, $\bar{\mathbf{m}}_j$ the weighted mean and $\mathrm{Var}_h\{\mathbf{m}|j\}$ the weighted variance of the prototype vectors: $\bar{\mathbf{m}}_j = \sum_k h_{jk} \mathbf{m}_k / H_j$ and $\mathrm{Var}_h\{\mathbf{m}|j\} = \sum_k h_{jk} \|\mathbf{m}_k - \bar{\mathbf{m}}_j\|^2 / H_j$ [18] [2]. Notice that the contribution of each variable and each map unit to the distortion measure can be easily measured. In the following, each of the three component terms is discussed in more detail.

E_{qx} : The first term, local data variance, measures the quantization quality of the SOM as the variance of the data vectors within each Voronoi set. If the neighborhood function values for each map unit are normalized to unity such that $H_{b_i} = 1, \forall i$, E_{qx} corresponds to classical vector quantization error. The E_{qx} is slightly different from the measure typically used to calculate quantization quality of a SOM, the average quantization error $\sum_i \|\mathbf{x}_i - \mathbf{m}_{b_i}\|/n$: in E_{qx} squared distances are used, and the data vectors are compared to the centroid of the Voronoi set, not to prototype vector, which only acts as the grouping criterion.

E_{nv} : The last term, neighborhood variance, measures the topological quality of the SOM. More specifically, since it is a measure of closeness of prototype vectors close to each other on the map grid, it measures the trustworthiness of the map topology. In [17] it was indeed noted that a strong point of the SOM as opposed to many other projection algorithms is its trustworthiness. Notice that the neighborhood variance only takes data vectors into account as weighting factors (N_i). Therefore, like the smoothness measure by Hämäläinen [7] and the topographic product [1], E_{nv} actually measures the quality of the map grid in representing the topology of the prototype vectors, not the topology of the data. In contrast, in the topographic function [19], topographic error [9], and preservation of kNN sets [17] measures the data vectors have an important role in measuring topological quality.

E_{nb} : The middle term, neighborhood bias, is perhaps the most interesting because it links the quantization and ordering together. It can be interpreted as the stress between those properties. To our knowledge, such measures have not been proposed earlier.

3 Experiments

Two sets of experiments are reported to illustrate the behavior of the terms in the decomposed distortion measure. In the first set of experiments, we control the degree of flexibility of the SOM by varying the final neighborhood radius in the training. As a simplified rule of thumb, using large neighborhood radii should result in stiff maps, which (over)stress topological ordering at the cost of quantization accuracy.

[2] Note that these elements are also utilized in [13].

In the second set of experiments, we introduce topological errors to the map configuration by altering the topological coordinates of map units, either locally or globally to create a twist in the map.

Controlling the stiffness of the map We control the stiffness of the map by training several maps by varying the final neighborhood radius in the training between 0 to 12. We evaluate the terms in the decomposition for each map. Figs. 1 and 2 show the average relative magnitudes of the decomposition terms together with confidence bounds over 20 repeated runs. Fig. 1 shows the results for a one-dimensional map with 100 map units trained with two-dimensional uniformly distributed data (see Fig. 3). Fig. 2(a) and (b) show the results for two-dimensional maps trained with uniformly distributed two- and five-dimensional data, respectively.

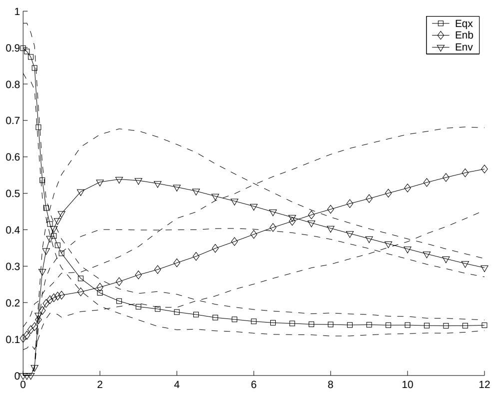

Figure 1: The relative magnitude of the terms in the decomposed distortion measure of maps trained and evaluated with varying final training radii shown on the x-axis. One-dimensional map is trained with two-dimensional data. The dashed lines indicate the 95% intervals of 20 repeated runs.

For large values of neighborhood radius the relative E_{nv} error decreases with increasing training radius. For small values there is a sharp drop at approximately $r = 0.5$. At this point, the values of the neighborhood function become so small that its effect on the distortion measure vanishes and SOM effectively becomes the k-means algorithm.

In Fig. 3, the local values of the distortion measure and each of the decomposition terms are shown for one instance of the 1-dimensional map trained with 2-dimensional data. The final training radius is 12. From the figures, the border effect can be clearly seen in the ends of the map. Of the decomposition terms, the border effect is most prominent in E_{qx} and E_{nb}. For E_{nb}, the areas of larger curvature get high values due to elongated Voronoi regions. For E_{nv} on the other hand, the biggest values can be seen where the map units are most sparsely distributed.

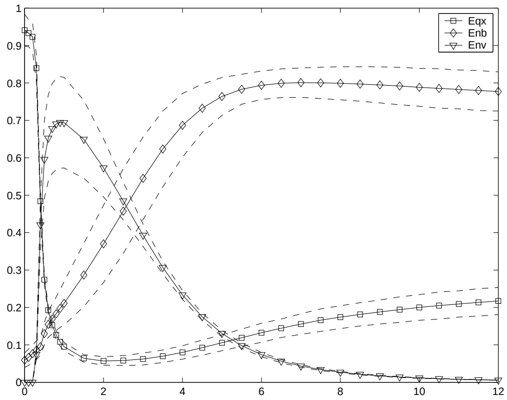

(a)

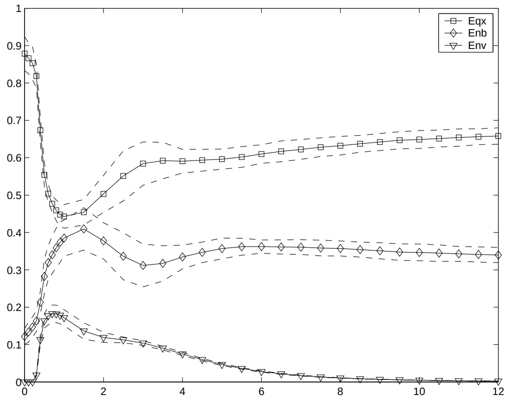

(b)

Figure 2: The same as in Fig 1, but two-dimensional maps were trained with two-dimensional and five-dimensional data in (a) and (b), respectively.

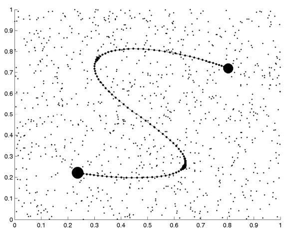

(a) E_d

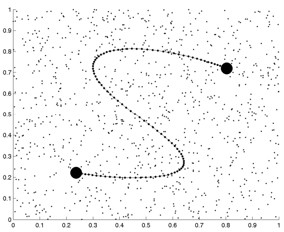

(b) E_{qx}

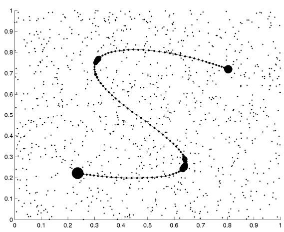

(c) E_{nb}

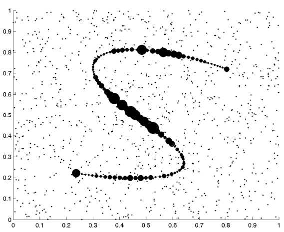

(d) E_{nv}

Figure 3: The normalized distortion measure evaluated locally for each map unit. Bigger circle indicates higher value.

Introducing topological errors A two-dimensional map was trained with two-dimensional uniformly distributed data. We introduced topological errors in the map by exchanging the prototype vectors of two map units (see Fig. 4). The distortion measure and its terms were evaluated in each map unit for both the correct and the distorted map. The results are shown in Fig. 5 with black bars corresponding to the correct map and white bars to the distorted map. The twist in the map can be observed from the distortion measure E_d, and its E_{nb} and E_{nv} terms.

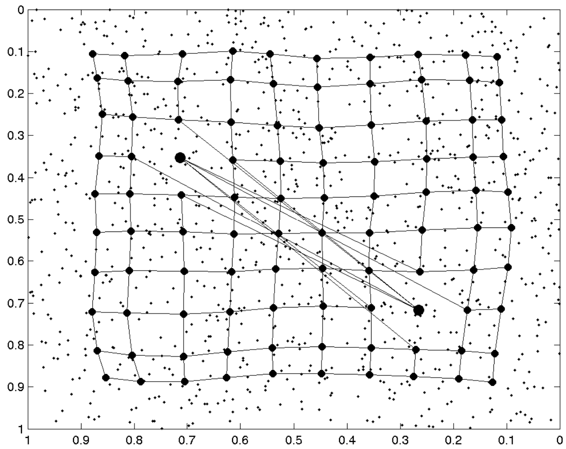

Figure 4: The map with an artificially introduced topological error. For evaluation of the terms in the distortion measure, see Fig. 5.

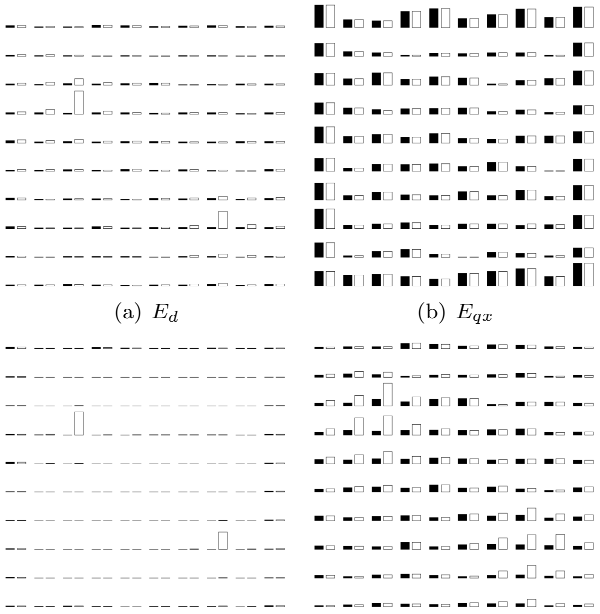

(a) E_d (b) E_{qx} (c) E_{nb} (d) E_{nv}

Figure 5: The distortion measure (a) with terms (b-d) evaluated for a map with a topological error (seen in Fig. 4). The bars in a-d are the distortion (component term) values for each map unit. The black bars indicate the values of the map with no topological defects and the white bars the values of the map with a topological error. The topological error and its location can be clearly observed from (a) and (c-d).

In the next experiment, a two-dimensional SOM was trained with two-dimensional data consisting of two disjoint uniformly distributed rectangles. After training, a twist was introduced by switching the prototype vectors for the lower part of the map (see Fig. 6). Again, the distortion measure and its terms were evaluated for each map unit for both the correct and the twisted map. Results are shown in Fig. 7.

Again, the twist can be seen from the distortion meausure, although it is somewhat masked by the non-uniform distribution of the data. The effect of the twist is most apparent in E_{nb} at the sides of the map, which is natural since this is where the prototype vectors have been moved the most. Also the border effect at top and bottom borders on the map can be clearly seen from E_{qx} and E_{nb}. The E_{nv} shows the change in the distribution of data between the two rectangles. Notice that it is closely related to the U-matrix [15]; also the U-matrix visualizes the local trustworthiness of the SOM. A benefit of the distortion measure is that it can be evaluated for each vector component separately. From Fig. 7(e-f) it can be clearly seen that the twist in the map takes place in the X-component.

4 Discussion

In this paper, the decomposition of the SOM distortion measure was introduced and some preliminary results were presented. It is clear that decomposition terms respond differently to different error scenarios in the map. The decomposition terms can be evaluated separately for each map unit, and even for each vector component, which makes it possible to locate the source of the distortion error. It may be possible to use the decomposition terms also to identify the type of the error. In [20] Li discussed different kinds of topological errors in SOMs and proposed that they could be made use of to tell something of the true topology of the data set.

One of the most important uses of error measures is comparision of different maps to each other. Not only this, but ideally the measures should be:

(1) comparable for different data sets: starting from different (sized) sample from the same distribution, to a completely different distribution
(2) comparable for different SOMs: different neighborhood radius, different number of map units, even different network architecture
(3) local, such that the source (location) of the error (on the map) can be identified

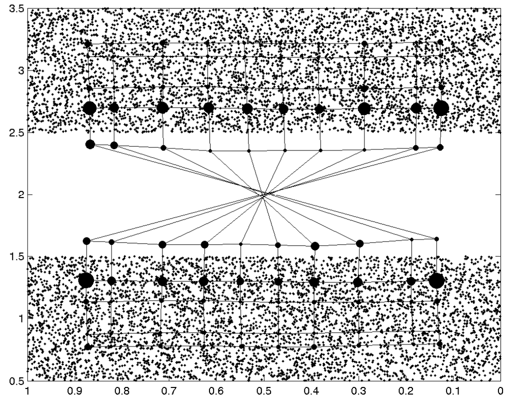

Figure 6: The twisted map has a topological error on a global scale. For the evaluation of the distortion measure and its terms, see Fig. 7. The size of the circles indicates the magnitude of the distortion measure.

As of yet, it is unclear whether the terms of the distortion measure, or some normalized version thereof, fits these requirements.

5 Summary

For reliable analysis of data, model assessment should reflect both the quantization and topology preservation aspects of the SOM. We have presented a decomposition of the SOM distortion measure that measures these properties locally for map units. A novel aspect of the work is that it introduces a measure for stress between the quantization and topology preservation properties; also, the evaluation is done locally for map units in order to locate possible anomalies in the map. Experiments illustrate the behavior of the component terms for a set of problems. We intend to study the decomposition further: we continue by making a full review and an experimental comparison of the proposed measures to evaluate a SOM. Ultimately, the measures should guide in optimizing the map topology and in the training to get a SOM with desired properties.

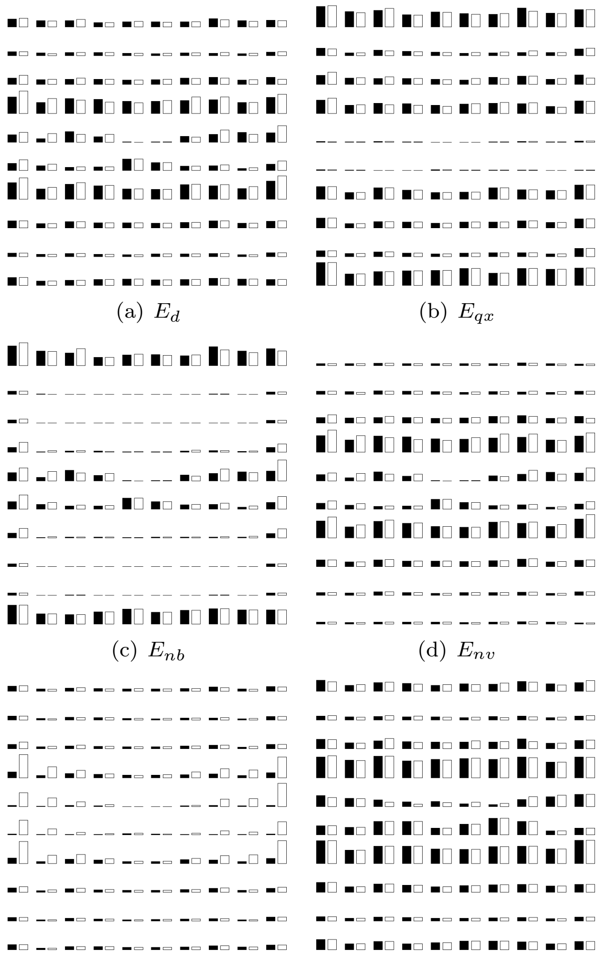

(a) E_d (b) E_{qx} (c) E_{nb} (d) E_{nv} (e) E_d X (f) E_d Y

Figure 7: The distortion measure (a), its component terms (b-d) and distortion measure measured in X and Y-direction separately (e-f) for a well ordered (black bars) and a twisted map (white bars).

References

[1] Hans-Ulrich Bauer and Klaus R. Pawelzik. Quantifying the neighborhood preservation of self-organizing feature maps. *IEEE Transactions on Neural Networks*, 3(4):570–579, July 1992.

[2] E. Erwin, K. Obermayer, and K. Schulten. Self-organizing maps: Ordering, convergence properties and energy functions. *Biol. Cyb.*, 67(1):47–55, 1992.

[3] Bernd Fritzke. Growing Cell Structures – A Self-Organizing Neural Network for Unsupervised and Supervised Learning. *Neural Networks*, 7(9):1441–1460, 1994.

[4] Bernd Fritzke. Growing Grid – a self-organizing network with constant neighborhood range and adaptation strength. *Neural Processing Letters*, 2(5):9–13, 1995.

[5] G. J. Goodhill and T. J. Sejnowski. A unifying objective function for topographic mappings. *Neural Computation*, 9:1291–1303, 1997.

[6] Thore Graepel, Matthias Burger, and Klaus Obermayer. Phase transitions in stochastic self-organizing maps. *Physical Review E*, 56:3876–3890, 1997.

[7] Ari Hämäläinen. *Self-Organizing Map and Reduced Kernel Density Estimation.* PhD thesis, University of Helsinki, 1995.

[8] S. Kaski and K. Lagus. Comparing self-organizing maps. In *Proceedings of International Conference*

on Artificial Neural Networks (ICANN) 96, pages 809 – 814, 1996.

[9] Kimmo Kiviluoto. Topology preservation in self-organizing maps. In *Proceedings of the International Conference on Neural Networks (ICNN'96)*, volume 1, pages 294–299, Piscataway, New Jersey, USA, June 1996. IEEE Neural Networks Council.

[10] Teuvo Kohonen. Self-Organizing Maps: Optimization Approaches. volume 2, pages 981–990. Elsevier Science Publishers, 1991.

[11] Teuvo Kohonen. *Self-Organizing Maps*. Springer-Verlag, 1995.

[12] M. A. Kraaijveld, J. Mao, and A. K. Jain. A nonlinear projection method based on kohonen's topology preserving maps. *IEEE Transactions on Neural Networks*, 6(3):548–559, May 1995.

[13] Jouko Lampinen and Timo Kostiainen. Recent advances in self-organizing neural networks. In U. Seiffert and L. Jain, editors, *Self-organizing neural networks: Recent advances and applications*, pages 75–94. Physica Verlag, 2002.

[14] Jouko Lampinen and Erkki Oja. Clustering Properties of Hierarchical Self-Organizing Maps. *Journal of Mathematical Imaging and Vision*, 2(2-3):261–272, November 1992.

[15] A. Ultsch and H. P. Siemon. Kohonen's Self Organizing Feature Maps for Exploratory Data Analysis. In *Proceedings of International Neural Network Conference (INNC'90)*, pages 305–308, Dordrecht, Netherlands, 1990. Kluwer.

[16] Akio Utsugi. Hyperparameter selection for self-organizing maps. *Neural Computation*, 9(3):623–635, 1997.

[17] Jarkko Venna and Samuel Kaski. Neighborhood preservation in nonlinear projection methods: An experimental study. *Lecture Notes in Computer Science*, 2130:485–, 2001.

[18] Juha Vesanto. *Data Exploration Process Based on the Self-Organizing Map*. PhD thesis, Helsinki University of Technology, 2002.

[19] Thomas Villmann, Ralf Der, Michael Herrmann, and Thomas M Martinetz. Topology preservation in self-organizing feature maps: exact denition and measurement. *IEEE Transactions on Neural Networks*, 8(2):256–266, 1997.

[20] J. Zupan X. Li, J. Gasteiger. On the topology distortion in self-organizing maps. *Biological Cybernetics*, 70:189–198, 1993.