Explicit Magnification Control of Self-Organizing Maps for "Forbidden" Data

Erzsébet Merényi, *Senior Member, IEEE*, Abha Jain, and Thomas Villmann

***Abstract*—In this paper, we examine the scope of validity of the explicit self-organizing map (SOM) magnification control scheme of Bauer *et al.* (1996) on data for which the theory does not guarantee success, namely data that are n-dimensional, $n \geq 2$, and whose components in the different dimensions are not statistically independent. The Bauer *et al.* algorithm is very attractive for the possibility of faithful representation of the probability density function (pdf) of a data manifold, or for discovery of rare events, among other properties. Since theoretically unsupported data of higher dimensionality and higher complexity would benefit most from the power of explicit magnification control, we conduct systematic simulations on "forbidden" data. For the unsupported $n = 2$ cases that we investigate, the simulations show that even though the magnification exponent α_{achieved} achieved by magnification control is not the same as the desired α_{desired}, α_{achieved} systematically follows α_{desired} with a slowly increasing positive offset. We show that for simple synthetic higher dimensional data information, theoretically optimum pdf matching $(\alpha_{\text{achieved}} = 1)$ can be achieved, and that negative magnification has the desired effect of improving the detectability of rare classes. In addition, we further study theoretically unsupported cases with real data.**

***Index Terms*—Data mining, high-dimensional data, map magnification, self-organizing maps (SOMs).**

I. Potential Benefits, and Known Limits of SOM Magnification Control

ONE theoretically interesting and powerful data analysis aspect of self-organizing maps (SOMs) is the *map magnification*, which relates the density Q of SOM weights in the input space, to the probability density function (pdf) P of the input data by the following power law:

$$Q(\mathbf{w}) = cP(\mathbf{w})^{\alpha}. \tag{1}$$

Here, α is the *magnification exponent* and c is a constant [1]. Certain values of α have been associated with particular quantization or information theoretical properties [1], [2]. An SOM with $\alpha = 1$ maximizes information theoretic entropy, therefore, such mapping produces the best approximation to the pdf of the data with the given number of codebook vectors (SOM weights). $\alpha = 1/3$ for 1-D data corresponds to minimum mean-squared error quantization, or, in general, $\alpha = d/(d+p)$ leads to minimum mean-squared error quantization of d-dimensional data in p-norm [2]. Thus, magnification control enables the realization of different similarity concepts in the underlying cost function (description error) based on the respective p-norms. For example, near-zero values for α approximate the maximum norm in the description error. As suggested in [1] on the basis of biological observations of the "perceptual magnet" effect that enlarges the cortical representation areas for rare events (e.g., events of particular danger that an organism must be alert to), $alpha < 0$ can enable better categorization by allocating larger-than-proportional areas of the SOM lattice for low-frequency inputs. Regions of lower data probability become more accurately represented by prototypes, which can lead to detection of rare clusters. That, in turn, increases the chance of discovery. In data mining, this mechanism could alert for very small, "interesting," or "suspicious" groupings in data such as caused by mineralogical deposits of tiny spatial extent that may be indication of past life in Mars imagery, or by terrorist activity in data gathered by security agencies. As is known, the inherent property of the Kohonen SOM (KSOM) [3] is a map magnification of $\alpha = 2/3$ (under certain conditions) [4], [5]. This value of α is optimal in neither minimum distortion nor maximum entropy sense. An SOM variant called conscience algorithm [6] is constructed to achieve $\alpha = 1$, but cannot induce any other value. The conscience algorithm is based on heuristics, and although it works well in practice, theoretical proof does not exist for the achieved map magnification.

Controlling the magnification of self-organizing neural maps is, therefore, an extremely attractive possibility because various values of the magnification exponent can affect desirable quantization properties and serve specific data mining purposes.

The explicit SOM magnification control introduced by Bauer *et al.* [1] (referred to as BDH from now on) provided a powerful principled approach to obtaining a desired magnification exponent for 1-D data and for n-dimensional data whose components are statistically independent. To briefly quote the method, let V and A denote the input data manifold and the SOM grid of processing elements (PEs, nodes, or neurons), respectively, P and Q the pdf of V, and the pdf of the SOM weights in V, respectively. The SOM PEs are indexed by their (potentially multidimensional) grid locations $\mathbf{r}$. The weight attached to node $\mathbf{r} \in A$ is $\mathbf{w_r}$. For any $\mathbf{v} \in V$ input, the KSOM learning algorithm [7] selects a winner node $\mathbf{s}$ by

$$\mathbf{s} = \arg\min_{\mathbf{r}} \|\mathbf{v} - \mathbf{w_r}\| \tag{2}$$

Manuscript received July 22, 2005; revised June 22, 2006; accepted August 18, 2006. This work was supported in part by the Applied Information Research Program of NASA, Science Mission Directorate, under the Grants NAG9-10432 and NNG05GA94G.

E. Merényi is with the Electrical and Computer Engineering Department, Rice University, Houston, TX 77005 USA (e-mail: erzsebet@rice.edu).

A. Jain was with the Electrical and Computer Engineering Department, Rice University, Houston, TX 77005 USA. She is now with the Microsoft Corporation, Redmond, WA 98052 USA.

T. Villmann is with the Clinic for Psychotherapy, University of Leipzig, Leipzig 04107, Germany (e-mail: thomas.villmann@medizin.uni-leipzig.de).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNN.2007.895833

and then, updates the weights according to

$$\Delta \mathbf{w_r} = \epsilon h_{\mathbf{rs}}(\mathbf{v} - \mathbf{w_r}) \tag{3}$$

where the *neighborhood function* $h_{\mathbf{rs}}$ defines the degree of update for each weight. $h_{\mathbf{rs}}$ is typically a Gaussian function centered over the winner node, which we will assume here. In general, the neighborhood can be defined in many different ways. In (3), the learning rate ϵ is globally defined, i.e., it is the same for all PEs for a given time step and its value is independent of any local properties of the map. The key idea of the BDH learning is to make the learning rate dependent on the local input density so as to effect

$$\epsilon_{\mathbf{r}} \propto \epsilon_0 P(\mathbf{w_r}). \tag{4}$$

This goal is achieved by changing ϵ in (3) to

$$\epsilon_{\mathbf{s}}(t) = \epsilon_0 \left(\frac{1}{\Delta t_{\mathbf{s}}} \left(\frac{1}{\|\mathbf{v} - \mathbf{w_s}\|^d} \right) \right)^m \tag{5}$$

where m is a free parameter, t is the time step, and $\Delta t_{\mathbf{s}}$ is the time difference since the PE $\mathbf{s}$ won last. ϵ_0 is a constant and d denotes the effective dimensionality of the receptive field of $\mathbf{w_s}$. Note that $\epsilon_{\mathbf{s}}(t)$ is locally determined, but then, applied to *all* weight updates in the current step. By doing so, the local property of the map is propagated to the neighbors. With a derivation similar to that of Ritter and Schulten [4] in which the $\alpha = 2/3$ property of the KSOM was proven, Bauer *et al.* showed that the learning rate in (5) modifies the power law under (1) to

$$Q(\mathbf{w}) = cP(\mathbf{w})^{\alpha'} = P(\mathbf{w})^{(2/3)(1+m)} \tag{6}$$

where the free parameter m can be used for controlling the value of α'. Therefore, to achieve a desired magnification exponent α', $m = 3/2 * \alpha' - 1$ needs to be used in (5). The reader is referred to [1] for further details.

Most real data, of course, do not obey the conditions stipulated for BDH learning (listed in Section II), yet it is real-data scenarios that would benefit the most from explicit magnification control. Extension of analytical proof of the BDH scheme for such higher dimensional cases hinges on analytical proof of the SOM as seen from, e.g., [4], with which there has not been much luck so far. This, however, does not necessarily mean that the BDH algorithm should not work for more complex data than it is proven for.

This paper examines and extends the limits of the BDH through carefully constructed numerical simulations. It also extends and makes more precise our preliminary results published in [8] and [9].

II. Behavior of the BDH Scheme on 1-D and 2-D Data

The theory put forward by Bauer *et al.* [1] proves that the BDH algorithm will successfully induce the intended value of α for the following cases:

- 1-D input data;
- n-D data, $\mathbf{v} = (v_1, \ldots v_n)$, if and only if $p_{\mathbf{v}}(\mathbf{v}) = p_{v_1}(v_1) * \cdots p_{v_n}(v_n)$ (i.e., the pdf factorizes into the marginals).

We examine the possible validity of the BDH algorithm beyond these limited cases. We specify various values of the magnification exponent α as the desired target value to be induced by the BDH. In each case, α achieved by the converged SOM is calculated by a histogram-based method, as used in [10] and compared to the desired α. Details of the histogram method are discussed in [11]. We provide a basic description in Appendix A after the discussion in Section IV.

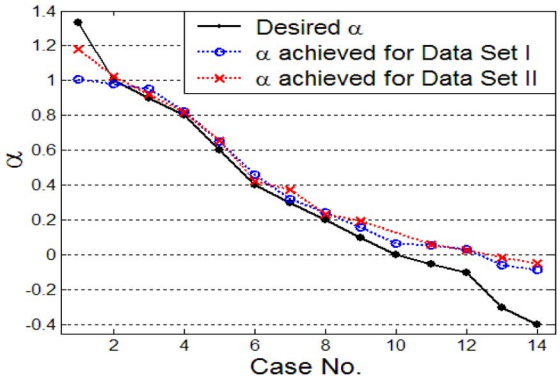

Fig. 1. Results of magnification control on 1-D data for data set I (open circles) and data set II (crosses). The graphs compare the values of the magnification exponent, α, obtained through the explicit magnification control scheme of Bauer *et al.* [1] with the desired values of α (filled dots) that were input to the BDH algorithm. The obtained values were derived by the evaluation of the converged SOM, similarly as described in [10], using the known pdfs of data sets I and II. Case number refers to separate runs of the BDH algorithm with different desired α values.

First, we confirm that the BDH works well for the previous, theoretically supported cases. Simulations with 1-D data, similar to those in [1], are shown in Fig. 1 for two 1-D data sets. Data set I is generated by $p(v) = 2v, v \in [0,1]$ and data set II is generated by $p(v) = 3v^2$, $v \in [0,1]$. As seen from Fig. 1, the values of the magnification exponent to be induced by the BDH are in tight correspondence with the actual values achieved, between $0 \leq \alpha_{\text{desired}} \leq 1$. The slight discrepancies can be attributed to the granularity of the quantization (100 SOM weights). The discrepancies are appreciable for $\alpha_{\text{desired}} < 0$ and for $\alpha_{\text{desired}} > 1$; however, it is worth noting that the deviations are not random: the α_{achieved} values seem to be attracted to 1 or 0 (for $\alpha_{\text{desired}} > 1$ and $\alpha_{\text{desired}} < 0$), respectively. The source of the obvious discrepancies outside of the $0 \leq \alpha_{\text{desired}} \leq 1$ range could be due to less accuracy in the calculation of α_{achieved} or instability of the magnification control. To our knowledge, this has not been investigated for the BDH. We know, however, from analyses for other magnification control approaches that instability regions may exists [12].

Next, we investigate the performance of the BDH algorithm on simple 2-D data. Experiment 1) will verify the BDH under the stipulated theoretical conditions of independency. Experiments 2) and 3) explore theoretically unsupported cases where data in 2-D are weakly and strongly correlated, respectively. The following cases, with different correlations (ρ) between the two dimensions of the data space, are evaluated.

1) **Data independent in the two dimensions:** $p_{\mathbf{v}}(\mathbf{v}) = p_{v_1}(v_1)p_{v_2}(v_2)$, generated according to the following pdf:

$$\begin{aligned} p_{v_1} &= 2v_1, \qquad v_1 \in [0,1] \\ p_{v_2} &= 2v_2, \qquad v_2 \in [0,1] \\ p_{\mathbf{v}}(\mathbf{v}) &= 4v_1v_2, \qquad v_1, v_2 \in [0,1]. \end{aligned} \tag{7}$$

As demonstrated in Fig. 2, the BDH magnification control works as advertised, for this data set.

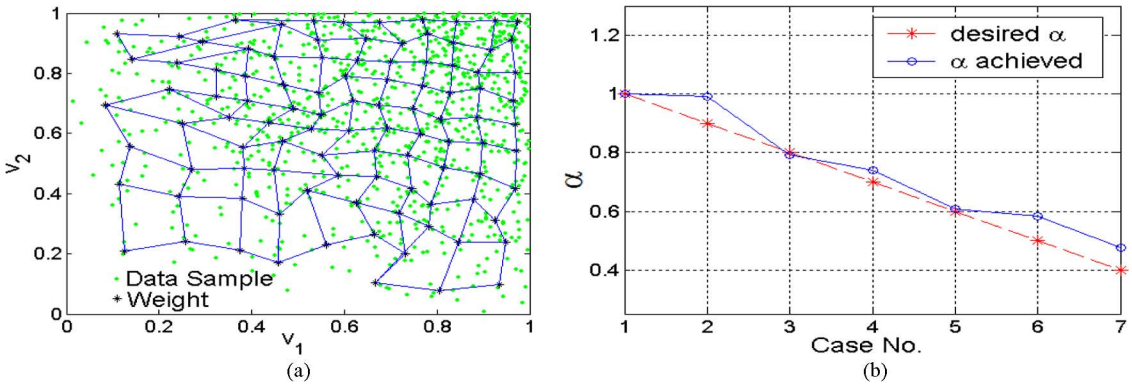

(a) (b)

Fig. 2. BDH simulation with the data set defined under experiment 1) in Section II, whose 2-D input samples $\mathbf{v} = (v_1, v_2)$ are such that v_1 and v_2 are independent. (a) 2-D input data samples (small green dots) and distribution of SOM weights (larger black dots), resulting from BDH with $\alpha_{\text{desired}} = 0.6$ after 2 000 000 learning steps. Weights adjacent in the SOM lattice are connected. (b) Comparison of the achieved α_{achieved} to α_{desired} shows very good agreement, as expected. The discrepancies are largely due to the fact that the theoretical results are asymptotic and we only have a finite number of PEs (100).

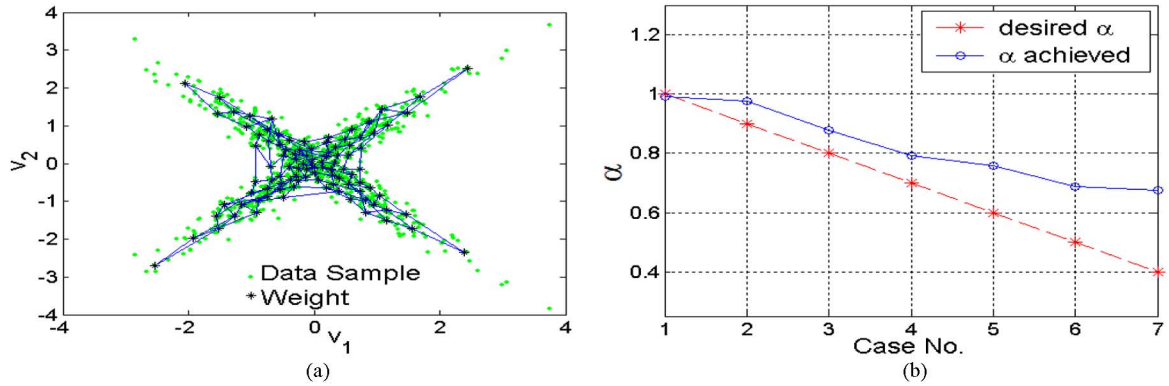

(a) (b)

Fig. 3. BDH simulation with the data set under experiment 2) in Section II, whose 2-D input samples $\mathbf{v} = (v_1, v_2)$ are such that v_1 and v_2 are weekly correlated, $\rho_{v_1 v_2} \ll 1$. (a) 2-D input data samples (small green dots) and distribution of SOM weights (larger black dots), resulting from BDH with $\alpha_{\text{desired}} = 1.0$ after 2 000 000 steps. Weights adjacent in the SOM are connected. (b) Difference between α_{achieved} achieved and α_{desired} increases in a predictable manner as α decreases from 1.

2) **Data weakly correlated in two dimensions:** $\rho \ll 1$. The data consists of two equal size subsets of 2-D samples, defined by

$$\begin{aligned} \mathbf{v} &= (v_1, v_2) \\ v_2 &= v_1 + n, \quad \text{for the first subset} \\ v_2 &= -v_1 + n, \quad \text{for the second subset} \end{aligned} \tag{8}$$

where $n = \mathcal{N}(0, 0.0625)$. v_1 and v_2 are weakly correlated with the correlation coefficient $\rho_{v_1 v_2} = 0.0044$. From Fig. 3, it can be seen that α achieved and the desired α are almost equal at $\alpha = 1$ and the two values differ increasingly but in a predictable manner as α decreases. This is a stronger result than available from the theory, as the theory only guarantees successful prediction if and only if v_1 and v_2 are independent.

3) **Data strongly correlated in two dimensions:** $\rho \approx 1$. This data set consists of 2-D samples generated by

$$\begin{aligned} \mathbf{v} &= (v_1, v_2) \\ v_2 &= v_1 + n \end{aligned} \tag{9}$$

where $n = \mathcal{N}(0, 0.25)$. The correlation coefficient is $\rho_{v_1 v_2} = 0.9026$. In this strongly correlated case, even though α achieved by the map differs from the desired α value, there is a clearly observable trend that the achieved values of α are systematically decreasing, following the desired values with a shift that increases slowly with decreasing α values (Fig. 4). This is again a stronger result than the theory provides.

Additional 1-D and 2-D cases are analyzed and more details given on analysis considerations in [11].

III. BDH on Higher Dimensional Data

The BDH algorithm has no analytical justification for data with two or more nonseparating dimensions. However, the systematic results from numerical simulations on 2-D "forbidden" data in Section II encourage investigation of the BDH behavior on higher dimensional data.

Our methodology to examine the performance of the BDH algorithm on 1-D and 2-D data sets was to induce a certain value of α, evaluate the α achieved by the converged SOM, and compare the two. However, evaluation of α is not an easy task in general, especially when the input pdf is unknown, as is most commonly the case. The evaluation of α involves the estimation of the pdf of the data and that of the SOM weights. So far, we were using a histogram-based method for α evaluation [10]. This method becomes inapplicable for high-dimensional data as

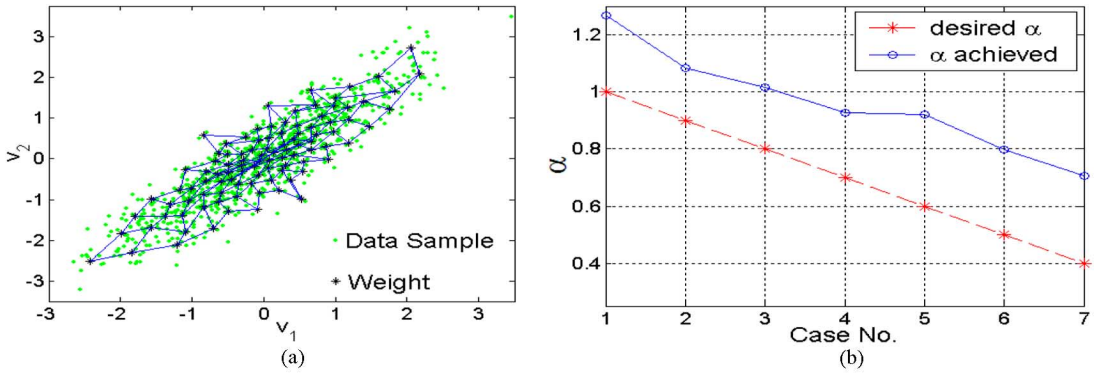

Fig. 4. BDH simulation with the data set under experiment 3) in Section II, whose 2-D input samples $\mathbf{v} = (v_1, v_2)$ are such that v_1 and v_2 are strongly correlated, $\rho_{v_1 v_2} \approx 1$. (a) 2-D input data samples (small green dots) and distribution of SOM weights (larger black dots), resulting from BDH with $\alpha_{\text{desired}} = 1.0$ after 1 000 000 steps. Weights adjacent in the SOM are connected. (b) Difference between α_{achieved} achieved and α_{desired} shows a clear trend of a shift that slowly increases with decreasing α_{desired}.

the number of samples required for the estimation of the pdf increases exponentially with dimensionality. For the 6-D and 8-D data in this paper the required number of samples for pdf estimation could easily be over ten million, a number that far exceeds the total number of our data samples in each case. A summary of requirements of pdf estimation for higher dimensional data is given in [11]. Here, we have to restrict ourselves to giving a brief introduction, along with further pointers to related issues and references in Appendix A.

There exist several neural approaches to estimate the pdf of the data on the prototype level. Prototype-based approaches are the DeSieno conscience learning [6], frequency-sensitive competitive learning [13], the magnification-controlled neural gas [14], kernelized variants of the SOM [15], and a Gaussian mixture approach based on SOM [16], to name just a few. However, comparison could be difficult. For example, VanHulle reports, based on 1-D experiments on artificial data, that conscience learning could be unstable for precise pdf estimation because the conscience may be taken up by only a few prototypes [15]. Also, there is no theoretical proof for conscience learning. Frequency-sensitive competitive learning is theoretically proven only for the 1-D case and we do not know of numerical verification for higher dimensions.

For higher dimensional cases, we evaluate the performance of the algorithm indirectly: by observing the resulting map and comparing to independently known properties of the input data. We concentrate on cases where we can evaluate the magnification results in some meaningful way, from an application's point of view. One special case is pdf estimation (forcing $\alpha = 1$), another is negative magnification. Both can be particularly useful for complex, high-dimensional data if systematic experiments indicate a *predictable* behavior of the BDH for such "forbidden data." For example, if successful forcing of negative magnification can be shown, the magnification effect is useful for data mining regardless of the exact value of the induced magnification exponent, which we cannot calculate.

A. BDH-Induced Negative Magnification on 6-D Synthetic Data

Finding very small classes in a data set is a challenging task. Input classes of rare occurrence may have little or no representation in the output map when KSOM is used. As shown for 2-D factorizing data in [1], application of BDH with $\alpha < 0$ results in negative magnification: The areal representation of low-frequency input samples becomes enlarged in the SOM. This is promising for detection of rare classes. Does it, however, work for higher dimensions?

We use two synthetic six-band (spectral) images with known cluster structure, described in Fig. 5, to show that negative magnification can be induced on this type of "forbidden" data. A spectral image consists of n coregistered image bands, each

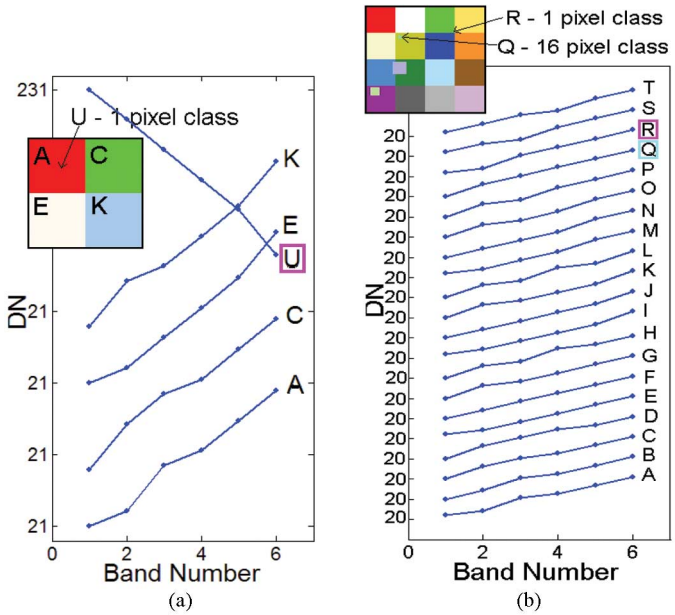

Fig. 5. The 6-D synthetic image data sets: the known class structure and spectral signatures of the classes. Each spectral signature is displayed with an offset for clear viewing. The data numbers (DN) on the y-axis indicate the values of the first band of each respective spectrum. The spatial area of both images is 128×128 pixels. (a) Five-class image data set with four large square classes (A, C, E, and K, with 4095, 4096, 4096, and 4096 pixels, respectively), and one one-pixel class U. The spectral signature of class U is dramatically different from the rest. (b) The 20-class image data set, with 16 nearly equal area (32×32 pixels) classes A–P, and four additional very small classes: R, Q, S, and T contain 1, 16, 64, and 128 pixels, respectively. Classes Q–T take away the corresponding number of pixels from those of the A–P classes in which they are embedded.

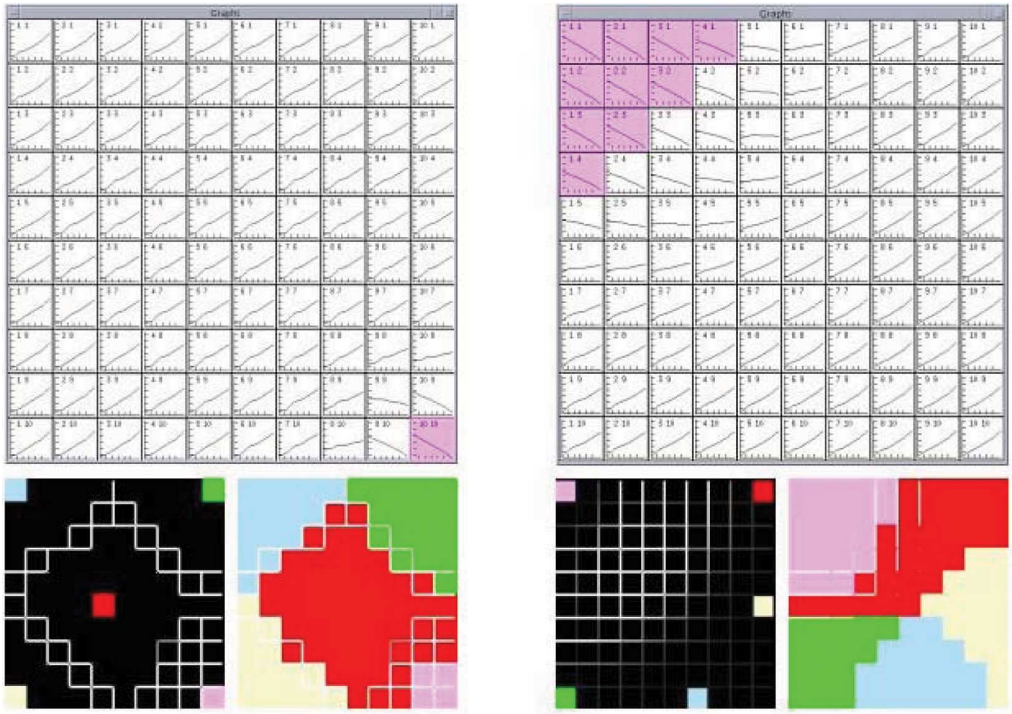

Fig. 6. Magnification results with the five-class data set. (a) Using KSOM. Top: Weight vectors of the 10×10 SOM. Only one PE represents the rare class U in the lower right corner. Bottom Left: Clusters identified in the map by visualizing the weight distances of adjacent PEs (the darker the fence between two PEs, the smaller the difference between the corresponding weights). Bottom Right: This figure shows which class each weight vector is closest to, which complements the information on the left. (b) Using BDH with $\alpha = -0.8$. Top: Weight vectors in the 10×10 SOM. The rare class U is now represented by ten PEs. Bottom Left: Clusters identified in the SOM. Bottom Right: The map of nearest class for each weight vector.

of which is taken at a different wavelength. Every pixel of the image, therefore, is characterized by an n-dimensional vector, called the spectrum, similar to the data in Fig. 5. In real spectral images, taken of material surfaces, the spectrum carries compositional information about the material in the respective pixel. See, for example, [17] for more detail. Spectral images are powerful information sources and are used in many areas of scientific research, business, industry, defense systems, etc. Detailed and precise exploitation of such data is of great interest. One especially valuable capability is the discovery of small, interesting groups of data.

The synthetic data sets used in this section are 128×128 pixel images, with a 6-D vector associated with each image pixel. These 6-D vectors are the input patterns for the SOM. Both images are artificially divided into rectangular areas within which the spectral signatures are the same. The spatial distribution of the various classes, along with their mean spectral signatures, are described in Fig. 5. The two images represent two levels of data complexity. The five-class data is illustrated in Fig. 5(a). Class U is a rare class with only one data point and with a spectral signature that is very different from the signatures of the other four classes. The rest of the classes have 4096 or 4095 data points each. In this data set, correlation coefficients between the different dimensions range from $0.004(\rho_{v_2 v_3})$ to $0.9924(\rho_{v_3 v_6})$, which renders it a "forbidden" case for application of BDH. When a 10×10 KSOM is used, the rare class U is represented by only one PE in the SOM [Fig. 6(a)]. BDH with $\alpha_{\text{desired}} = -0.8$ magnifies the rare class in the map: In Fig. 6(b), it is represented by ten PEs.

The choice of $\alpha_{\text{desired}} = -0.8$ comes from an extrapolation of the curves in Fig. 4, from which we estimated that a value of $0 > \alpha_{\text{desired}} > -0.5$ will likely induce $\alpha_{\text{achieved}} \approx 0$, so, in order to ensure a significant negative induced value, we need to specify $\alpha_{\text{desired}} \ll -0.5$. On the other end, Fig. 1 cautions us not to choose too large negative values. This extrapolation is admittedly only a hypothesis at present, not only because the simulations in Fig. 4 were not conducted for negative values but also because even if we had values for 2-D data assuming the same behavior for higher dimensional data would still remain a hypothesis. Evidence of successful forcing of negative magnification may motivate extension of the Fig. 4 simulations to the negative range in future work.

The synthetic image data set with 20 classes is similar to the five-class image data set in concept, but it has 20 classes, as shown in Fig. 5(b). Two of the classes, marked R and Q are very small, with only one and 16 data points, respectively. Here, the signatures of the rare classes are very similar to those of the larger classes, unlike the contrasting signature of the class U in the previous case. Trying to find these rare classes among the 20 classes using the same 10×10 SOM is a larger data mining challenge than finding the U class in the five-class image. Correlation coefficients between the different dimensions range from $0.0081(\rho_{v_5 v_4})$ to $0.5641(\rho_{v_4 v_2})$, so this too is a "forbidden" case for BDH according to the available theory. Clustering this data set using KSOM is depicted in Fig. 7(a). The rare classes are detectable but each is poorly represented, only by a single PE. Also, because of the tight quarters (the same SOM size for more classes), the separating fences (the contrasts between the

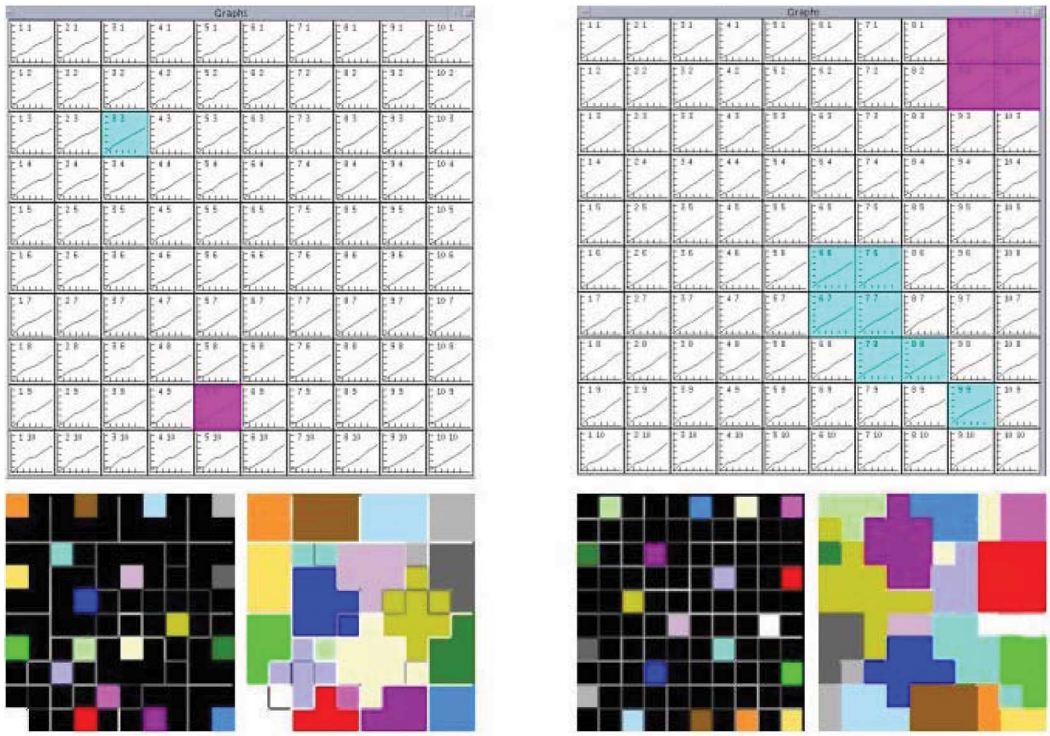

(a) (b)

Fig. 7. Magnification study with the 20-class image. (a) Using KSOM. Top: Weight vectors in the 10×10 SOM. Only one PE is allocated for each of the rare classes R and Q. Bottom Left: Clusters identified in the SOM. Fence intensity from black to white is proportional to the distance between the corresponding weights. Bottom Right: This figure shows which class each weight vector is closest to, which complements the information on the left. (a) Using BDH with $\alpha_{\text{desired}} = -0.8$. Top: Weight vectors in the 10×10 SOM. Now, four and seven PEs represent the rare classes R and Q, respectively. Bottom Left: Clusters identified in the SOM. Bottom Right: Map of closest class for each weight vector.

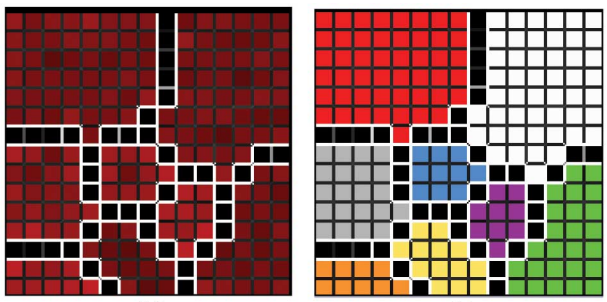

(a) (b)

Fig. 8. The 6-D eight-class data set as represented by an SOM that learned by the conscience algorithm. (a) Cluster boundaries, visualized as the distance of weights of adjacent PEs, similarly to that of the U-matrix [18] except we compute and display the distances to all eight neighbors separately, by the fences on the edges and corners of the PE grid cells. White is high fence (large dissimilarity), black is low fence (great similarity). The SOM has a 15×15 square grid. Each grid cell is shaded by an intensity of red proportional to the number of data points mapped to the PE in that grid cell. Black grid cells between the strong fences indicate that the receptive fields of the corresponding PEs are empty. The fairly uniform density over the nonempty PEs indicate a good approximation of maximum entropy mapping. The entropy of this map (relative to the theoretical maximum) is 0.998 for the active PEs and 0.962 for all PEs. (b) Known class labels superimposed over the PE grid cells. Both representations show that the PEs (and SOM weights) are divided among the classes in proportion to the sizes of the classes: A and B (red and white) contain 4096 data points each, C and O (green and grey) 2048, and D, H, I, and M have 1024 points. The corresponding number of designated PEs are A:48, B:49, C:25, O:21, D:13, H:9, I:10, and M:9. The deviations from the exact $4:2:1$ proportions can be due to the small size of the SOM, integer arithmetic, and the formation of intercluster gaps.

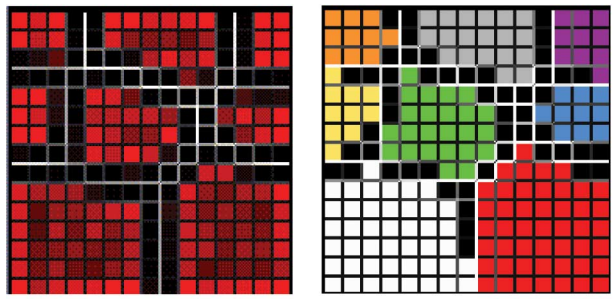

(a) (b)

Fig. 9. Eight known classes of the same 6-D patterns as in Fig. 8, as represented by an SOM that learned via the BDH magnification control, forcing $\alpha = 1$. Since this is a "forbidden" data set with strong interdimensional correlations, inducing $\alpha_{\text{desired}} = 0.7$ effectively produced $\alpha_{\text{achieved}} = 1$, as suggested by the experiments in Fig. 4. (a) Cluster boundaries and data density are visualized the same way as in Fig. 8. The entropy of this map (relative to the theoretical maximum) is 0.96 for the active PEs and 0.914 for all PEs. Note that many of the dark grid cells on either sides of single line fences have data points mapped to them albeit few. (b) Known class labels superimposed over the PE grid cells. Both representations show that the PEs (prototype vectors) are divided among the classes in proportion to the sizes of the classes, similarly to the conscience algorithm results in Fig. 8: A and B (red and white) contain 4096 data points each, C and O (green and grey) 2048, and D, H, I, and M have 1024 points. The corresponding number of designated PEs are A: 47, B: 44, C: 24, O: 19, D: 10, H: 10, I: 10, and M: 10. The deviations from the exact $4:2:1$ proportions can be due to the small size of the SOM, integer arithmetic, and the formation of intercluster gaps.

PE weights) are less pronounced, which makes the small classes less discernible. Fig. 7(b) shows that BDH with $\alpha_{\text{desired}} = -0.8$ magnifies the rare classes in the map: The 1-pixel class R is now represented by four PEs in contrast to one PE in the map formed by KSOM, and class Q is represented by seven PEs. The detectability of the rare classes has been increased by magnified

areal representation as well as by higher contrast in the weight differences.

The aim of negative magnification is to enhance the SOM areal representation of low-density data regions by forcing the allocation of more prototypes to those regions. This may improve the detectability of unknown rare classes. Naturally, the representation of the nonrare classes in such a map will be somewhat repressed as is apparent from the bottom of Fig. 7. The map obtained with $\alpha_{\text{achieved}} < 0$ and the one obtained with $\alpha_{\text{achieved}} \approx 1$ together provide a more complete picture of the structure of a data set.

In order to further chart the behavior of the BDH, we now present additional controlled experiments with synthetic as well as real data sets. Beside $\alpha < 0$, of special interest is the case of $\alpha = 1$ because it effects maximum entropy quantization, and thus helps faithful mapping of the input data structure.

B. *Inducing $\alpha = 1$ Magnification on 6-D Synthetic Data*

After demonstrating that negative magnification can be achieved on higher dimensional data, now, we show that $\alpha = 1$ can also be achieved by BDH, fairly accurately, on similar data. $\alpha = 1$ is a special case in that the heuristic conscience algorithm by DeSieno [6] is believed to achieve maximum entropy quantization, therefore, if we verify that, we can compare the properties of the SOM obtained by BDH magnification with an SOM obtained by conscience, to determine if the BDH realized $\alpha = 1$. In addition, we can compare the entropies of the two maps.

The simulations presented in this section serve two purposes. One is to demonstrate that the conscience algorithm indeed produces pdf matching, and the second is to show that by forcing $\alpha = 1$ the BDH can accurately model the pdf, at least on the level of the cluster structure of the data. We do this by a direct comparison of the clusters detected by the SOMs with the known labels of data points (the known clusters) in the artificially generated data. In many data mining pursuits, among them remote sensing and medical applications, the main objective is to find interesting, relevant groupings in the data, but not necessarily a finer approximation of the pdf. Such challenges can be met without a very precise evaluation of α.

The data set we use for this purpose is an eight-class synthetic image consisting of six image bands, similar to the data described in Fig. 5 except that here eight spectral types are distributed over subareas of the 128×128 pixel image in the following manner: Classes A and B each cover 4096 pixels, classes C and O are each 2048 pixels, and classes D, H, U, and M have 1024 pixels. Gaussian noise, about 10% on average, was added to create more realistic variations within the spectral classes. The conscience algorithm is expected to map each of these classes onto areas in the SOM that are proportional in size to the areas of the classes, namely classes D, H, I, and M should each occupy half as many PEs than either of class C or O, and A and B both should be represented by twice as many PEs as C or O, and by four times as many PEs as any of D, H, I, and M. Fig. 8 shows that indeed, this is the case within the accuracy allowed by the size of the 15×15 square SOM grid, by integer arithmetic, and taking into account the empty PEs that form dividing gaps between clusters. The clusters are captured using a somewhat modified version of the U-matrix [18], in that we compute and visualize the distances to the neighbor weights on either side of a PE separately, and do this also for the diagonal neighbors (we use an eight-neighbor neighborhood). The intensity of the "fences" between PEs is proportional to the distance between weights, on the black to white grayscale. White is large difference and black is great similarity. Clusters are clearly outlined by the white fences. Fences within the clusters are uniformly very low (virtually 0). Out of 225 SOM PEs, classes A and B cover 48 and 49 PEs, C and O cover 25 and 21 PEs, and the smallest four classes occupy 13, 9, ten, and nine PEs, respectively. The 41 black PEs belong to intercluster gaps. The largest deviation from the expected values occurs in the smallest classes, which is understandable considering that just one additional PE in each of H and M, taken away from D would even out the areas to ten and 11 PEs each. The SOM was run for two million steps to ensure convergence, but the cluster structure in Fig. 8 was already formed after 2–300 000 steps. It is noteworthy that the precision of the pdf estimation is better than the level of the cluster structure: The intensity of the red color in each PE cell in Fig. 8(a) is proportional to the relative winning frequency of the respective PE, and as is evident the red shades are fairly uniformly distributed within clusters.

Since a maximum Shannon entropy (equiprobabilistic) mapping provides the most faithful match of the pdf by the given number of prototypes, the entropy of the SOM is a good indicator of the quality of the learned quantization. The Shannon entropy of the SOM is given by

$$I = \sum_{i=1}^{N} P(j) \log_2 P(j) \tag{10}$$

where j is the PE index in the SOM ($j = 1, \ldots, N$), $P(j) = (I_j / I_N)$ is the probability that PE j wins an input sample, where I_j is the number of input samples mapped to (won by) PE j, and I_N is the total number of input samples. The higher the value of the map entropy, the closer the mapping is to equiprobabilistic. In an equiprobabilistic map $P(j) = (1/N) \; \forall \; j$, and in that case, I would have a maximum value of $log_2 N$. We compute a normalized entropy by dividing I by the theoretical maximum $log_2 N$ (thus, the maximum possible value of the entropy of any map is one), so that entropies can be compared across SOMs of different sizes. The (normalized) entropy of the conscience SOM in Fig. 8, is 0.998 for the active PEs, and 0.962 for all PEs.

After this verification of the expected performance of the conscience algorithm, we can evaluate the magnification performance of the BDH on the same data. The pairwise correlations of the various dimensions of this data set are typically strong, (most are between 0.3–0.8, and only two are less than 0.05). Based on the 2-D experiments summarized in Fig. 4, and lacking any guidance from theory or other works, we hypothesize that a similar trend (as in Fig. 4) may apply for higher dimensional forbidden data, i.e., in order to produce a certain α_{achieved} value one needs to request an α_{desired} value that is offset according to some observed function. Since this paper is the first one to chart such functions, the most reasonable assumption we can make is the functional relationship we charted for 2-D data with strong

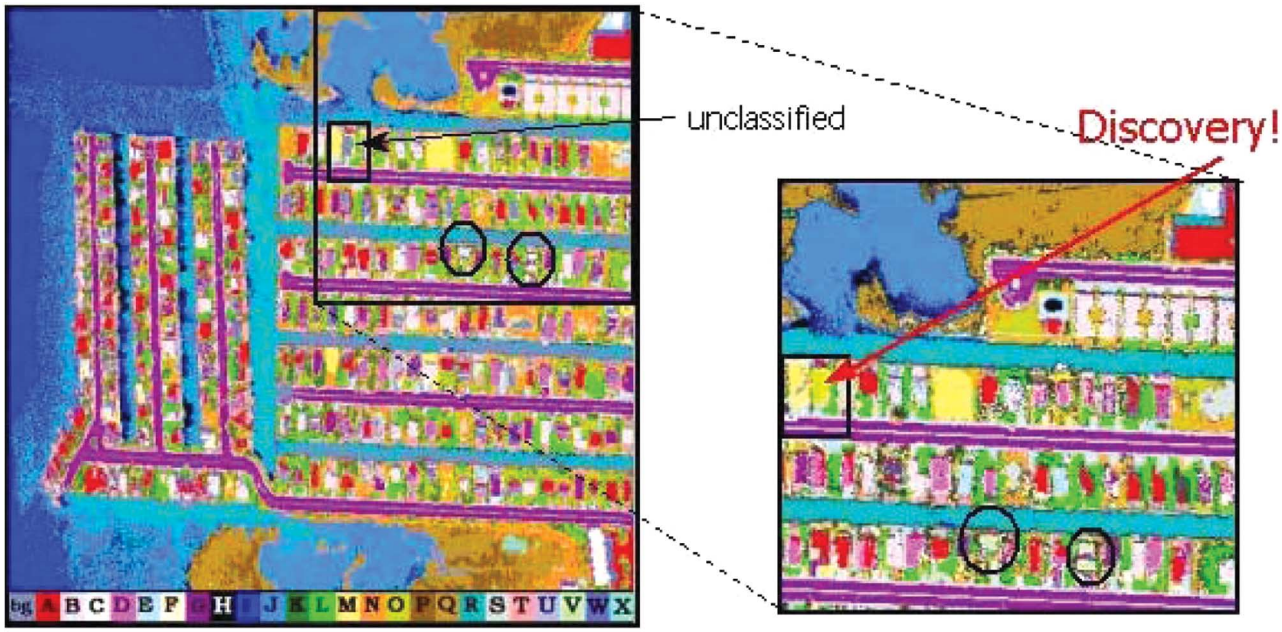

Fig. 10. Comparison of supervised classification and BDH clustering with $\alpha < 0$, on an eight-band spectral image of Ocean City, MD. (a) Earlier supervised classification that satisfactorily mapped 24 known cover types of interest, based on verification against ground truth. Shown centered in the small black rectangle within the framed upper right quadrant is an unclassified gray spot (the color of the background, "bg") apparently of the shape of a building, to the right of a yellow rectangular patch. He pale aqua spots in the black ovals belong to the rare class V, known at the time of the supervised classification. (b) SOM clustering using BDH magnification control with $\alpha_{\text{desired}} = -0.8$ on the upper right quadrant of the image. First, notice that the agreement between the supervised class map and this cluster map is striking, which inspires confidence in the clustering. Second, notice that the spot that remained unclassified in the supervised map is now filled exactly and with a color (greenish-yellow) that is different from all previous class colors: The spectral signature of this area is distinct. We discovered a new class. Moreover, this cluster only occurs at this location, and nowhere else: We discovered a small rare class. Fig. 11 shows the SOM view of this discovery. Note that the formerly known rare class V (pale aqua spots in black ovals) became better defined. Class spectral signatures are plotted in Fig. 12.

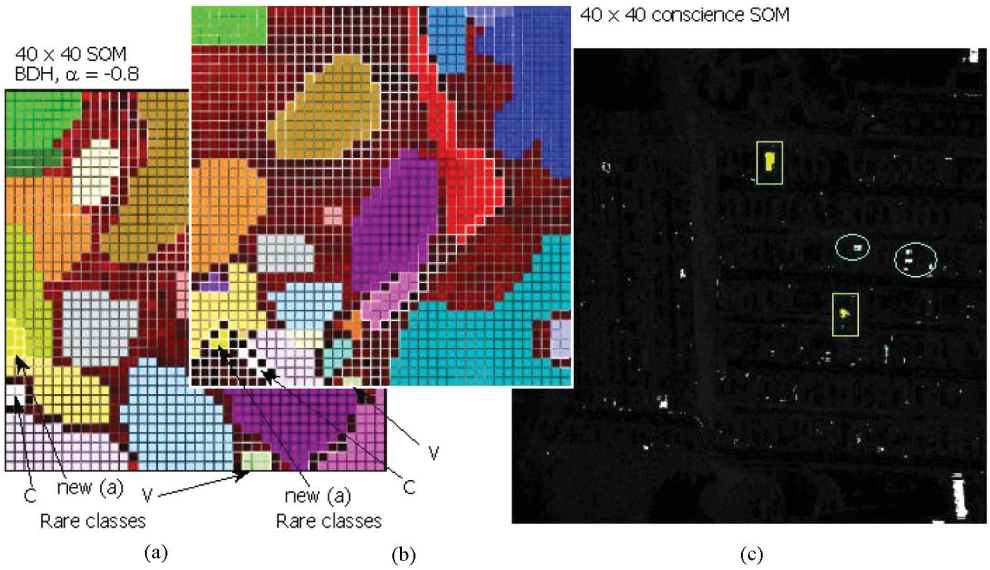

Fig. 11. Comparison of SOMs developed by BDH versus conscience learning. (a) SOM learned by BDH $\alpha_{\text{desired}} = -0.8$, using the upper right quadrant of the 512×512 pixel Ocean City image, shown framed in Fig. 10. (b) SOM learned by the Conscience algorithm ($\approx \alpha = 1$), using the entire Ocean City image. (c) Rare classes in the image. The newly fond cluster (indicated in the rectangles) only occurs in one other spot in the entire image, outside of the BDH-discovered location within the upper right quadrangle. It is apparent, as explained in Section III-C, that the rare clusters are magnified in the BDH SOM in comparison to the Conscience SOM.

dimensional correlations, and correct our assumption according to the outcome of the simulation, if needed. By extrapolation of Fig. 4, we assume that in order to achieve $\alpha_{\text{achieved}} = 1$ the BDH needs to be run with a choice of $\alpha_{\text{desired}} = 0.7$. The SOM, formed by BDH magnification forcing of $\alpha_{\text{achieved}} = 1$ in this way (Fig. 9), shows that it has a very similar area distribution over the eight classes as the conscience SOM in Fig. 8 (again, with similar accuracy considerations). This indicates that the BDH fairly closely achieved the desired maximum entropy mapping on this "forbidden" data set. The normalized entropy

of this map is 0.96 for the active PEs and 0.914 for all PEs. We note that the normalized entropy of the BDH SOM is lower than that of the conscience SOM. This is also obvious from the density of the cells that are close to the cluster boundaries in Fig. 9. By preliminary observation, the entropy of the BDH SOM very slowly increases over a very long run, however, it does not reach the entropy of the conscience SOM in approximately seven million steps. We think that this discrepancy could be due to other circumstances that the BDH may be more sensitive to, such as the combination of the number of clusters and the size of the SOM lattice, or the learning schedule including the decay rate of the neighborhood. This needs to be evaluated in subsequent work.

One difficulty with the BDH is the estimation of the data and weight density at the locations of the weights [see (4)] [19]. The estimation [see (5)] uses a rough approximation of the volume of the receptive field of $\mathbf{w_s}$, $\|\mathbf{v}-\mathbf{w_s}\|^d$, where d is the *effective dimensionality* of the receptive field. This dimensionality may vary from weight to weight and is usually unknown, therefore, any global value of d is a crude estimate. For that reason, we can expect the BDH to work well with a range of d values (within $[1, n]$). Therefore, we chose $d = 2$, which is the squared Euclidean distance, and thus, has the advantage of saving on computational cost. We add that the relative insensitivity of the BDH to the value of d agrees with our experiences with a limited number of trials.

C. Finding Rare Clusters in a Real Spectral Image

As we mentioned in the introduction, magnification control can lead to the detection of rare clusters. In particular, for inverted (negative) magnification, prototypes are preferentially placed in low-density areas of the data space, resulting in more accurate description of those low-probability regions. Low-density regions may contain meaningful, separate clusters, which are not detectable if there are not enough representing prototypes but may become "visible" through better representation using negative magnification. We will demonstrate this effect with a real–world data example, which is an urban remote sensing spectral image of Ocean City, MD. We use a 512×512 pixel, eight-band subset of the Ocean City image to study the effect of forced negative magnification. This data set also has high pairwise correlations, the magnitudes of which are mostly between 0.5 and 0.95. We cannot compute the value of α_{achieved}, but we can compare the appearance of known small classes in the BDH SOM and in an SOM that learned with the conscience algorithm to see if the rare classes occupy larger areas in the BDH SOM than in the conscience SOM. In addition, we look for previously unidentified clusters. An earlier supervised classification that was independently verified against ground truth provides the knowledge of existing clusters in the data. Fig. 10 demonstrates the discovery of one new—very small—cluster. It also shows another small cluster (pale aqua, class V) that was known at the time of the earlier supervised classification, but was more definitely outlined by BDH clustering. Fig. 11 compares the two SOMs. Shown in Fig. 11(a) is the 40×40 SOM formed by BDH learning with $\alpha_{\text{desired}} = -0.8$, using only the upper right quadrant of the image (framed in Fig. 10), i.e., 1/4 of the data. The newly

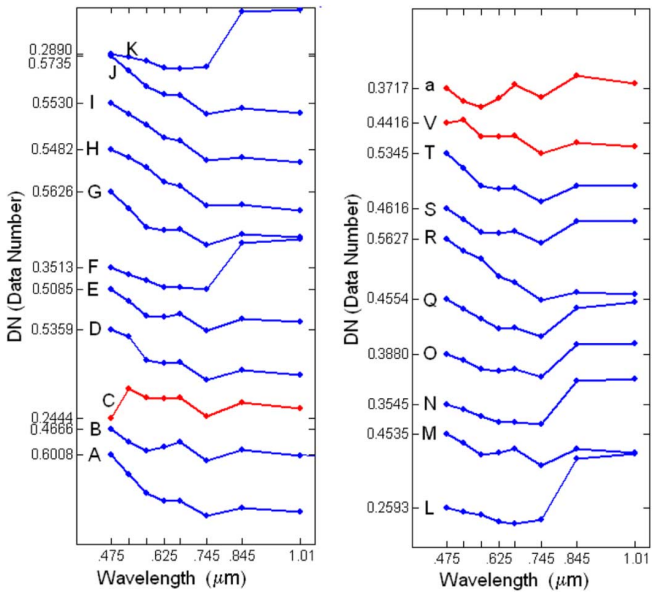

Fig. 12. Mean spectral signatures of the clusters identified in the Ocean City spectral image. Spectra are offset for clarity. The DN value on the vertical axis indicates the value in the first spectral band. The newly discovered cluster (greenish-yellow in Fig. 10) is labeled "a" on top of the right panel. The spectral dissimilarity of this cluster with the rest is obvious. The other rare classes discussed in Section III-C are class C and V.

discovered rare cluster (greenish-yellow) is indicated by the middle arrow. The spectral signature of this cluster is distinctively different from all other clusters, as seen in Fig. 12. Also indicated are two other small clusters that correspond to the previously known V (pale aqua) and C (white) classes from the supervised class map in Fig. 10. The 40×40 SOM produced by conscience learning, using the *entire* image, is in Fig. 11(b). The greenish-yellow cluster was hard to see in this map, and was only "discovered" because we looked for it based on the BDH discovery. This rare cluster covers only three PEs in the conscience SOM in contrast to seven PEs in the BDH SOM where it is also contoured by better developed "fences." Similarly, the previously known small V class is represented by four PEs in the conscience SOM versus six PEs in the BDH SOM, even though the conscience SOM was learned with four times as many data points, including more occurrences of the V class in the large image outside the upper quadrangle, and as seen from Fig. 11(c). The previously known white class occupies four PEs in both SOMs, in spite that within the 1/4 subimage used for BDH clustering, the white class only occurs in a small rectangle (not circled) in the upper right corner of Fig. 11(c), while there are many more white class pixels in the entire image used for the conscience SOM training [most notably the long vertical rectangle in the lower right corner of Fig. 11(c)]. These observations clearly indicate that, compared to Conscience the SOM, the BDH preformed negative magnification.

IV. Discussion of Problems and Conclusion

We mention here some issues related to this paper but not discussed previously. All of our evaluations assume topologically correct mapping by the respective SOMs (that the learning does not result in "twisted" maps). Ensuring and verification of this is not easy, especially for high-dimensional, complex

data. Although supporting theories exist [20], [21], those, similarly to the BDH itself, need systematic evaluation for higher dimensional data. In these applications, we used data of which we knew the cluster structure from independent investigations, thus we could verify the validity of the clusters identified by the SOM.

In our experience, finding "good" learning parameters is more difficult for the BDH algorithm than for KSOM or for the conscience algorithm. One especially important question seems to be the value of negative α: Can one "overdistort" the rest of the map while discovering rare clusters, by using a too large negative value? In several cases of real data, we saw that while extremely rare classes were beautifully magnified and further separated into meaningful subclusters such as in discovering rare Martian mineral types [22], the rest of the data were mapped onto thin filament-like groups of PEs leaving most of the SOM grid empty.

A related item is that it is hard to estimate what α_{desired} value should be used as input to the BDH, in order to achieve a specific α value. The experiments presented in Fig. 4 suggest a trend but we caution against straight extension of those values to higher dimensional data. Some preliminary experiments with the eight-class data seem to indicate that $\alpha = 1$ may successfully be induced within a range of α_{desired}. We found that to be the case for $0.6 \leq \alpha_{\text{desired}} \leq 0.9$, *with this particular data set*. It may be worth further simulations to get a firmer grasp on stability regions (that we suspect may exist), especially for negative values. The current lack of capability to evaluate α_{achieved} from the SOM of high-dimensional data makes this problem even harder.

The previous considerations suggest that magnification control may influence other properties of the SOM, namely, the ability of topographic mapping. Thus, if one is interested in optimizing both magnification and topographic mapping, one has to balance between these aspects and to prove the topography by appropriate tools. Alternatively, the magnification scheme can be integrated directly into the growing self-organizing map (GSOM) approach [23].

The estimation of the data and weight density at the locations of the weights [see (5)] is problematic in the BDH algorithm, as we mentioned in Section III-B. A more accurate assessment may be possible by the estimation of the volume of the Voronoi cells, which may be done through the size of the respective receptive fields (that the algorithm can record continuously) and may result in marked improvement of the algorithm. This is a follow-up task for us, worth pursuing in our opinion, especially since the BDH approach seems to be the most convenient method to control the magnification (see [12]).

Of course, the most interesting and outstanding issue is the theoretical justification of magnification control for higher dimensional cases. Kohonen proposed a theoretical approach for higher dimensions by modeling the local receptive field densities with simple hyperspheres of constant data density, which enables local factorization [19]. However, this assumption is generally not true for real data. In addition, Kohonen's considerations require that the SOM dimensionality match that of the data. A modification of the SOM winner determination by Heskes [24] leads to an energy function for the learning of SOMs, valid also for higher dimensions, but derivation of magnification properties seems to be very difficult because in the Heskes scheme the neighborhood function influences the winner selection. If we were to follow the derivation of Ritter and Schulten [4], a recursive equilibrium equation would result for which no solution exists at present. Another extension of SOMs is the neural gas paradigm whose learning dynamics also has an energy function [25]. For the neural gas, analytical proofs exist for several different magnification control paradigms including the BDH learning. However, since the energy function is obtained by the modification of the neighborhood function such that it is evaluated in the data space, the neighborhood structure among PEs is lost. The difficulties illuminated by these investigations further motivate the numerical simulations we described in this paper.

In conclusion, we presented systematic experiments with the map magnification control by BDH [1], on data for which the BDH scheme is not supported by existing theory. Based on our observations of the systematic BDH behavior on 2-D nonseparating "forbidden" data, we were able to induce maximum entropy ($\alpha = 1$) mapping and negative magnification on 6-D synthetic data. We also showed that negative magnification worked on 8-D real image data and helped enhance the areal representation, and thereby, the discovery of rare clusters. While the range of our studies is too limited to draw definite conclusions, the simulations indicate *consistent* behavior of the BDH on some set of "forbidden" data. This encourages further simulations to investigate the predictability of the BDH for potential analyses of complex, high-dimensional data. That, in turn, perhaps can inspire more theoretical studies and algorithm development.

APPENDIX
HISTOGRAM-BASED METHOD FOR EVALUATION OF α

The following power law relates the density of weights in the input space $Q(\mathbf{w})$ to the pdf $P(\mathbf{w})$ of the input samples

$$Q(\mathbf{w}) = P(\mathbf{w})^{\alpha} \times \text{const} \tag{11}$$

where $Q(\mathbf{w})$, the density of weights in the input space, is the number of reference vectors in a small volume $d\mathbf{w}$ of the input space. If both sides of (11) are divided by N_W, the total number of weight vectors, on the left, we get $Q(\mathbf{w})/N_W$, the pdf of the weight vectors, and on the right, const/N_W can be absorbed into a new constant $= \text{const}/N_W$. From now on, we will use the same power law as in (11), but with the understanding that $Q(\mathbf{w})$ now denotes the pdf of the weight vectors.

From (11), it is clear that finding the value of $\alpha\,(\alpha_{\text{achieved}})$ requires the estimation of the two pdfs. The following histogram-based method can be used to estimate the pdfs and evaluate α.

Let p_i and q_i be samples of the two densities $P(\mathbf{w})$ and $Q(\mathbf{w})$, respectively. p_i and q_i can be obtained from the input samples and weights by partitioning the input space into N_B bins and constructing frequency histograms in the following way:

$$q_i = \frac{\left(\frac{n_{W_i}}{N_W}\right)}{V_i}, \qquad i = 1, 2, \ldots, N_B \tag{12}$$

where n_{W_i} is the number of weights in the ith bin, N_W is the total number of weights, and V_i is the volume of the ith bin

$$p_i = \frac{\left(\frac{n_{I_i}}{N_I}\right)}{V_i}, \qquad i = 1, 2, \ldots, N_B \tag{13}$$

where n_{I_i} is the number of input samples in the ith bin, N_I is the total number of input samples, and V_i is the volume of the ith bin.

The value of α that best satisfies

$$q_i = \text{constant} \times p_i^{\alpha}, \qquad i = 1, 2, \ldots, N_B \tag{14}$$

in other terms, one that minimizes the error measure

$$E(\alpha) = \sum_{i=1}^{N_B} (q_i - \text{constant} * p_i^{\alpha})^2 \tag{15}$$

is the value of α_{achieved} by the SOM. α_{achieved} can be found by varying α in a range around the value of α_{desired}, determining the corresponding value of the error $E(\alpha)$, and then, choosing the α_{achieved} as that value of α that minimizes $E(\alpha)$.

The constant for a particular value of α can be determined by noting that $\sum_{i=1}^{N_B} q_i*(V_i) = 1$ (where $q_i*(V_i)$ is the probability that a weight lies in the ith bin; summing this over all bins yields one). Multiplying both sides of (14) by V_i and summing over N_B bins gives

$$\text{constant} = \frac{1}{\sum_{i=1}^{N_B} p_i^{\alpha}}. \tag{16}$$

This histogram-based method for determining the value of α_{achieved} seems simple but there are difficulties in extending it to higher dimensions. The *bin* in a d-dimensional space is a hypercube. In order to have an accurate estimation of the pdf, and hence, that of α_{achieved}, it is important to determine the correct size of the bin. As it turns out, this is not an easy task. The number of input samples required for estimating either of the pdfs for high-dimensional data is prohibitive. As a consequence of this, the evaluation of α_{achieved} in a general high-dimensional case is a difficult problem. We refer the reader to [26] for discussion of bin size selection and other issues involved in pdf estimation.

ACKNOWLEDGMENT

The authors would like to thank Dr. B. Csathó, Byrd Polar Institute, Ohio State University, for the Ocean City image and ground truth and P. Tracadas, Rice University, for great contribution to software development for the HyperEye© environment, in which the presented higher dimensional simulations were run.

REFERENCES

[1] H.-U. Bauer, R. Der, and M. Herrmann, "Controlling the magnification factor of self-organizing feature maps," *Neural Comput.*, vol. 8, no. 4, pp. 757–771, 1996.

[2] P. L. Zador, "Asymptotic quantization error of continuous signals and the quantization dimension," *IEEE Trans. Inf. Theory*, vol. IT-28, no. 2, pp. 139–149, Mar. 1982.

[3] T. Kohonen, *Self-Organization and Associative Memory*. New York: Springer-Verlag, 1988.

[4] H. Ritter and K. Schulten, "On the stationary state of Kohonen's self-organizing sensory mapping," *Biol. Cybern.*, vol. 54, pp. 99–106, 1986.

[5] D. R. Dersch and P. Tavan, "Asymptotic level density in topological feature maps," *IEEE Trans. Neural Netw.*, vol. 6, no. 1, pp. 230–236, Jan. 1995.

[6] D. DeSieno, "Adding a conscience to competitive learning," in *Proc. Int. Conf. Neural Netw.*, New York, Jul. 1988, vol. I, pp. I-117–I-124.

[7] T. Kohonen, *Self-Organizing Maps*. Springer-Verlag, Berlin, Germany: Springer-Verlag, 1997.

[8] A. Jain and E. Merényi, "Forbidden magnification? I," in *Proc. Euro. Symp. Artif. Neural Netw.*, Bruges, Belgium, Apr. 28–30, 2004, pp. 51–56.

[9] E. Merényi and A. Jain, "Forbidden magnification? II," in *Proc. Euro. Symp. Artificial Neural Netw.*, Bruges, Belgium, Apr. 28–30, 2004, pp. 57–62.

[10] H. Ritter, "Asymptotic level density for a class of vector quantization processes," *IEEE Trans. Neural Netw.*, vol. 2, no. , pp. 173–175, Mar. 1991.

[11] A. Jain, "Issues related to data mining with self-organizing maps," M.S. thesis, Electr. Comp. Eng. Dept., Rice Univ., Houston, TX, 2004.

[12] T. Villmann and J. C. Claussen, "Magnification control in self-organizing maps and neural gas," *Neural Comput.*, vol. 18, pp. 446–469, 2006.

[13] S. C. Ahalt, A. K. Krishnamurty, P. Chen, and D. E. Melton, "Competitive learning algorithms for vector quantization," *Neural Netw.*, vol. 3, no. 3, pp. 277–290, 1990.

[14] T. Villmann, "Controlling strategies for the magnification factor in the neural gas network," *Neural Netw. World*, vol. 10, no. 4, pp. 739–750, 2000.

[15] M. M. Van Hulle, "Faithful representations and topographic maps," in *Wiley Series and Adaptive Learning Systems for Signal Processing, Communications, and Control*. New York: Wiley, 2000.

[16] J. Lampinen and T. Kostiainen, "Generative probability density model in the self-organizing map," in *Self-Organizing Neural Networks*, ser. Studies in Fuzziness and Soft Computing, U. Seiffert and L. Jain, Eds. Heidelberg, Germany: Springer-Verlag, 2001, pp. 75–92.

[17] T. Villmann, E. Merényi, and B. Hammer, "Neural maps in remote sensing image analysis," *Neural Netw.*, vol. 16, pp. 389–403, 2003.

[18] A. Ultsch and H. P. Simeon, "Kohonen's self organizing feature map for exploratory data analysis," in *Proc. INNC-90-PARIS I*, Paris, France, 1990, pp. 305–308.

[19] T. Kohonen, "Comparison of SOM point densities based on different criteria," *Neural Comput.*, vol. 11, pp. 20 813/42 095–20 813/42 095, 1999.

[20] H.-U. Bauer and K. R. Pawelzik, "Quantifying the neighborhood preservation of self-organizing feature maps," *IEEE Trans. Neural Netw.*, vol. 3, no. 4, pp. 570–579, Jul. 1992.

[21] Th. Villmann, R. Der, M. Herrmann, and Th. Martinetz, "Topology preservation in self-organizing feature maps: Exact definition and measurement," *IEEE Trans. Neural Netw.*, vol. 8, no. 2, pp. 256–266, Mar. 1997.

[22] E. Merényi, A. Jain, and W. H. Farrand, "Applications of SOM magnification to data mining," *WSEAS Trans. Syst.*, vol. 3, no. 5, pp. 2122–2128, 2004.

[23] H.-U. Bauer and T. Villmann, "Growing a hypercubical output space in a self-organizing feature map," *IEEE Trans. Neural Netw.*, vol. 8, no. 2, pp. 218–226, Mar. 1997.

[24] T. Heskes, "Energy functions for self-organizing maps," in *Kohonen Maps*, E. Oja and S. Kaski, Eds. Amsterdam, The Netherlands: Elsevier, 1999, pp. 303–315.

[25] T. Martinetz, S. G. Berkovich, and K. Schulten, "Neural gas network for vector quantization and its application to time-series prediction," *IEEE Trans. Neural Netw.*, vol. 4, no. 4, pp. 558–569, Jul. 1993.

[26] D. W. Scott, *Multivariate Density Estimation-Theory,Practice and Visualization*. New York: Wiley, 1992.

Erzsébet Merényi (M'98–SM'05) received the M.Sc. degree in mathematics and the Ph.D. degree in computer science from Szeged (Attila József) University, Szeged, Hungary, in 1975 and 1980, respectively.

Currently, she is a Research Professor in the Electrical and Computer Engineering Department, Rice University, Houston, TX. She previously worked as a Staff Scientist at the Lunar and Planetary Laboratory, the University of Arizona, Tucson. Her interests include artificial neural networks, self-organized learning, manifold learning, segmentation and classification of high-dimensional patterns, data fusion, data mining, knowledge discovery, and application to information extraction from multi- and hyperspectral data for the identification of surface composition, including geologic, ecosystem and urban mapping from planetary remote sensing imagery, and analysis of biological tissues from multiband medical data. She has been analyzing data from space missions for over 20 years, including the development of custom algorithms for unique data returned by the (then) Russian Vega spacecraft from their close fly-by of Comet Halley in 1986. More recently, she has been working with data from the Imager for Mars Pathfinder, the Mars Exploration Rovers, Cassini VIMS, and from NASA's airborne hyperspectral sensor AVIRIS.

Abha Jain received the B.Tech. degree (with honors) from the Indian Institute of Technology (IIT) Kharagpur, India, in 2001 and the M.S. degree in electrical and computer engineering from Rice University, Houston, TX, in 2004. Her M.S. thesis was on data mining with self organizing maps.

In June 2004, she joined Microsoft Corporation, Redmond, WA, as a Software Design Engineer, where she is now a Program Manager.

Thomas Villmann received the Ph.D. degree and the habilitation with venia legendi in computer science from the University Leipzig, Leipzig, Germany, in 1996 and 2005, respectively.

Since 1997, he has been a Senior Researcher at the Medical Department of the University Leipzig. At the hospital for psychotherapy, he leads the computer science group and the research group of computational intelligence. Several research stays have taken him to France and the United States. His research area includes the theory and broad range of machine learning approaches such as neural maps, prototype-based clustering and classification, pattern recognition and evolutionary algorithms as well as applications in medicine, bioinformatics, satellite remote sensing, and other.

Dr. Villmann is a founding member of the German chapter of the European Neural Networks Society (GNNS).