In Proceedings of ESANN'98, 6th European Symposium on Artificial Neural Networks, ed. Michel Verleysen.
Brussels, Belgium: D-Facto, pp. 185-190.

Methods for Interpreting a Self-Organized Map in Data Analysis

Samuel Kaski, Janne Nikkilä, and Teuvo Kohonen

Helsinki University of Technology,
Neural Networks Research Centre,
P.O. Box 2200, FIN-02015 HUT, FINLAND

Abstract. The Self-Organizing Map (SOM) can be used for forming overviews of multivariate data sets and for visualizing them on graphical map displays. Each map location represents certain kinds of data items and the value of a variable in the representations can be visualized in the corresponding locations on the map display. Such component plane displays contain all the information needed for interpreting the map but information about the relations of the variables remains implicit. We have developed methods that visualize explicitly the contribution of each variable in the organization of the map at different locations. It is also possible to measure the contribution of each variable in the cluster structure within an area of the map to summarize, for instance, the characteristics of clusters.

1. Introduction

The SOM algorithm [2, 3] forms a mapping of a usually two-dimensional map lattice into the high-dimensional data space. There is a *model vector* connected to each point of the discrete lattice. The model vectors are situated in the data space; they act as an ordered set of models of different types of data items.

The map can be used as an ordered groundwork for illustrating different aspects of the data set. In addition to visualizing the values of the original variables as component planes (examples are shown in Fig. 2**a**) the map can be used to visualize the clustering tendency of the data in different regions of the data space. The model vectors follow the distribution of the data items and therefore the distances between the model vectors connected to neighboring points on the map lattice are shorter in clustered areas than in sparser regions. The so-called U-matrix display [4], an example of which is shown in Fig. 1, depicts the distances between neighboring model vectors as gray levels.

This work was supported by the Academy of Finland.

2. Methods

The methods presented in Secs. 2.1. and 2.2. aim at making explicit the contribution of the variables in the organization of a map at each location. In Sec. 2.3. we introduce a measure of how well a variable explains the organization within an area on the map, for instance around a cluster.

A data set for demonstrations. A relatively simple data set (cf. [3]) will be used for demonstrating the methods. A more realistic case study will be presented later in Sec. 3. The material consists of 13 properties of 16 animals. Each variable has the value one if the animal has the property and zero if it does not. A SOM of the animal data set is shown in Fig. 1. Different regions of the map represent different kinds of animals in an ordered fashion.

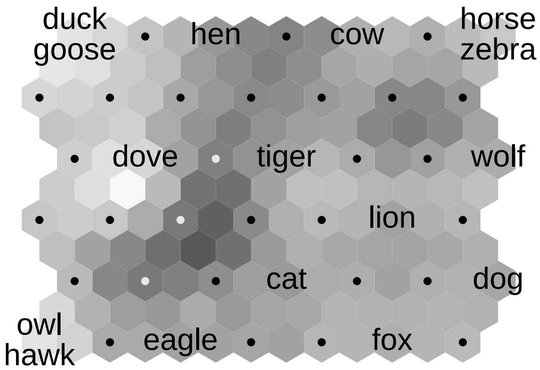

Figure 1: An overview of the relations of 16 animals generated by the SOM. The gray levels describe the clustering tendency in the data set: light shades correspond to clustered areas and dark colors to sparser regions in between the clusters. The shades have been smoothed spatially.

2.1. Local Factors

The SOM can be thought of as a nonlinear lattice of points that are determined by the model vectors in the high-dimensional data space. It is not possible to interpret the nonlinear lattice as simply as for example the set of linear factors obtained by factor analysis. The lattice can, however, be approximated *locally* by a linear hyperplane which is fitted to represent the model vectors within a certain radius on the map. The approximation can be computed with the principal component analysis algorithms resulting in two *local factors*.

The combined contribution of a variable on the local factors, computed as the sum of squares of the “factor loadings”, at each location of the map lattice can be visualized as a gray-level display that resembles a component plane (Fig. 2**b**). It can be seen in the figure that the variable “has hair” contributes strongly to the organization of the map along a stripe in the middle of the map where the representation changes from birds to other animals. The variable "has hooves" contributes to the organization in the top right corner.

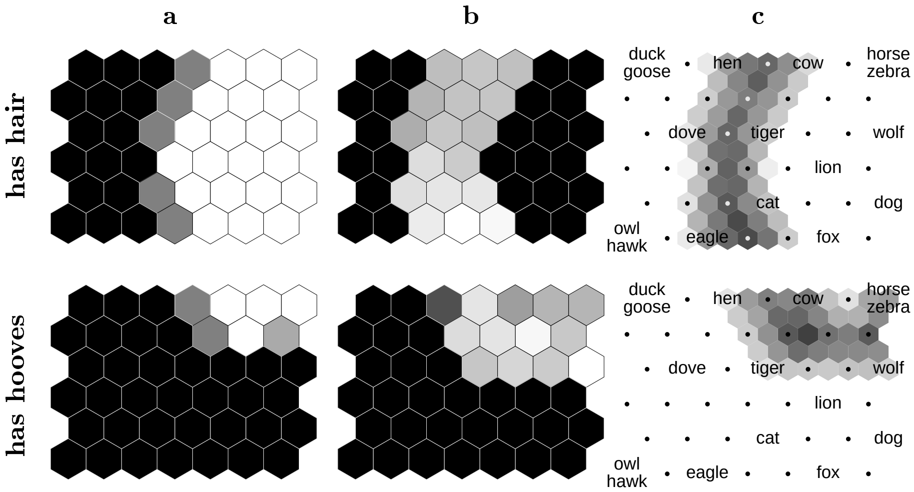

Figure 2: Sample illustrations of the methods applied to the animal data. The top row visualizes the variable "has hair" and the bottom row "has hooves", respectively. **a** The component planes. Each plane describes the values of one variable at each location on the map. **b** The contribution of the variables in the two local factors (white: maximal contribution, black: minimal contribution). **c** The (spatially smoothed) contribution of the variables in the local cluster structures (dark: large contribution, white: minimal contribution).

2.2. Contribution of a Variable in the Cluster Structures

There exists another method that is closely related to the analysis of local factors presented in the previous section but computationally much simpler. In the U-matrix display (cf. Fig 1) the cluster structures in the data are visualized as gray levels depicting the distances between model vectors connected to neighboring locations on the map lattice. The contribution of a variable in the U-matrix can be measured as the share of the distances stemming from the variable. If the share is large in a certain area the variable explains well the local cluster structure. A large share implies also that the types of data that nearby locations on the map represent differ predominantly in the values of the variable.

Examples of graphical displays showing the contribution of two variables in the local cluster structures of the animal map are shown in Fig. 2**c**. Dark stripes occur around the same locations as the light stripes in Fig. 2**b** indicating that both of the two methods can be used similarly.

2.3. Summary Generation

The methods described above aim at making the basis of organization of the SOM explicit. They do not, however, further reduce the amount of data, and we have therefore developed a method for generating briefer summaries of the important characteristics of the maps. In this study the method is used in a partly manual mode but most of the steps can be automated.

After the user has found some interesting area on the map, for example a cluster, we aim at summarizing *which of the original variables explain best the cluster structure around the area*. If the area is chosen to comprise a cluster and the sparser space around it, the best variables can be used to characterize the cluster. The measure that we use is the correlation between the cluster structure revealed by all the variables together (the U-matrix) and the structure revealed by each variable alone (cf. Sec. 2.2.). If the correlation is large the value of the variable changes strongly for example when crossing from one cluster to a nearby one and remains almost constant within a cluster. Therefore the variable can be used in explaining the local cluster structure around the cluster.

3. Case Study

We applied the interpretation methods to analyzing a data set that describes different aspects of poverty in 128 countries of the world. The set consisted of 39 indicators for each country, published by the World Bank [5].

A display generated by the SOM of the cluster structures in the data set is shown in Fig. 3**a**. It is relatively straightforward to understand the overall structure of the map based on the component planes. Poverty increases from left to right, which can also be seen by plotting the distribution of the gross national product (GNP) per capita on the map [1]. It is not as straightforward to interpret the fine structure of the map, however, and we have applied the method described in Sec. 2.3. to generate a summary of the local characteristics of the clusters in the data (Table 1). Although the table is a very reduced summary of the data set it is still impossible to fully interpret the results here, and we shall therefore concentrate on one highlight only.

Three of the four variables that explain best the cluster structure locally around cluster 6 deal with illiteracy and education. One example, the variable "% of household consumption spent on education" seems to have a local maximum at cluster 6 (cf. Fig. 3**b**). It can be seen clearly from the display showing the contribution of the variable on the local clustering tendency (Fig. 3**c**) that the value of the variable changes rapidly at the borders of the cluster.

4. Discussion

We have presented novel methods that aid in interpreting the characteristics of the data that different regions on a SOM represent. The methods concen-

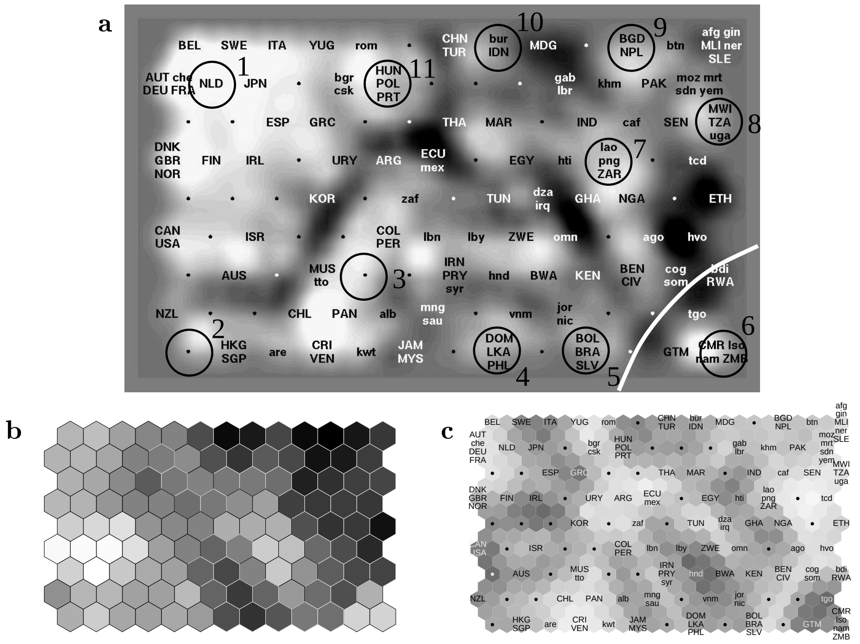

Figure 3: **a** A display generated by the SOM of the structures of poverty in the world. Shades of gray denote the cluster structure; representatives of 11 clustered (light) areas have been encircled manually. The area surrounding cluster 6 has been demarcated with the white line. One of the variables ("% of household income spent on education") that best explain the cluster structure in this area has been illustrated as a component plane in **b**. The relative contribution of the variable in the cluster structure is visualized in **c**.

trate on the local structure of the data, which can be contrasted with any straightforward statistical methods for comparing different clusters that could be extracted from the SOM display. The methods can be recommended to be used especially in exploratory tasks in which it is important to find novel, perhaps unexpected features from the data and to summarize the data set.

References

[1] S. Kaski and T. Kohonen. Exploratory data analysis by the self-organizing map: Structures of welfare and poverty in the world. In A.-P. N. Refenes, Y. Abu-Mostafa, J. Moody, and A. Weigend, eds., *Neural Networks in Financial Engineering*, pp. 498–507. World Scientific, Singapore, 1996.

[2] T. Kohonen. Self-organized formation of topologically correct feature maps. *Biol. Cybern.*, 43:59–69, 1982.

	1	2	3	4	5	6	7	8	9	10	11
Adult illiteracy (%) total						X			X		
% of hh consumption: food (cereals & tubers)									X		
% of hh consumption: clothing & footwear				X						X	
% of hh consumption: gross rents, fuel & power (total)				X							
% of hh consumption: gross rent, fuel & power (fuel & power)										X	
% of hh consumption: medical care		X									
% of hh consumption: education						X					
% of hh consumption: transport				X							
% of hh consumption: transport		X									
% of hh consumption: other consumption (other consumer durables)		X									
Population per physician	X							X			
Population per nursing person								X	X	X	
Births attended by health staff (%)											X
Babies with low birth weight (%)				X				X	X		
Daily calorie supply (per capita)											X
Primary school enrollment (% of age group) female					X	X					
Primary net enrollment (%)							X				
Primary pupil-teacher ratio										X	
Household income (% share by lowest 20 percent)			X		X		X				X
Household income (% share by 2nd quintile)			X		X		X				X
Household income (% share by 3rd quintile)	X		X								
Household income (% share by 4th quintile)	X	X									
Household income (% share by highest 20 percent)			X		X		X				
Household income (% share by highest 10 percent)	X										
Maternal mortality (per 100,000 live births)						X		X			

Table 1: A summary of the clusters, numbered in the top row according to Fig. **3a**, in the poverty data set. The four variables that explained best the cluster structure around each cluster are marked with X's.

[3] T. Kohonen. *Self-Organizing Maps.* Springer, Berlin, 1995. (Second extended edition 1997).

[4] A. Ultsch. Self-organizing neural networks for visualization and classification. In O. Opitz, B. Lausen, and R. Klar, eds., *Information and Classification*, pp. 307–313. Springer, Berlin, 1993.

[5] World Bank. *World Development Report 1992.* Oxford Univ. Press, New York, NY, 1992.