Hochdimensionale Datenräume – A Topology-Independent Similarity Measure for High-Dimensional Feature Spaces

Projektleiterin
Prof. Dr. Gabriele Peters

Zeitraum
2005 – 2007

Kontakt
Prof. Dr. Gabriele Peters
Fachbereich Informatik
Fachhochschule
Dortmund
Emil-Figge-Straße 42
44227 Dortmund
Tel.: (0231) 755-6796
E-Mail: gabriele.peters
@fh-dortmund.de

Abstract

In the field of computer vision feature matching in high dimensional feature spaces is a commonly used technique for object recognition. One major problem is to find an adequate similarity measure for the particular feature space, as there is usually only little knowledge about the structure of that space. As a possible solution to this problem we present a method to obtain a similarity measure suitable for the task of feature matching without the need for structural information of the particular feature space. As the described similarity measure is based on the topology of the feature space and the topology is generated by a growing neural gas, no knowledge about the particular structure of the feature space is needed. In addition, the used neural gas quantizes the feature vectors and thus reduces the amount of data which has to be stored and retrieved for the purpose of object recognition.

1 Introduction

In the field of computer vision objects are often represented by feature vectors describing local areas of. These local descriptors often are vectors of high-dimensional feature spaces. To identify equal or similar objects, for example for the purpose of object recognition, feature matching techniques are common, and for these matching techniques similarity measures for the feature vectors are needed. One major problem when choosing the similarity measure is often the lack of knowledge about the structure of the feature space. For example the features in the SIFT feature space are not uniformly distributed. Using the Euclidean distance leads to the problem, that the direct distance between two features cannot be used as an absolute measure of their similarity. This is a general problem of non-uniform feature spaces. Wrongly presumed uniformity can result in a classification of unsimilar features as similar and vice versa. Here we describe how to obtain a similarity measure which is suitable for the task of feature matching without knowledge of the particular structure of the high-dimensional feature space. Using a growing neural gas, the topology of the feature space is first learned and then used as a basis for the similarity measure. The similarity between two feature vectors will incorporate the length of the shortest path between those two nodes of the neural gas the feature vectors are mapped on. Besides the ability to adapt to non-uniformly distributed feature spaces the neural gas also quantizes the feature vectors. On the one and, this can be accompanied by a possible loss of information. But on the other hand, it also vastly reduces the amount of data which has to be stored and retrieved for feature matching purposes.

2 Defining a Similarity Measure on a Feature Space

The neural gas described in the previous section generates a topology of the feature space which can be used for a similarity measure. The generated topology is represented by a graph the nodes of which are the nodes of the GNG and the edges of which connect neighboring and thus similar nodes. Accordingly, we can describe the distance between two nodes (and later between two feature vectors of the high-dimensional space) by the number of edges on the shortest path between them. By doing so we utilize the ability of the neural gas to reflect the structure of the feature space.
To develop our similarity measure we need a distance matrix for the GNG graph. This distance matrix is derived by the calculation of paths of length n, where n is the number of edges connecting two nodes. How to determine nodes that are reachable on paths with a distinct length is described in [ICANN 2007], as well as the derivation of the distance matrix and the definition of the independent similarity measure (see figure 1 for an illustration).

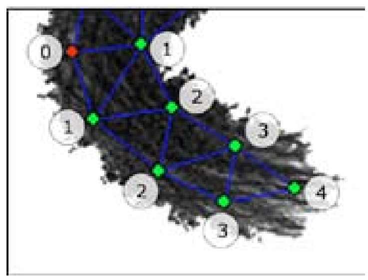

Figure 1. Topology-independent similarity measure. The gray areas represent the feature vectors of the high-dimensional space. Some nodes of the growing neural gas are depicted by red and green dots. The numbers they are labeled with are those values of the metric which express the distance between the red node and each of the green nodes. For example, the shortes distance between the red node and the upper green node labeled with "2" is a path of 2 edges.

3 Experiments

Figure 2. Database. A selection of 798 sample images on which we carried out our experiments.

We carried out our experiments on a database of 798 gray value images, a few of which are shown in figure 2. As features we consider patches of 18 x 18 pixels, thus our high-dimensional feature space has 324 dimensions. These features are not optimal descriptors for the purpose of object recognition. Nevertheless, we chose them for the evaluation of the similarity measure because they can be evaluated more easily by visual inspection than more advanced feature descriptors such as the SIFT vectors, for which a visual interpretation is much harder. Per sample image we extracted about 250 feature vectors, the positions of which have been determined with a KLT detector, resulting in a total of about 200,000 features. We ran the growing neural gas algorithm. After the processing of 300 feature vectors we added one node to the GNG and stopped the growing procedure after it consisted of 300 neurons. (All of the 200,000 features have been used for the generation of the GNG. After 300 nodes have been incorporated into the gas not many changes of the topology of the GNG were caused by the remaining features. Thus, those remaining features contributed to the stabilization of the GNG only, rather than to its overall topology.) Figure 3 shows exemplarily how some feature vectors have been quantized.

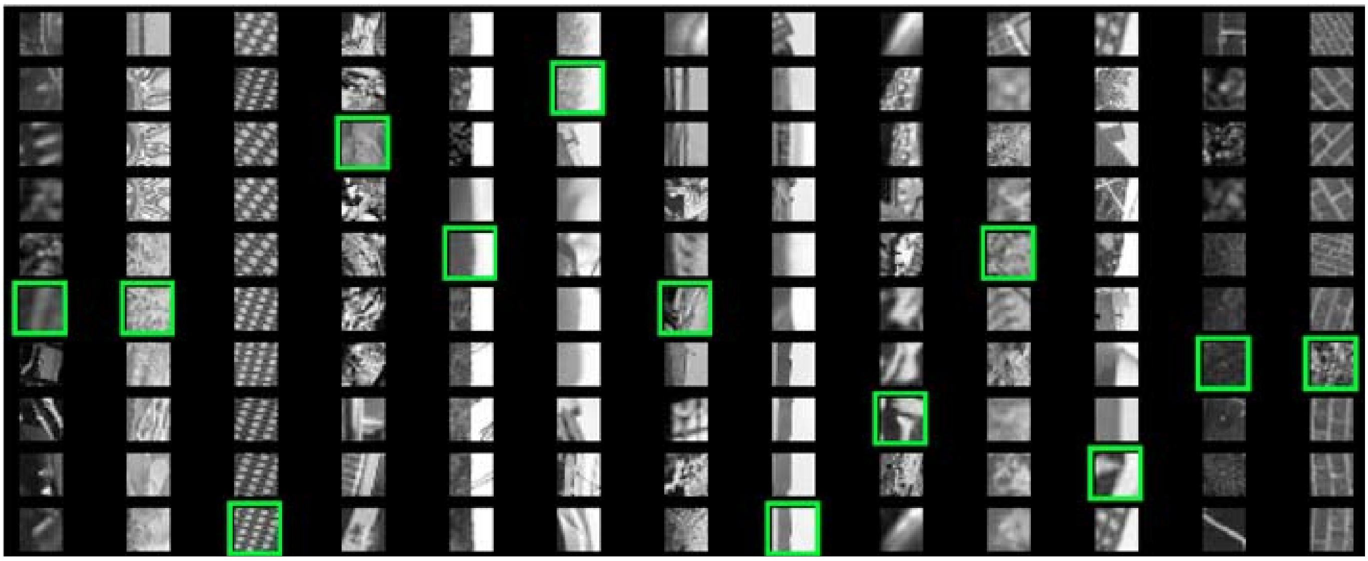

Figure 3. Quantization of feature vectors by the growing neural gas. Each column stands for one node of the GNG. Each entry (i.e., row) of a column shows one feature vector in form of a 18 x 18 gray value patch, which has been assigned to this node. In each case a column shows the last 10 feature vectors which have been mapped onto it. The features marked by green frames are those which are the last assigned.

5 Results

The final purpose of defining feature vectors and endowing their space with a suitable similarity measure, is the adequate encoding of the characteristics we intend to measure. In this case the application is an encoding of the visual characteristics of objects for purposes such as storage, classification, or recognition. Therefore, we have to analyze whether the features classified as being (mathematically) similar are also assessed by humans as being (visually) similar. In other words, the similarity (or difference) we determine by the proposed method must bear some correlation with the perceptual similarity (or difference) of two feature vectors. If our systems have to respond in an „intelligent" manner, they must use a similarity model resembling the perceptual similarity model of humans. Having these considerations in mind, we decided to assess the quality of the similarity measure by a visual inspection of the feature classification. Figure 4 shows again nodes of the growing neural gas, represented by 8 columns of the 10 last assigned feature vectors each.

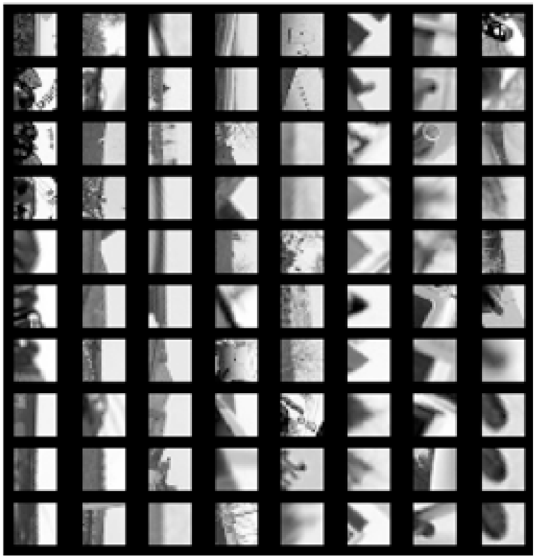

Figure 4. Classification of feature vectors into 8 neighbouring nodes of the growing neural gas. The columns represent the nodes. The entries of the columns are the last 10 feature vectors which were mapped onto them. The distance between the nodes is equivalent to the numbers of columns between them.

This time neighboring columns show adjacent nodes of the GNG, thus the number of columns between two nodes in the diagram is proportional to the distance between the nodes in the GNG. For example, the node represented by the first column and the node represented by the last column have a distance of 7 edges. We can summarize the results of the visual inspection as follows: Firstly, the similarity between features belonging to one node (i.e., features within one column) is, in general, larger than between features of different nodes. Secondly, one can observe a gradual decrease in similarity from the left to the right node for most of their assigned features. For example, the second column displays a larger overall similarity to the first column than the last column. Summarizing, one can say that the classification of features emerged from the proposed similarity measure corresponds to the assessment of the perceptual similarity by humans. Object recognition experiments remain to be done.

References

- Jochen Kerdels and Gabriele Peters, A Topology-Independent Similarity Measure for High-Dimensional Feature Spaces, 11th International Conference on Artificial Neural Networks (ICANN 2007), 2007.
- Jochen Kerdels, Dynamisches Lernen von Nachbarschaften zwischen Merkmalsgruppen zum Zwecke der Objekterkennun, diploma thesis, Technical University Dortmund, 2006.