The Empathic Visualisation Algorithm: Chernoff Faces Revisited

218

Contact
Andreas Loizides
Mel Slater
Department of Computer Science
University College London
A.Loizides@cs.ucl.ac.uk

These four pictures of faces represent the financial state of four different companies. In which one would you invest? Without knowing any details of what the facial expressions represent, it is very likely that you would choose the company represented by the bottom-right face as a good investment. In fact, more information can be gleaned from the faces by knowing that the degree of happiness represents profitability, the degree of fear represents liquidity, and so on.

The empathic visualization algorithm (EVA) is a fundamental extension of the type of data visualisation first introduced b y Chernoff,[1] who exploited the idea that people are hardwired to understand faces, and therefore can very quickly understand information encoded into facial features. In particular, it is very easy to cluster Chernoff faces into groups that represent similarities in the underlying data they represent.

Given an n k data matrix of n observations on k variables, the original Chernoff method assigned each variable to correspond to a particular facial feature like shape of the nose or shape of the eyes. The mapping from data to visual structure was arbitrary, and the resulting face had no correspondence to the underlying semantics of the data. Such faces are good for understanding pattern, but any individual face seen in isolation does not readily convey anything about the data without knowledge of the specific mapping used.

EVA provides an automatic mapping from semantically important features of the data to emotionally or perceptually significant features of the corresponding visual structure, such as a face. In other words, a single glance at the visual structure informs the observer of the global state of the data, since the visual structure has an emotional impact on the observer that is designed to correspond to the impact that would have been generated had the observer been able to analyse the underlying data itself. Finer details concerning interpretation of the visual structure a re then available through knowledge of the relationships between semantically important features of the data and emotionally significant aspects of the visual structure.

The Method

It is assumed that an n k data matrix is to be represented by a visual structure, ideally one that is naturalistic, in the sense that humans can immediately and transparently interpret the meaning of this structure at a high level. A human face is our paradigmatic example, and we will stay with this example from now on. There are a number of global characteristics that can be used to describe the emotional expressions of a face, such as its degree of happiness, sadness, calmness, fear, or anger. There are also a number of features of a face, such as muscle tensions, that have been used in the Park and Waters model[2] to determine the overall facial expression.

Correspondingly, there will be a number of important global characteristics that are of importance to the consumer of the data. For example, the magnitude of a particular combination of variables (representing overall company performance), thedifference between a variable and a threshold value, and so on. These global characteristics of the data correspond to the global characteristics of the visual structure. Finally, features of the data (particular combinations of the variables) will determine the features of the visual structure. Thus, if all these features of the data were known, all the features of the visual structure would be known, and the visual structure could be rendered. The global characteristics of such a rendered visual structure could then be measured.

A genetic program is used to determine features of the data so as to minimise a fitness function, which measures the difference between the global characteristics of the data and the corresponding global characteristics of the rendered visual structure. The goal is to minimise the difference, so that the ideal visual structure is one in which its global characteristics correspond exactly to the global characteristics of the underlying data. This can easily be achieved by choosing random functions over the set of k variables to form the first generation of features from the data. Successive generations are formed in the usual way by measuring the fitness of each rendered face and then selecting feature functions with probability proportional to fitness. In our experience, the genetic program typically converges after 75 generations. Initial experiments have shown that even non-expert users can use our technique to quickly interpret the significance of the data.

Summary

We have introduced a new method for constructing an automatic mapping from data to visual structure, which enforces a homomorphism between important characteristics of the data and the emotional or perceptual impact of the visual structure. Such visual structures are informative "at a glance" but can also reveal important detailed information about the data.

References

1. Chernoff, H. (1971). The use of faces to represent points in n-dimensional space graphically. RN NR-042-993, Department of Statistics, Stanford University, December 1971.
2. Waters, K., Parke, F. (1987). A muscle model for animating 3-dimensional facial expression. *SIGGRAPH 87 Conference Proceedings, Computer Graphics, 21* (4), 17-24.