Data Visualisation is concerned with techniques for visualising large scale multivariate data sets in order to help in understanding and exploration. Examples include financial data, business information and results from experiments. A visualisation method maps such data ideally into intuitive visual structures, but in most cases there is no obvious mapping from the data to the visual structure.

This thesis explores the use of naturalistic visual structures (for example, human faces) as representations of such multivariate data sets. An automatic mapping from the data to the visual structure is constructed through the use of the empathic visualisation algorithm (EVA) constructed for this purpose. EVA, is a fundamental extension of the type of data visualisation first introduced by Chernoff, who exploited the idea that people are hardwired to understand faces and therefore can quickly interpret information encoded into facial features. We use faces as our paradigmatic example, but the method is not limited to this case only.

Given an $n \times k$ 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 faces 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 human 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 are then available through knowledge of the relationships between semantically important features of the data and emotionally significant aspects of the visual structure.

EVA uses a Genetic Program (GP) to map the quantitative measurements from the multidimensional data set to the qualitative measurements of the visual structure. The genetic program typically converges after about 75 generations. Experiments using simulated and real financial data sets have shown that even non-expert users can use the constructed visualisation to quickly interpret the significance of the data. Subjects were able to cluster faces according to company performance, each being represented by a visual structure, in little time with a high success rate and without being told of the significance in the visual structure, compared to a time consuming procedure of going through raw spreadsheet data. Moreover, an experiment on Fear of Public Speaking data, has shown that EVA performs an accurate mapping from this data, to the emotions of the subjects giving the talk. Furthermore, the method requires being able to quantify emotional expressions in order to measure correspondence, fitness of data to the human face. Another set of experiments revealed that emotional

expressions can be measured based on movements of certain parts of a face model.

We introduce 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 structures are informative 'at a glance' but can also reveal important detailed information about the data.