Acume: A New Visualization Tool for Understanding Facial Expression and Gesture Data

Daniel McDuff - MIT Media Lab

March 24, 2011

Contents

Chapter 1

Introduction

1.1 What is Acume?

The Acume is an application that allows efficient exploration of facial action and gesture data. It provides a means to visualize and interact with behavioral data from large populations.

The open-source code is designed to directly address the needs of the face and gesture research community, while also being extensible and flexible for accommodating other kinds of behavioral data. Source code, application and documentation are available at http://affect.media.mit.edu/acume.

The application is written in the Processing language. If you wish to adapt the source code you will need to download a version of the Processing environment and install the ControlP5 GUI library. These can be found at the following addresses respectively:

Processing

http://processing.org/

ControlP5 GUI Library

http://www.sojamo.de/libraries/controlP5/

For greater detail about the applications facilities please see the paper that describes the application:

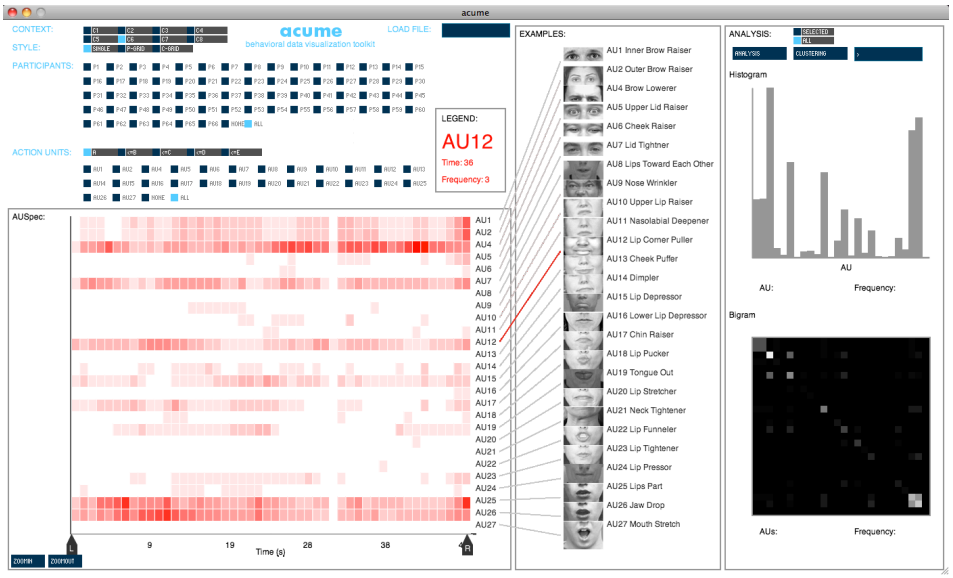

Figure 1.1: A screenshot of *Acume*, an open-source toolkit for exploring and visualizing behavioral data on multiple scales.

Acume: A New Visualization Toolkit for Exploring Facial Expression and Gesture Data. D. McDuff, R. el Kaliouby, K. Kassam and R. Picard, IEEE Conference on Automatic Face and Gesture Analysis, Santa Barbara, 2011.

Acume is designed to be adaptable and extensible such that other researchers can adapt the code for their own needs. Whilst this documentation is designed to serve the needs of those that use the software it may not cover all potential questions and problems. Please contact the author if you are would like to use Acume but are struggling with any aspect of it.

1.2 Examples

The application can be used to explore facial expression data across large populations. Here we show some examples of such an application. These images are screenshots taken from Acume. They show facial action unit data from multiple participant. Figure 1.1 shows the main application window this a set of data loaded.

For more information please refer to the two case studies that are explored in the publication described above.

Chapter 2

Application

2.1 Components

Acume currently consists of three linked panels. The spectrogram panel, example panel and analysis panel. The code has been designed such that supplementary panels could be added to suit the needs of the user.

2.2 Spectrogram Panel

The Spectrogram Panel has two views the AUSpec view, which represents the aggregated data from all participants, and the GRID view that displays side by side plots, one for each of the selected participants. In both cases action units can be turned on and off using the AU radio buttons. In the AUSpec view the user can zoom into a specific time segment and as the cursor is held over a region of the plot the AU, frequency and time instance are indicated in the legend above the plot. In the GRID view the AU and time instance information are displayed in the legend as the mouse hovers over each plot. The user can also select a particular GRID plot and blow it up.

2.2.1 AUSpec

Action Unit Spectrograms (AUSpec) are a method of aggregating dynamic behavioral data that is time aligned. Intensity of color at a particular point on the plot is proportional to the number of times that action unit was seen at that time instance across the population under consideration.

2.3 Examples Panel

The application provides the option to display the definition and/or a visual example, of the selected AUs in the Examples Panel. The user can load a definition text file and add a folder of example images specific to their task. Generic

file naming systems should be used for the stimuli and participant videos, as for the example images, these are documented in Chapter 3.

In the spectrogram mode the user has the ability to left-click on a particular time point and an image of the corresponding frame of the stimuli will appear to the right of the Examples Panel. In the GRID view a similar facility is available. This enables the user to click on the GRID plot of a particular subject at a particular time point and images of the corresponding frames of the videos of the subject and the stimuli will appear to the right of the Examples Panel.

2.4 Analysis Panel

The Analysis Panel is divided into two areas. In both cases the following representations can include all data or just data for the selected participants and AUs.

The frequency tab provides a histogram and bigram representation of the AU data. The color intensity of each square in the bigram represents the relative frequency of AU combinations within the particular dataset. If the cursor is held over a region of the histogram or bigram plot the relevant AU (or AUs) and the frequency associated with them will be displayed below the plot. The bigram allows for immediate identification of frequently occurring AU combinations, information unavailable from the AUSpec.

The clustering tab provides the ability to apply an unsupervised clustering algorithm to the data for each response. A distance matrix is calculated to provide a measure of the similarity between responses. To calculate the similarity between responses a hamming measure is used and the absolute difference in bits calculated between the binary GRID spectrograms. The rows within the distance matrix are then clustered using K-Means. The number of clusters can be set using the text box on the Analysis Panel.

The mode button on the Analysis Panel allows the user to view the mode spectrogram for each of these clusters. Indicating which action units at a particular time instance were seen across a majority of the members of that cluster.

Chapter 3

Input

In this section we describe the form that the input data should take and the naming scheme for supplementary data such as pictorial examples and stimulus and response videos.

3.1 Data

The data should be saved in .CSV format. The rows should correspond to unique samples. The first column should contain the numerical code of the participant for which that sample corresponds. The second column should contain the numerical code of the context for which that sample corresponds. The thirds column should contain the number of the time sample (e.g. frame number, second etc.). The remaining columns of the matrix should contain the corresponding action unit data.

3.2 Action Unit Labels

The action unit labels should be saved in a text file, named AULabels, with the action unit number in the first column followed by a space and then the action unit name. Each action unit should be on a different line of the file. The file should be saved in the 'Labels' folder.

3.3 Images

This section refers to the example images that appear on the examples panel. The user can provide images for any number of action units. If no image for an action unit in the input data file is provided then the corresponding label on the example panel will be blank.

The images files should be saved in .jpg format in the Images folder. The

title should be 'AU' + AU number + '.jpg' (e.g. AU1.jpg for AU1, AU2.jpg for AU2 etc.)

3.4 Videos

This section refers to videos of the stimuli and responses that may be matched with the data plots such that the user can click on the spectrogram or GRID plot and skip to that time point in the relevant video.

3.4.1 Stimulus Videos

The video files should be saved in .mov format in the Stimuli folder. The title should be 'c' + movie number + '.mov' (e.g. c1.mov for movie 1, c2.mov for movie 2 etc.)

3.4.2 Participant Videos

The participant video files should be saved in .mov format in the Subjects folder. The title should be 'p' + participant number + movie number + '.mov' (e.g. p12.mov for participant 1 and movie 2, p32.mov for participant 3 and movie 2 etc.)

Chapter 4

Acknowledgements and Contact Information

4.1 Acknowledgements

I would like to thanks Karim Kassam for providing data that inspired this software. Proctor and Gamble for funding a scholarship and the Media LAb consortium.

4.2 Contact Information

Daniel McDuff, Media Laboratory, MIT, Cambridge, 02139.
djmcduff (at) mit.edu