

MDE - Multimodal Data Explorer for Flexible Visualization of Multiple Data Streams

Isabelle Arthur¹, Jordan Quinn¹, Rajesh Titung¹, Cecilia O. Alm^{1,2}, and Reynold Bailey¹

¹Golisano College of Computing and Information Sciences, ²College of Liberal Arts

Rochester Institute of Technology

Rochester, USA

{isa6513,jaq1856,rt7331,coagla,rjbvcs}@rit.edu

Abstract—Static data visualizations are pervasive across STEM research disciplines. However, as the complexity of datasets continues to grow, often incorporating multiple modalities, the limitations of static visualization become apparent. Multimodal datasets require new tools for data exploration and visualization. In response to the limited capabilities of prior alternatives which tend to be less flexible or center on project-specific visualizations for data exploration, we are developing the Multimodal Data Explorer (MDE) - a novel application which is capable of exploring custom time-series multimodal data in context, providing a comprehensive solution that can enable a deeper understanding of the relationships between multiple streams of data. In addition, the client-server design of MDE facilitates multiple users' data exploration at the same time.

Index Terms—multimodal, visualization, time-series data

I. INTRODUCTION

Early research in affective science and computing mostly relied on unimodal datasets, such as images, speech audio, text transcriptions, and individual biophysical modalities [1]–[4]. However, relying on single modalities for affective modeling can be insufficient [5]. With advances in computational processing, incorporating multimodal data into affective modeling has become increasingly common, considering that humans tend to rely on information from multiple senses [6], e.g., when reacting to situational contexts and environmental changes or making decisions. Despite much research on affective modeling using multimodal datasets, there is still a need for additional research on broadly generalizable and adaptable multimodal data visualization. Existing applications are often project-specific or lack flexibility in the types of involved data streams or their user customization functionality. To address this gap, we have developed a novel application called the *Multimodal Data Explorer* (MDE) that enables users to load and visualize a range of multimodal datasets. Specifically designed to work with time-aligned multimodal data, MDE provides a generic and flexible solution for multimodal data visualization.

The MDE visualization system is designed to facilitate researchers' ability to spot trends in multimodal data through

This material is based upon work supported by the National Science Foundation under Award Nos. IIS-1851591 and DGE-2125362. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.




Fig. 1. A screenshot of a displayed page of the Multimodal Data Explorer (MDE) featuring a time-series view of an example multimodal dataset. This visualization includes videos, transcripts, facial expression features, galvanic skin response, and other time-series data streams. MDE allows users to add time-specific comments and customize the interface based on the specific data streams they wish to explore.

the use of a novel visualization interface. MDE is a Django-based web application [7], which allows the display of videos, transcripts, interactive data plots, and other time-synchronized data streams. The main video serves as the anchor for the remaining modalities time-wise which allows the data to be displayed in relevant segments. Additionally, MDE allows for in-place commenting and data annotation in the data timeline. Figure 1 shows a snapshot of several multimodal data streams from an experimental setting involving a dyadic task, where the data comprises modalities such as video, speech audio, dialogue transcripts, facial features, galvanic skin response features, and eye movements. The application can generate distinct pages for each multimodal data instance in the dataset. Users can further customize the viewing of the applications' display by manipulating the various window components during an exploration session. The outcome is a flexible system that aids researchers in exploring multimodal datasets which can lead to a deeper understanding of the relationships across the various data streams.

II. SELECT RELEVANT PRIOR WORK

Multimodal data analysis holds great importance in various scientific domains [8], [9]. One reason for this is that multimodality can help overcome the limitations of a single




Fig. 2. This diagram shows the behind-the-scenes interactions of a displayed page. When a user loads the page, the subject view (green) processes the data and sends it to the subject display which manages most user interactions and retrieves graphs to display.

modality [10], providing researchers with more complete information. Adequate exploration of multimodal data requires tools capable of displaying and processing different modalities concurrently. The development of MDE’s interface was informed by prior research. Although static data visualizations are widely used, there is a need for improved tools [11], [12], particularly for time-evolving data [13]. Projects such as M2Lens [14] or the Multimodal Learning Analytics platform [15] demonstrate the need for new applications to visualize multimodal data, and that interactive tools are preferred for better exploration and understanding of underlying information [16]. In addition, prior work has discussed the benefits of using interactive plots over static charts [11], especially pertinent to time-series data visualizations such as line graphs.

III. OVERVIEW OF THE MDE SYSTEM

The above discussions highlight the limitations of unimodal or static, non-interactive data visualization systems. In particular, the limitations of static visualization become even more evident when trying to explore multimodal, time-series data. The following is a description of the MDE system and how it accomplishes our goal of visualizing multimodal data streams.

The core of the system is the subject view, as highlighted in green in Figure 2. When generating the desired visualization of a given subject’s data, the subject view accesses a configuration file which details how the dataset is structured in order to identify the relevant data files before processing them. These data files may include any provided videos, transcripts, CSV files, etc. that the subject view needs to pass onto the subject display (an HTML page). The subject display will also call back to the subject view during the client’s session to get updated transcripts and save changes to a comments file.

The subject display mostly interacts with its own static files and a separate Django application [7] we created that generates charts on-the-fly using the *Bokeh* Python graphing library [17]. This application receives information about selected CSV file columns, current timestamp, and desired time segment width

and returns the corresponding chart for display. The static folder contains the subject display’s CSS and JavaScript files. The JavaScript files allow for dynamic widget creation, chart access, comment creation, and many other functionalities.

MDE’s user interface is designed to be user-friendly, intuitive, and flexible. To avoid clutter, MDE provides the option to split a dataset’s data visualizations across different pages. This keeps the visualizations organized and straightforward to parse. To ease the learning curve of using a new system, we carefully considered how to make MDE intuitive. A dashboard-style design, recognizable iconography, inviting buttons, and logical feature placement are common themes in the user interface’s design that guide a user in exploring data while using the system. Additionally, tooltips are included to further describe a given feature to the user. Flexibility is an important consideration, as researchers’ experiments may require different layouts and data streams. To this end, they can add or remove widgets as needed, and are able to switch between which CSV columns are being displayed in a widget. When uploading a dataset, researchers may set up a configuration file for their data, so they can utilize their own naming conventions.

IV. DEMO OUTLINE

The demo aims to show MDE’s capabilities and gather feedback for improvements. We will introduce MDE and provide an overview of its key capabilities and briefly introduce the example data that will be used for the demo. Attendees may also interact with MDE to attempt specific tasks using the example data and provide anonymous feedback indicating how well MDE facilitated exploration.

ETHICAL IMPACT STATEMENT

Data examples were collected in IRB-approved experiments with informed consent.

REFERENCES

- [1] Paul Ekman and Wallace V Friesen. Facial action coding system. *Environmental Psychology & Nonverbal Behavior*, 1978.
- [2] Ali Mollahosseini, Behzad Hasani, and Mohammad H Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing*, 10(1):18–31, 2017.
- [3] Cecilia O. Alm and Richard Sproat. Emotional sequencing and development in fairy tales. In Jianhua Tao, Tieniu Tan, and Rosalind W. Picard, editors, *Affective Computing and Intelligent Interaction*, pages 668–674, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [4] Sander Koelstra, Christian Muhl, Mohammad Soleymani, Jong-Seok Lee, Ashkan Yazdani, Touradj Ebrahimi, Thierry Pun, Anton Nijholt, and Ioannis Patras. DEAP: a database for emotion analysis; using physiological signals. *IEEE Transactions on Affective Computing*, 3(1):18–31, 2011.
- [5] Soujanya Poria, Erik Cambria, Rajiv Bajpai, and Amir Hussain. A review of affective computing: From unimodal analysis to multimodal fusion. *Information Fusion*, 37:98–125, 2017.
- [6] Shinsuke Shimojo and Ladan Shams. Sensory modalities are not separate modalities: plasticity and interactions. *Current Opinion in Neurobiology*, 11(4):505–509, 2001.
- [7] Django Software Foundation. Django (Version 4.2) [Computer Software]. <https://www.djangoproject.com/>, 2013.
- [8] Nicolle M. Correa, Tulay Adali, Yi-Ou Li, and Vince D. Calhoun. Canonical correlation analysis for data fusion and group inferences. *IEEE Signal Processing Magazine*, 27(4):39–50, 2010.
- [9] V.D. Calhoun, T. Adali, N.R. Giuliani, J.J. Pekar, K.A. Kiehl, and G.D. Pearlson. Method for multimodal analysis of independent source differences in schizophrenia: Combining gray matter structural and auditory oddball functional data. *Human Brain Mapping*, 27(1):47–62, 2006.
- [10] Felix Bießmann, Sergey Plis, Frank C. Meinecke, Tom Eichele, and Klaus-Robert Müller. Analysis of multimodal neuroimaging data. *IEEE Reviews in Biomedical Engineering*, 4:26–58, 2011.
- [11] Tracey L. Weissgerber, Vesna D. Garovic, Marko Savic, Stacey J. Winham, and Natasa M. Milic. From static to interactive: Transforming data visualization to improve transparency. *PLoS Biology*, 14(6):e1002484–e1002484, 2016.
- [12] Robert Newell, Ann Dale, and Celia Winters. A picture is worth a thousand data points: Exploring visualizations as tools for connecting the public to climate change research. *Cogent Social Sciences*, 2(1):1201885, 2016.
- [13] Kari-Jouko Räihä, Anne Aula, Päivi Majaranta, Harri Rantala, and Kimmo Koivunen. Static visualization of temporal eye-tracking data. In Maria Francesca Costabile and Fabio Paternò, editors, *Human-Computer Interaction - INTERACT 2005*, pages 946–949, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [14] Xingbo Wang, Jianben He, Zhihua Jin, Muqiao Yang, Yong Wang, and Huamin Qu. M2lens: Visualizing and explaining multimodal models for sentiment analysis. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):802–812, 2022.
- [15] René Noël, Diego Miranda, Cristian Cechinel, Fabián Riquelme, Thompsen P. Tiago, and Roberto Munoz. Visualizing collaboration in teamwork: A multimodal learning analytics platform for non-verbal communication. *Applied Sciences*, 12(15):7499, 2022.
- [16] Min Lu, Noa Fish, Shuaiqi Wang, Joel Lanir, Daniel Cohen-Or, and Hui Huang. Enhancing static charts with data-driven animations. *IEEE Transactions on Visualization and Computer Graphics*, 28(7):2628–2640, 2022.
- [17] Bokeh Development Team. *Bokeh: Python library for interactive visualization*, 2018. <https://bokeh.pydata.org/en/latest/>.