MUVTIME: a Multivariate Time Series Visualizer for Behavioral Science

Emanuel Sousa[1], Tiago Malheiro[1], Estela Bicho[1], Wolfram Erlhagen[2], Jorge Santos[1] and Alfredo Pereira[3]

[1]*Algoritmi Centre, University of Minho, Guimarães, Portugal*

[2]*Center of Mathematics, Department of Mathematics and Applications, University of Minho, Guimarães, Portugal*

[3]*CiPSI, School of Psychology, University of Minho, Braga, Portugal*

eafsousa@gmail.com, { tmalheiro, estelabicho}@dei.uminho.pt, wolfram.erlhagen@math.uminho.pt

{alfredo.pereira, jorge.a.santos }@psi.uminho.pt

Keywords: Multivariate Time Series, Visualization, Cognition.

Abstract: As behavioral science becomes progressively more data driven, the need is increasing for appropriate tools for visual exploration and analysis of large datasets, often formed by multivariate time series. This paper describes MUVTIME, a multimodal time series visualization tool, developed in Matlab that allows a user to load a time series collection (a multivariate time series dataset) and an associated video. The user can plot several time series on MUVTIME and use one of them to do brushing on the displayed data, i.e. select a time range dynamically and have it updated on the display. The tool also features a categorical visualization of two binary time series that works as a high-level descriptor of the coordination between two interacting partners. The paper reports the successful use of MUVTIME under the scope of project TURNTAKE, which was intended to contribute to the improvement of human-robot interaction systems by studying turn-taking dynamics (role interchange) in parent-child dyads during joint action.

1 INTRODUCTION

Like many other research fields, behavioral science is becoming increasingly more data intensive. Behaviors can now be captured, analyzed, and quantified semi-automatically and this is leading to a change in the way psychological science is made (Jaffe, 2014). A few trends are evident in the move towards "big data" in psychological research. One derives from the unexpected consequences of having previously collected data made available in the public domain for others to freely reuse and reanalyze. This kind of aggregation has already occurred for particular types of datasets and opened the way for synthetic modeling approaches to future studies – e.g. the NeuroSynth and NeuroVault projects combine functional magnetic resonance imaging (fMRI) data from multiple sources (Gorgolewski et al., 2015; Yarkoni, 2012). Another line of development is related to automatic sampling, sometimes with residual human investment in data acquisition, for instance biological data obtained from fitness trackers, used in studies of sleeping habits and how they interfere in people's general mood (Swan, 2013). Finally, a type of studies with a long tradition in psychological research, micro-analytic studies of social interaction – highly informative but notoriously challenging and labor-intensive, and thus limited in number, (Bakeman and Gottman, 1997; Burgoon et al., 2007) – are now at least more tractable. The term micro-analytic here refers to situations where social interaction is quantified, potentially across more than one modality, on a time-scale of milliseconds. Examples of this trend include the appearance of the field of multimodal corpora studies (Kipp et al., 2009), the first case of an almost exhaustive sampling of a single child's entire language input, conducted in the first two years (Roy et al., 2015), and novel developmental psychology studies of infant's sensorimotor dynamics and its effect on learning, during mother-infant social play (Smith et al., 2011; Yu and Smith, 2013; Pereira et al., 2014b). In this type of studies, data is analyzed so that complex relations between the variables coding different dimensions of human activity (e.g. gaze direction or body posture) can be determined.

Consider the case of a mother and an infant engaged in a joint task: Different dimensions of the mother's socially contingent activity, such as gaze direction, body posture, speed of motions, and frequency of vocalizations, among others, will affect similar dimensions of the infant's behavior and, in turn, also be affected by the infant's behavior in a well-coordinated social exchange (De Barbaro et al., 2013).

This recent trend of growing behavioral datasets bears a great potential for information extraction. This potential, however, is often limited by the tools available to behavioral scientists. Often, the sheer amount of data requires specific data visualization tools, not only because of the challenges introduced by dataset size but also by the very nature of the data itself. For instance, time series resulting from motion capture can contain both a macro and a micro-structure: e.g. average velocity vs. velocity variations produced by adapting in real-time to a social partner. Typically, the dataset is a multivariate time series and is difficult to visually explore and identify associations between variables. This is especially the case when one intends to interact with non-stationary time series data that rapidly changes depending on the time region of interest or when we need contrasting visualizations.

In order to assist with some of these data analysis tasks, we developed MUVTIME (MUltivariate Visualization of TIME Series), an interactive tool for multivariate time series visualization for social interaction studies, intended to facilitate the process of interleaving visualization and numerical analysis. MUVTIME was developed in the context of project TURNTAKE (Lisboa et al., 2014; Pereira et al., 2014a), a project aimed at improving Human-Robot Interaction design by studying developmentally parent-child turn-taking dynamics with the ultimate objective of improving robots' ability to adjust to individual rhythms in interaction, a factor known to influence the quality of a social interaction (Jaffe et al., 2001). Data analysis was both hypothesis-driven – motivated by studies of mother-infant attachment style (Jaffe et al., 2001) – and exploratory – the relationship between vocal and motor coordination is unclear.

Our main goal was applying data visualization techniques to the specific question of understanding real-time social coordination as it happens in mother-child interactions. Functional requirements and visualization design rationale for the tool derived from the research questions, our previous experience in using behavior coding and annotation tools (the authors include developers and end-users), and the specific requirements of visualizing turn-taking dynamics. This produced a minimum set of functional requirements we selected for development.

MUVTIME implements a time brushing tool allowing a user to focus on specific time periods, analyzing the data with different visualizations. It also integrates a video-playing tool where current time is marked in a time cursor and data views are all linked by time. This is particularly critical in behavioral research since it allows immediately contrasting the actual data of an interaction episode with the judgment made by the researchers when visualizing behavioral events. The data visualization techniques include a minimum core set of time series visualizations plus a novel visualization of turn-taking, a high-level categorical visualization, automatically extracted from the data, and that works as a visual descriptor of the coordination between two interacting partners, following the work of (Jaffe et al., 2001).

The paper describes the MUVTIME most important features and then reports and discusses its utilization on the TURNTAKE project.

2 RELATED WORK

In this section we review some computer applications for interactive exploration of time oriented data. A more comprehensive survey can be found in Aigner et al. (2011)

One domain with an intensive use of this tool is medical care and research where it is used for medical diagnosis, treatment management, and data exploration. (Catarci et al., 2003) presented one of the first interfaces for dealing with biomedical data that also included some limited capabilities of time series visualization. A more time oriented graphic tool was proposed by Bade et al. (2004) for comparing data of patients in treatment with corresponding medical guidelines It featured methods for displaying qualitative and quantitative temporal information in the same graph. Another project, the CareGiver (Brodbeck et al., 2005), displayed both categorical and numerical data in different bands while maintaining a unique timeframe.

Other tools, more focused on data exploration, were KNAVE-II (Shahar et al., 2006) which allowed to visualize and explore both raw data and temporal abstractions derived from it (e.g. episodes of high blood pressure). PatternFinder (Fails et al., 2006) featured a graphical interface for performing queries (formulated as sequences of events) on a database of patient histories and produce visualizations. Lifelines 2 (Wang et al., 2009) allows temporal alignment of event-based time series data of different patients according to the occurrence of a specific event (e.g. the application of a therapy) and compare posterior events. Similan (Wongsuphasawat and Shneiderman, 2009), was designed to query and visualize medical records, when searching for specific sequences of categorical events. CareCruiser (Gschwandtner et al., 2011) organizes multiple treatment plan algorithms in a tree like or flow chart view. The user can then chose one plan and visualize the history of the patients that received that particular line of treatment. VisuExplore (Rind et al., 2011) uses some of the techniques from Bade et al. (2004) but more specifically turned to data exploration.

More general tools for interactive explorations of time series data also exist. The multiple iterations of TimeSearcher (Buono et al., 2007; Hochheiser and Shneiderman, 2004) were built around the concept of time box, a graphical object that a user can draw with the mouse on a two-dimensional plot of time series data (where the horizontal and vertical axis represent time and the measured dimension, respectively). This acts as a filter by defining a period of interest (the time interval that fits in the box) and a range of values (values that range between the vertical limits of the box). The result of the filter is the set of data items whose values fit inside those ranges. GeoTime (Kapler and Wright, 2005) is a 3D visualization tool for space-time data exploration where a ground plane represents spatial information while variation in time is displayed in the third dimension. VIS-STAMP (Guo et al., 2006) also allows exploration of geo-temporal data using self-organizing maps (SOM) for data clustering. EventViewer (Beard et al., 2007) is a framework for visualization of sensor-based data acquired on multiple locations during long time spans. Exploration of both temporal and spatial patterns is supported. FacetZoom (Dachselt et al., 2008) allows a user to navigate time oriented data using different time granularities.

In the behavioral science domain, the analysis of categorical time series is also common. These time series are typically obtained through annotations of events occurred in observation sessions recorded on video and normally performed by trained researchers in an exploratory fashion or using structured coding schemes. One highly influential computer tool, MacSHAPA (Sanderson et al., 1994), was developed for fast annotation of these recordings and geared towards exploratory sequential data analysis (ESDA). Annotations were entered as values within the cells of a spreadsheet, where the categorical variables were represented as columns. It included some graphic visualization capabilities, but most importantly it offered digital frame-by-frame video playback in VCR tape recordings. As computers evolved, similar tools, free and commercial, became available. Some are concentrated on speech analysis, for example ELAN (Wittenburg et al., 2006) and EXMARaLDA (Schmidt and Wörner, 2009) or multimodal corpora like ANVIL (Kipp, 2012), TASX-annotator (Milde et al., 2001) or MacVisTa (Young and Bann, 1996). Recently, Yu and colleagues (Yu et al., 2012; Yu et al., 2009) proposed a tool for interactive exploration of time series with video playback. The tool includes several timeline visualizations methods for visual exploration of both numerical and categorical time series. Two other tools, ChronoViz (Fouse et al., 2011) and BEDA (Kim et al., 2013) were also developed for visualization and annotation of time series data with video playback. Notables (Lee et al., 2013) is an online platform featuring a visualization called “plexlines” where categorical events occurred during an interaction session between clinicians and children are displayed as circles in a timeline.

3 MUVTIME

MUVTIME was designed to assist on the process of multivariate time series data analysis, both for direct data exploration and in the process of prototyping processing algorithms. It was implemented in Matlab because this computational platform is our primary tool for complex analysis. Thus, we can apply processing methods on multivariate time series and rapidly visualizing them without having to export the data to csv or other format and import it on other tool such as TimeSearcher or ChronoViz. MUVTIME is built around two Matlab objects designed to work with time series data: *timeseries* and a collection of timeseries (*tscollection*). *Timeseries* objects contain a univariate time vector and a multivariate data vector. They also contain metadata, e.g.: units, starting and ending time, data quality measures, as well as a structure for events description. The *tscollection* is formed by a set of *timeseries* objects that share a common time vector but may refer to different data types. When performing time series processing in Matlab it is

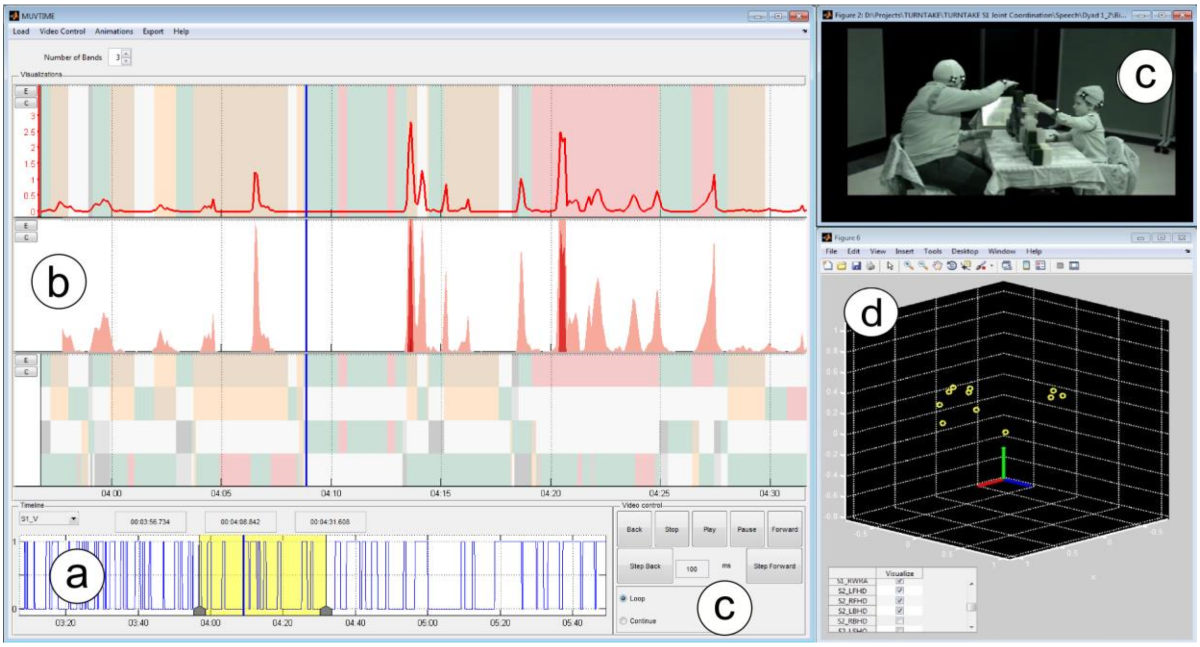

Figure 1: Depiction of MUVTime main interface. a) Timeline navigator. The yellow panel defines the time period under visualization. b) Visualization panels. Each panel displays one or more time series with a specific visualization method. c) Video player and control. Video time is represented by the vertical blue line on the visualization panels. d) Point light display. 3D visualization of the motion capture data represented as points in a 3D space.

useful to use these objects since they offer a set of standard methods for manipulation (e.g. interpolation, concatenations) that ensure data consistency and they also allow basic feature extraction like averages and standard deviation. MUVTIME can be called from the Matlab command prompt directly with a *tscollection* object as input argument. It can also be called with no input arguments and a collection stored on a ".mat" file can be loaded through the menu options. A depiction of the application's main interface is presented in Figure 1, where its main elements are visible: The timeline navigator, the visualization panels, the video window and control and the Point light walking display. Next we describe each of the elements in detail.

3.1 Main Interface

3.1.1 Timeline Navigator

The timeline navigator controls the time window of the data currently in visualization in the band graphs (Figure 1.a). Graphically, it consists of a standard time series line plot (with data referring to one of the time series of the collection), where time runs horizontally from left to right. An overlaid semi-opaque yellow panel defines the time period under visualization. Initially, when a collection is loaded, the time window covers the complete time series but its width can be adjusted with the computer mouse by click-and-drag of the cursors on the right and left bottom corners of the window. The complete window can also be displaced through click-and-drag of the yellow region. The time series plotted on the background (the user can control which one it is) serves as cue for the adjustment of the brushing window on a particular time region (Yu et al., 2012).

3.1.2 Visualization Panels

The visualization panels (Figure 1.b) are the main component of MUVTIME where the data is visualized. The total number of panels can be defined on the edit box on top of the visualizer. Within each panel, the visualization can be changed by clicking on the "Edit button" (top left corner of each band). This opens an options window where a particular visualization can be selected and its parameters adjusted. The time interval visualized is defined by the sliding panel on the timeline navigator. When this time panel is adjusted the visualization panels are automatically updated so that the time window under visualization is the same on all of them. The graphs are always kept aligned so that each horizontal coordinate is vertically aligned with the same coordinate in the other graphs. This allows visual comparison of different visualizations, and different time series, for the same time region.

3.1.3 Video Player and Control

MUVTIME can optionally load and play a video file on a separate visualization window on the screen. The underlying video player is the Microsoft windows media player that is part of the free set of tools available for Microsoft windows. This "outsourcing" of the video playing decreases the computational effort for Matlab, allowing it to run fluidly even with the video playing. The player can be controlled either through a standard video interface (Figure 1.c) or through keyboard shortcuts. Two different reproduction modes are available. In the first one, the video loops continuously in the time window of visualization. In the second the video controller automatically shifts the time window, running the video until its end. A vertical blue line is always drawn on both the time brushing panel and in the visualization panes, marking the current video time. A double click on any location of a visualization panel will make the video current time jump to the instant defined by horizontal position of the mouse cursor when clicked.

3.1.4 Point Light Display

The point light display tool (Figure 1.d) is a visualization window that can display motion capture data represented as point lights if the time series collection includes time series of 3D data point. Point lights refers to using a small sphere to represent a motion capture marker and is a type of motion visualization originally derived from studies of biological motion perception, in particular point light walkers (Johansson, 1973; Johnson and Shiffrar, 2013). In MUVTime it allows contrasting live video with motion data, mostly for controlling data reliability. The tool includes a menu for choosing which points the user wants to visualize (MUVTIME will automatically search time series with three dimensional data). Point motion is synchronized with the current time.

3.2 Data Visualizations

For each visualization panel, the user can decide which data and type of visualization to use. Next we describe the four types of visualization available.

3.2.1 Overlapping Line and Categorical Graphs

Line graphs are probably the most common type of visualization for time series and are also included in MUVTIME. The interface allows multiple time series to be plotted as lines in a single visualization panel (figure 2a). For each line that is appended to the visualization a new vertical axis, colored the same as the corresponding line, is added on the left side of the pane. The scale on the axis is automatically set to fit the range of the data, within the time range of the current time interval of visualization. With this visualization of multiple axes, data of different dimensions can be compared in terms of their trends (e.g. one rises when other falls). The color of each line and corresponding axis is automatically set when the line is added but can be altered by the user. When the mouse cursor hovers the graph, a vertical line appears also displaying the values of the variables on the instant defined by the mouse horizontal position on the graph.

Overlapped on the line plots, the user can add a categorical visualization that will be displayed as a set of shaded areas. This allows inspecting the relations between numerical and categorical variables. MUVTIME allows that any time series can be defined as categorical, as long as its values are all integers. By default, the color map for the categorical time series is automatically defined depending on the total number of categories. MUVTIME can generate two types of color maps: one defines the colors in order to facilitate visual differentiation, maximizing their distance in the CIELAB color space (Schanda, 2007); the other assumes that the categorical space is scaled. Thus the color map is defined as a heat map of the categories, based on their numerical value. The user can also define a color map by loading .xls file where each line is formed by four numbers, the first one referring to the category numerical identifier and the other ones describing the colors in the RGB space.

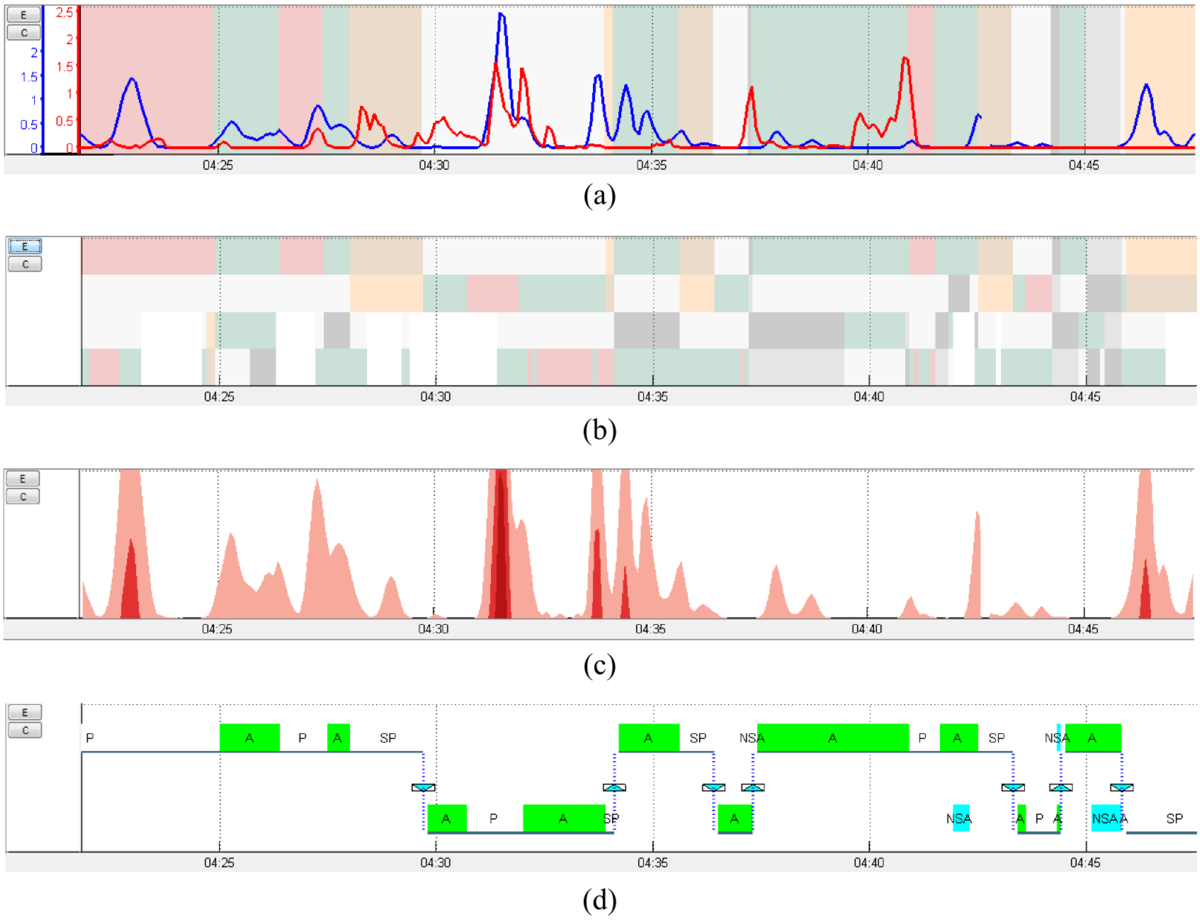

Figure 2: Available visualization types: (a) Overlapping line and categorical graphs. (b) Stacked categorical map. (c) Horizon graph. (d) AVTA diagram.

The choice of using colors to represent categories raises the question of scalability considering that color differentiation becomes increasingly difficult as the number of colors increases. However, considering the application domain it is unlikely that a large number of categories will be required for same plot.

3.2.2 Stacked Categorical Maps

This type of visualization (figure 2b) allows the comparison of multiple categorical time series as a set of stacked colored bands. Like in the previously described visualization, any time series whose values are exclusively integers can be chosen as categorical and displayed in this form. The color map can also be automatically defined to maximize differentiation or to act as categorical heat map.

3.2.3 Horizon Graphs

Horizon graphs (figure 2c) (Few, 2008; Saito et al., 2005) are a type of plot idealized to minimize the space necessary to display a time series. They are based on area plots but employ a few techniques to minimize space while facilitating perception. First of all the negative part of the graph is drawn mirrored to the x-axis, but with a different color to allow distinction. Second, the vertical direction of the area plot is divided into bands where the color tone is darker as the values fit into increasingly higher ranges. Thirdly, instead of being stacked, the bands are overlaid, reducing the necessary space for display. This type of visualization reduces the height of the graph compared to a normal line graph by relation of $1{:}n_{bands}$. We choose $n_{bands} = 3$ thus reducing the height 6 times without compromising the ability of the users to distinguish the bands (Heer et al., 2009).

3.2.4 AVTA Diagram

The AVTA diagram (figure 2d) is a data abstraction that can be obtained from two binary time series that quantify the activity of the two agents by some measure. It is inspired by the Automated Vocal Transaction Analysis system (Cassotta et al., 1964; Jaffe et al., 2001), initially proposed to represent all possible dyadic states of a speech interaction between two partners. It defines a behavioral dimension called a turn and a turn rule, such that each turn is unambiguously attributed to one of the partners (i.e. similar to our informal notion of who is "holding the floor" in conversation). A turn begins at the instant that any participant vocalizes alone, and it is held until the other vocalizes alone, at which point the turn is exchanged. The turn itself is at a different conceptual level than the vocal states, as it can be a composite of all of them. Beyond the turn, the system also defines states of conversation in terms of absence/presence of vocalization, duration of pauses and activity. In the TURNTAKE project, we generalized this system assuming a broader definition of activity that also includes body motion of any kind. While in the initial works the diagram as used for schematic demonstrations, MUVTime is able to automatically generate this visualization based on input data permitting its use for interaction analysis. It receives as input two time series representing either vocal activity (e.g. speech transcriptions) or motion (e.g. some measure of the physical motion of a tracking marker). The input signals are binarized to zero or one by thresholding. The turn attribution rule is similar to the one of Jaffe et al. (2001). The possible activity states attained by each of the partners are the following ones:

1. **Active (A)** – A continuous interval of time where one individual is the turn holder and is active alone, containing no period of inactivity greater than *x ms* where x is defined by the user.
2. **Pause (P)** – A joint period of no activity greater than or equal to *x ms* bounded by the Active periods of the turn holder.
3. **Switching Pause (SP)** – A joint silence greater or equal to *x ms*, initiated by the turn holder but terminated by unilateral activity of the partner, that gains the turn.
4. **Interruptive Simultaneous Activity (ISA)** – It is a period of time that begins with the activity of the partner that does not hold the turn, while the turn holder is still active, and ends when the turn holder stops being active, and at that point the partner that initiated the interruption gains the turn.
5. **Noninterruptive Simultaneous Activity (NSA)** – It begins with the activity of the partner that does not hold the turn and ends when the partner who holds the turn is active continuously.

MUVTIME performs the binarization of the two time series according with a user-defined threshold. It then runs the time series to determine the turns and the dyadic states according with the rules defined. The diagram itself is formed by two horizontal axes, each one describing the state of a partner. The turn holder is marked by a thin horizontal gray bar drawn on its timeline. The other states are represented by colored boxes stacked on the gray bar representing the turn (A - green, P and SP - white, ISA - yellow, NSA - blue). Vertical blue dashed lines linking the timelines mark the turn changes and blue arrows distinguish their direction of change. While this visualization is not as compressed as, for instance the stacked categorical plots, the fact that it displays the activity of the individuals in two separate streams facilitates the understanding of who has the turn and the identification of dyadic states in which the individuals have simultaneous activity (e.g. NSA).

4 TURN-TAKING DYNAMICS DURING JOINT PLAY: A CASE STUDY OF MUVTIME

Next, we summarize a few key findings from using MUVTIME with the multivariate time series dataset generated in the TURNTAKE project. A main goal was to study turn-taking dynamics, when children are engaged in open-ended joint play with a parent. Visualization was critical since we needed a tool that allowed the researcher to replay time series data (e.g. who is active/inactive in voice or motion signals) synchronized with the video of the interaction.

We conducted a set of experimental studies where mother-child dyads had to engage in a joint action task that varied in level of difficulty; we captured each dyad's: vocalizations, head and wrist movements using a motion capture system, and recorded a video of the interaction. Each interaction section lasted approximately 8 minutes resulting in a (multimodal) multivariate time series dataset ranging from 16×10^4 to 26×10^4 data samples (movements were captured at 200Hz).

4.1 Methods

4.1.1 Participants

Ten adult-child dyads participated in the study; children's age ranged from 52.1 months to 78.6 months and the adult was always either the child's mother or father.

4.1.2 Procedure

Dyads were tested in two tasks: one was more demanding and it required the parent to teach the child how to build an object (the object was too complex for the child alone to succeed) and the rules forced both partners to engage in the construction; a second task still included a joint goal – the task was to build the tallest tower possible with blocks – but placed no restrictions on the dyad in terms of how they could achieve the task. Thus, the second task of constructing a tower was not free play but was clearly less constrained and demanding than the first task.

Participants sat across each other on a small table and wore a sports headband and one wristband on both wrists. Reflective markers, 14 mm in size, were then attached to the bands using Velcro placed. There were four markers in the head and one in each wrist. A video camera pointed at the table workspace area recorded video and audio of the interaction.

4.1.3 Data Coding and Processing

Motion data was coded and processed using Vicon Nexus software package (Vicon, 2015); speech data was transcribed and time-coded using ELAN (Wittenburg et al., 2006). Transformation of coded motion and speech transcription data into a dataset of suitable for time series analysis involved five main steps: data reduction of the motion data, using the MSV (mean-square velocity) approach (Gray et al., 2005) – this computes a unidimensional signal from 3D point data; calculation of vocalization on/off binary time series; registering motion and speech data; calculation of the AVTA model of turn-taking for all possible pairs of binary time series (speech or markers); and exporting datasets to a .mat file.

The final processed dataset, one per dyad and task, is a *tscollection* with a large number of time series: vocalization on/off per partner; head and wrists' MSV per partner; individual AVTA states per pair of signals (e.g. adult's vocalization with child's vocalization; adult's binarized head MSV with child's binarized head MSV). Each AVTA pair generates 12 binary time series.

4.2 Data Visualization Findings

Figure 3 shows a set of time series plots produced with MUVTIME. The plots are from a single time period of one dyad, while they were engaged in the most difficult task – building an object by copying it from an image. This example is representative of the entire sample; the equivalent plots in the majority of dyads were similar to this one.

The first two panels contain a visually detailed version of an AVTA diagram; the plot depicts the dyadic turn-taking state, labels each one and shows two running tracks, one per participant, and it is possible to see who is the turn holder at each moment. The first panel, a), shows an AVTA diagram for adult's vocalization paired with child's vocalizations; the second panel, b), shows an AVTA diagram for adult's right hand movements paired with child's head movements – this is a characteristic pair since head stabilization is often associated with paying attention to the social partner in social play (Yu and Smith, 2013). In panel b), the diagram was calculated on the fly by specifying a threshold on the MSV signals (any value above the threshold counts as partner active) and we used a conservative value. Both plots immediately show the coordinated nature of these social exchanges: interruptions are infrequent and partners smoothly alternate turns. It is also clear that speech modality and movements have different time scales.

The information in panels a) and b) can be compressed and depicted by stacked categorical maps. Panel c) shows the same information of panel a) but using this type of visualization and panel d) shows the same information but with a highly conservative activity threshold. This comparison shows full detail of the AVTA diagrams is not mandatory: panel c) still shows clearly smooth turn-taking. Panel d) is similar to panel b) but with the threshold change, any small movement counts as activity and this affects the turn-taking state (e.g. there are more interruptions visible for instance). The motion data still shows evidence of turn-taking but also shows how the decision of what counts as making a movement or standing still is problematic and changes the conclusions on the dynamics of the interaction.

A third set of panels, e) to g), shows an effort to further compress the full multidimensional dyadic state. What is shown are stacked categorical maps of binary variables namely: speech on/off of each

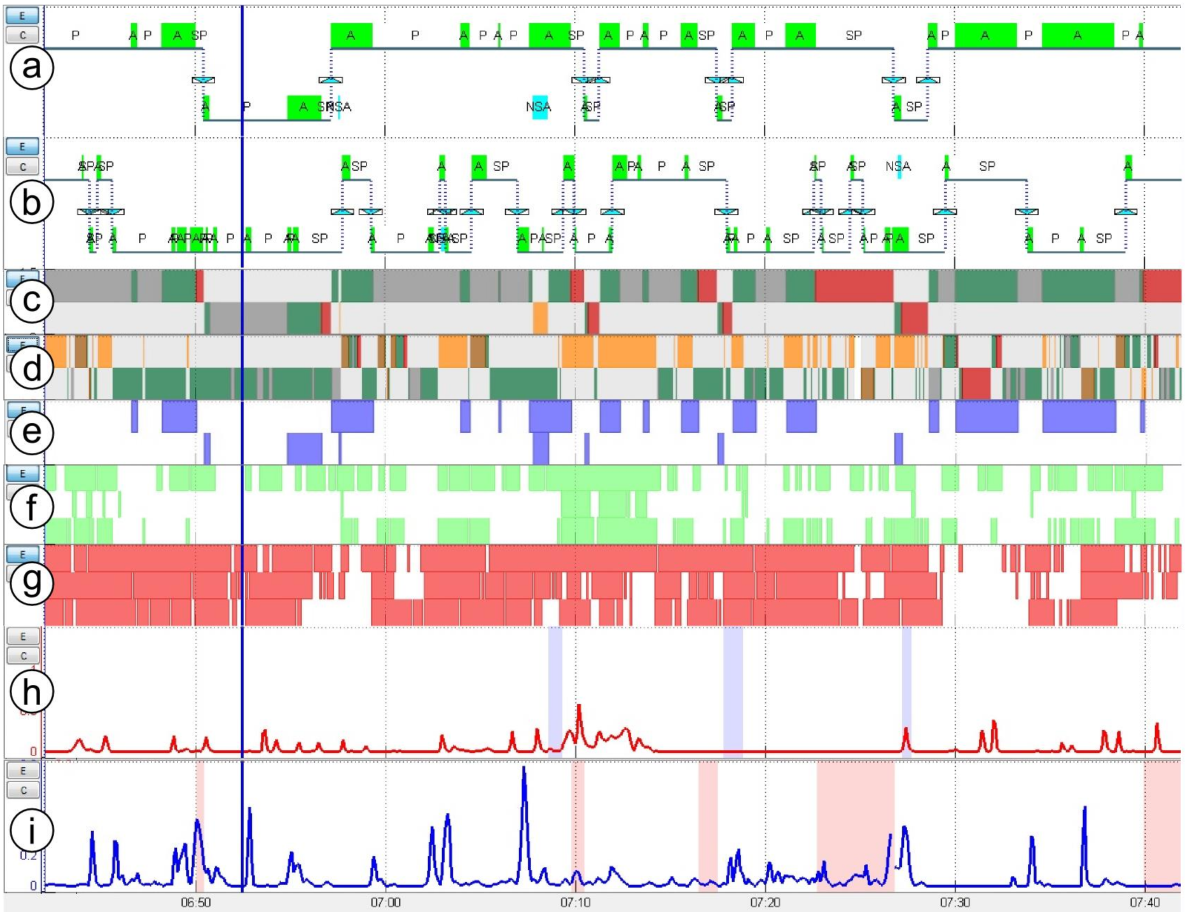

Figure 3: Example of multivariate visualization from the same time period of a particular dyad while they were engaged in a joint construction task. a) AVTA diagram for adult's vocalization and child's vocalizations; b) AVTA diagram for adult's right hand movements and child's head movements; c) same information as (a) but as a stacked categorical view; d) same information as (b) but with a different threshold and with a stacked categorical view; e-g) stacked categorical maps of binary variables, namely (e) speech on/off of each partner in panel (f) binarization of head, right and left wrist of adult's MSV and (g) binarization of head, right and left wrist of child's MSV; h) adult's switching pause (categorical) and MSV value for the child's head (line); i) reverse of (h).

partner in panel e); head, right and left wrist of adult's binarized MSV in panel f); and head, right and left wrist of child's binarized MSV in panel f). Again, the compression shows that for some signals, the key information is preserved: in panel e) the smooth, well-coordinated nature of turn-taking when vocalizing is still evident. Panels f) and g) show motion data as on/off. Although panel g) shows the child as constantly moving, this visualization was still highly informative: first, even with a strict threshold, the adult has periods of complete stillness, a signal known to be used by adults in infant-directed demonstrations of novel actions, e.g. (Rolf et al., 2009); second the constant activity of the child calls into the question if this corresponds to constant activity, natural differences in terms of motor stability with adults or the threshold decision itself – an example of visualization assisting quantitative modeling as a preliminary phase.

Finally, panels h) and i) show an example of contrasting speech with motion data, in this case an important specific combination. A critical aspect of turn-taking is the switching pause (the moment of silence when the two partners switch) – it is a marker of coordination strength (Jaffe et al., 2001). Panel h) shows the adult's switching pause as a categorical map and overlaid, the quantitative MSV value for the child's head of coordination; panel i) shows the reverse.

5 CONCLUSIONS

In recent years, behavioral science has followed a general trend in science and is becoming more data driven. As a consequence, a greater emphasis is now placed on tools and processes for data exploration and visualization of large datasets. This paper described MUVTIME, an interactive graphical tool for multivariate time series visualization, developed in the context of TURNTAKE, a social interaction research project whose purpose is to study the interaction dynamics of parent-infant dyadic pairs and apply the resulting conclusions to Human-Robot interaction design.

MUVTIME was developed to allow fluent interaction of the user with the data. The user can control the time window of visualization, contrast different visualizations of the same or different time series, and also compare the video recordings of the experiments with the data. Also, it features an abstract visualization named AVTA, automatically extracted from the data and that works as high-level descriptor of the interaction between the dyadic pairs.

During the course of the project we have applied MUVTIME extensively to visualize the time series resulting from the interaction studies. The tool allowed us to make fast visual evaluations of interaction performance of the dyads, determining key moments of the interaction and contrasting the time series with the real action recorded in video. The overlapping of line and categorical plots was particularly useful in suggesting correlations between numerical and categorical variables while the strictly categorical visualizations (stacked categorical maps and AVTA) provided insightful visual footprints of the interactions. The horizon graph, initially implemented due to its high information/space ratio has not been particularly explored in the context of the TURNTAKE analysis because of our current focus on categorical variables. Future studies will include examining quantitative measures of motion like mean square velocity and in this case the horizon plots might be particularly useful.

ACKNOWLEDGEMENTS

This research was supported by: Marie Curie International Incoming Fellowship PIIF-GA-2011-301155; Portuguese Foundation for Science and Technology (FCT) Strategic program FCT-UID/EEA/00066/2013; FCT project PTDC/PSI-PCO/121494/2010. AFP was also partially funded by the FCT project (IF/00217/2013).

REFERENCES

Aigner W, Miksch S, Schumann H, et al. (2011) Visualization of Time-Oriented Data. *Human-Computer Interaction*: 69-103.

Bade R, Schlechtweg S and Miksch S. (2004) Connecting time-oriented data and information to a coherent interactive visualization. *Proceedings of the 2004 conference on Human factors in computing systems CHI 04* 6: 105-112.

Bakeman R and Gottman JM. (1997) *Observing interaction: An introduction to sequential analysis*: Cambridge university press.

Beard K, Deese H and Pettigrew NR. (2007) A framework for visualization and exploration of events. *Information Visualization* 7: 133-151.

Brodbeck D, Gasser R and Degen M. (2005) Enabling large-scale telemedical disease management through interactive visualization. *European Notes in Medical Informatics* 1: 1172-1177.

Buono P, Plaisant C, Simeone A, et al. (2007) Similarity-Based Forecasting with Simultaneous Previews: A River Plot Interface for Time Series Forecasting. *2007 11th International Conference Information Visualization (IV '07).* IEEE, 191-196.

Burgoon JK, Stern LA and Dillman L. (2007) *Interpersonal adaptation: Dyadic interaction patterns*: Cambridge University Press.

Cassotta L, Feldstein S and Jaffe J. (1964) AVTA: a Device for Automatic Vocal Transation Analysis. *Journal of the experimental analysis of behavior* 7: 99-104.

Catarci T, Santucci G and Silva SF. (2003) An Interactive Visual Exploration of Medical Data for Evaluating Health Centres. *Journal of research and practice in information technology* 35: 99-119.

Dachselt R, Frisch M and Weiland M. (2008) FacetZoom: A Continuous Multi-Scale Widget for Navigating Hierarchical Metadata. In: ACM (ed) *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* Florence, Italy: ACM Press, 1353.

De Barbaro K, Johnson CM, Forster D, et al. (2013) Methodological considerations for investigating the microdynamics of social interaction development. *IEEE Transactions on Autonomous Mental Development* 5: 258-270.

Fails JA, Karlson A, Shahamat L, et al. (2006) A visual interface for multivariate temporal data: Finding patterns of events across multiple histories. *IEEE Symposium on Visual Analytics Science and Technology 2006, VAST 2006 - Proceedings*: 167-174.

Few S. (2008) Time on the Horizon. *Visual Business Intelligence Newsletter*: 1-7.

Fouse A, Weibel N, Hutchins E, et al. (2011) ChronoViz: a system for supporting navigation of time-coded data. *CHI'11 Extended Abstracts on Human Factors in Computing Systems.* ACM, 299-304.

Gorgolewski KJ, Varoquaux G, Rivera G, et al. (2015) NeuroVault. org: A web-based repository for collecting and sharing unthresholded statistical maps of the human brain. *Frontiers in neuroinformatics* 9.

Gray J, Breazeal C, Berlin M, et al. (2005) Action parsing and goal inference using self as simulator. *Robot and Human Interactive Communication, 2005. ROMAN 2005. IEEE International Workshop on.* IEEE, 202-209.

Gschwandtner T, Aigner W, Kaiser K, et al. (2011) CareCruiser: Exploring and visualizing plans, events, and effects interactively. *IEEE Pacific Visualization Symposium 2011, PacificVis 2011 - Proceedings*: 43-50.

Guo D, Chen J, MacEachren AM, et al. (2006) A Visualization System for Space-Time and Multivariate Patterns (VIS-STAMP). *IEEE Transactions on Visualization and Computer Graphics* 12: 1461-1474.

Heer J, Kong N and Agrawala M. (2009) Sizing the horizon: the effects of chart size and layering on the graphical perception of time series visualizations. *CHI '09Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, April 4–9, 2009*: 1303-1312.

Hochheiser H and Shneiderman B. (2004) Dynamic query tools for time series data sets: Timebox widgets for interactive exploration. *Information Visualization* 3: 1-18.

Jaffe E. (2014) What Big Data Means For Psychological Science. *APS Observer.*

Jaffe J, Beebe B, Feldstein S, et al. (2001) Rhythms of dialogue in infancy: Coordinated timing in development. *Monographs of the society for research in child development*.

Johansson G. (1973) Visual perception of biological motion and a model for its analysis. *Attention, Perception, & Psychophysics* 14: 201-211.

Johnson K and Shiffrar M. (2013) *People watching: Social, perceptual, and neurophysiological studies of body perception*: Oxford University Press.

Kapler T and Wright W. (2005) GeoTime information visualization. *Information Visualization* 4: 136-146.

Kim JG, Snodgrass M, Pietrowicz M, et al. (2013) Visual Analysis of Relationships between Behavioral and Physiological Sensor Data. *Workshop on Visual Analytics in Healthcare.* Washington, USA.

Kipp M. (2012) Multimedia Annotation, Querying, and Analysis in Anvil. *Multimedia Information Extraction: Advances in Video, Audio, and Imagery Analysis for Search, Data Mining, Surveillance, and Authoring*: 351-367.

Kipp M, Martin J-C, Paggio P, et al. (2009) *Multimodal corpora: from models of natural interaction to systems and applications*: Springer.

Lee J, Lin S and Karahalios K. (2013) Visualizing Patterns of Social and Communicative Behavior in Children Using Plexlines. *Workshop on Visual Analytics in Healthcare.*

Lisboa IC, Sousa E, Santos JA, et al. (2014) Parent-child vocal coordination during joint construction tasks: a replication of the AVTA model of turn-taking. *9th Meeting of the Portuguese Association of Experimental Psychology APPE 2014.* Covilhã, Portugal.

Milde J-T, Milde J-T, Gut U, et al. (2001) The TASX-environment: an XML-based corpus database for time aligned language data. *Proceedings of the IRCS Workshop On Linguistic Databases, 11-13 December 2001*: 174-180.

Pereira AF, Lisboa IC, Sousa E, et al. (2014a) Vocal and motor coordination in joint-construction tasks during parent-child social interactions. *Workshop Dynamic Interactions Between Visual Experiences, Actions and Word Learning, part of the Fourth Joint Conference on Development and Learning and on Epigenetic Robotics.* Genoa, Italy.

Pereira AF, Smith LB and Yu C. (2014b) A bottom-up view of toddler word learning. *Psychonomic bulletin & review* 21: 178-185.

Rind A, Aigner W, Miksch S, et al. (2011) Visual Exploration of Time-Oriented Patient Data for Chronic Diseases: Design Study and Evaluation. *Information Quality in e-Health. Lecture Notes in Computer Science.* 301-320.

Rolf M, Hanheide M and Rohlfing KJ. (2009) Attention via synchrony: Making use of multimodal cues in social learning. *Autonomous Mental Development, IEEE Transactions on* 1: 55-67.

Roy BC, Frank MC, DeCamp P, et al. (2015) Predicting the birth of a spoken word. *Proceedings of the National Academy of Sciences*: 201419773.

Saito T, Miyamura HN, Yamamoto M, et al. (2005) Two-tone pseudo coloring: Compact visualization for one-dimensional data. *Proceedings - IEEE Symposium on Information Visualization, INFO VIS*: 173-180.

Sanderson P, Scott J, Johnston T, et al. (1994) MacSHAPA and the enterprise of exploratory sequential data analysis (ESDA). *International Journal of Human-Computer Studies.* 633-681.

Schanda J. (2007) CIE Colorimetry. *Colorimetry.* Hoboken, NJ, USA: John Wiley & Sons, Inc., 25-78.

Schmidt T and Wörner K. (2009) EXMARaLDA—Creating, analyzing and sharing spoken language corpora for pragmatics research. *Pragmatics-Quarterly Publication of the International Pragmatics Association* 19: 565.

Shahar Y, Goren-Bar D, Boaz D, et al. (2006) Distributed, intelligent, interactive visualization and exploration of time-oriented clinical data and their abstractions. *Artificial Intelligence in Medicine* 38: 115-135.

Smith LB, Yu C and Pereira AF. (2011) Not your mother's view: the dynamics of toddler visual experience. *Developmental Science* 14: 9-17.

Swan M. (2013) The quantified self: Fundamental disruption in big data science and biological discovery. *Big Data* 1: 85-99.

Vicon. (2015) *Nexus 2*. Available at: http://www.vicon.com/products/software/nexus.

Wang TD, Plaisant C, Shneiderman B, et al. (2009) Temporal summaries: Supporting temporal categorical searching, aggregation and comparison. *IEEE*

Transactions on Visualization and Computer Graphics 15: 1049-1056.

Wittenburg P, Brugman H, Russel A, et al. (2006) ELAN: a professional framework for multimodality research. *Proceedings of LREC.* 5th.

Wongsuphasawat K and Shneiderman B. (2009) Finding comparable temporal categorical records: A similarity measure with an interactive visualization. *VAST 09 - IEEE Symposium on Visual Analytics Science and Technology, Proceedings*: 27-34.

Yarkoni T. (2012) Psychoinformatics new horizons at the interface of the psychological and computing sciences. *Current Directions in Psychological Science* 21: 391-397.

Young FW and Bann CM. (1996) ViSta: The visual statistics system. Technical Report 94–1 (c).

Yu C and Smith LB. (2013) Joint Attention without Gaze Following: Human Infants and Their Parents Coordinate Visual Attention to Objects through Eye-Hand Coordination. *PLoS ONE* 8.

Yu C, Yurovsky D and Xu TL. (2012) Visual data mining: An exploratory approach to analyzing temporal patterns of eye movements. *Infancy* 17: 33-60.

Yu C, Zhong Y, Smith T, et al. (2009) Visual data mining of multimedia data for social and behavioral studies. *Information Visualization* 8: 56-70.