Visual Exploration of Common Behaviors for Developmental Health

Yi Han, Agata Rozga, John Stasko and Gregory D. Abowd
School of Interactive Computing
Georgia Institute of Technology
yihan@gatech.edu, agata@gatech.edu, stasko@cc.gatech.edu, abowd@gatech.edu

***Abstract*—Detecting early signs of developmental issues in a child is critical for successful interventions. Using behavior capture technologies such as video cameras, microphones and wearable sensors, it is increasingly possible to record and analyze human behavior in unprecedented detail. However, the analysis of such rich and complex dataset poses new challenges for developmental psychologists. We present the initial design of a visualization tool aimed at supporting the exploration of digitally captured behavioral data for studying developmental health.**

I. Introduction

Developmental psychologists seek to understand how people's behaviors change over the human lifespan. Developmental problems in childhood are particularly important as they could influence the rest of a child's life. One way to assess problems by psychologists is to observe how a child interacts with others. For remote and retrospective assessments, today these interactions are often captured on video. However, reviewing video is time-consuming and one video camera can only capture behaviors from one angle. As a result, as part of a multi-site, interdisciplinary project, we are exploring how to use additional sensing technologies such as multiple high-resolution video cameras and wearable sensors to capture more subtle, non-obvious behaviors [3]. The richer, more complex captured behaviors make it more challenging for data annotation and analysis. As a result, we are also creating new computational methods to help automatically detect and label behaviors of interest to support behavior analysis for psychologists. Nevertheless, along with the vast amount of automatically labeled behaviors comes new challenges in how to effectively explore them to study typical and atypical behavioral patterns.

Through discussions with developmental psychologists, we learned that they are particularly interested in finding typical or atypical behaviors among and between groups of children. There are two reasons: 1) they need to know what is considered a typical behavior among a large group of children in order to tell when they see an atypical behavior, 2) when they find an atypical behavior in a group of children, they are interested in seeing how this behavior correlates with the diagnosis of developmental disorders. The task of finding typical and atypical behavior patterns is normally achieved through human observations. However, with today's new behavior capturing technologies, psychologists want to investigate if they can find subtle, non-obvious behavior patterns that are difficult to identify through human observation, especially across a large group of children. The current tools being used for exploring behavior patterns, such as video annotation tools [6], are only effective in looking at behaviors of a small number of children. Therefore, we see an opportunity for a new visualization tool that can effectively support the exploration of typical and atypical behavior patterns exhibited by groups of children. Information visualization (infovis) represents a natural fit to the problem because it leverages humans' superb visual cognitive ability in finding insights from a large information space [2]. By interactively exploring the data, new interesting questions may arise from discovering unexpected visual patterns.

The application of infovis has been widely studied in understanding physiological health but largely overlooked in psychological health. Shneiderman highlighted three overlapping application domains in improving healthcare: personal, clinical and public health [5]. Specifically, extensive work has been done in analyzing the temporal events in electronic health records [4]. However, little work in infovis has focused on psychological health, not to mention developmental health that emphasizes behavior pattern exploration. As part of a project that aims to understand human behaviors through novel computational methods, we believe infovis can play a key role in supporting the exploration of the digitally captured behaviors for developmental health[1].

II. Behavior Data

We used a short social play protocol to collect social and communicative behaviors for analysis. To date we have collected data from over 140 children aged 15-30 months (>190 sessions, each child could have up to two sessions recorded at different ages.) The protocol is a five-minute interaction between an examiner and a child at a table [3]. There are five stages in the protocol: 1) greeting the child, 2) rolling a ball back and forth with the child, 3) reading a book with the child, 4) placing the book on one's head and pretending it's a hat, and 5) tickling the child. The child's reaction to the examiner's speech or intentional pause are scored by the examiner based on the expected responses such as whether the child gazed back at the examiner when his or her name is called.

We instrumented a lab to capture the behaviors exhibited by the child, the parent (who is often in the same room) and the examiner. The lab contains 13 video cameras, an overhead Microsoft Kinect sensor and two microphones. During the sessions, we additionally equipped the child and examiner with a lapel microphone and Affectiva Q sensors on both wrists that capture temperature, eletrodermal activity and accelerometry.

[1]Another paper (Lee et al.) submitted to this workshop also describes a visualization tool for the same dataset. This paper considers the problem of being able to create understanding across multiple sessions while the other focuses on individual sessions

The onsets and offsets of the child's behaviors during the play sessions are annotated frame-by-frame by multiple research assistants in ELAN [6], an open-source data annotation tool, on four key behavior variables: (1) Gaze (examiner's face, examiner's hands, parent's face, ball, book); (2) Speech (vocalizations, verbalizations), (3) Gestures (e.g., pointing, reaching); (4) Vocal affect (crying/whining, laughing). The specific subcategories such as "Gaze at Examiner Face" and "Gaze at Ball" in the Gaze behavior variable are the basic unit for annotation. These annotations are being used as ground truth by the rest of the research team that are designing computational techniques to automatically identify the behaviors. Note that at this point the number of fully annotated sessions ($\sim$40) is smaller than the total number of recorded sessions ($>$190).

In this visualization prototype, we used the human annotation data for the exhibited behaviors. In the future, these annotations will be replaced with behaviors automatically detected by computational methods which better capture subtle, non-obvious behaviors [3].

III. Understanding Needs in Psychology Research

We spoke with developmental psychologists informally to learn about their interest in studying social and communication behaviors. We discovered two interests that guided our design: 1) Responsive behaviors in social interactions and 2) Modes of communication used in the response.

Psychologists look for responsive behaviors in social interactions when studying developmental health in children. As a result, we used a unit of analysis called an "Examiner Bid (EB)" which is a time window that bins the responsive behaviors. This time window is similar to what is commonly used in developmental psychology research for coding responsive behaviors. In the social play protocol, the examiner uses specific phrases to introduce activities to the child (e.g., "Look at my ball!"), to guide the child through each stage (e.g., "Ready, set, go!), and even to prompt the child to engage in specific behaviors (e.g., "Can you turn the page?"). A complete list of these speech bids can be found on the x-axis of the Multiples View in Figure 1 and Figure 2. The examiner also engages in specific behaviors to elicit a response from the child, such as pausing the ball and book activity to see whether and how the child will react. Therefore, we defined an EB as a window of time that includes the examiner's speech/pause and a configurable amount of seconds afterwards. For example, the EB "look at my ball" includes the time when the examiner utters the words, holds up the ball, and a few seconds afterwards when the expected response should occur.

Psychologists look at four modes of communication in a child's response to assess social behaviors: 1) Attention 2) Speech 3) Gesture and 4) Affect. The four key variables in the human annotation (gaze, speech, gesture, vocal affect) are designed to code these modes of communication. It is important to understand how the different modes are used together in social interactions.

We created our design based on these research interests. Additionally, we also found that psychologists typically only use statistical visualizations in their research for communication purposes. Therefore, we thought leveraging their familiarity with bar charts would be a good place to start as it may lower the barrier of entry in using the visualizations.

Fig. 1. Subgroup Selection in the Distribution View. (a) Using the slider to focus on the aggregation of behaviors in the Ball stage EBs (b) A group of sessions that showed laughing behavior is selected (c) None showed verbalization but it is not uncommon (d) Few showed reaching behavior while many others do (e) The six sessions in the selected group are highlighted.

IV. Visualization Design

The visualization tool contains three views as shown in Figure 1: 1) Multiples View that shows how common a behavior is among a group of children over the social play protocol, 2) Distribution View that shows how common the frequency of a behavior is over a consecutive set of EBs and 3) Sessions View that shows how individual sessions are selected and visualized. In addition, there is a control panel for filtering the data and configuring the views. Using a user-centered design, we present the visualization design progressively based on the needs of developmental psychologists. The visualizations are implemented with the D3 visualization library [1].

Multiples View

Psychologists want to see which behaviors are commonly exhibited in each EB over the course of the social play protocol. The Multiples View (Figure 2) shows this information. The four modes of communication are displayed in four vertically-stacked bar charts. The x-axis of the view is the list of EBs. They are in the order of when they should appear in the play protocol. The y-axis of each bar chart shows the number of children that exhibited the behavior. In Figure 2 (a), for example, at a glance a psychologist can immediately notice that there are more sessions showing laughing behaviors at the tickling EBs. This behavior is expected because the examiner is engaging the child in a tickling game. Such expected typical behavior patterns were very useful in helping us demonstrate

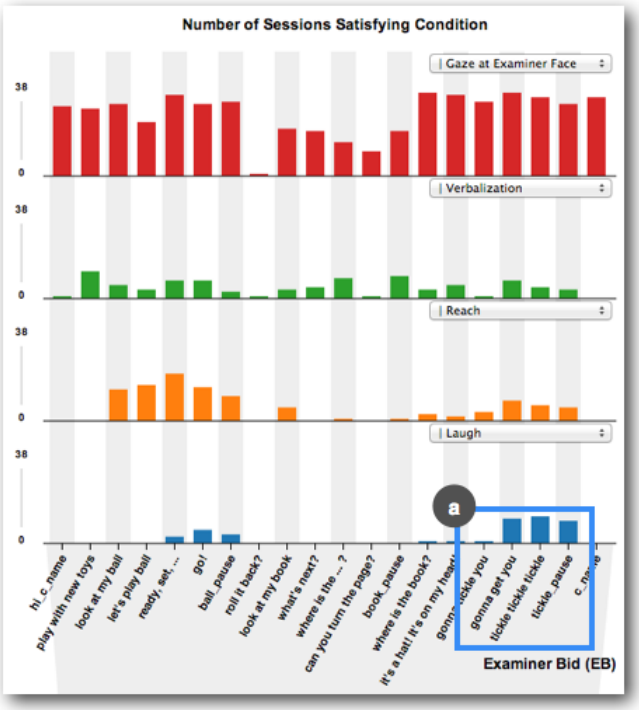

Fig. 2. MultiplesView. (a) More sessions showing laughing behavior at the tickling EBs.

the utility of the visualization prototype and establish trust in the visual representations. To inspect other subcategories of behaviors such as "vocalization," a psychologist can use any of the dropdown boxes above the bar charts to select them. If a psychologist would like to know how many sessions exhibited a behavior more frequently, they can tune a threshold with a slider in the control panel to find children that exhibited more behaviors in each EB.

Distribution View

Psychologists also want to find subgroups of children that show different amounts of behaviors over a group of EBs in the social play protocol. This information is shown in the Distribution View that is located beneath the Multiples View as shown in Figure 1 and Figure 3 (a). This view shows the distribution of sessions over the total amount of behaviors across a number of consecutively selected EBs. The EBs can be selected from the range slider above the Distribution View. The four selected behaviors in the Multiples View are shown on the x-axis. The y-axis shows the aggregated behavior amount. The sessions are distributed among 10 equal-sized bins over the y-axis accordingly. An additional bin at the bottom holds all the sessions that did not exhibit the specified behavior.

Sessions View

After finding a group of sessions of interest from the Multiples or Distribution View, psychologists will need to identify those specific sessions for further investigation. The Sessions View shows the list of sessions that are loaded into the system as shown in Figure 1 and Figure 3 (b)(c). Each session includes two visual components next to a session label: (1) A bar showing the amount of behaviors of the session and (2) A circle indicating whether the session is selected.

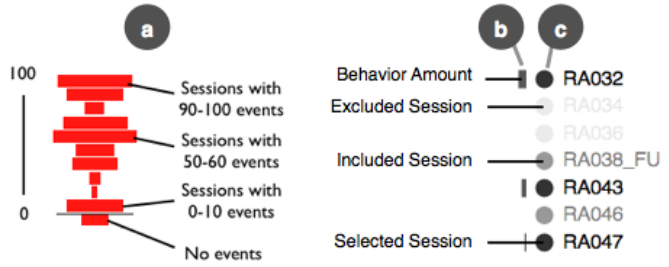

Fig. 3. Distribution and Sessions View. (a) Session distribution illustration (b) Bar: behavior amount of a session in a selection (c) Circle: selection status of a session.

V. Usage Scenarios

We present two usage scenarios of the visualization tool that explores how different groups of children exhibit common typical or atypical behaviors.

A. *Is a group of children behaving atypically?*

When psychologists see an atypical behavior exhibited by a group of children, they are interested in learning about who these children are and what other behaviors they have in common. In Figure 1 (b), a group of six sessions has been selected in the Distribution View after observing this group of children is behaving differently than the others. This group of sessions have children showing laughing behavior in the Ball stage EBs as selected by the range slider in Figure 1 (a). As there is only a small subset of sessions showing this behavior, we assume it is an atypical behavior that should be further investigated. When the six sessions in the bar are selected, these sessions are highlighted everywhere in the visualization tool while keeping the prior visualization faded behind in place. This design allows the psychologist to see if the selected subgroup of children in these sessions is behaving differently then the rest of the children. If more sessions are highlighted in a longer bar (typical behavior), it means the selected children behave similarly to the larger group. From Figure 1 (c) in the Distribution View, it seems that all these children that laughed did not verbalize at all over the entire social play session! However, when they are highlighted in a longer bar, the visualization implies that it seems typical for children to not verbalize in the social play session. On the other hand, few of them showed reaching behaviors that are common in the Ball stage as shown in Figure 1 (d). Are these children behaving atypically? From the highlighted six sessions in Figure 1 (e), the psychologist can see exactly which sessions they are in and either further the exploration by looking at different behavior variables, tuning other parameters or examining the video data of those sessions to better understand what might have caused these children to behave differently from others.

B. *How do children at different ages behave in the social play session?*

To assess developmental progress, it is important for developmental psychologists to understand how children should behave at different ages. By adjusting the age range slider on the control panel in two side-by-side windows as shown in Figure 4 (a), we can easily see the differences in common behaviors exhibited by younger and older children. Notably, (1)

Fig. 4. Comparing Younger to Older Children. (a) Different age ranges selected (b) More older children verbalized in their sessions (c) More younger children vocalized in the ball session (d) More younger children whined and cried.

more older children verbalized over the course of the session (Figure 4 (b)), (2) more younger children vocalized at the Ball stage (Figure 4 (c)) and (3) more younger children whined and cried (Figure 4 (d)). (1) and (3) can be expected but (2) might be interesting in exploring further. Why do children tend to vocalize more when they are younger and playing ball?

VI. Discussion

One key challenge in our work is that psychologists are not accustomed to using visualization to explore data. They typically only use statistical visualizations for communicating their research findings. Even though occasionally they use video annotation tools with a timeline visualization, it is not part of their standard workflow. As a result, we have encountered some initial reluctance toward using visualization for exploration. Nevertheless, after seeing an interactive demonstration of the visualization and how it provides additional opportunities for finding new research questions through visual exploration, psychologists started to become intrigued by the idea of using visualization in this new way in their work.

VII. Conclusion

Infovis is a promising new approach in supporting psychologists explore social and communicative behavior patterns for studying developmental health. Currently we are only using the human annotation data to create the visualizations. In the future, we would like to add computationally generated labels to the visualization because they could capture subtle behaviors that are not currently coded by the human annotators [3]. With more uncertainly and complexity in the these behavior labels, we foresee a great deal of opportunities for infovis to make an impact in developmental health down the road.

Acknowledgment

This work was supported by the NSF Expedition Award number 1029679. We would like to thank Opal Ousley, Lauren Adamson, Daniel Messinger, Karrie Karahalios, John Lee, Sanny Lin and the Ubicomp research group at Georgia Tech for their valuable feedback to this work.

References

[1] M. Bostock, V. Ogievetsky, and J. Heer. D3: Data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, 2011.

[2] J.-D. Fekete, J. J. v. Wijk, J. T. Stasko, and C. North. The value of information visualization. In A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North, editors, *Information Visualization*, number 4950 in Lecture Notes in Computer Science, pages 1–18. Springer Berlin Heidelberg, Jan. 2008.

[3] J. M. Rehg, G. D. Abowd, A. Rozga, M. Romero, M. Clements, S. Sclaroff, I. Essa, O. Y. Ousley, Y. Li, C. Kim, H. Rao, J. C. Kim, L. Lo Presti, J. Zhang, D. Lantsman, J. Bidwell, and Z. Ye. Decoding children's social behavior. In *2013 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.

[4] A. Rind, T. D. Wang, W. Aigner, S. Miksch, K. Wongsuphasawat, C. Plaisant, and B. Shneiderman. Interactive information visualization to explore and query electronic health records. *Foundations and Trends in HumanComputer Interaction*, 5(3):207–298, 2013.

[5] B. Shneiderman, C. Plaisant, and B. Hesse. Improving healthcare with interactive visualization. *Computer*, 46(5):58–66, 2013.

[6] P. Wittenburg, H. Brugman, A. Russel, A. Klassmann, and H. Sloetjes. Elan: a professional framework for multimodality research. In *In Proceedings of Language Resources and Evaluation Conference (LREC*, 2006.