Final Report for 3R01DC007428-04S1 'Collaborative Research Supplement and Infrastructure and Research Equipment for Advancement of Science

a. The three original Specific Aims are unchanged from the funded application. They are as follows:
   (1) Do our initial findings of comprehension preceding production with some grammatical constructions and use of some principles of word learning generalize to a larger sample of children with autism?
   (2) How well does the early language development of these children, including both comprehension and production, predict their language abilities when they reach ages 6-8 years?
   (3) How well does children's on-line efficiency in understanding language and language principles predict the individual variation characteristic of ASD?

The ARRA Supplement contributes to and extends all three Specific Aims by providing an ecologically-valid, densely sampled, and extremely efficiently analyzed audio-visual corpus of the speech and home environments of 3 participants, 1 of whom was diagnosed with an ASD and 2 who are typically developing.

b. Studies and Results:

This report covers the period from August 1, 2009 to July 31, 2011, including both years of the ARRA supplement. This report includes 2 major parts; Part 1 focuses on the technical aspects of the Speechome Recorder itself, and Part 2 focuses on the collection and preliminary analyses of the corpus.

During the first 1.25 years of the project, and then continuing through data collection, transcription, and analysis, the MIT team was focused on refining and developing the Speechome Recorder hardware and software. In particular,
1. The specific hardware components were refined and updated from previous prototypes.
2. Numerous software designs were developed and implemented, including those involved in remote upkeep, Speechome-human interfaces, serverside data auditing and management, transcription pipeline, and the front-end website.
3. Development of a smaller, less expensive Speechome Light was begun.

During Autumn 2010, the Speechome Recorders themselves became under construction. Therefore, at this time the UConn team focused on participant recruitment. Some recruitment had been conducted during the Spring of 2010; however, because the Speechome Recorders were not ready for deployment at that time, none of those families were enrolled in the study. A total of four families were contacted during Autumn 2010, and three agreed to participate in the study. Parents signed consent forms for themselves and their minor children; moreover, consent was obtained from frequently visiting grandparents, therapists, and friends. Pictures were taken of all consented participants, and data auditing revealed 0 incidence of errors, in which recordings were made of unconsented individuals.

Part 1: Technical Report on the Speechome Recorder

Speechome Recorder:

The Speechome Recorder (Figure 1) is a portable version of Speechome (Roy et al., 2006) audio/video recording technology. Its compact design enables swift, cost-effective deployment in clinics and homes. In these contexts, the Speechome Recorder was created to capture recordings of child/caretaker interaction and other behavior occurring in the course of daily life.

Figure 1: The Speechome Recorder

Hardware System Schematic:

Three different prototypes of the Speechome Recorder were created until we arrived at our current design (Figure 1). Figure 2 shows the schematics of the Speechome Recorders that were deployed.

As you can see in the schematics, the recorder has a dual camera system: One overhead panoramic camera and a frontal camera to capture facial expressions. The cameras used in this design are Lumenera Le165c video cameras which transmit data over Ethernet. Both cameras are outfitted with Fujinon FE185C057HA-1 lenses which have a 185 degree angle-of-view. This allows the overhead camera to capture most of the room from the top. The cameras are configured to record at 15 frames per second at a resolution of 960 by 960 pixels. Figure 3 shows frames (downsized) taken from the overhead and frontal cameras from one of the recorders. As you can see, the resolution of the video is sufficient to identify human participants and surrounding objects.

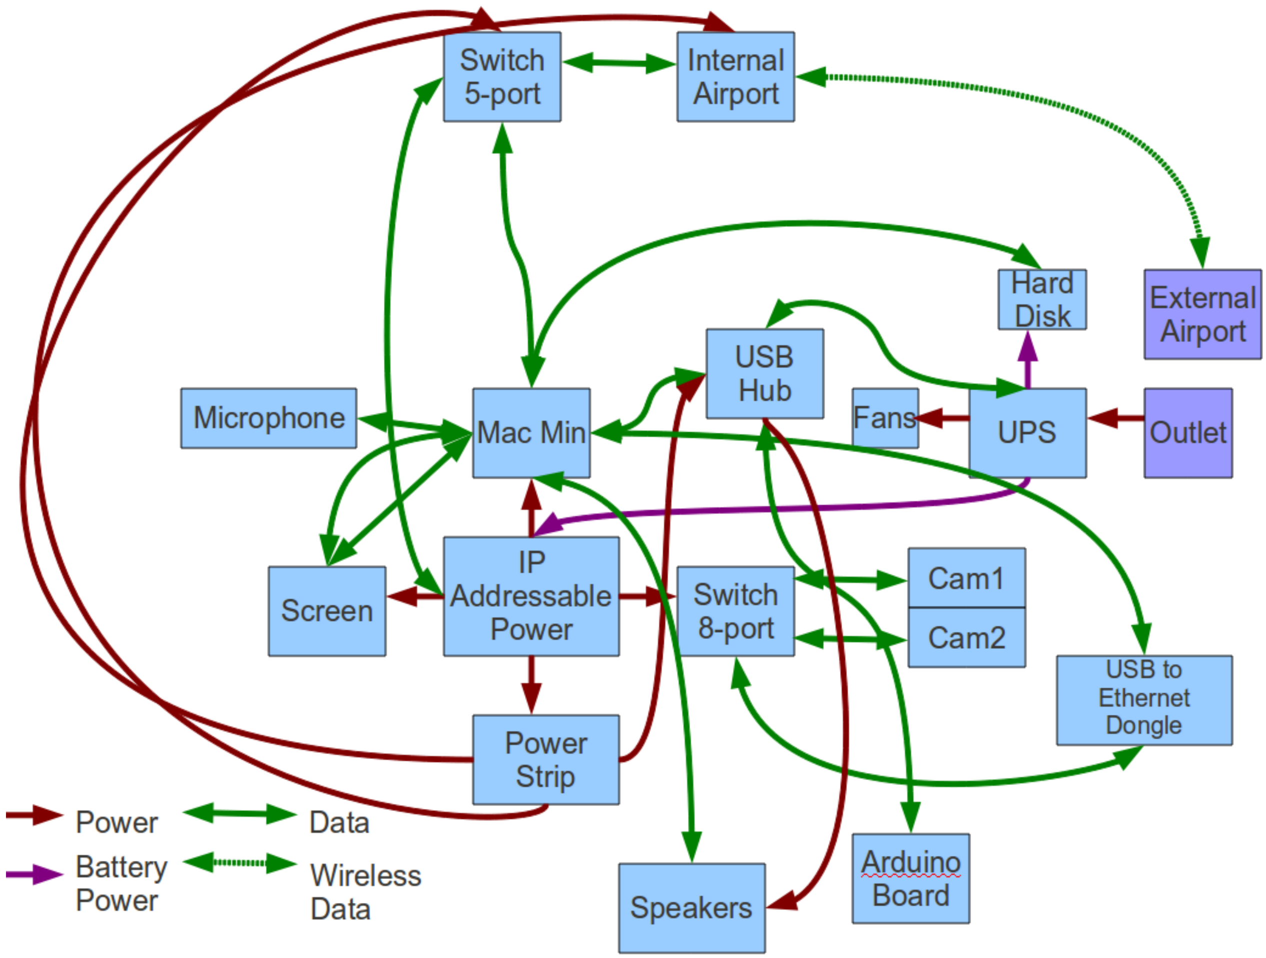

Figure 2: Schematic of the Speechome Recorder

The audio sensor in the recorders are AKG C562CM boundary layer microphones. These microphones use the surface in which they are embedded as a pickup. This allows a microphone placed in the head of the recorder to pick up speech in any corner of the room.

The Speechome Recorders are designed to be able to run for months without the need for maintenance and/or technical support. To that effect, all recorders are outfitted with cooling systems, voltage regulators, UPS battery backup, and 4 TB storage capabilities. Additionally, the recorders are outfitted with an IP-Addressable power supply which allows us to remotely turn power on or off to any of the hardware components in the recorders and thus remotely debug them.

Figure 3: Frames taken from the overhead and frontal cameras

Software on the Speechome Recorders:

In this section we will describe the software running on the Speechome Recorders.

Remote Upkeep:

As mentioned previously, the Speechome Recorders were designed to be able to run for months without the need for on-site visits. To enable that, we wrote a comprehensive diagnostic software suite that runs on the recorders. This software automatically checks the operational status of all hardware components in the recorders every hour and sends a report to our servers at MIT. This way we are alerted to all issues in a recorder almost instantaneously. Moreover, at the end of each day all data recorded that day is transmitted over the Internet to our servers at MIT. This allows the recorders to use that space if need be (i.e. if the 4TB drive is full). This removes the need for us to manually replace the drives on the recorders every few months. Additionally, this grants us instant access to the recorded data for transcription, analysis, and backup.

User Interface:

Figure 4 shows the user interface on the Speechome Recorders. As you can see there are five buttons in the interface. Here we will describe what each button does.

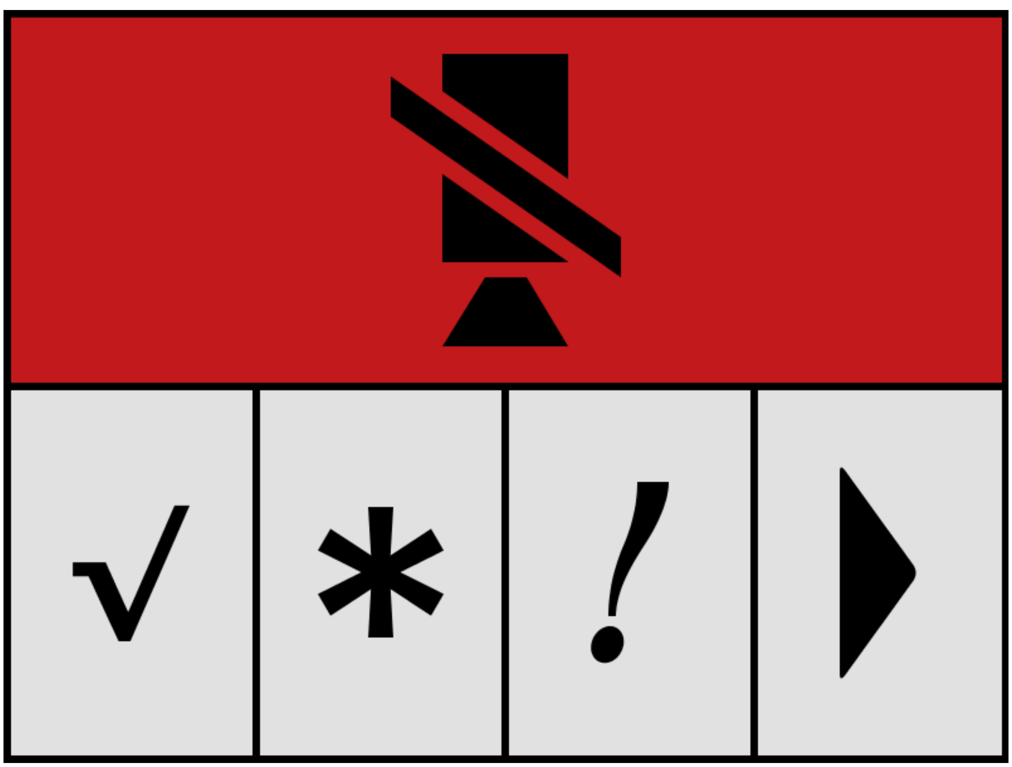

Figure 4: User Interface of the Speechome Recorder

On/Off button
The top red button is the on/off button. It allows the users to turn recording on or off (the button turns green when recording).

Consent button
The left most button at the bottom is the consent button. This button takes a picture of a person who's consenting to be recorded. The picture is sent to our servers at MIT. This picture is then used by our auditing officer to verify that only people who have consented are recorded. This process is explained in more detail later in this report.

Ooh button
The second button from the left is the “Ooh” button. This button allows the user to event mark (i.e., time stamp in the video and audio record) anything interesting or important that was recently recorded. This allows the researchers and/or the family to easily access interesting and/or important recorded material.

Oops button
The third button from the left is “Oops” button. This button allows the user to mark any recently recorded segment for deletion. Any segment marked by the “Oops” button will be automatically deleted without being seen by anyone else.

Playback button
The last button is the playback button. This button allows the user to playback and review any recorded data on the recorder's screen.

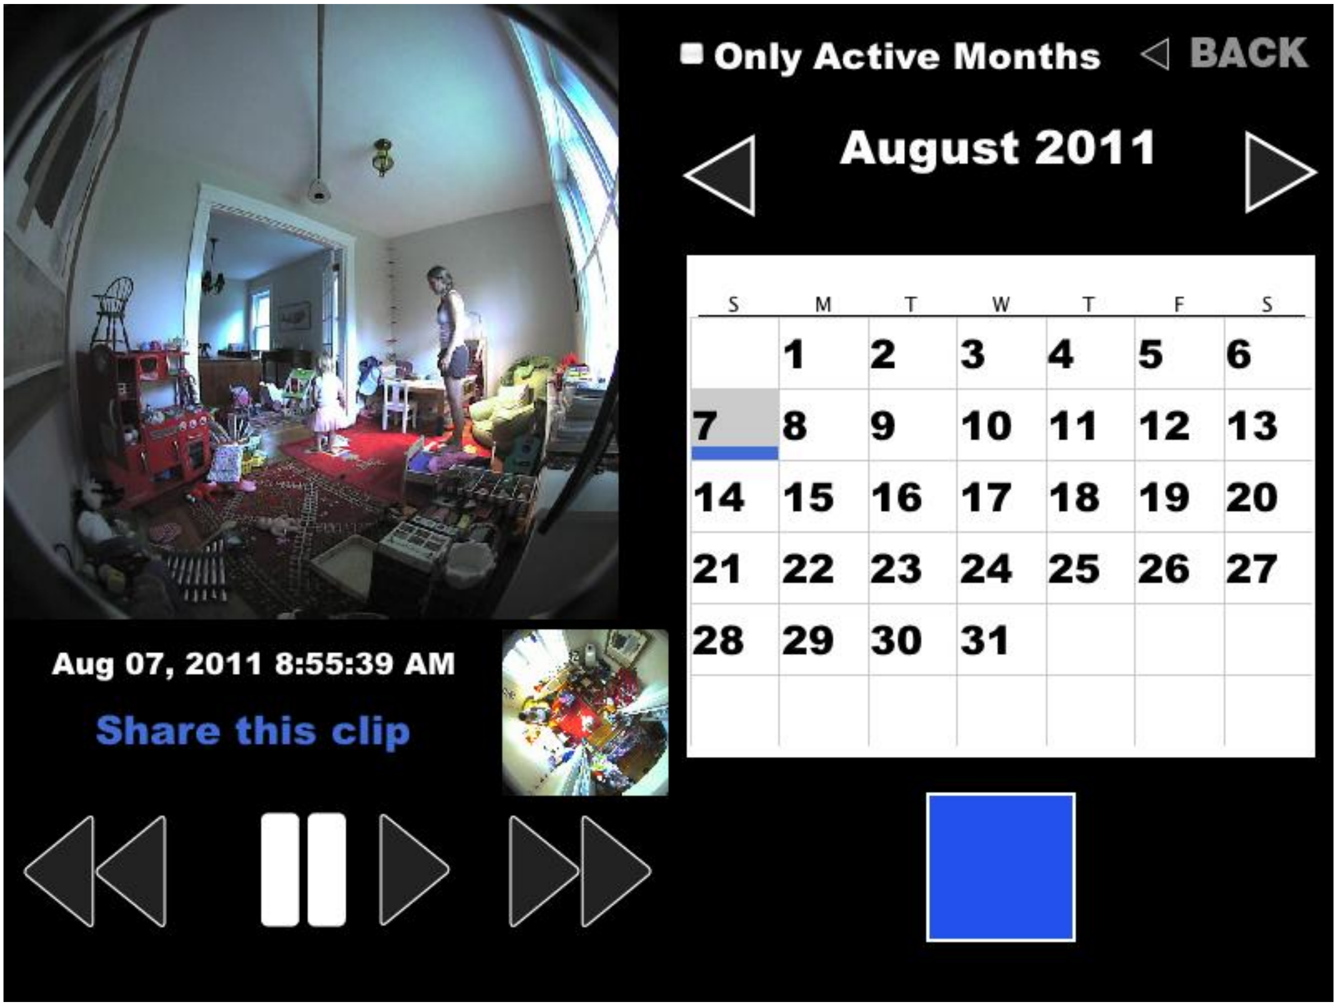

Figure 5: On-board SR playback software

Figure 5 shows the playback interface. The calendar on the right highlights the days where data was recorded. Once a day is selected, the user can playback the clips from that day using the movie player on the left. The user can also switch between the overhead and the frontal views. If the user for any reason wishes to share the clip with family members (or keep it themselves), he/she can select the “Share this clip” button which marks the clip for sharing. Per the user's request, we then turn the clips into DVDs and ship them to the user. As we will describe later in this report, users can also use the Speechome Recorder's website to view, download, and share all recorded clips from their house.

Server-Side Software:

In this section we will describe the software running on our servers at MIT.

Data Management:

As mentioned earlier, data is uploaded from the recorders to our servers at MIT on a daily basis. All data uploaded from the recorders are also backed up daily. Moreover, the data is down-sampled and converted into mpeg movies viewable by any media player (as described later in this report).

Data Auditing:

We developed software that allows our auditing officer to efficiently review recorded data to insure that only people who have consented have been recorded. Since there is too much data to inspect manually, we first pass the data to our automatic face recognition software. This software picks out frames that have human faces in them. Our software then randomly picks one of these frames for every 15 seconds of recorded video. These frames, along with the picture of the people who have consented to be recorded, are passed to our auditing software which is used by the auditing officer for inspection.

Figure 6 shows our auditing software. The numbered buttons on top represent each of the active recorders (the fifth recorder is a test recorder in our lab). Once a recorder is selected, frames from that recorder are shown (the left image) along side a list of pictures of people who have consented to be recorded (the picture scroll-panel is on the right). As mentioned in previous sections, these pictures are taken on the Speechome Recorders upon consent. The auditor can then compare the faces in the frames to the consent pictures. If a face is not found, the auditor can use the giant red “X” button to mark that picture. When complete, a report is sent to us about all marked images. Once the persons who have not consented are identified we either try to get the persons' consent or if that fails, delete all recorded data that includes those persons.

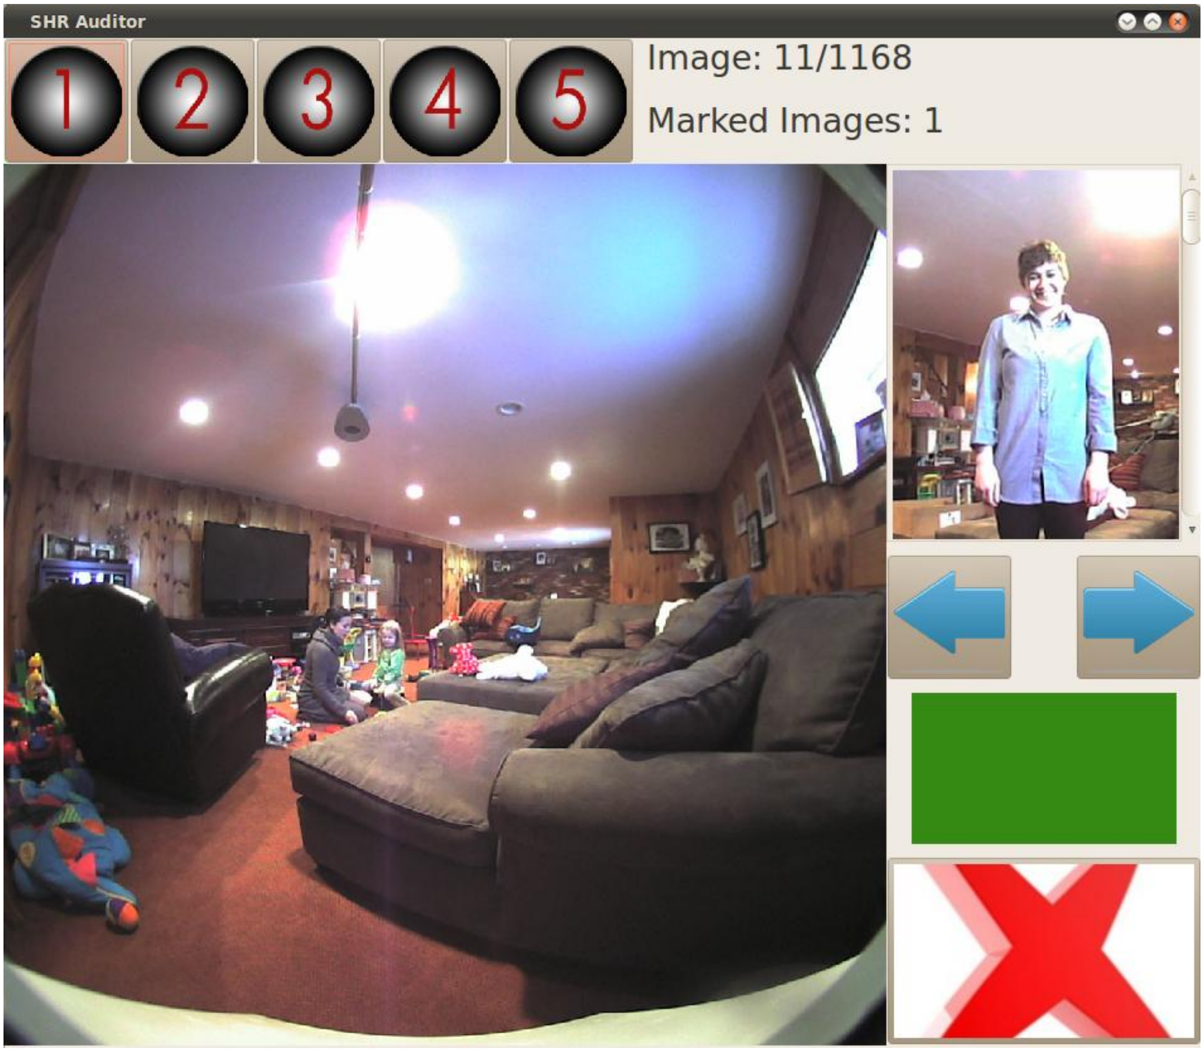

Figure 6: Data Auditing Software

Transcription Pipeline:
Once the data has been audited, we then begin preparing data for transcription. First, we pass the audio through our noise-reduction (band-pass filter) software to make the audio as clean as possible. Next, we run the audio through our speech detector which differentiates human speech from other audio signals. The program then takes all the human speech and divides it up into short utterances.

These utterances are then passed to our transcription system, BlitzScribe, which facilitates fast and accurate speech transcription (Roy & Roy, 2009). Human transcribers at the University of Connecticut transcribe speech data using BlitzScribe. It's important to note that all the steps up until the human transcription are completely automatic and require no human intervention. Figure 7 illustrates this audio pipeline.

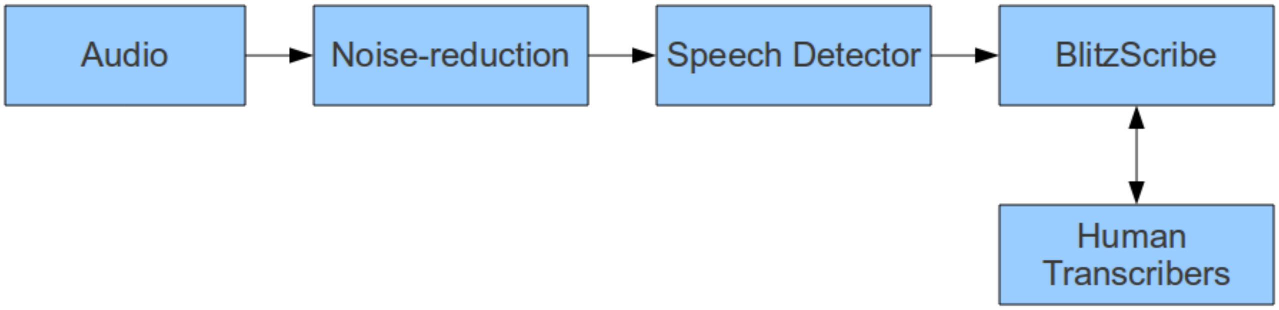

Figure 7: Transcription Audio Pipeline

Front-end Website:
Finally, a front-end website was created to allow researchers to easily browse through and view recorded data. The website is hosted at the MIT Media Lab and is highly secure as well as password protected. Figure 8 shows the front page of this website. As you can see, the website allows you to browse the data from any of the deployed recorders. Once a recorder is selected, you are presented with several options, seen in Figure 9.

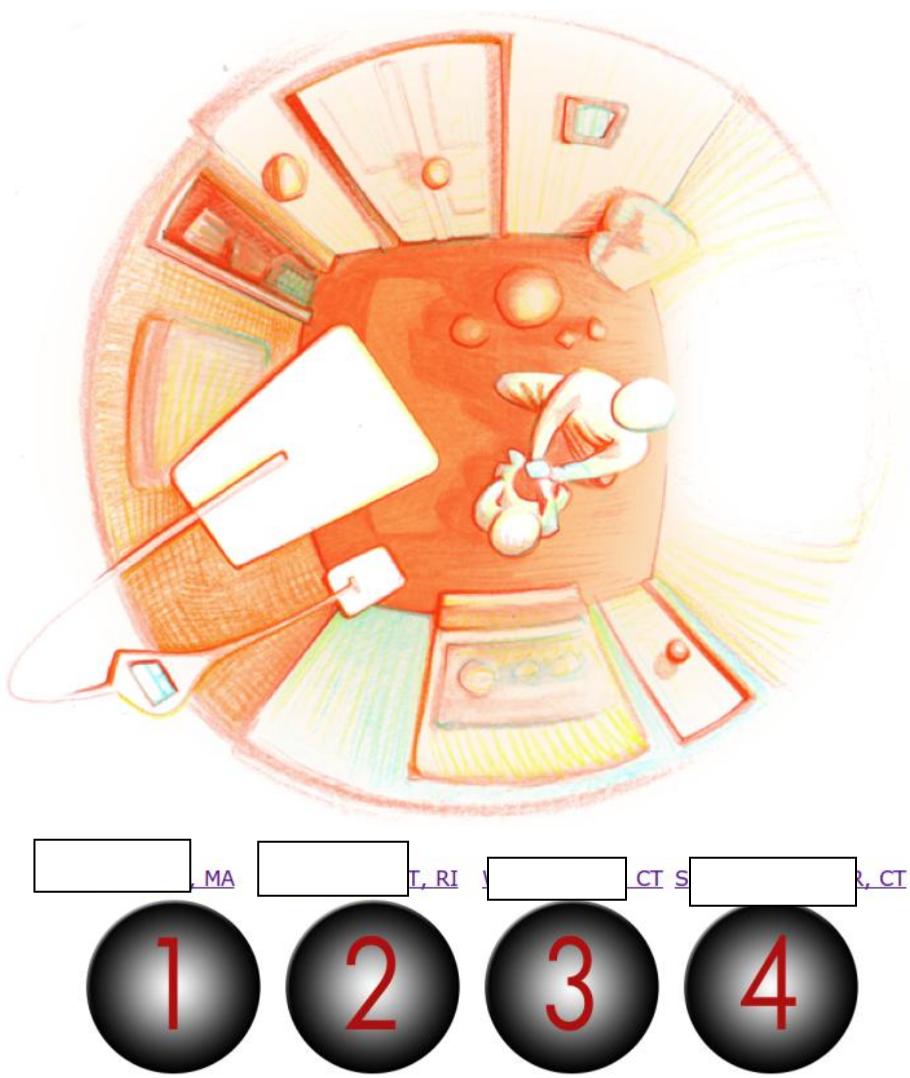

Figure 8: The Speechome Recorders' Front-end Website

The “Data Reports,” “Audit Reports,” “Transcription Progress,” and “Recorder Status” links allow us to monitor different aspects of the selected recorder, everything from the hardware status of the recorder to the amount of data recorded and transcribed from that recorder.

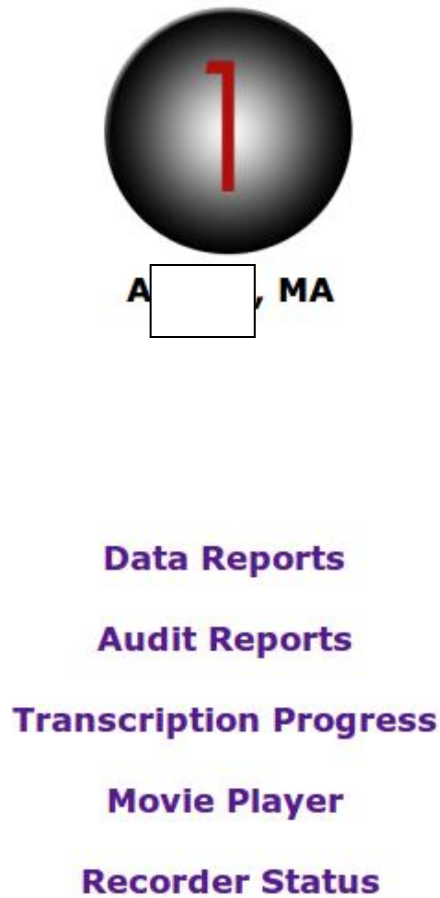

Figure 9: Speechome Recorder Website

The “Movie Player” link points to another password page that allows the researchers (and the families) to browse and view the movies and the transcriptions (if available) of any segment of the recorded data.

Figure 10, shows the movie player interface along with a sample movie segment and its corresponding transcription. It should be noted that because of our fast automatic data processing pipeline, it takes on average only 48 hours for any recorded data to be available on the website for browsing (excluding transcription).

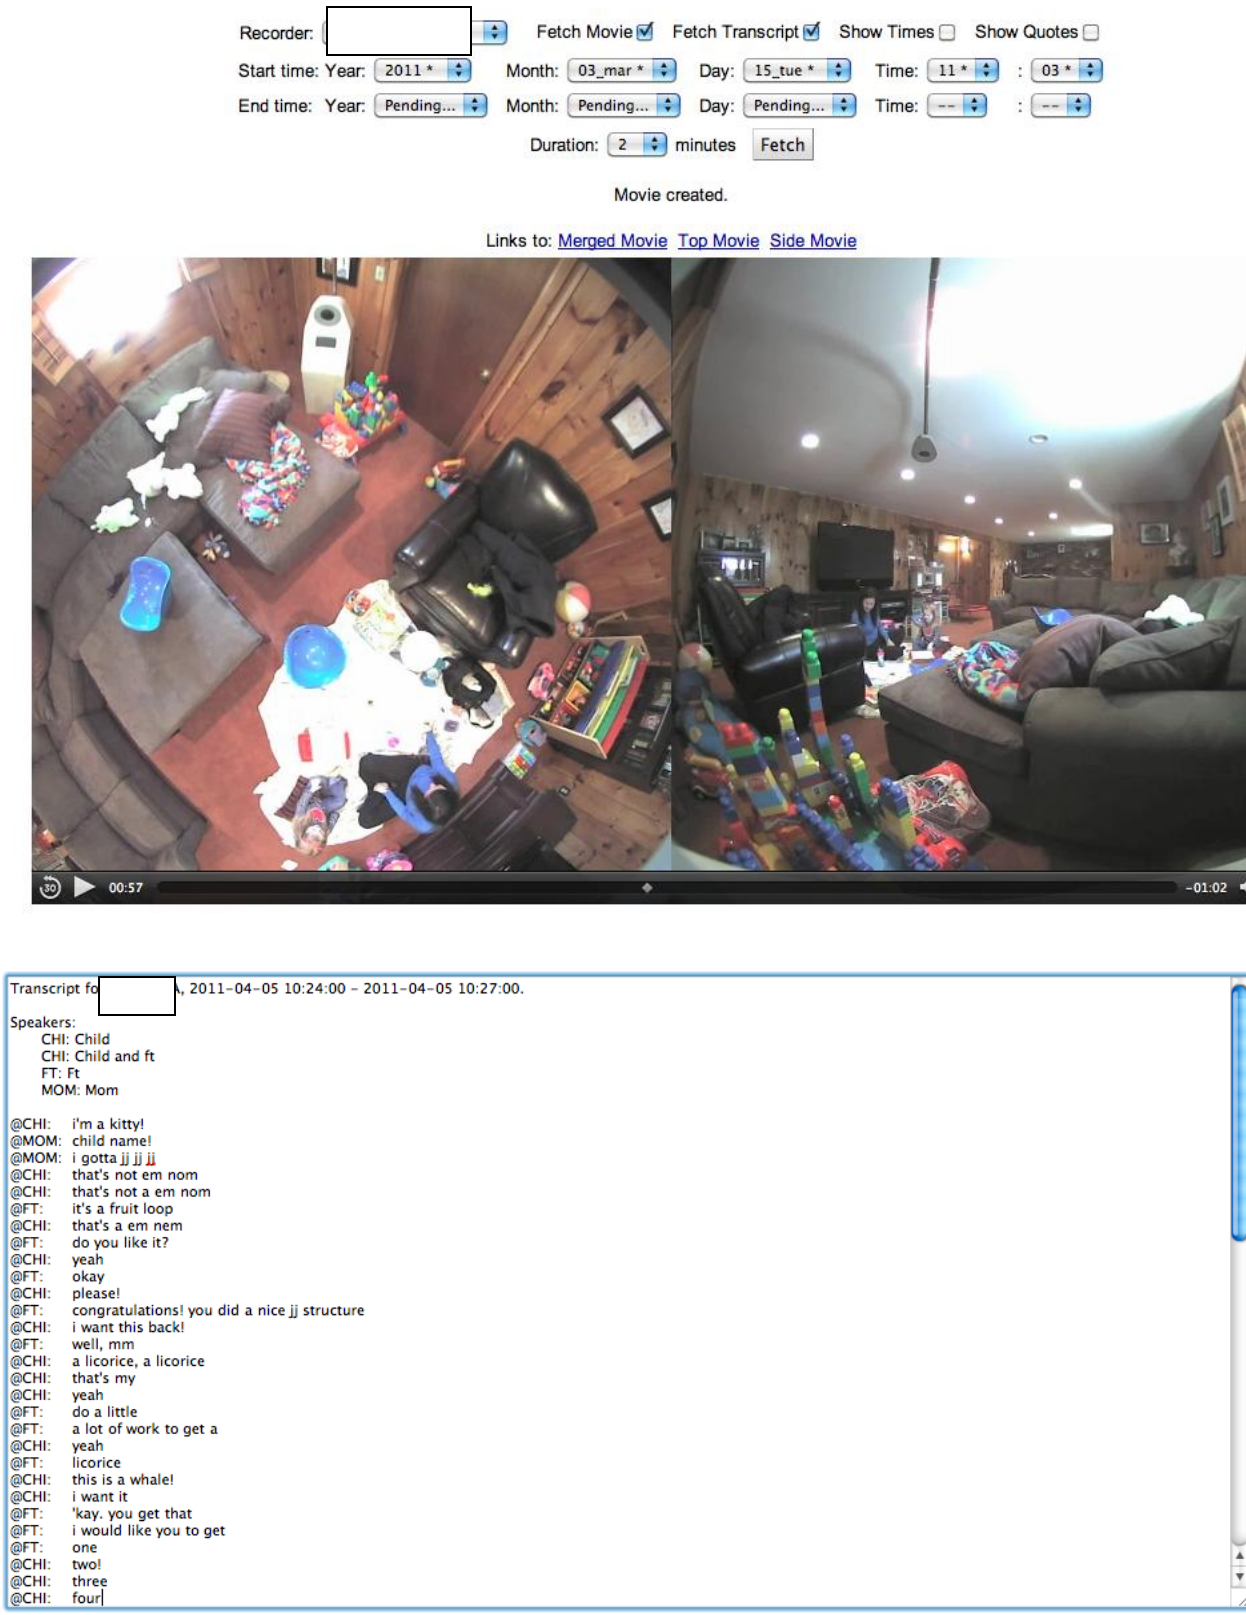

Figure 10: Movie Player and Transcription Viewer

Data:

The following table shows the amount of data recorded and transcribed from our deployed recorders as of September 1st, 2011.

Table 1

Recorder	Total Data Recorded (Hours)	Percent Data Transcribed	Number Of Words Transcribed
A, MA	51.5	100.00%	145729
W, CT	10.3	72.00%	25586
SW, CT	34.2	5.00%	4144
Pawtucket, RI*	40.53	0.00%	0
Total	136.53	44.00%	175459

- Still actively recording.

Speechome Recorder Light:

The main design of the current Speechome Recorder was done about four years ago. Though the recorders were designed to be as efficient and compact as possible at that time, we think that with new technology we can redesign the recorders to be considerably smaller, faster, cheaper, and with greater capabilities. In this section we will describe the design of the new Speechome Recorders.

Major Updates:

First, we intend to use iPads as the sole display and interface device for the new recorders, eliminating the sometimes problematic touch screen. Using iPads has the following advantages:

- Through WiFi, the families can operate their recorder from anywhere in the house;
- The viewing screen will be 2X bigger than the touch-screens with a much greater resolution;
- Can use iPad for clip viewing, clip sharing and viewing recording statistics from anywhere in the house; and
- Remote viewing through the SR's cameras (i.e. checking on the kids playing from another room).

Second, the new recorders will use standard off the shelf web-cams, which have become extremely cheap with very good quality and resolution in recent years. Replacing our current cameras with web-cams not only saves a significant amount of money and space, it will also enable the new recorders to record directly to QuickTime, eliminating the need for transcoding our videos as well as eliminating issues with audio/video synchronization.

Third, the new recorders will be outfitted with Microsoft Kinect sensors. These sensors allow for very sophisticated 3D mapping. These 3D sensors are capable of collecting up to 60,000 points of data and can track the motion of the body, limbs, and joints at around 30 frames per second. This allows for real time detection and analysis of the human form and its movement. The 60,000 points of data collected per frame allows objects as small as fingers to be detected (Figure 11). Moreover, the data generated by the Kinect can potentially be used for automatically identifying and tracking humans. Figure 12 shows an example of using the Kinect to capture 3D depth information from the environment. Even with state-of-the art computer vision algorithms, it is very difficult to automatically and reliably track a human body, let alone limbs and joints using data from our current recorders. However, as shown here adding a Kinect 3D sensor to the recorders greatly simplifies this task.

Figure 11: Capturing fingers using the Kinect's 3D sensors

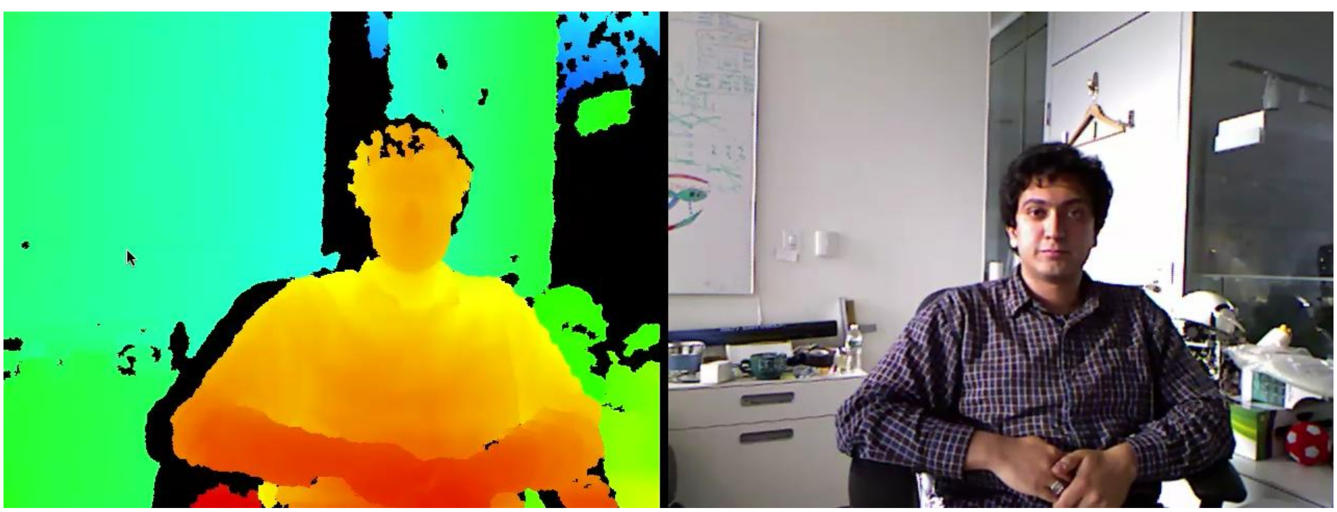

Figure 12: Kinect's 3D Capture

Fourth, we intend to store very little data locally on the recorders. All recorded data will be uploaded to our servers at the earliest possible time, therefore eliminating the need for massive storage devices on-board the recorders. This will make the new recorders smaller, lighter, and cheaper.

Fifth, by leveraging recent breakthroughs in low-cost cameras, sensors, and storage devices the new recorders will be significantly cheaper than the current ones. We estimate the new recorders to cost about $3,000 per unit compared to the more than $15,000 that each of the current recorders cost. This cost saving will be a major contribution to scaling SR use for scientific and everyday purposes.

Finally, the new design will be significantly simpler, lighter, and aesthetically pleasing than the previous design. The body of the new recorders will be made out of wood, making the recorders look more like furniture. The next section explains our current design.

Design:

Figure 13 shows our current (and very crude) approximation of the new design. We intend to have the recorders installed in a corner (which is in most cases the ideal location) of rooms. Thus, we are using a trapezoidal design which allows the recorders to “lean” into the corner for stability. The trapezoid provides clearance for corner obstructions, and the central support leg is far enough from the actual corner to allow for obstructions like baseboard heaters. The goal is to have the main web-cam at about 4' from the floor and the top web-cam at about 7'.

Experiments to Verify Design:

In order to compare the performance of the new design with that of the old design, we intend to install a prototype of the new recorder next to the current recorder in a lab setting. We will also set up software which will enable recording to be synchronized on the two recorders such that whenever one records, the other will also automatically start recording, enabling us to have data captured at identical times and conditions from both recorders. Once installed, we can test and compare the following features of our new design versus the old design:

- *Software stability*
- *Software ease of use*
- *Video/Audio quality*
- *Transcription*
- *Automatic speech detection*
- *Automatic speaker identification*
- *Transportability*
- *Ease of assembly*
- *Durability*
- *Aesthetics*

Ideally, the new design will surpass or at the very least have similar performance as the old design in all of the aforementioned categories.

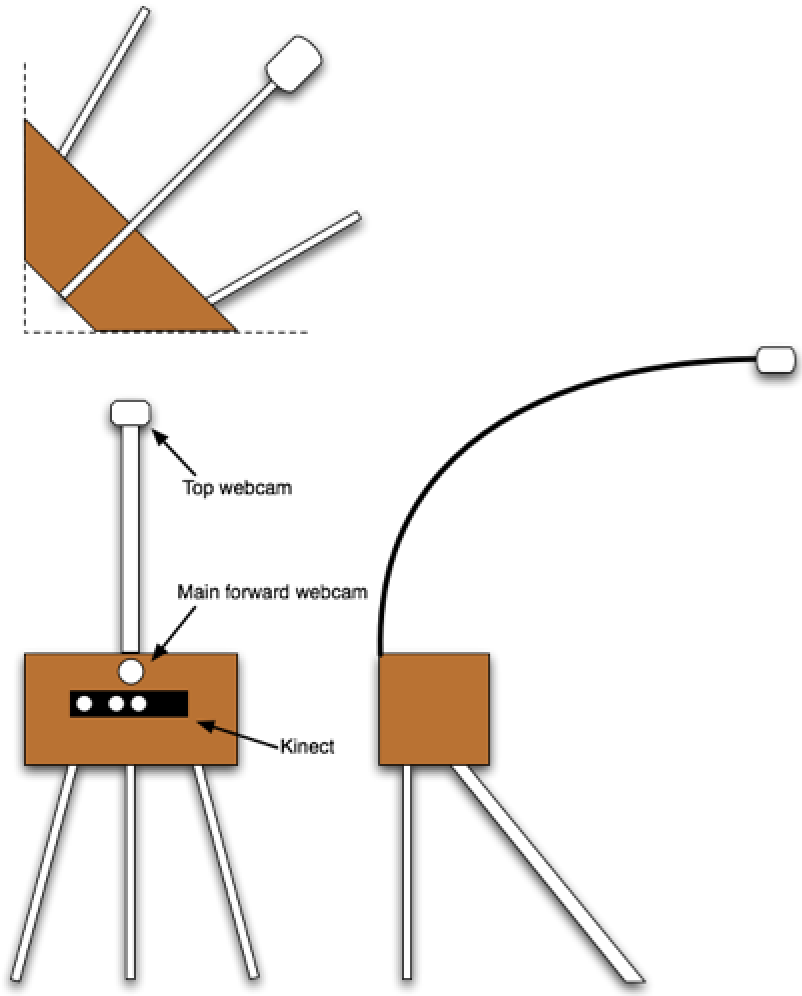

Figure 13: New Speechome Recorder Design

References:

Roy, D., Patel, R., DeCamp, P., Kubat, R., Fleischman, M., Roy, B., et al. (2006). The Human Speechome Project. In *Proceedings of the 28th Annual Cognitive Science Conference* (pp. 2059–2064). Mahwah, NJ: Lawrence Earlbaum.

Roy, B. C., & Roy, D. (2009). Fast transcription of unstructured audio recordings. In *Proceedings of Interspeech*. Brighton, England.

Part 2: Data Collection, Transcription and Preliminary Analyses

The Speechome Recorders (SRs) were installed in recruited participants' homes in February 2011 (Acton SR) and March 2011 (Willington SR, South Windsor SR). SRs were installed in family rooms, in the basement or on the 1st floor. Families were visited several times during the installation and consent process, and given extensive training in the SR-human interfaces. Subsequently, families' use of their SR was monitored via the front-line website, and families were contacted by phone and/or e-mail (as preferred) on a weekly basis to encourage use of the SR and to address any questions.

Transcription training began in May 2011, and transcription using Blitzscribe began in earnest in June 2011. Two undergraduates have been trained in Blitzscribe and have been transcribing from 5-10 hours/week through October 2011. The A corpus and the W corpus have been fully transcribed; the SW corpus is close to completion.

Preliminary analyses have been carried out with some of the data from the A corpus. This child was given an ASD diagnosis during her 2nd year, and was recorded for our study from 33 through 39 months of age. A total of 13 sessions during February 2011 have been analyzed; these comprise all of the sessions, lasting longer than 20 minutes, which were recorded during this month. Basic findings about A's speech are in the following table:

Table 2:

Session	MLU words	word types	word tokens	tokens/minute
1	2.328	117	530	31.18
2	2.592	165	418	12.48
3	2.373	154	468	22.83
4	2.461	205	699	17.48
5	2.612	186	611	23.50
6	2.652	392	2537	20.22
7	2.322	176	740	16.26
8	2.644	407	2259	22.15
11	2.723	437	3332	19.83
12	2.549	186	897	25.27
13	3.128	181	604	22.37
14	2.442	154	526	18.46
15	2.641	142	366	12.20

A's speech output is actually remarkably similar to that of a same-age typically developing speaker of English. For example, the TD children in Cohort 2 of the longitudinal parent grant, when they were 32 months of age (Visit 4), had average MLUs of 2.80 (range 1.74 to 4.16), and produced an average of 20.2 word tokens per minute. Thus, A has made great strides in her language development since diagnosis and the onset of intervention.

Four of the 13 sessions, designated in yellow in Table 2, comprised home therapy sessions; the rest of the sessions involved the 'free play' of A and her family. As the table shows, A's speech output is not markedly different during therapy vs. free play sessions. The therapy sessions are frequently longer than the free play sessions, resulting in more word types and tokens produced, but the MLU and token/minute measures are quite consistent. A slight increase in MLU from early to later sessions might be seen; however, it is not yet statistically reliable.

Because A's speech output is relatively advanced for her diagnostic group, our preliminary analyses focused on her verb use, rather than on her mapping of words to objects as originally planned. A's verb uses were extracted and coded for occurrence during present and past contexts.

Our first analyses serve to demonstrate how similar was A's verb use compared with that of TD children. A total of 1,260 verb tokens were extracted, of which 128 referred to past events (10%) while the rest referred to present events. This overwhelming preference for talking about the 'here and now' rather than the 'there and then' is also seen with TD children (Sachs, 1983). Moreover, of the 1132 verb tokens referring to present events, 43 were correctly marked with the 3rd person present indicative (e.g., *He needs a banana*). Of the 1086 verb tokens referring to present events that were unmarked, 1021 were nonetheless correct, because these were either declarative sentences with A talking about herself (e.g., *I want it*), or imperative sentences with A directing her father or sister (e.g., *close it*). Again, the facts of English make such pervasive use of unmarked present tense entirely typical for this age and context.

Many (67%) of the errors of omission were also quite typical, including missing 3rd person –s *(birdy come_)*, missing auxiliaries (*Where ___he go*?), missing progressive suffix –ing (*I'm stay__*), and missing 'to' indicating the infinitive (*I want__ hold this*).

Only 3 errors of 'commission' (Snyder, 2007), in which the wrong agreement marker was used, were observed: (1) *I don't wanna has socks*. (2) *People has to pick up.* (3) *They all wears.* Such a low error rate for commission errors is also typical; however, these utterances will also be re-checked to make sure they were transcribed correctly.

A's references to past events are shown in Table 3. These include the verb types which were produced in unmarked form (e.g., *You brush it yesterday*), those produced with the correct regular past suffix –ed (*What happened?*) and those produced with the correct irregular past marker (*That broke*). I find it highly significant that A also produced one token with an overgeneralized past tense marker (e.g., *I throwed*), as this is the first documented evidence of a child diagnosed with an ASD producing such an overgeneralization. These overgeneralizations, while errors, provide compelling evidence that the child is able to go beyond her input (i.e., no adult would say this) to produce a rule-based form. Such overgeneralizations are highly significant for children with ASDs because of their well-known difficulties with rule abstraction (e.g., Minshew, Meyer & Goldstein, 2002).

A also produced 16 verb tokens that were unmarked for past tense, 19 that were marked for regular past tense, 82 that were marked for irregular past tense, and 10 that included an auxiliary marked for past tense (e.g., *I was using it*). As Table 3 shows, only 1 verb was produced in the regular past in more than one session (*happened*), whereas 10 verbs (those in the red font) were produced in the irregular past in more than one session. This, too, is typical: verbs that take the irregular past tense are among the most frequent in adult English, and thus are the most likely to recur across sessions. Interestingly, only 3 verbs appeared in both marked and unmarked past tense forms (green font: *have/had, do/did/done, bump/bumped*). These findings suggest that A is at the beginning of acquiring the regular past tense, and we expect to observe these forms increasing in frequency across the next few months.

Table 3

Session	Unmarked	Regular Past	Irregular Past	Overgeneralized
1		happen	was	throwed
			did	
			made	
			stuck	
			threw	
2	brush	walk	broke	
			took	
			got	
			put	
3	have	check	had	
		punch		
		jump		
4		want	drew	
5	rip	whip	went	
			found	
			knew	
			shut	
6	spill		fell	
	miss		done	
	roll			
7	no new past tense forms			
8	kill		read	
	cry			
	spill			
	step			
11	do	clean	brought	
	eat	look	rang	
	bump			
	mix			
12	drop	bumped	hurt	
		moved		
		turned		
13	carry	called		
14			said	
			flew	
15	no new past tense forms			

Thus far, A's use of verb tense and agreement appears to be developing entirely typically. However, we also observed a set of errors that were quite unexpected and thus far unexplained. These include 21 tokens in which the verb was referring to present events, was unmarked, and yet appeared in a noun frame; for example, *I am a get.* Table 4 presents the verb types that were used in this frame; across the 13 sessions comprising about 0.5 months, this frame was quite productive (i.e., appeared with a number of different verbs).

The video context was scrutinized for possible interpretations of this novel frame; it appears that it refers to ongoing (I am walking) or imminent (I am going to show you) events. However, what is puzzling is why A did not use her already-attested correct constructions for these construals, with these verbs. For example, A produced *I am gonna X* at Session 1 and many subsequent sessions, so why not insert these other verbs into that frame? Moreover, as Table 4 shows, A produced some instances of verbs in the 'I am verbing' frame after using them in the 'I am verb' frame, but other such instances appeared before: For example, *I am going* was produced at session 5, so why not produce this frame again with *go* at session 7? And *I am doing* was produced at session 5, so why not produce this frame again with *do* at session 8? One other such instance was observed, in which a verb was used in a noun context: *I don't have a cry.*

Interestingly, only 2 instances were observed in which A used this frame appropriately, with nouns: *I am a mermaid* and *I am a elephant*.

Table 4

Session	Uses of 'I am a V'	1st sesson of use with the 'be+Ving' form
1	walk	
1	play	3
5	pay	
5	put	11
6	show	
7	go	5
7	cut	8
8	do	5
11	draw	
11	get	14
13	make	5
13	open	11

Even at this preliminary stage, these 'I am a V" errors point to several exciting aspects of A's speech and language. First, they reveal that she is not just imitating the adult speech in her environment, because no adult native speaker of English would produce such an error. Therefore, A is going beyond her input, a process that is accepted for TD children (e.g., Chomsky, 1965; Goldin-Meadow, 2003; Naigles et al., 2009) but thus far unexplored (and perhaps unpredicted) for children with ASDs. Children with ASD are frequently characterized as producers of routine and rigid language formats, and the very fact that A has produced a novel form—and uses it with a variety of different verbs—belies this characterization as true for all children with ASD. Second, the fact that this frame *is* productive for A, that she uses it multiple times and it appears to convey similar interpretations each time, points to the implication that A derived this frame from her own (albeit incorrect) analyses of her input. That is, she (implicitly, of course) performed some analyses of her input speech and contexts of use, and came up with a frame that is unattested in English, but functions adequately (at least for these 2 weeks) to communicate. In other words, A is not just slavishly memorizing the words and sentences she hears, but also analyzing them for underlying patterns, and creating new patterns as well. She is both a rule-user and a rule-creator, two attributes not commonly applied to individuals with ASD. Of course, she has also come up with an incorrect rule for English, so additional analyses will be needed of her input, to attempt to reconstruct her mis-parsing (i.e., such that she mistook a noun frame for a verb frame).

Future analyses will investigate A's continued use of this frame, and if, when, and/or how she eventually discards it. Moreover, the speech of the other two children, both TD, will be scrutinized for this and/or similar novel frames of production.

It is very important to point out that, without the highly 'dense' sampling of data from the SR, we might only have observed 1-2 of the 'I am a V' utterances, leading to the hypothesis that these were idiosyncratic rather than productive errors. The sheer volume of speech in A's corpus has thus allowed us to establish her predominantly correct usage of verb tense and agreement, along with the intriguing errors that demonstrate her use and creation of rules.

c. Significance:

Autism and related disabilities are severe disorders of language development. This study has enabled the development and piloting of the most valid assessments of children's levels of speech production thus far, by recording daily samples in the home environment over a period of 2-6 months. As more data are analyzed, we will be able to ascertain each child's speech sophistication on a variety of lexical, grammatical, and pragmatic levels, capture the shape of his/her developmental change at these levels, and compare these with his/her levels of language comprehension. This project is in line with recent studies of young typically-developing children's motor (Adoph et al., 2008) and language (Naigles et al., 2009) development, which have revealed that such 'dense' sampling vastly increases sensitivity to the occurrence and non-occurrence of words/motor behaviors, thus rendering more accurately the patterns of development involved in their use. Moreover, as more data are analyzed we will be able to further ascertain the early predictors of later language outcome and individual variation in ASD, by providing additional extremely detailed measures of how early speech production changes (or not). These findings will thus continue to reveal when and where children with autism diverge from typically developing children, providing detailed and specific information for service providers and parents. Finally, the development of the much less expensive Speechome-Light will enable parents, researchers, and service providers to utilize the technology to assess many more children, expanding the positive consequences of this tool to many more people affected by autism.

d. Publications/Presentations:

Roy, D. (2011, July) A Study of Language Development in Context. Plenary presentation, *International Association for the Study of Child Language,* Montreal, Canada.

Naigles, L. (2011) Not sampling, getting it all. In E. Hoff (Ed.) *Research Methods in Child Language: A Practical Guide* (pp. 240-253). Oxford:Wiley-Blackwell.

e. Project-Generated Resources:

None at this time.