

Quantifying Narrative Ability in Autism Spectrum Disorder: A Computational Linguistic Analysis of Narrative Coherence

Molly Losh · Peter C. Gordon

© Springer Science+Business Media New York 2014

Abstract Autism is a neurodevelopmental disorder characterized by serious difficulties with the social use of language, along with impaired social functioning and ritualistic/repetitive behaviors (American Psychiatric Association in Diagnostic and statistical manual of mental disorders: DSM-5, 5th edn. American Psychiatric Association, Arlington, 2013). While substantial heterogeneity exists in symptom expression, impairments in language discourse skills, including narrative (or storytelling), are universally observed in autism (Tager-Flusberg et al. in Handbook on autism and pervasive developmental disorders, 3rd edn. Wiley, New York, pp 335–364, 2005). This study applied a computational linguistic tool, Latent Semantic Analysis (LSA), to objectively characterize narrative performance in high-functioning individuals with autism and typically-developing controls, across two different narrative contexts that differ in the interpersonal and cognitive demands placed on the narrator. Results indicated that high-functioning individuals with autism produced narratives comparable in semantic content to those produced by controls when narrating from a picture book, but produced narratives diminished in semantic quality in a more demanding narrative recall task. This pattern is similar to that detected from analyses of hand-coded picture book narratives in prior research, and extends findings to an additional narrative context that proves particularly challenging for individuals with autism. Results are discussed in terms of the utility of

LSA as a quantitative, objective, and efficient measure of narrative ability.

Keywords Autism · Endophenotype · Language · Narrative · Phenotype

Introduction

Impaired social communication is a core feature of Autism Spectrum Disorder (ASD; American Psychiatric Association 2013). While substantial heterogeneity exists in the language abilities of individuals with ASD, impairments in language discourse skills, including conversation and narrative are observed across the spectrum of ability in ASD (Tager-Flusberg 2005). Discourse is the facet of language most closely tied to interpersonal interaction, and therefore most related to the social impairments seen in ASD (Tager-Flusberg 2000). Observations of discourse impairment in ASD appear in the earliest clinical descriptions by Leo Kanner and Hans Asperger, who separately described among their patients highly formulaic language and lack of language used for social purposes, often despite “astounding vocabularies” and intact structural language (Asperger 1991; Kanner 1943). Following these early clinical accounts, systematic empirical investigations have documented widespread difficulties with discourse in ASD, including problems initiating and maintaining conversational interactions (Capps et al. 1998; Tager-Flusberg and Anderson 1991), trouble abiding by common discourse conventions (e.g., turn taking, register) (Paul et al. 2009; Reichow et al. 2008; Tager-Flusberg and Caronna 2007; Volden et al. 2009), and a lack of inclination or ability to construe and convey experiences through narrative (Bruner and Feldman 1993; Loveland et al. 1990; Loveland and Tunali 1993).

M. Losh (✉)
Roxelyn and Richard Pepper Department of Communication
Sciences and Disorders, Northwestern University, Evanston,
IL 60208, USA
e-mail: m-losh@northwestern.edu

P. C. Gordon
Department of Psychology, University of North Carolina,
Chapel Hill, NC, USA

Narrative discourse (or storytelling) is a ubiquitous form of communication, used to derive meaning from experiences and construct shared understanding of events (Bruner 1986, 1990). Prior research has revealed a host of narrative difficulties in ASD. In semi-structured conversational interactions, individuals with ASD appear disinclined to narrate (Capps et al. 1998; Losh and Capps 2003), and when prompted tend to produce stories lacking coherence (Capps et al. 2000; Diehl et al. 2006; Gallese and Goldman 1999; Losh and Capps 2003, 2006; Loveland and Tunali 1993; Tager-Flusberg 2000). Relative to controls, individuals with ASD also display trouble enriching stories with psychological and social significance through the use of narrative evaluation (e.g., causal language and explanations of protagonists' thoughts and feelings) (Capps et al. 1998, 2000; Diehl et al. 2006; Losh and Capps 2003; Tager-Flusberg and Sullivan 1995).

These prior studies have together provided highly descriptive, clinically valid characterizations of the narrative difficulties encountered by individuals with ASD across language contexts. Yet the detailed hand-coding methods employed have some well-acknowledged drawbacks: hand coding can be labor intensive and subjective; achieving reliability across coders requires extensive training which may be difficult to implement across different research sites; finally, protocols for coding often must be adapted for use with different populations, tasks, and dialects. And while measures of narrative ability would appear to be useful measures for tracking response to language and communication intervention, the time, training, and labor involved in narrative coding makes it difficult to implement such a metric in treatment studies, particularly in more open-ended contexts that are most sensitive to the narrative impairments in ASD.

This study aimed to build on prior research on narrative in ASD by evaluating the utility of tools from computational linguistics for quantifying the narrative abilities of high functioning individuals with ASD across different discourse contexts. Techniques from computational linguistics are now being applied to a range of clinically-relevant tasks, such as assessing the quality of interactions of couples (Black et al. 2010; Simmons et al. 2005), language-function by individuals with schizophrenia (Elvevag et al. 2007) and classification of emotions (Metallinou et al. 2010). These techniques have the potential to complement and perhaps replace the detailed hand-coding methods employed in prior studies.

Latent Semantic Analysis

This study investigates narrative in ASD by using Latent Semantic Analysis (LSA; Landauer and Dumais 1997) as a quantitative tool to compare the meanings of narratives

produced by ASD and typically-developing individuals. The *vector-space* techniques behind LSA were originally developed to tackle problems of information retrieval in computerized databases (for reviews see Jurgens and Stevens 2010; Turney and Pantel 2010), but through LSA they have been applied successfully to problems in psychology. LSA has been used to model vocabulary acquisition (Landauer and Dumais 1997), comprehension of metaphor (e.g., Kintsch 2002) and textual coherence (Foltz et al. 1998; Ledoux et al. 2006). It has also been used in educational applications for text selection and essay scoring (Franzke et al. 2005). Of greater relevance to the present study, LSA has been used successfully to quantify anomalous patterns of the language produced by individuals with schizophrenia with thought disorder, showing that it can successfully capture the atypical nature of these individuals' responses in a word association task as well as their incoherence in the telling of narrative stories (Elvevag et al. 2007). The successful application of LSA to schizophrenia suggests that it can be used as an objective, consistent measure of the language-related symptoms of a complex psychiatric disorder, which like in ASD can be characterized by off-topic, idiosyncratic language.

Whereas the general rationale behind LSA is straightforward, the manner in which it operates is relatively complex. Generally speaking, LSA determines the meaning of a word through its patterns of statistical co-occurrence with other words, and in this way it captures Firth's dictum that we can "... know a word by the company it keeps" (Firth 1957, p. 11). More specifically, LSA characterizes the meaning of a word as a point in a vector space. Words with similar meanings are close together in this space while words with dissimilar meanings are far apart. Figure 1 gives a schematic outline of the general steps involved in creating a vector space model and also notes some important specifics about LSA as it was used in this study. Creation of a vector space begins with selection of a corpus for training. For present purposes this corpus was a selection of general readings through college level in English. General readings through lower grade levels are also available as are specialized readings on topics such as psychology and biology, with additional corpora available in French. The selected corpus is then processed so that it is suitable for analysis; that involves steps such as eliminating very high frequency words and eliminating affixes so that only the stem (root word) remains. Then a unit of "context" is chosen, in this case the paragraph, allowing a word-by-context distribution to be obtained as a critical step toward characterizing the meaning of a word by the kinds of contexts it appears in. The next step is feature extraction, which in LSA consists of word co-occurrence within contexts (i.e., paragraphs in this case). Thus, two words have similar meanings if they appear in similar




Fig. 1 A schematic of the steps involved in creating a vector-space model of meaning

contexts (i.e., collections of similar words). Finally, there is a global operation across the collections of word-in-context co-occurrences. In LSA singular value decomposition, a method that is similar to factor analysis, is used to determine dimensions that capture the systematic variation in word meaning. The resulting vector space for the (default) version of LSA that we used has 300 dimensions.

The dimensions that make up the space cannot be understood directly in the sense that they cannot be mapped onto intuitively-clear, lexical-semantic features such as those employed in hand-coding schemes (e.g., causal connectors, emotion terms). However, the LSA high-dimensional space captures important aspects of the meaning of words because words with similar meanings are closer within the space than are words with dissimilar meanings, with the distance of words within the LSA space strongly related to human judgments about similarity in meanings of words (Landauer and Dumais 1997; Landauer et al. 2007).

Once the LSA space has been created from the source corpus it can be used to evaluate the semantic similarity of different types of language samples ranging from individual words, to sentences, to complete narratives. In LSA the meaning of a narrative is represented as a point in the vector space with its position determined by the positions of its component words. Thus, the similarity in meaning between pairs of narratives is given by their proximity in this high-dimensional space just as it is for pairs of words. It is important to note that LSA is incapable of representing any meaning that depends on the order of the words in a narrative and that its reliance on co-occurrence statistics means that very high frequency words (such as articles, prepositions and pronouns) do not affect LSA's determination of the meaning of a narrative because such very high frequency words co-occur with nearly all other types of words. While these characteristics limit LSA's ability to fully capture all meanings, the technique remains useful as a quantitative measure of the meaning of content words.

Here we provide a first evaluation of the usefulness of LSA as an automatic tool for assessing the quality of narratives produced by individuals with ASD. LSA was used to test whether the semantic consistency between narratives of controls and those from individuals with high-functioning ASD varied with the demands of the

communicative situation, using tasks that differed in cognitive and interpersonal demands placed on the narrator (i.e., picture book narratives with illustrative stimuli to support narrative production, and narrative recall tasks). Success in characterizing narrative performance in terms of semantic consistency across narratives produced by different individuals would allow LSA to be used both when narratives can be compared to a "true standard", as when individuals attempt to retell a spoken story (narrative recall tasks), and when there is no true standard, as when individuals respond to a prompt consisting of a series of pictures but no words (picture book task). Successful implementation of LSA in this study would support the potential of LSA as a quantitative assessment of narrative performance across a range of communicative environments, that could complement more comprehensive but labor intensive hand coding methods. Here, we applied LSA to narrative samples obtained from a narrative recall task not previously examined, and a wordless picture book context previously studied with this sample. Analyses addressed group differences in narrative ability indexed by LSA across these different discourse contexts, and also examined associations between LSA and key hand coded measures of narrative coherence previously conducted.

Methods

Participants

The subjects included 22 high-functioning individuals with ASD (referred to as the ASD group) and 26 typically-developing controls who participated in the narrative study conducted by Losh and Capps (2003). All children demonstrated IQs within the normal range and groups were matched on verbal IQ using the WISC-III (Wechsler 2003), and chronological age. Children ranged in age between 8 and 14 years. Children with ASD were evaluated using gold standard diagnostic instruments, the Autism Diagnostic Interview, Revised (Lord et al. 1994) and Autism Diagnostic Observational Schedule (Lord et al. 2000), and diagnosed according to DSM-IV criteria (American Psychiatric Association 1994). Based on ADOS and ADI algorithm scores, all children also met current DSM-5

criteria for Autism Spectrum Disorder (American Psychiatric Association 2013). See Losh and Capps (2003) for further details on participant characteristics.

Tasks

Narrative performance was compared in two types of tasks: narrative recall and spontaneous narrative production from a picture book.

Narrative Recall

Narrative recall was assessed using two oral fairy tale texts used in prior research to investigate the narrative abilities of children with intellectual disabilities (Dennis et al. 1994): “The Frog Prince” (Retelling #1, Opie and Opie 1974) and “The Practical Princess” (Retelling #2, Williams 1978). Text of the stories is provided in “Appendix”. The texts were read to children by the experimenter, with accompanying demonstration with puppets depicting the main characters. After hearing each story, children were asked to tell the story to the experimenter.

Picture Book Narration

The 24-page wordless picture book, *Frog, Where Are You?* (Mayer 1969) was used to elicit spontaneous narratives from subjects. The story is about a boy and his adventures searching for his missing pet frog. This storybook has been used in many prior studies of narrative discourse in both typical and atypical development [including ASD (Capps et al. 2000; Losh and Capps 2003; Tager-Flusberg and Sullivan 1995)]. Children were asked to narrate the story as they looked at the pictures.

Prompting during all tasks was kept to a minimum and included only prompts for elaboration or clarification when children paused excessively or made confusing or incoherent remarks. Tasks were presented in random order and were video- and audio-taped for verbatim transcription. Transcribers were blind to group status and were trained to greater than 80 % reliability. Fifteen percent of transcripts were also assessed for reliability, with greater than 80 % agreement for point-to-point word agreement and utterance boundary marking. Any disagreements were resolved by a third, senior transcriber.

Analysis Techniques

The semantic similarity of the narratives was measured by entering them into the LSA website maintained at the University of Colorado (<http://lsa.colorado.edu>) a process that simply involves cutting and pasting the transcribed narratives. Similarity was measured using the default

semantic space derived by LSA from a large sample of texts meant to approximate reading experience through college. This default semantic space was chosen because its range of training texts means that it is likely to capture semantic distinctions that would be missed by simpler or more specialized selections of text. In addition, while the current study compares language samples produced by children with ASD or typical development, we aimed to use a semantic space that in the future could be used to analyze adult narratives as well. Using LSA “pairwise similarity” option, similarity between narratives was measured separately for each of the three narrative productions: Retelling 1 (*Frog Prince*), Retelling 2 (*Practical Princess*) and picture book narration (*The Frog Story*), using LSA’s matrix comparison function. For the two retellings, this involved measuring the similarity of every participant narrative with each of the other participants’ narratives and with the text of the story as read prior to the retelling. For the wordless picture book there was no text version and so this simply involved measuring the similarity of participants’ narratives against one another.

Results

Narrative Similarity Between Individuals and to a Standard Text in the Narrative Recall Tasks (Retellings 1 and 2)

The similarity between pairs of narratives, and between each narrative and the story text for the retellings, can be used to determine the degree to which those narratives are centered on a common meaning defined by overlap in semantic content. In performing this analysis, LSA is given no information about which of the text samples is the original story, but it recovers that information through the overall pattern of similarity of the retellings and the original. This can be seen by analyzing the similarity matrix produced by LSA of retellings and the original text using multi-dimensional scaling (MDS) where the LSA similarities are treated as proximities.

The left panel of Fig. 2 shows the results for Retelling 1 (*Frog Prince*) and the right panel shows results for Retelling 2 (*Practical Princess*). The number of dimensions in the MDS solution was fixed at two for ease of interpretation; results are not dependent on the number of dimensions used or their interpretation. The original story (indicated for each story by the red square) falls in the center of the proximity space, the point at which the summed distance to all the other stories is minimized. Further, there is a cluster of narrative retellings (primarily produced by typically-developing children) that is very close to the original story in the MDS proximity space.

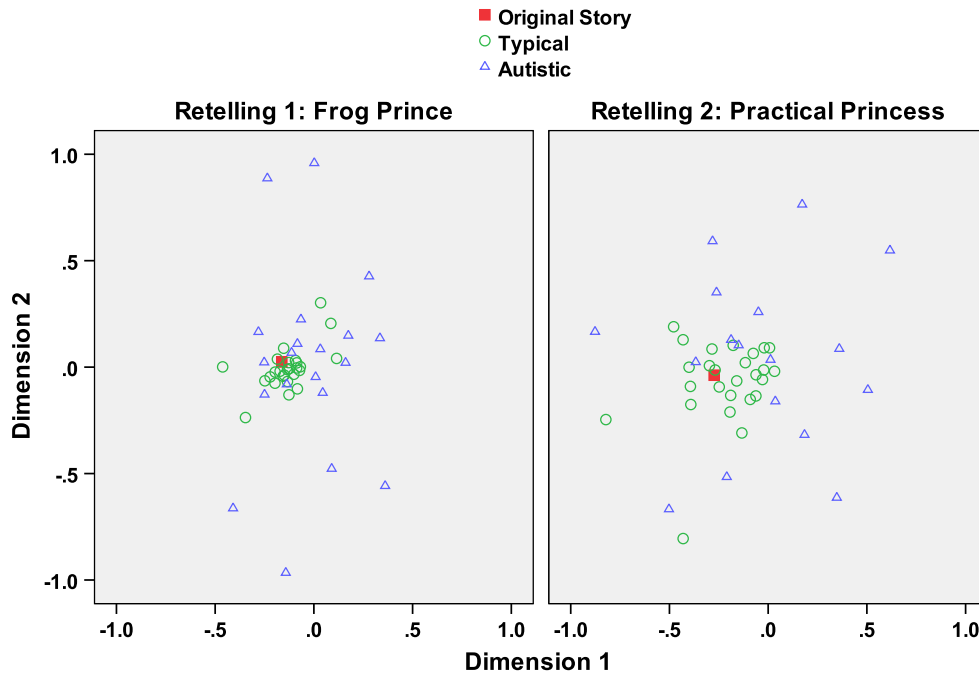


Fig. 2 Multidimensional scaling solution with two dimensions for LSA similarities of narrative retellings to each other and to the original story. The *left panel* shows the retelling of Story 1 (*Frog Prince*) and the *right panel* shows the retelling of Story 2 (*Practical Princess*). The critical results are: (1) that the text of the original story (denoted by the *red square*) is in the center of the space, (2) most narrative retellings from controls (denoted by *green circles*) are also close to the center of the space, and (3) retellings from the autistic participants (denoted by *purple triangles*) are more widely dispersed.

The ranges for the similarity in these graphs were set between -1 and 1 so that the central patterns would be visible. However, these limits meant that results for some narratives that were highly dissimilar to the others are not displayed. This occurred for one narrative in Retelling 1 and four in Retelling 2. In all cases these extreme values were due to narratives produced by autistic individuals. Thus, the variability seen in narratives by autistic as compared to typically-developing children is even greater than what is shown in the graphs (Color figure online)

Retellings in this cluster have the greatest proximity to both the original story and to the other retellings, indicating that they have the greatest overlap in semantic content with both the original story and the retellings by other children. The substantial dispersion in the similarity space for narrative retellings by ASD children indicates that their narrative retellings differ from the original story and from the retellings of the typically-developing children. In addition, the retellings that have the highest overall similarity to other retellings can be treated as a standardized definition of the meaning of the story as defined by having the greatest overlap in meaning with the other narrative retellings. Defining a standard using the similarity of narratives becomes important in cases where no objective standard for a narrative is available, such as when subjects were asked to narrate from the wordless picture book.

Narrative Performance Across Retelling and Picture Book Tasks

For the retelling tasks, the semantic accuracy of a participant’s narrative can be scored as its similarity to the original text. The left panel of Fig. 3 shows this measure of

narrative accuracy for each group. For both retelling tasks, similarity to the text standard was higher for the typically-developing group than for the ASD group (*Frog Prince*, $t(46) = 4.2, p < .001$, Cohen’s $d = 1.24$; *Practical Princess*, $t(46) = 5.5, p < .001$, Cohen’s $d = 1.62$). These results show that LSA appears sensitive to differences between ASD and typically-developing children in the accuracy with which they can retell a story that they have just heard in a communicative situation. Further, the plots for these data (left panel of Fig. 3) show that there is a great deal more variability in the narratives produced by the ASD group than by the control group, with somewhat more variability observed in the second, longer recall task (*Practical Princess*).

Analyses comparing individuals’ narratives to the original text are not possible for the picture book narration (*Frog Story*) because there is no original text. However, the results shown in Fig. 2 offer a path to performing directly-comparable analyses on the two retellings and the picture book narration by showing that the narratives that were most similar to all other narratives were also highly similar to the original text. Accordingly, the ASD group’s narrative performance was evaluated in the following way for




Fig. 3 Boxplots of the similarity of participants' narratives in the two retelling tasks and in the picture book narration. The *edges of the boxes* represent the 25th and 75th percentiles with the *center line* representing the median. The whiskers represent boundaries that are one and one half times the estimated inter-quartile range from the edges of the *box*. Observed values beyond the whiskers are plotted individually. The *left panel* shows similarity to a standard consisting of the text of the story that the participant heard and then retold. The *right panel* shows similarity to a standard consisting of the narrative when compared to the average of the four narratives produced by typically-developing children that were most similar to the other narratives produced by all typically-developing children. The individual narratives used to create this average standard were not included in boxplots

both retellings and the picture book narration. For each story, the four narratives from control subjects that were most similar to the other control narratives were selected and their LSA scores averaged to create an empirically derived standard that reflected the center of shared meaning across the different individual narratives. Narratives from children whose stories were not used for the averaged standard are shown in the right panel of Fig. 3. Just as when the original text was used as the standard, there were highly significant group differences in LSA similarity for both retellings (*Frog Prince*: $t(41) = 3.5$, $p = .001$, Cohen's $d = 1.08$; *Practical Princess*: $t(42) = 4.2$, $p = .001$, Cohen's $d = 1.30$) with the ASD group diverging from standards significantly more than controls for each recalled story. However, for the Picture Book Narration (*Frog Story*) the ASD group did not differ significantly from controls in their difference from the average standard ($t(33) = 0.56$, $p > .5$, Cohen's $d = 0.19$).

Correlations Between Hand Coding and LSA

To explore the validity of LSA we examined correlations between semantic similarity scores across conditions and several key hand coding variables indicative of narrative quality, conducted previously on transcripts of the picture book context as well as personal narratives elicited through conversation (Losh and Capps 2003): (1) the frequency of experimenter prompting (an index of children's reliance on the experimenter to scaffold the narrative and clarify confusing information presented); (2) the use of evaluative devices (e.g., explanations of characters' internal states, causal language to bind narrated events, etc.), and (3) the use of complex syntax, including adverbial, coordinate, and relative clauses as well as verb complements to combine sentential elements. Each of these factors was related to LSA narrative scores across groups. More extensive experimenter prompting was associated with lower semantic similarity (i.e., lower quality narrative skill as measured by LSA), $r = -.48$, $p < .005$. More sophisticated use of complex syntax (indexed by the range of different types of complex syntactic devices deployed) and narrative evaluation were also associated with higher (i.e., better) semantic similarity scores across narratives ($r = .43$, $p < .005$, $r = .25$, $p = .097$, respectively).

Discussion

Important insights into narrative discourse abilities in ASD have come from prior studies employing detailed hand-coding schemes (Capps et al. 2000; Diehl et al. 2006; Losh and Capps 2003; Loveland et al. 1990; Tager-Flusberg and Sullivan 1995). In the current study we attempted to build on this prior work by applying a computational linguistic tool, LSA, to objectively and efficiently characterize narrative performance among high-functioning individuals with ASD and typically-developing controls, across two different narrative tasks. In an earlier study we found that this high-functioning ASD group showed striking differences from controls on a number of narrative discourse measures (frequency and diversity of complex syntax, thematic coherence, explanations of protagonists' internal states) during semi-structured conversation, yet did not show many of these differences when narrating from a wordless picture book (Losh and Capps 2003). Here, using LSA, we confirmed the ASD group's relatively intact performance in the picture book task, and identified narrative recall as an additional context that proves particularly challenging for children with ASD. That is, the ASD group's semantic similarity scores were comparable to controls' in the picture book task, but in both narrative recall tasks they diverged significantly in mean semantic

similarity scores, reflecting significantly diminished narrative quality relative to controls.

Good correspondence was found between LSA semantic similarity scores and previously conducted hand coding variables from Losh and Capps (2003), where lower LSA scores across narrative contexts were associated with several indices of narrative impairment—heavy reliance on experimenter prompting, minimal use of narrative evaluation, and restricted use of complex syntax. In line with these findings, inspection of children’s recalled narratives revealed that without the aid of the picture book, the ASD group produced narratives with an abundance of off-topic and irrelevant remarks, departing considerably from the stories’ themes, and often resulting in incoherent stories. As testament to the challenges posed by narrative recall for the ASD group, and their apparent reliance on illustrations to facilitate narrative activity, many children displayed evidence of distraction and frustration in the narrative recall tasks (e.g., one child exclaimed, “I need a book ... I told you, I need a book”). Thus, while basic narrative skills may be evident under certain circumstances, the critical ability to recall and recount narratives in social interaction remains a significant challenge in ASD.

What might account for the profile observed here and in other reports (Capps et al. 1998; Losh and Capps 2003; Goldman 2008), in which children with ASD demonstrated notable difficulty producing coherent narratives when not provided with illustrated stimuli? This difference could be due to the attentional and interpersonal differences between narrating a story from a picture book and other discourse tasks such as retelling a story to an interviewer or engaging in a semi-structured conversation. While picture book narration still requires an extensive, coherent linguistic response, the presence of pictures and the sequence of pages serve to control attention and to organize language production in a way that retelling, which depends on memory and involves no overt cues, does not. Picture books may in fact capitalize on relatively strong visual spatial skills common in ASD, and an inclination towards visual stimuli noted among many individuals with ASD (Grandin 1995). This is consistent with prior reports that individuals with ASD often rely on visually salient features when discussing thoughts and emotions in narrative (Capps et al. 2000; Losh and Capps 2006). Further, a picture book provides an appropriate target for looking, whereas retelling a story creates more of an interpersonal demand to attend to and engage with an interlocutor. Of note, the more taxing demands of memory and interpersonal interaction entering into play when narrating based on recall of a story enacted by another, as in the fairy tale recall tasks, did not appear to be challenging for typically developing children, whose semantic similarity scores in the recall tasks were comparable or better than those for the picture book.

Findings together provide demonstration that LSA can be used to objectively and sensitively capture narrative performance differences between tasks and groups, and thus may prove to be an effective analytic tool for measuring narrative performance in ASD. In characterizing narrative coherence as overlap in meaning shared across different speakers, results also show how LSA can measure narrative coherence both in cases where an objective standard exists and in cases where that standard is induced from a collection of different narratives. This is important given that, as noted above, narrative recall and other less structured tasks appear to be more sensitive to the problems that children with ASD have with narration than wordless picture book tasks.

Nonetheless, practitioners may be hesitant to use unstructured tasks in assessment given inherent challenges employing controlled scoring when more variability is present. Because LSA provides an empirically derived standard of narrative quality, it may offer a way to implement more varied narrative tasks in clinical assessments and in measuring responses to intervention. Before this can be done, however, additional research will be needed to evaluate and refine the implementation of LSA across different ages and ability levels to establish developmental norms for different types of narratives (e.g., picture book, conversational, recall, etc.) as a flexible and accessible tool. Indeed, a major limitation of LSA in its current form is that it cannot be readily applied to discourse contexts such as conversational narrative, where topics may vary substantially across individuals. Because conversational interaction has been demonstrated to be among the most challenging discourse contexts for individuals with ASD, a goal for future work should be to continue to refine computational tools for use in less structured contexts that are most challenging for children with ASD, and that pervade daily social interactions. The increasing availability of openly accessible shared databases including language samples from both typically developing children and adults [e.g., The TalkBank Project (MacWhinney 2007)] could provide valuable data for the refining LSA for such purposes.

In sum, whereas additional study and refinement of computational linguistic tools is needed, LSA appears to show promise as an efficient method for quantifying narrative discourse performance that avoids many of the pitfalls of hand-coding systems. In the future, its efficiency may be increased substantially if transcripts of recorded speech can be generated using automatic speech recognition techniques, as is now being done with recordings from other clinically significant domains (Georgiou et al. 2011). As an analytic method LSA also provides a quantitative metric of narrative performance, and is therefore well suited for genetic and neurobiological research, where the

use of such quantitative phenotypic measures provides more statistical power to detect gene-brain-behavior relationships than categorically defined clinical constructs (Almasy and Blangero 2001; Gottesman and Gould 2003; Sykes and Lamb 2007). LSA also holds potential for application in large-scale studies, where time and labor intensive hand-coding methods are not feasible. Finally, LSA could serve as an important tool for measuring response to intervention, providing a reliable quantitative metric of language coherence that could be applied and tracked over time. By demonstrating the utility of an experimentally derived computational linguistic measure of narrative discourse performance, findings together suggest that LSA may be an important tool for future studies of ASD and other neurogenetic developmental disorders of language.

Acknowledgments We wish to acknowledge Patrick Plummer for his assistance with the preparation of transcripts for computational linguistic analysis. This work was supported by the National Science Foundation (BCS-0820394) and the National Institute of Deafness and other communication Disorders (1R01DC010191-01A1).

Appendix

Frog Prince Story (Opie and Opie 1974)

Once upon a time there was a princess. She went one day to play at the pond. She took her golden ball with her, which was her favorite toy. She played with the ball at the pond, until-oops-the ball fell in the water. She began to cry, and said, "Oh dear, I would give anything to have my lovely ball back again!" Suddenly a frog stuck his head out of the water and said, "Why are you crying?" The princess said, "What use are you, you nasty frog? My ball has fallen into the pond." The frog said, "I will get your ball for you, if you will be my friend and let me live with you and eat from your plate and sleep on your bed." "What nonsense," the princess thought, "He will never be able to get out of the pond, but I will promise what he asks because perhaps he can get my ball back." So she said to the frog, "If you will get my ball for me I will do all the things you ask." So the frog dived into the water, picked up the ball with his mouth, and threw it on the ground in front of the princess. When she saw it, she was so happy that she ran off home without thinking of the kind frog. Next day the princess heard a knock just as she was having her dinner. She opened the door and there was the frog! The princess did not like the frog but she remembered her promise. The frog said, "Let me come in and sit near you, and eat your dinner with you." So the frog ate dinner with the princess. After they had eaten, the frog said, "Carry me upstairs and put me on your bed." And the princess did what he asked, even

though she still thought he was a wet and nasty frog. She secretly hoped that he would be gone when she woke up. But the next morning she was amazed to find that the frog was gone. Instead there was the most handsome prince she had ever seen. He told her that a bad fairy had changed him into a frog until he could find a princess who would be his friend and let him eat with her and sleep on her bed. And the prince and the princess were married and lived happily ever after.

Practical Princess Story (Williams 1978)

Princess Bedelia was as lovely as the moon shining upon a lake full of water lilies. She was as graceful as a cat leaping. And she was also extremely practical. When she was born, three fairies had come to her cradle to give her gifts. The first fairy had given her beauty. The second had given her grace. But the third, who was a wise old creature, had said, "I give her common sense." "I don't think much of that gift," said King Ludwig, raising his eyebrows, "What good is common sense to a princess? All she needs is charm." Nevertheless, when Bedelia was eighteen, something happened which made the king change his mind. A dragon moved into the neighborhood. He settled in a dark cave on top of a mountain. He sent a message to the king. "I must have a princess to devour," the message said, "or I shall breathe out my fiery breath and destroy the kingdom." Sadly, King Ludwig called together his councilors and read them the message. "Perhaps," said the Prime Minister, "we had better advertise for a knight to slay the dragon. That is what is generally done in these cases." "I'm afraid we haven't time," answered the king. "The dragon has only given us until tomorrow morning. There is no help for it. We shall have to send him the princess." Princess Bedelia had come to the meeting because, as she said, she liked to mind her own business and this was certainly her business. "Rubbish!" she said, "Dragons can't tell the difference between princesses and anyone else. Use your common sense. He's just asking for me because he's a snob." "That may be so," said her father, "but if we don't send you along, he'll destroy the kingdom." Bedelia said, "I see I'll have to deal with this myself!" She left the council chamber. She got the largest and gaudiest of her state robes and stuffed it with straw, and tied it together with string. She packed 100 pounds of gunpowder into the center of the bundle. She got a strong young man to carry it up the mountain for her. She stood in front of the dragon's cave and called. "Come out! Here's the princess!" The dragon came peering out of the darkness. Seeing the bright robe covered with gold and silver embroidery, and hearing Bedelia's voice, he opened his mouth wide. When Bedelia gave the signal, the young man heaved the robe down the dragon's throat. Bedelia threw

herself flat on the ground, and the young man ran away. As the gunpowder met the flames inside the dragon, there was a tremendous explosion. Bedelia got up, dusting herself off. “Dragons,” she said, “are not very bright.” She left the young man sweeping up the pieces, and she went back to the castle to have her geography lesson.

References

- Almasy, L., & Blangero, J. (2001). Endophenotypes as quantitative risk factors for psychiatric disease: Rationale and study design. *American Journal of Medical Genetics*, *105*, 42–44.
- American Psychiatric Association. (1994). *Diagnostic and Statistical Manual of Mental Disorders (DSM-IV)* (4th ed.). Washington, DC: American Psychiatric Association.
- American Psychiatric Association. (2013). *Diagnostic and statistical manual of mental disorders: DSM-5* (5th ed.). Arlington, VA: American Psychiatric Association.
- Asperger, H. (1991). ‘Autistic psychopathy’ in childhood (U. Frith, Trans.). In U. Frith (Ed.), *Autism and Asperger syndrome*. London: Cambridge University Press.
- Black, M., Katsamanis, A., Lee, C.-C., Lammert, A. C., Baucom, B. R., Christensen, A., et al. (2010). Automatic classification of married couples’ behavior using audio features. In *11th annual conference of the International Speech Communication Association 2010 (Interspeech 2010)* (Vols. 3, 4, pp. 2030–2033).
- Bruner, J. (1986). *Actual minds, possible worlds*. Cambridge, MA: Harvard University Press.
- Bruner, J. (1990). *Acts of meaning*. Cambridge: Harvard University Press.
- Bruner, J., & Feldman, C. (1993). Theories of mind and the problem of autism. In S. Baron-Cohen, H. Tager-Flusberg, & D. J. Cohen (Eds.), *Understanding other minds: Perspectives from autism*. New York: Oxford University Press.
- Capps, L., Kehres, J., & Sigman, M. (1998). Conversational abilities among children with autism and developmental delay. *Autism*, *2*, 325–344.
- Capps, L., Losh, M., & Thurber, C. (2000). The frog ate a bug and made his mouth sad: Narrative competence in children with autism. *Journal of Abnormal Child Psychology*, *28*, 193–204.
- Dennis, M., Jacennik, B., & Barnes, M. A. (1994). The content of narrative discourse in children and adolescents after early-onset hydrocephalus and in normally developing age peers. *Brain and Language*, *46*(1), 129–165. doi:10.1006/brln.1994.1008.
- Diehl, J. J., Bennetto, L., & Young, E. C. (2006). Story recall and narrative coherence of high-functioning children with autism spectrum disorders. *Journal of Abnormal Child Psychology*, *34*(1), 83–98.
- Elvevag, B., Foltz, P. W., Weinberger, D. R., & Goldberg, T. E. (2007). Quantifying incoherence in speech: An automated methodology and novel application to schizophrenia. *Schizophrenia Research*, *93*(1–3), 304–316. doi:10.1016/j.schres.2007.03.001.
- Firth, J. R. (1957). *Papers in linguistics*. Oxford: Oxford University Press.
- Foltz, P. W., Kintsch, W., & Landauer, T. K. (1998). The measurement of textual coherence with latent semantic analysis. *Discourse Processes*, *25*, 285–307.
- Franzke, M., Kintsch, E., Caccamisa, D., Johnson, N., & Dooley, S. (2005). Summary Street[®]: Computer support for comprehension and writing. *Journal of Educational Computing Research*, *33*(1), 304–316.
- Gallese, V., & Goldman, A. (1999). Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Science*, *2*, 493–500.
- Georgiou, P. G., Black, M. P., & Narayanan, S. S. (2011). *Behavioral signal processing for understanding (distressed) dyadic interactions: Some recent developments*. Paper presented at the third international workshop on social signal processing (SSPW’11), ACM Multimedia’11, Scottsdale, AZ.
- Goldman, S. (2008). Brief report: Narratives of personal events in children with autism and developmental language disorders: Unshared memories. *Journal of Autism and Developmental Disorders*, *38*(10), 1982–1988. doi:10.1007/s10803-008-0588-0.
- Gottesman, I. I., & Gould, T. D. (2003). The endophenotype concept in psychiatry: Etymology and strategic intentions. *American Journal of Psychiatry*, *160*(4), 636–645.
- Grandin, T. (1995). *Thinking in pictures*. New York: Doubleday.
- Jurgens, D., & Stevens, K. (2010). *The S-Space package: An open source package for word space models*. Paper presented at the proceedings of the ACL 2010 systems demonstrations, Uppsala, Sweden.
- Kanner, L. (1943). Autistic disturbances of affective contact. *The Nervous Child*, *2*, 217–250.
- Kintsch, W. (2002). The potential of latent semantic analysis for machine grading of clinical case summaries. *Journal Biomedical Informatics*, *35*(1), 3–7.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*(2), 211–240.
- Landauer, T. K., McNamara, D. S., Dennis, S., & Kintsch, W. (Eds.). (2007). *Handbook of latent semantic analysis*. Mahwah, NJ: Lawrence Erlbaum.
- Ledoux, K., Camblin, C. C., Swaab, T. Y., & Gordon, P. C. (2006). Reading words in discourse: The modulation of lexical priming effects by message-level context. *Behavioral and Cognitive Neuroscience Reviews*, *5*, 107–127.
- Lord, C., Risi, S., Lambrecht, L., Cook, E. H., Leventhal, B. L., DiLavore, P., et al. (2000). The Autism Diagnostic Observation Schedule-Generic: A standard measure of social and communication deficits associated with the spectrum of autism. *Journal of Autism and Developmental Disorders*, *30*, 205–223.
- Lord, C., Rutter, M., & Le Couteur, A. (1994). Autism Diagnostic Interview-Revised: A revised version of a diagnostic interview for caregivers of individuals with possible pervasive developmental disorders. *Journal of Autism and Developmental Disorders*, *24*(5), 659–685.
- Losh, M., & Capps, L. (2003). Narrative ability in high-functioning children with autism or Asperger’s syndrome. *Journal of Autism and Developmental Disorders*, *33*(3), 239–251.
- Losh, M., & Capps, L. (2006). Understanding of emotional experience in autism: Insights from the personal accounts of high-functioning children with autism. *Developmental Psychology*, *42*(5), 809–818.
- Loveland, K. A., McEvoy, R. E., & Tunali, B. (1990). Narrative story telling in autism and Down’s syndrome. *British Journal of Developmental Psychology*, *8*, 9–23.
- Loveland, K. A., & Tunali, B. (1993). Narrative language in autism and the theory of mind hypothesis: A wider perspective. In S. Baron-Cohen, H. Tager-Flusberg, & D. J. Cohen (Eds.), *Understanding other minds: Perspectives from autism*. Oxford: Oxford University Press.
- MacWhinney, B. (2007). The TalkBank Project. In J. C. Beal, K. P. Corrigan, & H. L. Moisl (Eds.), *Creating and digitizing language corpora: Synchronic databases* (Vol. 1). Houndmills: Palgrave-Macmillan.
- Mayer, M. (1969). *Frog, where are you?*. New York: Dial Press.

- Metallinou, A., Lee, S., Narayanan, S., & IEEE. (2010, March 14–19). *Decision level combination of multiple modalities for recognition and analysis of emotional expression*. Paper presented at the 2010 IEEE international conference on acoustics, speech, and signal processing, Dallas, TX.
- Opie, I. A., & Opie, P. (1974). *The classic fairy tales*. London: Oxford University Press.
- Paul, R., Orlovski, S. M., Marcinko, H. C., & Volkmar, F. (2009). Conversational behaviors in youth with high-functioning ASD and asperger syndrome. *Journal of Autism and Developmental Disorders*, 39, 115–125.
- Reichow, B., Salamack, S., Paul, R., Volkmar, F. R., & Klin, A. (2008). Pragmatic assessment in Autism Spectrum Disorders: A comparison of a standard measure with parent report. *Communication Disorders Quarterly*, 29(3), 169–176. doi:10.1177/1525740108318697.
- Simmons, R., Gordon, P. C., & Chambless, D. (2005). Pronoun use in marital interaction: What do “you” and “I” say about marital health? *Psychological Science*, 16, 932–936.
- Sykes, N. H., & Lamb, J. A. (2007). Autism: The quest for the genes. *Expert Reviews in Molecular Medicine*, 9(24), 1–15.
- Tager-Flusberg, H. (2000). Understanding the language and communicative impairments in autism. *International Review of Research in Mental Retardation*, 23, 185–205.
- Tager-Flusberg, H. (2005). Designing studies to investigate the relationships between genes, environments, and developmental language disorders. *Applied Psycholinguistics*, 26, 29–39.
- Tager-Flusberg, H., & Anderson, M. (1991). The development of contingent discourse ability in autistic children. *Journal of Child Psychology and Psychiatry*, 32, 1123–1134.
- Tager-Flusberg, H., & Caronna, E. (2007). Language disorders: Autism and other pervasive developmental disorders. *Pediatric Clinics of North America*, 54(3), 469–481. doi:10.1016/j.pcl.2007.02.011.
- Tager-Flusberg, H., Paul, R., & Lord, C. (2005). Language and communication in autism. In F. Volkmar, R. Paul, & A. Klin (Eds.), *Handbook on autism and pervasive developmental disorders* (3rd ed., pp. 335–364). New York: Wiley.
- Tager-Flusberg, H., & Sullivan, K. (1995). Attributing mental states to story characters: A comparison of narratives produced by autistic and mentally retarded individuals. *Applied Psycholinguistics*, 16(3), 241–256.
- Turney, P. D., & Pantel, P. (2010). From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research*, 37, 141–188.
- Volden, J., Coolican, J., Garon, N., White, J., & Bryson, S. (2009). Brief report: Pragmatic language in autism spectrum disorder: Relationships to measures of ability and disability. *Journal of Autism and Developmental Disorders*, 39(2), 388–393. doi:10.1007/s10803-008-0618-y.
- Wechsler, D. (2003). *Manual for the Wechsler Intelligence Scale for Children—Fourth edition*. San Antonio: The Psychological Corporation.
- Williams, J. (1978). *The practical princess and other liberating fairy tales*. New York: Parents’ Magazine Press.