Unraveling the Dyad: Using Recurrence Analysis to Explore Patterns of Syntactic Coordination Between Children and Caregivers in Conversation

Rick Dale and Michael J. Spivey
Cornell University

Recurrence analysis is introduced as a means to investigate syntactic coordination between child and caregiver. Three CHILDES (MacWhinney, 2000) corpora are analyzed and demonstrate coordination between children and their caregivers in terms of word-class *n*-gram sequences. Results further indicate that trade-offs in leading or following this coordination reflect individual differences between children at varying levels of development. Further analyses characterize the syntactic patterns that are coordinated, and results are consistent with recent language acquisition research on syntax acquisition. Overall, recurrence analysis reveals that there is a process of child-caregiver coordination taking place in ongoing conversation at the level of syntactic description.

Despite decades of research, there is no consensus on the appropriate characterization of the environment in which children acquire their first language. For example, there remains the perennial perspective on the irrelevance of the exact nature of the input, provided there simply is some (e.g., Chomsky, 1965; Newport & Gleitman, 2002; Pinker, 1994). Others approach child-directed speech as a complex but catered input stream

---

Rick Dale and Michael J. Spivey, Department of Psychology.

Correspondence concerning this article should be addressed to Rick Dale, Department of Psychology, Cornell University, Ithaca, NY, 14853. Internet: rad28@cornell.edu

tailored to the properties of language that children need to acquire (e.g., infant-directed speech; Broen, 1972; Moerk, 1992; Snow, 1972). Numerous researchers have additionally proposed that the input need not necessarily be specially tailored, but is far from irrelevant: Statistical properties of language input are crucial for its acquisition (e.g., Gomez & Gerken, 2000; Maratsos & Chalkley, 1980; Redington, Chater, & Finch, 1998; Saffran, 2003). The study of syntax acquisition has been particularly illustrative of this theoretical diversity. Proposals for the origin of grammatical knowledge range from innate knowledge, such as that of pronominal structures (Lidz, Waxman, & Freedman, 2003), to tailored input that underlies learning of syntactic patterns, such as when caregivers contrast correct forms with a child's error (e.g., Saxton, 1997). Some have proposed that children can abstract grammatical categories from extensive exposure to word distribution patterns (e.g., Mintz, 2003; Lewis & Elman, 2001; Reali & Christiansen, 2003).

Methodological diversity goes hand in hand with theoretical diversity. The last 20 years have seen an amazing upsurge of methods, from brain imaging to sophisticated preverbal behavioral tests. A particularly promising and now well-established methodology is the analysis of records of child-caregiver interaction. This methodological approach has cast light on the theoretical dimensions of syntax acquisition. For example, some early work sought to identify the social or structural cues for aiding the child's language development. Hirsh-Pasek, Treiman, and Schneiderman (1984) conducted an early analysis of interaction between 40 mother-child pairs demonstrating differential maternal responding contingent on the grammaticality of a child's utterance. Debate opened anew on the nature, availability, and sufficiency of this *implicit* negative evidence (Bohannon & Stanowicz, 1988; Demetras, Post, & Snow, 1986; Marcus, 1993; Morgan, Bonamo, & Travis, 1995; Morgan & Travis, 1989; Moerk, 2000). More recently, research has suggested that caregivers issue contrastive responses to a child's ungrammatical utterances, serving both to model how a structure is used and as evidence that

the child has erred (Chouinard & Clark, 2003; Saxton, 2000). Accompanying this is the growing evidence that positive statistical information available in the input is sufficient to drive considerable generalization for learning grammatical structure (e.g., Lewis & Elman, 2001; Mintz, 2003; Reali & Christiansen, 2003). These examples of corpus-analytic research make the obvious assumption that the syntax of children and caregivers is following a path of alignment or coordination: Statistical input and contingent responses in a conversation, if effective, shape the language of both the child and caregiver toward "syntactic coordination."

There are weak and strong interpretations of this syntactic coordination (Dale & Spivey, 2005). A "weak" interpretation merely refers to the child coming to use the particular language spoken by the caregiver. A "stronger" interpretation suggests that in ongoing individual interactions, there is a process of coordination taking place. The child (and/or caregiver) is inclined to produce sequences of words or syntactic phrases, during a conversation, that match those being heard. Research on adult conversation has lately suggested that a wide variety of behavior is coordinated in this way during social interaction. For syntactic structures particularly, for example, Branigan, Pickering, and Cleland (2000) recently demonstrated that in a picture-description dialogue, participants often repeat syntactic structures employed by another member of the conversation. Also, Cleland and Pickering (2003) found that participants can be primed to use certain noun-phrase structures given what occurs in ongoing dialogue (see Garrod & Pickering, 2004, for a review).

Even earlier in development, research on preverbal vocalization suggests that coordination might be an important characteristic of the language-learning task. For example, Goldstein and colleagues (Goldstein, King, & West, 2003; see also Bloom, Russell, & Wassenberg, 1987) recently demonstrated experimentally that maternal responses contingent on infant vocalization increase the quantity and quality of those vocalizations. Maternal

modeling contingent on vocalization contributes to the complexity of these vocalizations within an individual interaction. Tamis-LeMonda and Bornstein (2002) also demonstrated, in extensive analysis of mother-child interactions, that maternal responses that consistently and closely follow a child's utterances correlate strongly with later language development. These timely caregiver interactions perhaps provide clues that, as grammar learning proceeds, there might also be a process of syntactic coordination.

The strong interpretation thus has important theoretical implications (Clark, 1996). Recently, Garrod and Pickering (2004) have argued that dialogue is such a fluid and seemingly simple task for us because it is steeped in coordination mechanisms found in many cognitive processes during social interaction. These coordination interpretations about language input during acquisition, both weak and strong versions, complement this discussion by pursuing the extent to which coordination occurs and changes at a syntactic level in child-caregiver interaction. Although the weak assumption can be corroborated easily by observing any child and caregiver interaction, the strong version remains a tricky issue to quantify. In this article, we demonstrate strong coordination in the child's grammar learning environment. By subjecting three corpora of child-caregiver interaction to extensive analysis, we show that this pattern of syntactic coordination holds during development. Our approach adapts an analytic technique, used in a variety of disciplines, called *recurrence analysis* (e.g., Church, 1993; Eckmann, Kamphorst, & Ruelle, 1987; Von Heijne, 1987; Zbilut & Webber, 1992; see Dale & Spivey, 2005, for a review, and Webber & Zbilut, 2005, for an excellent technical introduction). The analysis reveals global structural patterns concerning how child and caregiver language aligns during interaction. By "global," we mean drawing general quantitative measures, with minimal dependence on statistical assumptions, describing the extent to which a bout of child-caregiver interaction involves language structures that are more or less similar to each other. Doing so provides a quantification of syntactic coordination in

transcripts of naturalistic dialogue. The method is based on analyzing sequences of syntactic elements, time series of grammatical usage, allowing comparison of two such sequences, and revealing patterns of *recurrence*. The ordered sequences of concern here are time series of syntactic class usage by child and caregiver. The approach therefore provides a window on how structures used by the child "recur" in those used by caregiver (and vice versa).

In a study similar to the present article, Sokolov (1993) made use of a program dubbed CHIP to investigate patterns of morphosyntactic usage between child and caregiver within a particular utterance window (see also Sokolov & MacWhinney, 1990). Results revealed concomitant morphosyntactic usage, which Sokolov argued supported a process of fine-tuning in child-directed speech. Sokolov also offered a number of thoughts about who might be leading this concomitant usage: Is it the child or caregiver during development offering up the coordinated structures in conversation? Sokolov speculated that it might, indeed, be both. Recurrence analysis might also shed light on this question.

In this article, we add to current research on syntactic coordination by further analysis of temporal patterns in child-caregiver interaction. Recurrence analysis is introduced below and is directly compared to some simple natural language processing (NLP) models of document analysis. Following this, we present an analysis of three corpora drawn from the CHILDES database (MacWhinney, 2000). The current article therefore has two primary goals. The first is to present a method by which grammatical coordination might be explored in real-time, naturalistic corpora. This method can aid in pursuing a number of questions concerning the structural and temporal patterns in child-caregiver interaction. The second goal is to apply the method toward answering two specific questions: Is there strong syntactic coordination in child-caregiver interaction, and does one speaker follow another, or neither? In what follows, we first consider some simple measures that might address the first of these questions. These measures are taken from well-known analyses in NLP and

other disciplines. We then lay out recurrence analysis, define its computed structures, and address how questions concerning time and structure can be extracted from them. Finally, we present an application of these methods to three CHILDES corpora.

Batch Comparisons

One way of assessing coordination between child and caregiver is to employ NLP techniques. Transcripts of child-caregiver interaction can be treated as documents and subjected to common NLP analyses. There now exists a vast literature involving the comparison of documents and their word-distribution patterns, with diverse applications (e.g., text clustering, topic-change detection, among many others; Manning & Schuetze, 1999). Many of these methods might be termed "batch," because they extract word or word-co-occurrence frequency distributions from documents, thus discarding data on how these words or co-occurrences are ordered in the document. In other words, if the documents in question are transcripts of a conversation (e.g., CHILDES transcripts), the batch methods ignore information about time, treating each transcript like a "bag of words." Despite this, they do offer a means of comparing word-class usage of child and caregiver.

The most common basis for comparing documents is often termed the *vector space model*. A document is seen as a vector whose elements represent the frequency or probability of a particular word or phrase occurring in it. For example, legal documents can cluster together, and be separated by other documents such as movie reviews, by computing the distance between the document's vectors involving words like *herein* or *shall*. Because the words or word co-occurrences can be seen as dimensions or events in a distribution, legal documents and movie reviews can be shown to have significantly different probability distributions among the words that compose them.

In this article, we are concerned with the coordination of syntax between child and caregiver. Because syntax inherently

involves ordering of elements, the simplest basis for comparing the usage of child and caregiver is to consider bigram co-occurrences of grammatical elements. Merely considering individual items (e.g., *nouns*) in isolation is not sufficient to establish syntactic coordination. Instead, bigrams afford the smallest window in which an ordering of syntactic elements can be compared (e.g., *noun verb*, as in "Anne kissed").

Comparing probability or frequency distributions of bigrams used by child and caregiver is a very simple extension of existing NLP techniques and widely applied information-theoretic measures (see Manning & Schuetze, 1999, for an excellent introduction and review). Although it might be a promising direction to apply the more sophisticated techniques currently available, we consider a few simple batch methods here. The first step in NLP applications is to obtain the documents for analysis—in our case, to extract the frequency distribution over syntactic bigrams. For a given transcript in CHILDES, we calculate the frequency of syntactic bigrams used by child and caregiver separately. The child and caregiver thus have individualized frequency distributions. Consider Figure 1, in which we overlay a frequency distribution for child and for caregiver. Here we have computed the frequency for each syntactic bigram for child and caregiver from one sample of Abe's corpus (Kuczaj, 1976). We have ordered these bigrams along the *x*-axis using the mother's bigram frequencies and then entered the corresponding frequencies for the child. Although these distributions look very similar, the distribution from the child's usage in a different transcript shows a slightly different pattern (dotted line). All distributions are derived from syntactic patterns of an English conversation, so, unsurprisingly, there will generally be similar patterns in them. In order to tell whether coordination in such simple terms as syntactic bigrams is occurring between child and caregiver, we can quantify the closeness or similarity of these distributions over multiple transcripts. To do this, for any given distribution of the child drawn from a sample we compute two corresponding distributions for the caregiver. We compare the bigram distribution of child and caregiver
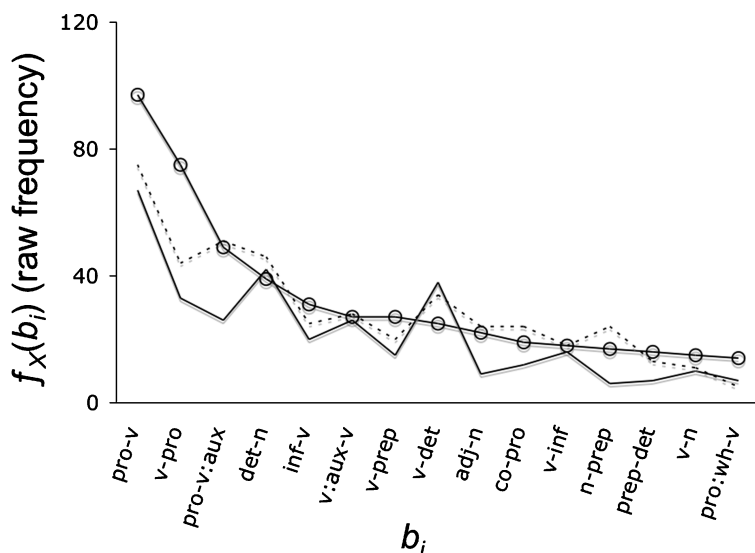
Figure 1. Example frequency distributions from one of Abe's transcripts. Caregiver is shown with open circles and Abe's distribution from the same transcript is shown in a solid line. Abe's distribution from a separate transcript is shown in a dotted line. Bigram patterns (*b*) shown along the *x*-axis, with corresponding frequencies on *y*-axis ($f_X(b)$). (*n* = noun; *prep* = preposition; *det* = determiner; *v* = verb; *wh-pro* = wh-pronoun; *adj* = adjective; *v:aux* = auxiliary verb; *pro* = pronoun; *inf* = infinitive marker "to").

in the same conversation and then compare the same distribution from the child with the caregiver's distribution from a transcript one step ahead in the corpus. The first comparison is measuring the similarity of bigram distributions of child and caregiver engaged in conversation, and the second is simply measuring usage from the child in one conversation and caregiver usage from another (i.e., the next sample in the corpus). This was done for Abe's entire corpus (208 samples). We present three basic measures frequently used in NLP document or word-distribution analyses. Each reveals significant coordination in raw syntactic bigram distributions.

Using the analogy of a vector space, these distributions can be compared using a distance metric. For example, Euclidean

distance (*ED*) can be computed, where each bigram frequency is considered a dimension along which caregiver and child have a value for a given transcript. The distance between two distributions (or vectors) is simply the square root of the squared distances along each dimension, given by

$$ED(m, k) = \sqrt{\sum^{B_T} [f_m(b) - f_k(b)]^2}$$

where *b* represents a given bigram occurring in one or the other distribution and f_X represents the frequency of that bigram in child or caregiver, with *k* representing the child's, and *m* representing the caregiver's. B_T is simply the set of all bigrams occurring in one or the other distribution (or both) extracted from an individual sample (or two, for the other-conversation comparison). If *ED* is close to zero, this means that child and caregiver distributions are very similar; greater values indicate disparity in bigram distribution. This measure reveals a difference in the expected direction, with *ED* smaller with same-sample distributions, $M = 62.5$ vs. 67.1, $t(207) = 2.5$, $p < .05$. Although the numerical difference seems small, this indicates that over the samples, there is a significant tendency for within-conversation bigram distributions to be more similar than a comparison of the child's and caregiver's from two separate conversations.

A commonly used information-theoretic measure is the Kullback-Leibler (KL) divergence, or relative entropy. Although this is not strictly a metric (as is *ED*), it has been used extensively to compare word distributions (e.g., Dagan, Lee, & Pereira, 1999). Each bigram frequency is converted into a probability by dividing it by the number of bigrams in the child's (or caregiver's) syntactic usage $P_X(b) = f_X(b)/\sum f_X(b)$, and

$$KL(m, k) = \sum_{B_m} P_m(b) \log \frac{P_m(b)}{P_k(b)}$$

This measure is often used to compare two distributions; it is zero when the distributions are exactly the same, and it exhibits increasing values when distributions differ. This measure in fact generates a highly reliable difference between same- and other-conversation comparisons, $M = .22$ vs. $.25$, $t(207) = 4.0$, $p < .0001$.

A third and final commonly used measure known as L1-norm simply takes the absolute value of the differences across the distribution: $\text{L1}(m, k) = \sum |P_m(b) - P_k(b)|$. Using this measure again obtains a significant difference between distribution types, $M = .75$ vs. $.82$, $t(207) = 9.9$, $p < .0001$.

Application of NLP measures shows that, at least in such simple terms as raw coordination of syntactic bigrams, child and caregiver are coordinated in conversation. Nevertheless, these batch measures cannot approach the strong version of the interpretation discussed in the introduction to this article: These batch measures compare distributions independent of how the bigrams are occurring in time. A more thorough investigation into syntactic coordination therefore requires measures sensitive to the temporal ordering of the bigrams. In the next section, we adapt a technique that has been used in both NLP and other contexts (e.g., heart rate, seismology, and postural adjustments, among others), called recurrence analysis. This method can be similarly applied to bigrams of syntactic usage, although inherently temporal in nature. The method will therefore permit more detailed quantitative hypotheses regarding the coordination between child and caregiver in ongoing conversation.

Recurrence Analysis

The batch measures just reviewed see child and caregiver syntactic usage as sequence-free bundles of bigrams with corresponding frequencies or probabilities. These distributions are then subjected to vector-based or information-theoretic analyses. Although they reveal a tendency for similar bigrams to be employed by child and caregiver, they do not yet address the problem of strong syntactic coordination. Instead of massing distributions of bigram usage, consider separating the usage by child

and caregiver from a transcript and forming two time series representing syntactic elements:

$$caregiver : s_{m,1}, s_{m,2}, \ldots s_{m,N_m}$$

$$child : s_{k,1}, s_{k,2}, \ldots s_{k,N_k}$$

Here, $s_{X,i}$ represents the ith syntactic element used by speaker X (k for child, m for caregiver). N_X is the total elements composing the time series for speaker X. Once two time series have been extracted from a transcript, we can seek bigrams that match between caregiver and child. If syntactic element i and $i + 1$ of the caregiver match j and $j + 1$ of the child, we have a matching bigram between the two time series: $(s_{m,i}, s_{m,i+1}) = (s_{k,j}, s_{k,j+1})$. This match can simply be represented by the time index pair (i, j), with the understanding that it represents a bigram "recurrence." By doing a comparison of all bigrams occurring in the transcript, we can construct a large grid, the rows of which represent bigrams of the caregiver and the columns represent that of the child. In each cell of this grid, we place a 0 or 1 depending on whether the cell's corresponding coordinates involve a recurring bigram: 1 when recurring, 0 otherwise.

For the sake of clarity, we present two examples of constructing this grid. First, consider the following very short and imaginary time series of child and caregiver syntactic usage, with an example recurrence in bold:

*caregiver (m): determiner **noun verb** determiner noun*

*child (k): **noun verb** determiner adjective noun*

Figure 2 shows the grid of bigram positions ordered in time, with caregiver along rows and child along columns. These row and column bigrams are placed by moving a window of size 2 along each speaker's time series. Each time a bigram used by the child is the same as that by the caregiver, a 1 has been placed in the corresponding cell. This structure is referred to as a recurrence plot (RP) and can become a very large grid of 1's and 0's when we build the structure from full transcripts. For example, and as a
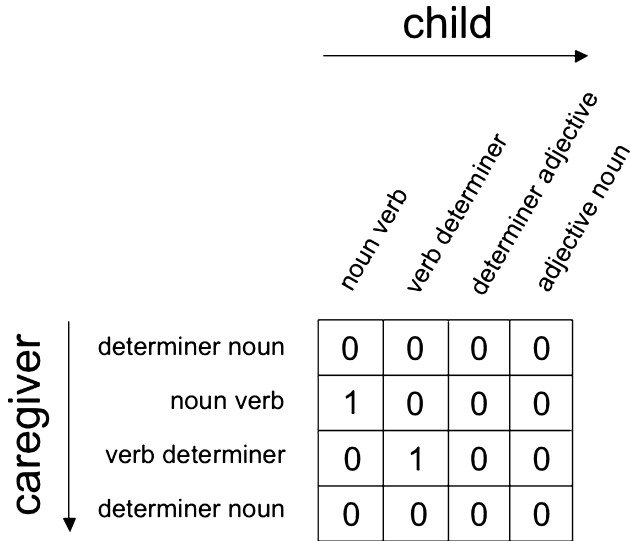
Figure 2. Example recurrence plot constructed from simple time series. 1's are recorded in cells corresponding to matching bigrams, with 0's otherwise.

fuller illustration, a plot from one of Abe's samples is shown in Figure 3. To translate this grid into a familiar Cartesian format, we place the caregiver's time series along the x-axis (time index i) and the child's along the y-axis (time index j). In addition, 0's are not plotted on this vast grid, and 1's are represented simply by filled pixels where recurrent bigrams are occurring.

In the sample that generates this plot, Abe is 2 years, 6 months of age. He is engaging in conversation with his caregiver and they are discussing a variety of toys and activities in their environment:

Father: Do you like those toys Abe?

Abe: Uhhuh like them see those donkeys down there

Abe: And what else?

Father: What's that?

Father: What are you doing Abe?

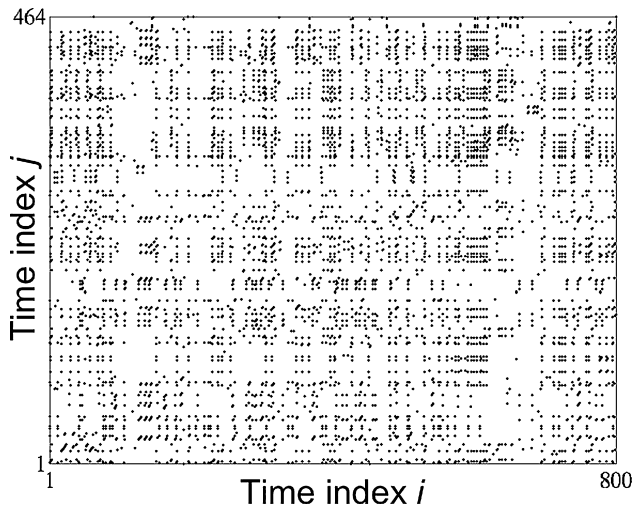Abe: I want to beat you up in the head!

Figure 3. A full recurrence plot from one of Abe's samples. Filled pixels in the plot represent points at which matching syntactic bigrams are occurring in the two time series of child and caregiver.

To begin building this full recurrence plot, we first extract the sequence of word-class usage by Abe and his caregiver. For example, the first usages by Abe's caregiver are

"do you like those toys abe"

The CHILDES coding makes available word class information on the morphosyntactic tier of the transcripts, and the syntactic sequence of this sentence is represented by

"verb pronoun verb determiner noun proper-noun # . . ."

All sentences used by one speaker are strung together in a time series of syntactic usages, thus maintaining the temporal ordering within conversation. If we were simply conducting a batch analysis as described earlier, we would initially compute the frequencies or probabilities of bigrams.

To conduct recurrence analysis, we instead build the large grid representing pairs of time indexes at which child and caregiver used the same syntactic bigram. Each point (i, j) will have as its elements the time index, i, at which caregiver used the

bigram, and time index j, at which Abe did. This actual visualization of the points reveals where matching syntactic bigrams are occurring in conversation (relative to when the child and when the caregiver produced the sequence of two syntactic elements).

In batch analyses, quantitative measures of syntactic coordination were computed by analyzing the frequency of word-class bigrams used in conversation. As already mentioned, these batch analyses do not retain temporal information about the usage of bigrams. An RP is in fact a raw structure that can capture the bigram distributions while also showing where they are occurring in time in the language of both members of the conversation. It is very easy to show that RP contains much of the information needed for batch analysis (see the Appendix). Nevertheless, what is more interesting is that measuring the temporal patterning of recurrent bigrams becomes possible: As child-caregiver interaction unfolds in time, we can inspect the relative temporal incidence of recurrent syntactic bigrams. From these RPs, quantitative measures might thus be extracted by analyzing the number and nature of points in the plot (Zbilut & Webber, 1992). This amounts to analyzing the temporal indexes i and j and computing measures based on their distributions. We introduce three measures here and will put them to use in full CHILDES analyses below.

The simplest and "rawest" measure drawn from RPs is the *recurrence rate* (RR). This measure is simply the percentage of the plot that is filled with points. This crude measure represents the overall extent to which child and caregiver are using same word-class patterns and has been used in other domains in psychology (e.g., posture; Shockley, Santana, & Fowler, 2003).

$$RR = \frac{\|RP\|}{N_m N_k}$$

An important issue discussed in the introduction to this article is whether child and caregiver engage in syntactic coordination in dialogue. Pursuant to this, another intuitive quantity can be

obtained by calculating the extent to which recurrence points occur in temporal proximity: Abe and caregiver tend to use similar bigrams at around the same time in conversation. In the simplest case, where the sequences of word-class usage by child and caregiver are of the same length ($N_m = N_k$), this question can be answered by noting the number of recurrence points (i, j) such that $|i - j| \leq w$. In other words, the patterns of syntactic usage that recur tend to fall near the diagonal "line of incidence" of the recurrence plot (where $i = j$) within some distance w (or width around the line of incidence).

For our purposes, because the sequence of usages of caregiver and Abe might differ in length, we cannot extract the points around the *simple* diagonal on the plot (see Figure 4). Instead, we generate a hypothesized line of incidence by drawing positive integers that approximate the line $j = N_k/N_m \times i$, the "true" diagonal from $(1, 1)$ to the point (N_m, N_k). A band of size w is extracted by translating this line along the axes for i and j. The band along the line of incidence is therefore a set of points approximating a temporal coincidence of syntactic bigrams. A measure from this set of points is straightforwardly defined. The *diagonal recurrence rate* of width w (RR_w) is number of points (i, j) falling in that band, divided by the total coordinates satisfying the approximate line and its translations from $-w$ to w. With D representing the number of *possible* point coordinates on the plot that are inside the integer-approximated band $j = N_k/N_m \bullet i \pm w$,

$$\mathrm{RR}_w = \frac{\|\mathrm{RP} \cap D\|}{\|D\|}$$

Finally, we are concerned with the question about temporal ordering of the recurrent patterns. The RP might reveal whether the same pairs tend to follow or precede Abe's usage. To illustrate in the simplest case, allow $N_m = N_k$, where the line of incidence is $i = j$. If a recurrence point falls *above* this line of incidence, its coordinates relative to that line of incidence are $(i, j + a)$, where a is the vertical distance from the line of incidence. In
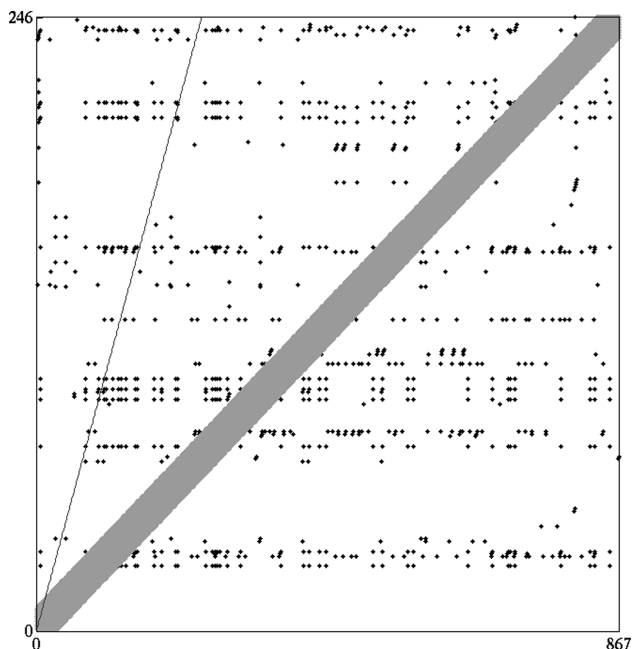
Figure 4. When $N_m \neq N_k$, the line of incidence with $i = j$ only crudely approximates temporal coincidence (thin black line). Instead, points are sampled along the line from $(1, 1)$ to (N_m, N_k) (thick gray line). Points within this band are the basis for diagonal recurrence measures.

other words, this recurrence point in the upper triangle of the plot indicates the pattern generating that same pair along the axis for j is occurring *later in time* relative to the axis for i. This would mean that if caregiver usage is always represented on the i-axis, the child's use of the pattern is occurring *after,* and therefore following, the caregiver usage. An intuitive grasp of this pattern is possible by looking at Figure 5. By shifting points around the line of incidence, it is clear that shifting away from the i-axis (to above the line) moves the recurrence point farther ahead in time. RR_w can then be divided in terms of recurrence contributed by the upper and lower triangles. If the caregiver sequence is always represented along the i-axis, the first element in points (i, j), the *speaker diagonal recurrence rate* can be simply defined as (for the caregiver):
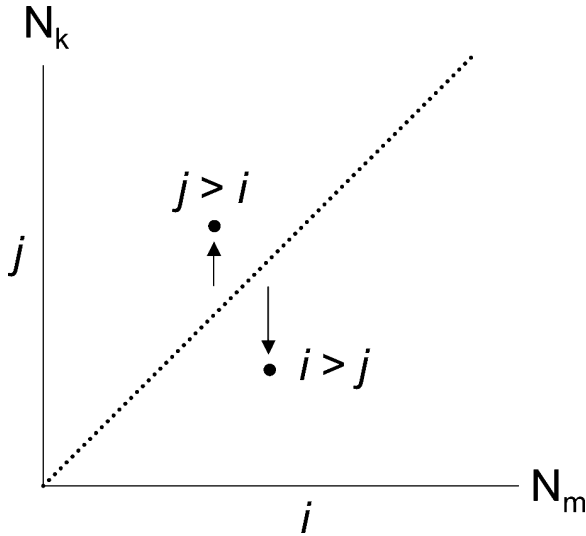
Figure 5. When a recurrence point is above the line of incidence, the time index *j* is larger, or farther ahead in the transcript, than index *i*. The reverse is true when the point falls below the line of incidence. When caregivers are always represented as *i*, this temporal pattern might reveal from whom the recurrent sequence originates.

$$\mathrm{RR}_{m,w} = \frac{\|\mathrm{RP} \cap D_{w+}\|}{\|\mathrm{RP} \cap D\|}$$

This measure is a percentage, representing the contribution to the diagonal recurrence rate from the upper ($w+$) or lower ($w-$) portions of the grid. Within a sample, if $\mathrm{RR}_{m,w} > \mathrm{RR}_{k,w}$, then, on average, the caregiver is "leading" the conversation (in terms of recurrent bigrams).

The RP in Figure 3 has a RR value of 1.30%, representing a set of 4,911 points. The diagonal recurrence with $w = 50$ is 1.58%, and with 150, it is 1.40%. For this plot, we can compute the extent to which points are being contributed by caregiver or child along the diagonal. In this particular sample, with $w = 50$, we have $\mathrm{RR}_{m,w}$ of 48.7%, and $\mathrm{RR}_{k,w}$ at 51.3%. Because $\mathrm{RR}_{k,w} > \mathrm{RR}_{m,w}$, it seems that Abe has a lead on the caregiver (although very weakly).

These simple example analyses serve to illustrate how an RP is constructed and analyzed. They permit comparison of syntactic usage by child and caregiver while providing the tools to measure temporal organization of this usage. Although batch analyses are capable of showing distributional similarity in the conversation, the RP now permits investigation of the stronger coordination hypothesis: Is there stronger syntactic coordination going on along the line of incidence than outside? Also, does the child lead or follow?

We introduced three measures that will be used in subsequent analyses to answer these questions. A few points are in order before continuing. First and most importantly, the values generated by these measures are small. As in batch analyses, we can therefore only expect meaningful patterns to emerge as we observe these measures across different transcripts (as Abe and the other children develop). This is, indeed, what is typically expected in the literature (e.g., Sokolov, 1993). Second, there are a number of parameters that can be explored in the analysis. Because we have very powerful analyses given the many transcripts of our chosen corpora (see next section below), it will be useful to test for robustness of any patterns by using multiple parameter values. We do this below by varying w and using patterns larger than just bigrams (trigrams and quadrigrams). As is customary, we will refer to this as the size of the "window" n-grams and vary the n parameter in these analyses. With bigrams, for example, $n = 2$.

Analysis

Materials

We selected three English corpora from the CHILDES database (MacWhinney, 2000): Brown's Sarah (Brown, 1973), Kuczaj's Abe (Kuczaj, 1976), and Sachs's Naomi (Sachs, 1983). These three corpora were used recently by Chouinard and Clark (2003) for the same reasons we choose them here: The sample

sizes are relatively large, numerous, and drawn at regular intervals. We selected transcripts from these corpora in which the child consistently used approximately two morphemes per utterance. The age range of the resulting transcript sets is presented in Table 1. Each sample was turned into two separate time series of word-class usage: one for usages by the child and another for usages by the caregiver.[1] Each was composed of the *syntactic class* to which the particular word belonged. For example, in one of Abe's samples, the child's usage of syntactic classes (gleaned from CHILDES' morphosyntactic tier) were strung together and evaluated according to some window size n (with bigrams, $n = 2$); likewise, all caregiver usage in that sample (any language issued to the child, whether father, mother, or other) composed another time series of usage. Table 1 presents these three corpora. Figure 6 presents further information about the transcripts. Abe is at a relatively greater level of development than Sarah and Naomi, who take more time to achieve longer length of sentences compared to caregivers—here, in terms of morpheme number per sentence (cf. mean length of utterance). To construct recurrence plots, we used the Matlab CRP Toolbox (Marwan & Kurths, 2002).

Procedure

For each child individually, we constructed RPs for each sample as outlined earlier.[2] Once again, this was done by recording the time indexes at which child and caregiver used the same sequence of size n (bigrams, trigrams, or quadrigrams). For window

Table 1

Total age range of the corpora, and samples used in recurrence analysis

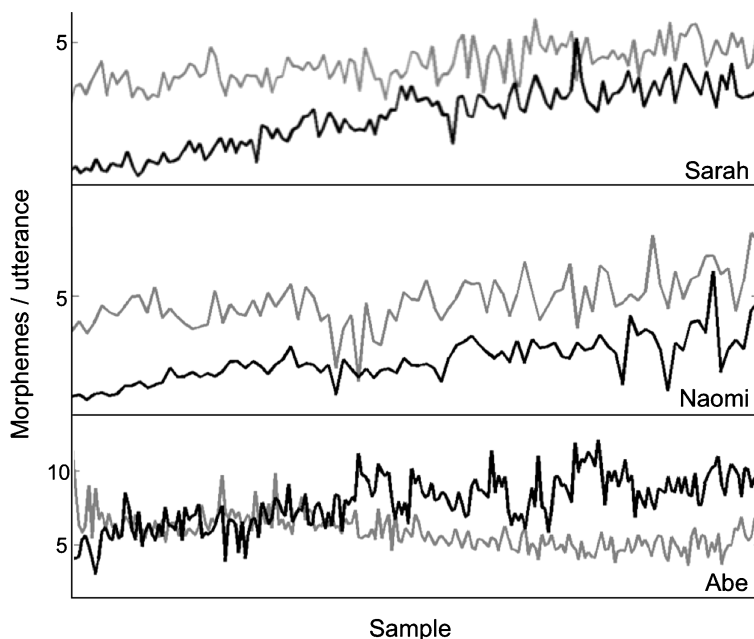| Corpus	Age of first sample	Age of last sample	Number of transcripts
Abe	2; 5	5; 0	208
Sarah	2; 9	5; 1	109
Naomi	1; 11	4; 8	76

Figure 6. Abe, Sarah, and Naomi (black lines) exhibit different levels of development in terms of the number of morphemes per sentence in the transcripts (gray lines present these data for caregivers). Values were calculated by dividing the total number of syntactic elements used by the number of sentences in a transcript.

size n and size of diagonal w, we chose three parameter values to explore. We built RPs with window sizes of $n = 2, 3$, and 4. In addition, diagonal recurrence was computed by using three values of w: 50, 100, and 150. Given patterns of temporal organization between caregiver and child syntactic usage, these should appear robustly with all parameter values. Finally, because $n = 2$ generates the largest amount of recurrence points (largest sets RP) and the bandwidth w of 50 the most stringent temporal context, we chose these two parameter values to explore speaker diagonal recurrence.

Three separate analyses were therefore conducted over each transcript's recurrence plot. First, we simply computed the total

recurrence (RR) of each plot. This results in 3 (values of n) $\times$ 3 (children) = 9 separate analyses for RR. We then computed the diagonal recurrence (RR_w) for each of the parameter values, resulting in 3 (n) $\times$ 3 (w) $\times$ 3 (children) = 27 separate diagonal analyses. Finally, we computed speaker diagonal recurrence ($RR_{m,w}/RR_{k,w}$) for each sample using $n = 2$ and $w = 50$, giving three (children) separate analyses.

The same analyses were conducted with two "control" conditions. These conditions were defined by creating different plots by either (a) manipulating the sequence of word-class usage in the *same* sample or (b) using sequences from two *different* samples (e.g., *sample* vs. *sample* + 1, as in above). For the former control, in a *shuffled-sample* condition, we disordered the child's utterances, keeping each utterance intact, but shuffling their order in conversation. Therefore, "batch" results (such as RR) should remain exactly the same because they are independent of the temporal organization of word-class usage. However, diagonal measures should be directly influenced by shuffling the sequence of word-class usages. If there is indeed temporal organization, it should be lost when we randomize the order of the sentences used by one of the speakers.

For the latter control, which we will refer to as *next-sample*, we conducted recurrence analysis (building RPs) between the child's sequence for a given sample and the sequence of usage by caregivers *one sample ahead in the corpus*. This was done in our batch-analysis examples presented earlier. In the next-sample control condition, we would still expect to find many recurrence points because the same language is being used across samples. However, we should observe differences in both recurrence rate *and* diagonal measures.

Each child and parameter set was considered separately for statistical analysis. Only recurrence results for a particular parameter set and child were collapsed. For example, an analysis would be built using bigrams ($n = 2$) and three separate conditions (same-sample, shuffled, and next-sample). These three conditions were compared in terms of recurrence rate, RR. Next, for this

bigram analysis, three values of w (50, 100, 150) were used to compute diagonal recurrence measures. Once again, the three conditions were compared in the diagonal scores they generate. This was done for all values of n and all three children.

Predictions

When translating these conditions into specific predictions, we should expect a few basic patterns to emerge. These can be phrased in terms of expected results in our three recurrence measures, RR, RR_w, and $RR_{m,w}/RR_{k,w}$. As mentioned earlier, RR represents the raw extent to which patterns of size n are used similarly by caregiver and child (not unlike batch measures). By definition, there will be no differences between same-sample and shuffled-sample conditions. However, if there is an overall tendency in conversation to use similar word-class sequences, there should be higher values of RR in same-sample (and shuffled-sample) conditions than in the next-sample condition.

If child and caregiver tend to use similar syntactic n-grams at the same time, guided by the temporal context of ongoing conversation, then we should expect RR_w to be higher than the value of RR in the whole plot, but only for the *same-sample* RP condition, because only it preserves the temporal organization of discourse. Indeed, as w becomes larger in this condition, we should expect that RR_w will drop, because we are widening our temporal scope and allowing greater asynchronies to be included. In other words, $RR_{50} > RR_{100} > RR_{150}$.

Finally, it is not immediately obvious what should be expected in terms of leading and following, $RR_{m,w}/RR_{k,w}$. Sokolov (1993) offered some discussion on what pattern of leading should emerge in conversation between caregiver and child. He remarked that there is a historical tendency to imply that early in grammatical development, maternal leading takes precedence. As already mentioned, debate occurred long ago over what structural correlations mean in child-caregiver interaction (Newport,

Gleitman, & Gleitman, 1977). This measure is therefore, to some extent, exploratory in nature. Importantly, it might be that each child will lead or follow depending on their level of syntactic development. A child at a younger age, and not yet mastering specific structures, might tend to follow the mother's usage, whereas children who are flexible grammarians might tend to take the lead in conversation. On the other hand, neither might offer the lead in conversation. The results concerning this recurrence measure are therefore a first step toward uncovering these patterns of following and leading, whether there is developmental change, and individual differences across children. A result involving caregiver/child trade-offs in leading would strengthen a coordinative perspective on human communication—and shed some light on its developmental origins. This theoretical perspective might recommend that it is neither the mother consistently guiding nor the child urging "I'd rather do it myself" (Newport et al., 1977), but, rather, "Let's do it together."

Results

As in the foregoing subsection, we consider each measure separately in this subsection. As already mentioned, we consider each child separately for statistical analysis. This in fact results in a more conservative test of the patterns and a potential window onto any individual differences or unfulfilled temporal organization across all children. Within each subsection, we conduct statistical tests by comparing the three conditions for each parameter set: same-sample, shuffled-sample, and next-sample plots.

RR: Overall Recurrence

The recurrence rate was computed for each of the three conditions and at each chosen value of n, 2–4. As expected by definition, RR was not different for same- and shuffled-sample conditions. To compare same and next conditions, we ran a 2 (same vs. next) × 3 (size of window, n) repeated-measures

ANOVA for *each* child. All three revealed a significant main effect of plot type ($p < .05$). All children and n-gram sizes revealed a greater RR for the same-sample plot condition (see Figure 7) compared to the next-sample condition. The recurrence rate is therefore significantly higher when we compare syntactic usage of caregivers and children in the same conversation—there seems to be an overall tendency to use similar syntactic sequences. As expected, as the window size is increased from 2 to 4, the probability of a recurrent sequence becomes very low and RR diminishes. Although these RR differences are small in magnitude, they are consistent across all parameter values and children.

RR_w: Diagonal-Windowed Recurrence

RR_w reveals the amount of recurrence occurring along a band of width w around the "true" diagonal by building a line from $(1, 1)$ to (N_m, N_k), as described earlier. There are two parameters to explore here. We assess RR_w at three lexical widths: 50, 100, and 150 words around the approximated diagonal. Once
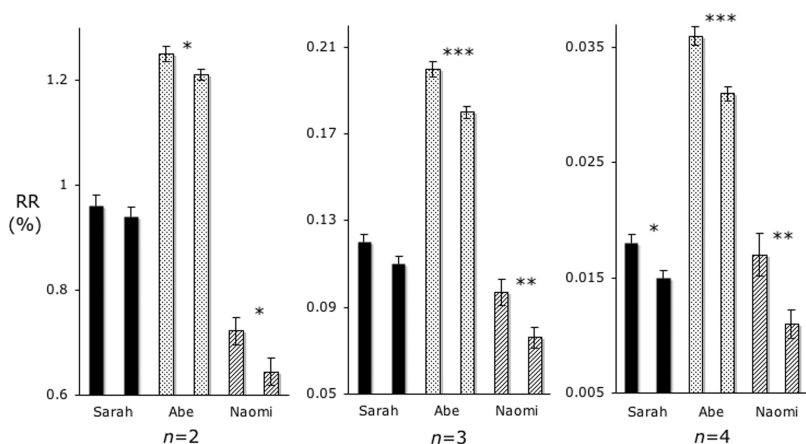

Figure 7.   Mean recurrence rate results for the three children, across n-gram sizes. For each pair of bars, the left bar represents same-sample recurrence and the right bar represents next-sample recurrence.

again, as earlier, we do this for all values of window size n. We ran a 3 (same-, shuffled-, and next-sample) × 3 ($w = 50$, 100, 150) × 3 ($n = 2, 3, 4$) repeated-measures ANOVA. Once again, all children showed a significant main effect of sample type ($p < .0001$). In addition, we should expect to find a significant interaction between sample type and size of band, w. All children exhibited a significant interaction ($p < .0001$). Figure 8 shows overall results. The same-sample plot, which maintains the temporal ordering of conversation between caregivers and children, consistently shows higher recurrence at all parameter values. In fact, all children also reveal a diminishing of RR_w as w increases in the same-sample plot conditions. Only Abe, however, shows this trend to be significant.
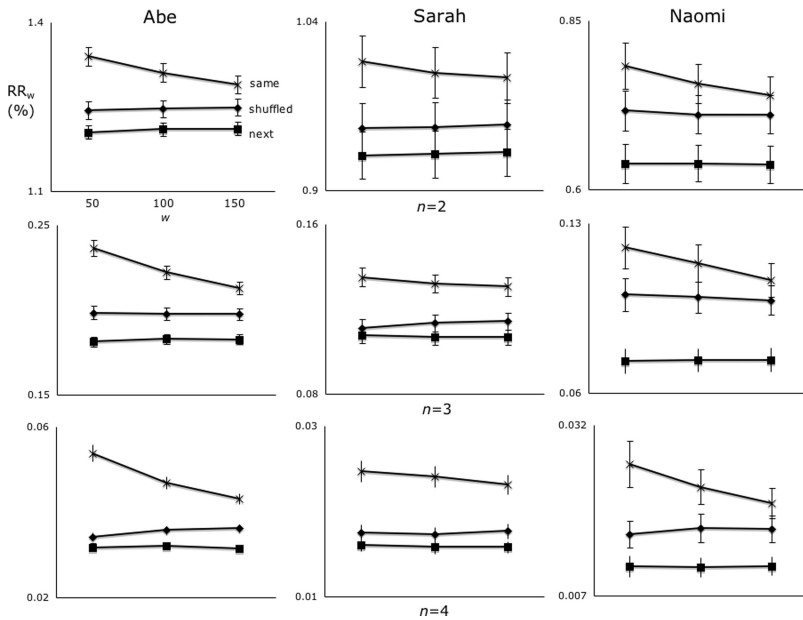

Figure 8.   Mean diagonal recurrence rates across all children, n values, and band width (w) around the sampled line from $(1, 1)$ to (N_m, N_k).

$RR_{m,w}/RR_{k,w}$: Leading Versus Following

In a first planned analysis, a straightforward way to test for a tendency to lead conversation is simply to consider the percentage of points contributed ("led") by child compared to caregivers across all samples. Such a broad analysis across samples only shows weak results. Paired *t*-tests on the three children only showed a significant effect for Abe and Sarah. Interestingly, Abe revealed a tendency to lead his caregivers, average $RR_{k,50}$ of 50.5% vs. $RR_{m,50}$ of 49.5%, $t(207) = 2.8$, $p < .01$. In other words, only 1% of recurrence points along the band of width 50 are contributed more by Abe than by caregivers when assessed across samples. Sarah, instead, revealed a small tendency to follow, average $RR_{k,50}$ of 49.5% vs. $RR_{m,50}$ of 50.5%, $t(109) = 2.7$, $p < .01$. Once again, this reveals only a 1% difference. Naomi showed no significant difference. Nevertheless, as shown in Figure 6, Abe is farther ahead in development than both Sarah and Naomi, and this seems to suggest that an individual difference might be revealed in these separate corpora.

To get a more interesting picture of patterns of leading or following, we conducted a different planned analysis. To what extent, over periods of development, does the child or caregiver lead conversation? If $RR_{k,50}$ in a transcript is greater than 50%, this is coded as a "score" for the child; if $RR_{m,50}$ is greater than 50%, the caregivers lead, and a score is counted for them. If children exhibit consistent patterns across samples of leading or following, then cumulative scores across transcripts should reveal a developmental pattern or trajectory; that is, if the process of leading or following is indeed small enough to recommend only a pure Bernoulli process with probability of .5 (there is only approximately a uniformly random chance that the child or the caregiver might lead), then cumulative scores for the children should not differ from what is expected from this random process. In Figure 9, we present two solid gray lines representing a 95% confidence interval based on 100 runs of such a Bernoulli process, meant to simulate the null condition if caregivers and child are
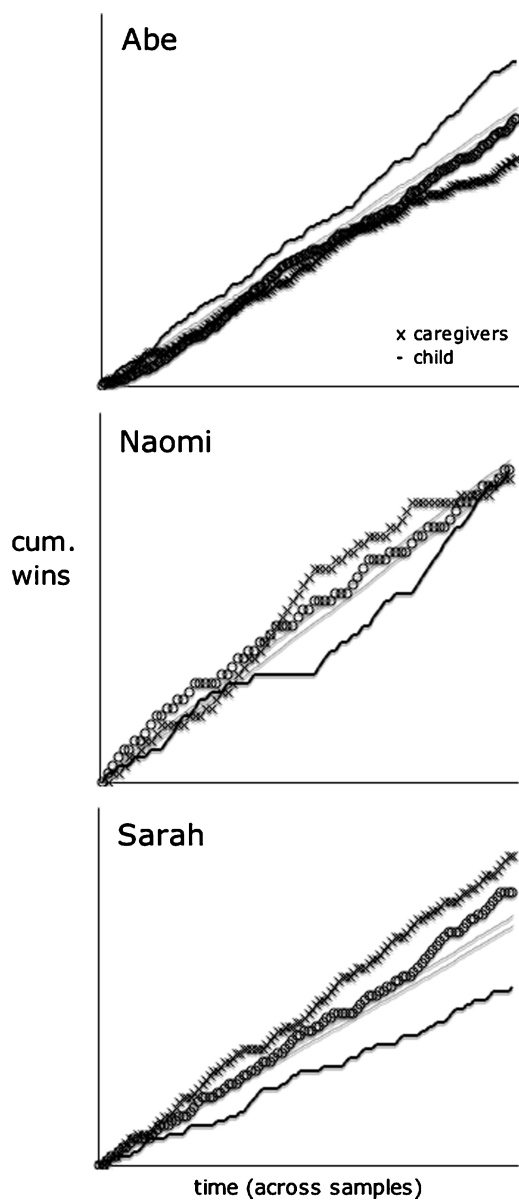
Figure 9. Cumulative wins across selected transcripts. See text for details.

equally likely to lead. In addition, Figure 9 shows cumulative scores across samples for each child. Caregiver (lines with "x" markers) and children (solid lines) reveal that cumulative scores consistently exceed this baseline interval. This demonstrates a first look at individual differences in contribution to coordinative usage. Abe, who is rather farther ahead in development than Naomi or Sarah, shows a consistent rise in his score. Sarah, on the other hand, is consistently behind the caregiver, and Naomi spends some time behind, but subsequently catches up. To make sure that these patterns are based merely on the quantity of language contributed by the child, we conducted a regression analysis between number of word classes used by the child and the proportion of $RR_{k,50}$ to $RR_{m,50}$. No children revealed a significant relationship. This can be further supported by seeing cumulative scores based on the shuffled-sample plots. The child's score is presented as the open "o" line in Figure 9. This shuffled-sample child's score does not consistently escape the 95% confidence interval. This preliminary look might therefore offer the possibility that caregivers and children are indeed taking turns leading recurring usage, and individual differences might coincide with level of development.

More importantly, both $RR_{X,50}$ tests should be considered together. Although there seems to be a developmental trend in leading and following, the initial statistical comparisons reveal that this leading is marginal. Therefore, although there might be a small tendency for children to become leaders or followers over developmental time, it is still the case that both sides of the conversation are contributing comparable amounts of recurring bigrams.

Discussion

These results further suggest that there is syntactic coordination in conversation between child and caregiver. Furthermore, the diagonal RR measure (Figure 8) suggests that this coordination is stronger in ongoing conversation—children and caregivers are more likely to use recurrent patterns of word-class usage in

that temporal context. Finally, the speaker diagonal RR provides hints that caregiver and child might trade off the lead during development. Abe, a skilled grammarian early on, consistently leads his caregivers, whereas Sarah and Naomi exhibit the reverse.

Although these results are promising, they suggest a crucial further analysis to gain an understanding of what underlies them: What is the nature of the recurrent patterns? The following analysis targets the specific *n*-gram patterns and their relative frequencies in the same-sample plot conditions. Although the foregoing results, analyzing only patterns within sentences, are compelling, it would be illuminating to gain insight into the specific patterns being coordinated. The following subsection seeks an answer.

Pattern Analyses

What specific syntactic sequences are recurrent between child and caregivers? By tracking the syntactic *n*-grams associated with each child-caregiver recurrence point we can assess the distribution of patterns and which specific part of speech is playing a greater or lesser role in guiding coordination. Although the current analyses are only preliminary in this direction, they offer some interesting suggestions about the organization of recurrent syntactic patterns. First, we demonstrate a Zipf-like distribution in the same-pair patterns. This provides an overall characterization of the pattern distribution within each transcript. Next, we target the trigram patterns to explore differential contributions of nouns versus verbs. This reveals that recurrence points, or same pairs, might be reflective of other discussion in the language acquisition literature.

Zipf-like n-*Invariant Distribution*

Before we get to pattern identity, one interesting property that emerges is a Zipf-like distribution in the *n*-grams. Zipf distributions have been explored extensively across numerous

disciplines (see Li, 2002, for a review). Although numerous explanations exist for this and similar power laws (e.g., see Mitzenmacher, 2003, and Newman, 2005, for reviews; see Van Orden, Holden, & Turvey, 2003, for a similar discussion in the psychology literature), the purpose of examining this distribution in the current transcripts is simply to characterize the distribution of the patterns producing the recurrence points. The simplest expression of a Zipf relationship is that the frequency of a word or pattern is inversely proportional to its rank to the exponent of α.

$$f(b_i) \propto \frac{1}{\mathrm{rank}^\alpha}$$

To examine this relationship here, for each *n*-gram in each transcript we compute the number of matches between child and caregiver. The patterns are then ranked within that transcript. The log transformation of the rank and bigram recurrence count is then subjected to linear regression. An example plot is presented in Figure 10, illustrating the strong linear relationship
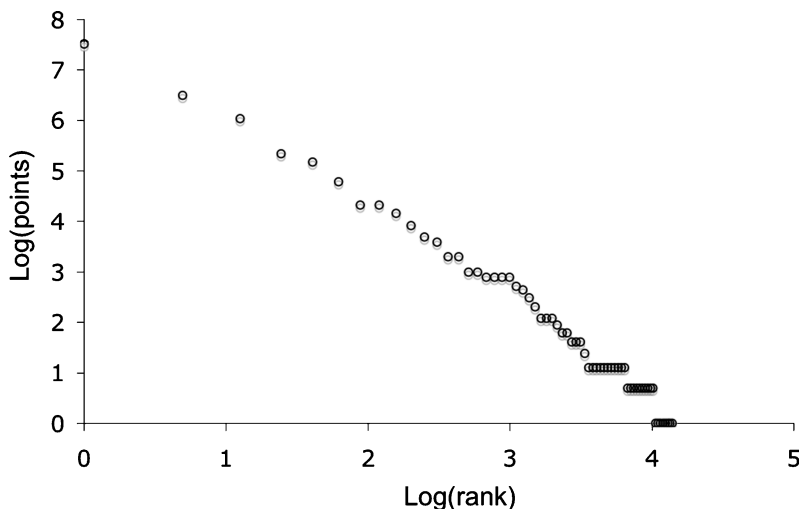

Figure 10.   Log-log graph of recurrence points (*y*-axis) and pattern ranking (*x*-axis).

between matches (recurrence point "frequency") and rank. The log-log regression scores across n values are presented in Table 2. Across all n-gram lengths, this power-law model produces highly reliable r^2 values. Although there is some debate about the applicability of Zipf when other distributions might be a better fit (e.g., Egghe, 2000; Li, 2002; Mitzenmacher, 2003), these regression scores suggest that this is a particularly good characterization of the pattern distributions here. The patterns that are coordinated are thus ranked in a "heavy tailed" distribution. This means that, similar to other properties in language (e.g., Zipf and word frequency) and nature in which Zipf applies (Li), there are highly frequent sequences of word classes guiding the recurrence patterns in conversation. This frequency drops off considerably, according to a power-law relationship, as we consider less frequent sequences. The next analysis begins to look at what word classes play a role in such highly frequent sequences.

Verbs and Nouns

Although the Zipf analysis broadly characterizes transcript pattern distribution, one might wonder which grammatical elements are contributing to the recurrent patterns. As in the presentation of batch analyses, we focus on a specific window size— here, on trigrams ($n = 3$). By targeting the patterns that contain verbs and nouns, we can investigate the extent to which these

Table 2

r^2 *values for log-log linear regressions of rank and recurrence points averaged across samples*

Corpus	$n = 2$		$n = 3$		$n = 4$	
	r^2	α	r^2	α	r^2	α
Abe	.96	.51	.96	.70	.94	1.00
Sarah	.95	.46	.95	.63	.94	.87
Naomi	.91	.53	.91	.71	.81	.89

different parts of speech contribute to the trigram recurrences in each plot. Each set of recurrent patterns containing a noun, for example, is counted. Likewise, all patterns with a verbal element are tallied. We can directly compare these values across plots. In all three children, the participation of verb elements in coordination between child and caregiver is greater than that of nouns ($p < .01$). This interesting result echoes extensive current discussion concerning the central role that verbal structures play in developing grammar (e.g., see Tomasello, 2003, and Clark, 2003, for a review).

The first results in recurrence measures indicate syntactic coordination between interlocutors during child language acquisition. These secondary results suggest that such recurrence is not simply a matter of trivial structural repetition. Instead, these recurrence patterns might be reflective of grammatical acquisition patterns.

General Discussion

In all three corpora, transcripts exhibit coordination between child and caregiver, particularly in the context of ongoing conversation. The pattern analyses provide some clues about these recurrent patterns and suggest that they might be reflective of ongoing debate concerning the characteristics of early grammatical behavior in children learning English (see below). It therefore appears that syntactic sequences are coordinated between children and their caregivers. In addition, the results from speaker diagonal recurrence have permitted a peek into whose patterns are guiding the temporal coordination. Results suggest a basis in individual differences, wherein advanced children are often leaders, whereas children earlier in development might be guided by caregivers.

These results are consistent with a strong coordinative interpretation on mechanisms in social interaction (Garrod & Pickering, 2004). At the syntactic level, coordination of structure seems to be occurring early in development. A number of

speculative remarks might be relevant. First, further analysis is required to discern whether this coordination is consistent across other children and other languages. It is interesting to suppose that coordinative processes, like syntactic coordination, are an important part of language acquisition. Second, just as Sokolov (1993) regarded concomitant morphosyntactic usage by child and caregiver as a process of fine-tuning, these results suggest that children and caregivers are trading off over development in a broader coordinative way. "Fine-tuning," in this sense, suggests a unidirectionality, when in fact the child can be actively involved in shaping the conversational context, particularly when this child reaches a higher level of grammatical development. Whereas Sokolov did not insist that parents are necessarily providing all of the guidance in this coordinative process, the current results further attest to a rich dynamic between children and their caregivers during syntax acquisition. In fact, this dynamic might reflect earlier coordination, as mentioned in the introduction to this article: Goldstein and colleagues (Goldstein, King, & West, 2003) have noted that preverbal infants are sensitive to contingent responses by caregivers and that caregivers are sensitive to responses issued by children (Goldstein & West, 1999).

There are, of course, a number of important limitations of the current results that should be acknowledged. First, because this is a "global" kind of analysis (see also Hart & Risley, 1995), it has yet to delve into specific structural recurrence, such as *aux*-questions or transitive verb constructions, and how they might be organized in time during conversation between child and caregiver (e.g., Fey & Loeb, 2002). Differing grammatical particles might exhibit stricter temporal or contextual organization in child language at earlier ages, and over development, it might grow into flexible usage. Although this is merely speculative, the method outlined here might be applied to such structures by targeting the specific patterns containing them. In fact, the kind of method Sokolov (1993) described might be integrated with recurrence analysis. Although recurrence analysis provides a generalized means of

exploring temporal organization of syntactic patterns, the nature of the patterns and how they relate might adopt similar comparison algorithms such as Sokolov's.

A second related limitation is that the recurrent patterns are simply sequential in nature and do not yet address more complex "structure-dependent" features of natural human language. Although these results might be preliminary in this direction, recurrence analysis might be applied to more abstract descriptions of the syntactic patterns occurring between child and caregiver. One way in which this can be addressed is to incorporate such structures as Treebank analyses (Marcus et al., 1994) within utterances, and again subject these structures to recurrence analysis. The current results are at least a first step toward such an analysis.

Finally, we have targeted just syntactic patterns for measuring coordination. An important direction for future application of this technique is to discover the contribution of lexical sequences in generating this syntactic coordination. The current results suggest that there is coordination in syntactic n-grams while leaving open the possibility that lexically organized structures might contribute to this coordination early on. Numerous perspectives on language structure suggest this might be the case (see Tomasello, 2003, chap. 4, for a review). Dale and Spivey (2005) performed a lexical version of recurrence analysis and found that syntactic coordination is not a trivial form of lexical repetition. Nevertheless, it should be acknowledged that structures organized around lexical items could account for early coordination. Indeed, it would be interesting to further apply recurrence analysis to this problem by engaging argument structure organized around particular items to find coordination at the "locus" of syntactic patterns (Chouinard & Clark, 2003) and guided early on by "lexical islands" (Tomasello).

Despite these limitations, the current results are quite robust across parameter values, and the leading versus following results might offer future directions in studying the nature of child-caregiver interaction at differing levels of grammatical

development. The results are consistent with other discoveries of strong coordinative patterns in human communicative behavior. Whether eye movements (Richardson & Dale, 2005), postural adjustments in conversation (Shockley et al., 2003), or syntactic coordination between interlocutors (Branigan et al., 2000), there are many levels at which humans "synchronize" while communicating. Neurophysiological bases might also be in sight, such as mirror neurons underlying the ability to perceive and generate parallel action or goal sequences (Gallese, Keysers, & Rizzolatti, 2004), sometimes hypothesized as the evolutionary basis for human communication (e.g., Rizzolatti & Arbib, 1998). In any case, at the level of syntactic description, our results take a first step toward establishing the prevalence of strong coordination in language development.

Notes

[1]Markers for hesitations, pauses, trailings off, among other "nonsyntactic" elements were removed from the sequences. To focus on syntactic usage, errors uniquely coded (with a "$*$" marker) were also modified to appear as proper usage.

[2]Only patterns *within* sentences were evaluated. This was done to avoid having bigrams with end-of-sentence markers wash out other syntactic bigrams. There might be interesting patterns in end-of-sentence recurrence, however. See Dale and Spivey (2005), in which strong coupling across sentences is observed in a similar analysis.

References

Bandt, C. (2005). Ordinal time series analysis. *Ecological Modelling*, *182*, 229–238.

Bloom, K., Russell, A., & Wassenberg, K. (1987). Turn taking affects the quality of infant vocalizations. *Journal of Child Language*, *14*, 211–227.

Bohannon, J., & Stanowicz, L. (1988). The issue of negative evidence: Adult responses to children's language errors. *Developmental Psychology*, *24*, 684–689.

Branigan, H. P., Pickering, M. J., & Cleland, A. A. (2000). Syntactic coordination in dialogue. *Cognition*, *75*, B13–B25.

Broen, P. (1972). The verbal environment of the language-learning child. *Monographs of the American Speech and Hearing Association*, *17*.

Brown, R. (1973). *A first language: The early stages*. Cambridge, MA: Harvard University Press.

Casdagli, M. C. (1997). Recurrence plots revisited. *Physica D*, *108*, 12–44.

Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.

Chouinard, M. M., & Clark, E. V. (2003). Adult reformulations of child errors as negative evidence. *Journal of Child Language*, *30*, 637–669.

Church, K. W. (1993). Char_align: A program for aligning parallel texts at the character level. *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*, *31*, 1–8.

Clark, E. (2003). *First language acquisition*. Cambridge, UK: Cambridge University Press.

Clark, H. (1996). *Using language*. Cambridge, UK: Cambridge University Press.

Cleland, A., & Pickering, M. J. (2003). The use of lexical and syntactic information in language production: Evidence from the priming of noun-phrase structure. *Journal of Memory and Language*, *49*, 214–230.

Dagan, I., Lee, L., & Pereira, F. (1999). Similarity-based models of word cooccurrence probabilities. *Machine Learning*, *34*, 43–69.

Dale, R., & Spivey, M. J. (2005). Categorical recurrence analysis of child language. In *Proceedings of the 27th Annual Meeting of the Cognitive Science Society* (pp. 530–535). Mahwah, NJ: Lawrence Erlbaum.

Demetras, M., Post, K., & Snow, C. (1986). Feedback to first-language learners. *Journal of Child Language*, *13*, 275–292.

Eckmann, J.-P., Kamphorst, S. O., & Ruelle, D. (1987). Recurrence lots of dynamical systems. *Europhysics Letters*, *5*, 973–977.

Egghe, L. (2000). General study of the distribution of N-tuples of letters or words based on the distribution of the single letters of words. *Mathematical and Computer Modelling*, *31*, 35–41.

Fey, M. E., & Loeb, D. F. (2002). An evaluation of the facilitative effects of inverted yes-no questions on the acquisition of auxiliary verbs. *Journal of Speech, Language, and Hearing Research*, *45*, 160–174.

Gallese, V., Keysers, C., & Rizzolatti, G. (2004). A unifying view of the basis of social cognition. *Trends in Cognitive Sciences*, *8*, 396–403.

Garrod, S., & Pickering, M. J. (2004). Why is conversation so easy? *Trends in Cognitive Sciences*, *8*, 8–11.

Goldstein, M. H., King, A. P., & West, M. J. (2003). Social interaction shapes babbling: Testing parallels between birdsong and speech. *Proceedings of the National Academy of Sciences USA*, *100*, 8030–8035.

Goldstein, M. H., & West, M. J. (1999). Consistent responses of human mothers to prelinguistic infants: The effect of prelinguistic repertoire size. *Journal of Comparative Psychology*, *113*, 52–58.

Gomez, R. L., & Gerken, L. (2000). Infant artificial language learning and language acquisition. *Trends in Cognitive Sciences*, *4*, 178–187.

Hart, B., & Risley, T. R. (1995). *Meaningful differences in the everyday experience of young American children*. Baltimore, MD: Brooks.

Hirsh-Pasek, K., Treiman, R., & Schneiderman, M. (1984). Brown & Hanlon revisited: Mothers' sensitivity to ungrammatical forms. *Journal of Child Language*, *11*, 81–88.

Kuczaj, S. (1976). *-ing, -s and -ed: A study of the acquisition of certain verb inflections*. Unpublished doctoral dissertation, University of Minnesota.

Lewis, J. D., & Elman, J. L. (2001). *Learnability and the statistical structure of language: Poverty of stimulus arguments revisited*. Paper presented at the Proceedings of the 26th Annual Boston University Conference on Language Development (pp. 359–370).

Li, W. (2002). Zipf's law everywhere. *Glottometrics*, *5*, 14–21.

Lidz, J., Waxman, S., & Freedman, J. (2003). What infants know about syntax but couldn't have learned: Experimental evidence for syntactic structure at 18 months. *Cognition*, *89*, B65–B73.

MacWhinney, B. (2000). *The CHILDES project: Tools for analyzing talk*. Mahwah, NJ: Erlbaum.

Manning, C., & Schuetze, H. (1999). *Foundations of statistical natural language processing*. Cambridge, MA: MIT Press.

Maratsos, M., & Chalkley, M. A. (1980). The Internal Language of Children's Syntax: The ontogenesis and representation of syntactic categories. In K. Nelson (Ed.), *Children's language*: *Vol. 2*. New York: Gardner Press.

Marcus, G. (1993). Negative evidence in language acquisition. *Cognition*, *46*, 53–85.

Marcus, M., Kim, G., Marcinkiewicz, M. A., MacIntyre, R., Bies, A., Ferguson, M., et al. (1994). The Penn Treebank: Annotating predicate argument structure. In *Proceedings of the 1994 Human Language Technology Workshop* (pp. 110–115). Princeton, NJ: Morgan Kaufman.

Marwan, N., & Kurths, J. (2002). Nonlinear analysis of bivariate data with cross recurrence plots. *Physics Letters A*, *302*(5–6), 299–307.

Mintz, T. (2003). Frequent frames as a cue for grammatical categories in child directed speech. *Cognition*, *90*, 91–117.

Mitzenmacher, M. (2003). A brief history of generative models for power law and lognormal distributions. *Internet Mathematics*, *1*, 226–251.

Moerk, E. (1992). *A first language taught and learned*. Baltimore, MD: Brookes.

Moerk, E. (2000). *The guided acquisition of first-language skills*. Norwood, NJ: Ablex.

Morgan, J. L., Bonamo, K. M., & Travis, L. L. (1995). Negative evidence on negative evidence. *Developmental Psychology*, *31*, 180–197.

Morgan, J. L., & Travis, L. L. (1989). Limits on negative information in language input. *Journal of Child Language*, *6*, 531–552.

Newman, M. E. J. (2005). Power laws, Pareto distributions and Zipf's Law. *Contemporary Physics*, *46*, 323–351.

Newport, E., & Gleitman, L. (2002). The invention of language by children: Environmental and biological influences. In D. J. Levitin (Ed.), *Foundations of cognitive psychology: Core readings* (pp. 685–704). Cambridge, MA: MIT Press.

Newport, E., Gleitman, H., & Gleitman, L. (1977). Mother I'd rather do it myself: Some effects and non-effects of maternal speech style. In C. A. Ferguson (Ed.), *Talking to children: Language input and interaction*. Cambridge, UK: Cambridge University Press.

Pinker, S. (1994). *The language instinct*. New York: Harper.

Redington, M., Chater, N., & Finch, S. (1998). Distributional information: A powerful cue for acquiring syntactic categories. *Cognitive Science*, *22*, 425–469.

Reali, F., & Christiansen, M. H. (2003). *Reappraising poverty of the stimulus argument: A corpus analysis approach*. Online supplement to the Proceedings of the 28th Boston University Conference on Language Development.

Richardson, D. C., & Dale, R. (2005). Looking to understanding: The coupling between speakers' and listeners' eye movements and its relationship to discourse comprehension. *Cognitive Science*, *29*, 39–54.

Rizzolatti, G., & Arbib, M. (1998). Language within our grasp. *Trends in Neuroscience*, *21*, 188–194.

Sachs, J. (1983). Talking about the there and then: The emergence of displaced reference in parent–child discourse. In K. E. Nelson (Ed.), *Children's language*: *Vol. 4*. Hillsdale, NJ: Lawrence Erlbaum.

Saffran, J. R. (2003). Statistical language learning: Mechanisms and constraints. *Current Directions in Psychological Science*, *12*, 110–114.

Saxton, M. (1997). The contrast theory of negative input. *Journal of Child Language*, *24*, 139–161.

Saxton, M. (2000). Negative evidence and negative feedback: Immediate effects on the grammaticality of child speech. *First Language*, *20*, 221–252.

Shockley, K., Santana, M.-V., & Fowler, C. A. (2003). Mutual interpersonal postural constraints are involved in cooperative conversation. *Journal of Experimental Psychology: Human Perception and Performance*, *29*, 326–332.

Snow, C. E. (1972). Mother's speech to children learning language. *Child Development*, *43*, 549–565.

Sokolov, J. (1993). A local contingency analysis of the fine-tuning hypothesis. *Developmental Psychology*, *29*, 1008–1023.

Sokolov, J., & MacWhinney, B. (1990). The CHIP framework: Automatic coding and analysis of parent-child conversational interaction. *Behavioral Research Methods, Instruments, and Computers*, *22*, 151–161.

Tamis-LeMonda, C. S., & Bornstein, M. H. (2002). Maternal responsiveness and early language acquisition. *Advances in Child Development and Behavior*, *29*, 89–127.

Tomasello, M. (2003). *Constructing a Language: A usage-based theory of language*. Cambridge, MA: Harvard University Press.

Van Orden, G. C., Holden, J. G., & Turvey, M. T. (2003). Self-organization of cognitive performance. *Journal of Experimental Psychology: General*, *132*, 331–350.

Von Heijne, G. (1987). *Sequence analysis in molecular biology: Treasure trove or trivial pursuits?* San Diego: Academic Press.

Webber, C. L., Jr., & Zbilut, J. P. (2005). Recurrence quantification analysis of nonlinear dynamical systems. In M. A. Riley & G. C. Van Orden (Eds.), *Tutorials in contemporary nonlinear methods for the behavioral sciences* (pp. 33–101). Retrieved March 1, 2005, from http://www.nsf.gov/sbe/bcs/pac/nmbs/nmbs.jsp

Zbilut, J. P., & Webber, C. L., Jr. (1992). Embeddings and delays as derived from quantification of recurrence plots. *Physics Letters A*, *171*, 199–203.

Appendix

Recurrence Analysis and Batch Measures

Although perhaps straightforward to readers familiar with NLP or recurrence techniques, it is of value to demonstrate the direct relationship between recurrence structures (i.e., RPs) and the batch measures considered in the text. Recurrence analysis can be treated as a "peeling" of batch measures *in time*. Consider again the set B_T of all syntactic bigram *types* occurring in

one transcript between child and caregiver. Rather than simply counting the occurrences of a bigram b_i in caregiver speech, imagine building sets $t.i._m$ and $t.i._k$ such that each element represents a point in time at which this bigram occurs:

$$t.i._m(b) = \{i \mid b = (s_{m,i}, s_{m,i+1})\}$$

The same can be done for the bigram b in the child's time series. The relevant RP of these time series can be defined as the union of all Cartesian products of $t.i._m$ and $t.i._k$ over B_T:

$$\text{RP} = \bigcup^{B_T} t.i._m(b) \times t.i._k(b)$$

Any such bigram b has the frequency $f_k(b)$ in the child's usage and $f_m(b)$ in the caregiver's usage. All batch measures are the sum over some function of these frequency values (or their corresponding probabilities): $\sum F[f_k(b), f_m(b)]$. Recurrence analysis is a "peeling" in time, as $f_m(b) = \|t.i._m(b)\|$, and any such batch measure can be expressed simply as a function over the collapsed time indices in which bigrams of B_T occur. In this case, sometimes referred to as "ordinal" or "categorical" (Bandt, 2005; Dale & Spivey, 2005), the situation is relatively simple. In a continuous time series of complex systems, the informational richness of the RP has been given a detailed technical treatment (Casdagli, 1997).