A longitudinal analysis of children's communicative acts

Claire Augusta Bergey (cbergey@wisc.edu)

Misha O'Keeffe (meokeeffe@wisc.edu)

Robert Hawkins (rdhawkins@wisc.edu)
Department of Psychology, University of Wisconsin-Madison
1202 W. Johnson Street
Madison, WI 53706 USA

Abstract

Children rapidly learn to use language to effect a variety of *communicative acts*, such as proposing actions, asking questions, and making promises. While prior work has characterized this development in cross-sectional corpora, these analyses have been unable to comprehensively track individual differences in children's acquisition of communicative acts. We analyzed a longitudinal corpus of parent-child interactions from ages 14 to 58 months. We find that children's repertoires of communicative acts diversify over this period, with stable individual differences in the diversity of children's communicative act repertoires. Further, the diversities of parents' and children's communicative act repertoires are correlated. Children with more diverse communicative act repertoires also have larger vocabularies and use more diverse syntactic frames, suggesting links between discourse development and lexical and syntactic knowledge. Taken together, this work provides new insight into individual trajectories of communicative development and connections between communicative act use and other levels of language structure.

Keywords: language development, discourse, pragmatics, corpus studies, natural language processing

Introduction

Natural conversation involves an intricate dance of different *communicative acts*. We request and grant permission, ask and answer questions, and make and refuse suggestions (Searle, 1969; Austin, 1975). The ability to appropriately initiate and respond to a full repertoire of communicative acts is crucial for getting things done in social interaction and therefore plays a foundational role in theories of language development (Bruner, 1975, 1981; Tomasello, 2010; Zufferey, 2014). Some capacity for initiating different communicative acts may be present pre-verbally (e.g. proto-imperatives to request objects or proto-declaratives to establish attention; Bates, Camaioni, & Volterra, 1975) and continues to develop into early childhood (Cameron-Faulkner, 2014; Casillas & Hilbrink, 2020). For example, an adult-like understanding of *promising* may not emerge until relatively late (Astington, 1988; Bernicot & Laval, 2004).

While considerable effort has been put toward characterizing variability and consistency in lexical and grammatical development (Fenson et al., 1994; Frank & Yurovsky, 2021; Kidd & Donnelly, 2020), it has been more difficult to characterize trajectories of *discourse*-level development. One classic longitudinal study (Snow, Pan, Imbens-Bailey, & Herman, 1996, which examined children from 14 to 32 months) quantified the growth of children's repertoires over development, but was limited to few, short sessions over a constrained age range. Thus, they were unable to measure more precise trajectories of children's communicative act repertoires.

Several recent efforts have aimed to scale up analyses of discourse-level development using recent breakthroughs in natural language processing (NLP). By automating communicative act labeling for parent-child conversations, these studies have allowed analysis of many more transcripts across many more children and ages (Nikolaus, Maes, Auguste, Prévot, & Fourtassi, 2022; Bergey, Marshall, DeDeo, & Yurovsky, 2022). They report finer-grained findings, such as estimates of the precise ages at which different communicative acts are acquired and the precise expansion of children's communicative act repertoires. However, these studies have relied on CHILDES, a primarily cross-sectional corpus of parent-child conversations (i.e., most children in the corpus are not followed longitudinally). Cross-sectional data make it challenging to measure individual differences in children's and dyads' communicative act use, as differences at single timepoints may reflect session-level factors.

In this paper, we extend the large-scale study of children's communicative acts to a large longitudinal corpus that includes naturalistic interactions from the same group of children as they develop from 14 to 58 months old. In Part I, we establish basic quantitative facts about parents' and children's communicative act use over development. First, we replicate findings from cross-sectional studies, showing that children's communicative act repertoires rapidly expand over this age range. We also show that parents' and children's communicative act repertoires become more similar as children grow. Further, leveraging the longitudinal nature of these data, we show that there are individual-level and dyad-level differences in the diversity of communicative acts children and parents use: individuals' communicative act use has some stability over time, and parent-child dyads tend to use similarly diverse communicative acts.

In Part II, we examine how communicative act use corresponds to other levels of language structure. Prior work has suggested that the size of one's communicative act repertoire correlates with productive vocabulary measures (Snow et al., 1996). We replicate and extend this finding, showing that the diversity of children's communicative act repertoires is correlated with their productive vocabulary in naturalistic speech, an independent vocabulary test, and their syntactic frame di-

In L. K. Samuelson, S. L. Frank, M. Toneva, A. Mackey, & E. Hazeltine (Eds.), *Proceedings of the 46th Annual Conference of the Cognitive Science Society.*

versity in naturalistic speech. Next, we examine how syntactic frames and communicative acts correspond in parents' and children's speech. Prior work has suggested that there is a tight form-function relationship in parents' speech to children, enabling children to more easily infer communicative intent from utterance form (Shatz, 1979; Ninio, 1992; Cameron-Faulkner & Hickey, 2011). We show that parents' speech has a high degree of correspondence between syntactic frame and communicative act type, constraining this inference problem, and that children approach adult-like levels of this correspondence over development. Taken together, our findings shed light on the individual paths children trace in learning communicative acts and their relationship to other levels of language structure.

Methods

LDP corpus. The Language Development Project (LDP) is a multi-year longitudinal corpus of caregiver-child interactions in the home (Goldin-Meadow et al., 2014). The LDP includes 64 typically-developing children in English-speaking households. Families were selected to be representative of the Chicagoland area in terms of race and socio-economic status. Researchers visited participants' homes and recorded a 90-minute session of their behavior at home every four months from when the child was 14 months to 58 months old (for a total of 727 sessions, with missing sessions from some participants; all families completed at least four sessions). We included only sessions in which the child and a caregiver each spoke at least 10 utterances, totaling 959,956 utterances across 692 sessions among 64 children. The LDP corpus is currently undergoing anonymization and is not yet publicly available. Communicative act label sequences and analysis code are available at `https://github.com/cbergey/communicative-acts`.

Annotating communicative acts. Recent advances in natural language processing (NLP) have made it possible to characterize children's communicative act development at scale. Progress labeling communicative acts in adult conversation (Stolcke et al., 2000; Boyer et al., 2010; Kumar, Agarwal, Dasgupta, & Joshi, 2018) laid the groundwork for automatically identifying communicative acts in children's conversations. However, classifying children's communicative acts comes with unique challenges given that the utterance forms children use to express different communicative acts are rapidly changing over development, and are often different from those used by adults.

Two recent papers have introduced models that classify children's communicative acts. Bergey et al. (2022) introduced an unsupervised model that automatically discovered a set of 15 communicative act types, and Nikolaus et al. (2022) introduced a supervised model based on the INCA-A coding system (Ninio, Snow, Pan, & Rollins, 1994). These papers characterized caregiver-child interactions in CHILDES, a primarily cross-sectional corpus, and provided converging evidence that children's communicative act repertoires expand rapidly over early childhood.

We classify communicative acts in the LDP using the model from Nikolaus et al. (2022). This model is preferable for our purposes because it classifies communicative acts according to INCA-A, allowing the model's classifications to be verified against human annotations and to be more easily compared to those in other work. The INCA-A includes 67 communicative act categories, ranging from common acts (*declarative statement, propose action, ask yes/no question*) to much more infrequent ones (*threaten to do, criticize nonverbal act, give intentionally non-satisfying answer*).

The tagging model developed by Nikolaus et al. (2022) is a Conditional Random Field model trained on the Snow et al. (1996) corpus, a corpus of parent-child interactions with hand-labeled communicative acts from children at 14, 20, and 32 months old. Specifically, it was trained to classify communicative act type based on the following features: the speaker (adult or child), the unigrams and bigrams (words and two-word combinations) in the utterance, the number of words in the utterance, the part of speech tags of the utterance, and the number of repeated words from the last utterance. It also captures transition probabilities between act types, constraining predictions of each utterance's act type using the act type of the prior utterance. The 40 most common communicative act types children produced and the distribution of ages at which they first produced them in the corpus are shown in Figure 1.

Validation. Nikolaus et al. (2022) report that their model has 72% communicative act classification accuracy on held-out portions of the Snow corpus, its training corpus. However, it is possible that the model would not generalize to other corpora—in particular, to children outside the training set's age range, which the LDP corpus includes.

We tested the model's performance on the LDP by sampling 10 utterances from each included session of each child (3,360 utterances) and having two human raters hand code them. The two human coders had 58% label agreement. The model's agreement with each coder was 40% and 44%; considering whether the model chose a label that *either* human coder chose, the model agreed with at least one of the coders 53% of the time. Model–human coder agreement did not vary substantially by child age: agreement with either coder ranged from 51% to 57%.

Though these agreement values are well above chance given the large label space of communicative acts, they are well below the validation accuracy found by Nikolaus et al., 2022 and the 80% human interrater reliability reported by Snow et al., 1996. This may be due to the fact that several labels are often appropriate for a given utterance, and reflect challenges in recreating the precise coding method of Ninio et al. (1994); it may also reflect that some aspects of the Snow corpus do not generalize to other corpora. In sum, we see lower model accuracy when generalizing to this new corpus, though the accuracy is not far from human agreement levels and the drop in agreement is not attributable to the wider age range of the LDP corpus.

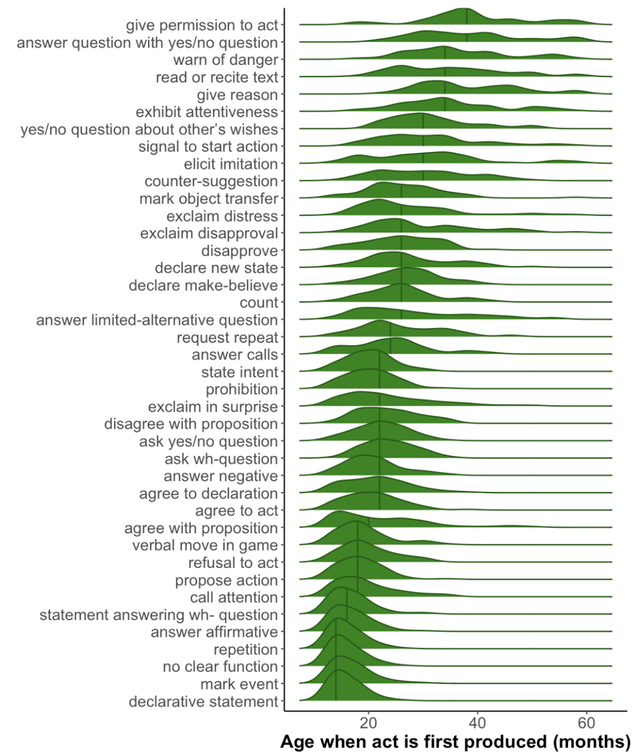

Figure 1: The first age at which children produce communicative acts. The 40 most common acts produced by children are shown here, ranked by median age of first production on the y-axis. Ridges show density of children first producing that act at that age; vertical lines indicate the median age at which an act is first produced.

Part I: Characterizing communicative act trajectories

Children's repertoires diversify over development. To measure the expansion of children's communicative act repertoires over development we calculate their *entropy*. Entropy measures the variability of children's communicative act use. For instance, a child who primarily uses one communicative act type and rarely uses any others would have low communicative act entropy, while a child who uses several communicative act types in more equal proportion would have higher communicative act entropy. Children's communicative act entropies increase significantly over development, demonstrating that children's repertoires diversify over the 14-month to 58-month range (Figure 2). We confirm this observation with a mixed-effects linear model predicting children's entropies from (log) age with a by-subject intercept, finding a significant effect of age ($\beta = 0.66, p < 0.001$). Meanwhile, parents' communicative act entropies are relatively stable and change only slightly over development; we confirm this observation with a mixed-effects linear model predicting parents' entropies from (log) child's age and a by-subject intercept ($\beta = 0.05, p = 0.006$). Modeling children's and parents' data together, we find a significant interaction

Utterance	Communicative act
C: Wash dolly's hair.	Declarative statement
P: You going to wash her hair?	Ask yes/no question
C: Yeah.	Answer affirmative
C: It's all dirty.	Declarative statement
P: Mhmm.	Agree w/proposition
P: Let's get a bowl.	State intent
P: If we're going to do that.	State intent
P: And let's put a towel down.	State intent
P: So we don't have water all over.	Give reason

Table 1: An example conversation between a 34-month-old and parent from the corpus. Communicative act types are labeled using the model from Nikolaus et al. (2022), an automated classifier based on the INCA-A coding scheme. C designates child speaker, P designates parent speaker.

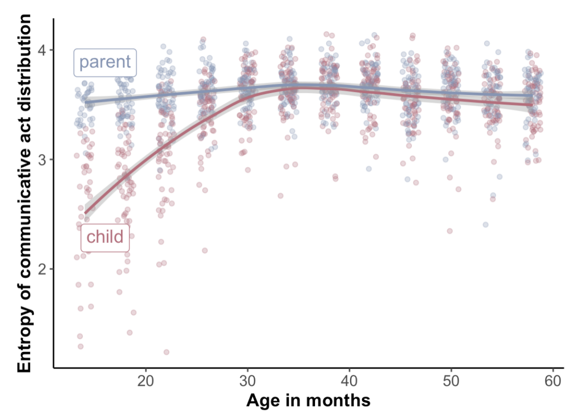

Figure 2: Children's and parents' communicative act entropies over development. Children use an increasingly diverse distribution of communicative acts while parents' distributions are relatively stable.

between age and speaker type, verifying that children's entropies change more over development than parents' do ($\beta = 0.60, p < 0.001$). To put these repertoires in more concrete terms, out of the 67 act types in the INCA-A, the median 14-month old is using 10 of them, the median 22-month old is using 20, and median 30-month-olds and older are using 26 or more.

Another way to characterize the change in children's communicative act repertoires is to examine the difference between children's distributions of communicative acts and their parents'. Using Jensen-Shannon distance (JSD), we calculated the difference between children's and parents' distributions at each timepoint. Using a mixed-effects linear model predicting JSD from (log) age and a by-subject intercept, we find that JSDs significantly decrease over development ($\beta = -0.23, p < 0.001$). That is, children's repertoires not only diversify but become more similar to their parents' over development (Figure 3).

Individual differences in communicative act repertoires. Prior cross-sectional work on parents' and children's communicative act repertoires has been unable to disentangle session-to-session variability in communicative act repertoires from stable individual differences. To examine whether children evince individual differences in their repertoires, we calculated the JSD between a child's communicative act distribution at each timepoint and their own distribution at each other timepoint. We then generated a null distribution of JSDs by calculating the JSD between random pairs of children at ages matched to those in the self-paired set. If children have stable individual differences in their communicative act distributions, the JSDs of self-paired child distributions should be lower than the JSDs of randomly paired child distributions.

We find that this is true: children's communicative act use is more similar to their own communicative act use at other timepoints than to other children's at those same other timepoints (mean self JSD = 0.178, mean random JSD = $0.200, p < 0.001$ by t-test). We also find that parents' communicative act use is more similar to their own communicative act use at other timepoints than to other parents' at those same other timepoints (mean self JSD = 0.081, mean random JSD $= 0.108, p < 0.001$). Both parents and children have small but stable individual differences in communicative act use over development.

To make individual differences in children's communicative act use more concrete, we selected two children whose communicative act distributions are most different on average across all sessions to compare (the average JSDs between their distributions at each age was highest); call them Child A and Child B. On average across all sessions, Child A more often has utterances with no clear function (26% of utterances), more often proposes actions (15%), and more often states intent (11%) than Child B does (8%, 8%, and 5% respectively). Child A also produces act types that Child B never does, such as expressing approval and giving permission to act. Child B more often makes declarative statements (33%), answers yes/no questions (11%), and answers affirmatively (6%) than Child A does (20%, 2%, and 1% respectively). Child B also produces act types that Child A never does, such as exclaiming distress and completing a statement on request.

Examining the parents of Child A and Child B, we find that Parent A more often proposes actions (27%), makes prohibitions (4%) and states intent (4%) compared to Parent B (13%, 1%, and 1% respectively). Parent B more often asks yes/no questions (17%), agrees with propositions (7%) and marks events (8%) than Parent A (10%, 3%, and 7% respectively). Parent B also produces act types that Parent A never does, such as counting, expressing sympathy, and praising for nonverbal acts. These dyads represent extreme examples, but give a tangible sense of individual differences in communicative act use.

Children with more diverse repertoires have parents with more diverse repertoires. What role do caregivers play in children's acquisition and deployment of communicative acts? Bergey et al. (2022) found that the entropies of parents' and children's communicative act repertoires are correlated—parents who use more diverse communicative acts are talking to children who use more diverse communicative acts. However, this finding used primarily cross-sectional data, meaning they could not establish whether these correlations reflected session-level factors (e.g., the activity dyads were engaged in during a session may constrain the diversity of communicative acts both caregiver and child deployed) or reflected stable, dyad-level differences.

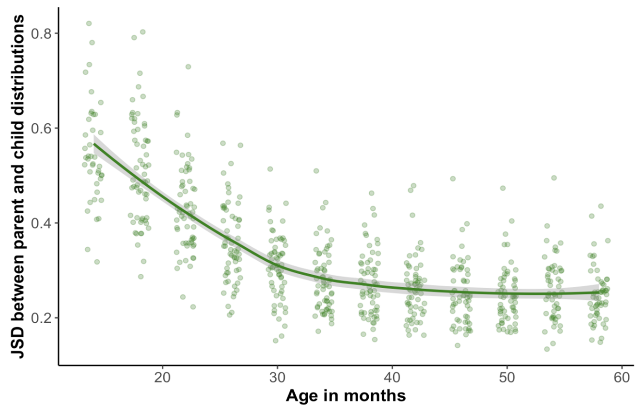

Figure 3: Plot of the Jensen-Shannon distance between children's and parents' distributions of communicative acts over development. Children's and parents' communicative act distributions become more similar as children grow.

First, replicating prior results on cross-sectional data, we find that caregiver entropy and child entropy tend to be correlated within sessions (significantly correlated in 9 out of 12 sessions; lowest correlation at 14 months, $r = 0.02, p = 0.88$, highest correlation at 38 months, $r = 0.56, p < 0.001$; average correlation $= 0.35$; correlations throughout are Pearson's correlation coefficient). Leveraging this longitudinal corpus, we can ask whether this holds even when we examine the correlations between caregivers' and children's average entropies across all sessions. It does: children's average entropies across all sessions are correlated with caregivers' average entropies across all sessions ($r = 0.41, p < 0.001$). Thus there is a stable dyad-level correspondence in diversity of communicative act use, such that parents-child pairs tend to have similar communicative act entropies.

There are several plausible causal explanations for why caregivers' and children's communicative act entropies are consistently correlated. Caregivers who use more diverse communicative acts may allow children to *learn* a more extensive communicative act repertoire over time, or give them more opportunities to *deploy* diverse communicative acts in the moment. Children who use more diverse communicative acts also may elicit more diverse communicative acts from their parents. Finally, parents and children may have correlated communicative act repertoires due to outside factors (e.g., genetic or environmental factors they share but that do not involve a causal relationship between their repertoires).

Though this observational dataset cannot definitively adjudicate between these causal explanations, because we have several observations of the same dyads, we can gain insight into which explanations are most likely by examining the relationship between caregivers' and children's repertoires over time.

To do this, we fit a mixed effects linear model predicting children's communicative act entropies at each timepoint from several factors: the child's own communicative act entropy in the prior session, the parent's communicative act entropy in the prior session, the parent's communicative act entropy in the current session, (log) child's age, and a by-subject intercept. The factors that significantly predicted children's communicative act entropy were their own communicative act entropy in the prior session ($\beta = 0.40, p < 0.001$), their parent's entropy in the current session ($\beta = 0.48, p < 0.001$), and (log) age ($\beta = 0.11, p = 0.004$). Parents' entropy in the prior session was not significantly predictive ($\beta = -0.05, p = 0.40$). That is, we primarily find evidence that children have consistent individual differences in communicative act entropy over time and that their communicative act entropy corresponds to that of their current conversational partner, and do not find that the parent's communicative act distribution in the prior session is predictive of the child's current communicative act distribution over and above those factors. However, as demonstrated above, all of these factors are correlated and compete to account for the same variance in children's repertoires. Our findings suggest child-level differences and within-session elicitation as key places to look to explain communicative act use, but in no way eliminate prior parent input as a causal factor.

Part II: Correspondence to other levels of language structure

Taken together, the findings from Part I show that children's communicative act use becomes more diverse over development, that there are stable individual- and dyad-level differences in communicative act use, and that children with more diverse repertoires tend to have parents with similarly diverse repertoires. Having isolated communicative act use to trace its trajectories, we now turn to its integral role in a developing language system with multiple levels of structure. Children are rapidly gaining competence across several aspects of language form during the preschool years, including their knowledge of words and syntactic structures. How does communicative act ability relate to these other language abilities?

Communicative act diversity correlates with vocabulary. We begin by considering the relationship between communicative act diversity and vocabulary size. We calculated children's total productive vocabulary size (the number of word types they produced) across all sessions. Children's total productive vocabulary across all sessions was significantly correlated with their average communicative act entropy across all sessions ($r = 0.37, p = 0.003$). As a stronger test, we fit a mixed effects linear model predicting children's productive vocabulary size at each session from their own communicative act entropy at that session, their parent's communicative act entropy at that session, their (log) age, and a by-subject intercept. Children's communicative act entropies, parents' communicative act entropies, and (log) age were all significant predictors of child productive vocabulary size ($\beta_{child-entropy} = 16.87, p = 0.04$, $\beta_{adult-entropy} = 68.03, p < 0.001$, $\beta_{age} = 289.93, p < 0.001$).

Given that communicative act entropy and productive vocabulary are measured from the same transcripts, one may worry that aspects of the conversation (the ongoing activity or other factors) are driving the correspondence between productive vocabulary and communicative act diversity, or that the productive vocabularies children demonstrate in these conversations are not representative of their true vocabulary size. The LDP corpus affords another, independent measure of vocabulary: children were tested on the Peabody Picture Vocabulary Test (PPVT) at 54 months old (Dunn & Dunn, 1965). This provides an independent and validated measure of vocabulary, late in our age range, as an outcome. We find that children's PPVT scores are significantly correlated with their average entropy over all sessions ($r = 0.29, p = 0.03$). Thus, both on a within-conversation measure and an independent one, children's vocabularies and communicative act entropies are significantly correlated.

Communicative act diversity correlates with syntactic ability. How does children's communicative act ability relate to their syntactic ability? To answer this question, we transformed each utterance into a frame and slot pattern by retaining any word in each utterance that is in the top 150 most common words in the corpus, and transforming all other words in each utterance into part of speech tags. This results in syntactic frames such as "Where's the [noun]?", "I [verb] it," and "Are you [verb]?" These are meant to capture syntactic structures that children have at their disposal, lumping over more particular semantic content.

Do children with more diverse communicative act repertoires also have more diverse syntactic frame repertoires? We find that at young ages they do: within session, children's syntactic frame entropy is significantly positively correlated with their communicative act entropy at all sessions from 14 to 38 months (highest correlation at 18 months, $r = 0.73, p < 0.001$; mean correlation over all sessions, $r = 0.29$) though they are not significantly correlated at older ages. Children's mean communicative act entropy and mean syntactic frame entropy across all sessions are correlated ($r = 0.40, p = 0.001$). Though we do not have independent measures of syntactic ability to strengthen this claim (as with vocabulary), this supports the idea that syntactic ability and communicative act ability are related, particularly at younger ages.

Correspondence between communicative acts and syntactic frames. A single communicative act can be achieved with many utterance forms: for instance, one can propose an action by saying "Get up," "Let me see it," or "Would you go?" The ways children correspond communicative acts with

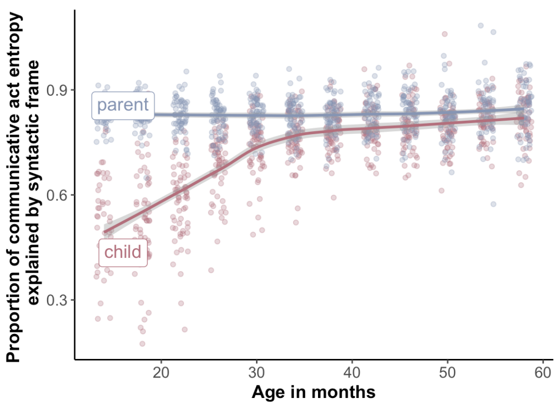

Figure 4: Plot of the proportion of communicative act entropy explained by syntactic frame (mutual information between communicative acts and syntactic frames divided by communicative act entropy).

syntactic frames tell us about how they are learning to express communicative intents with different utterance forms. The ways caregivers correspond communicative acts with syntactic frames tell us about the inference problem children face when trying to understand the communicative intent of an utterance. This correspondence is not always straightforward: for instance, "Can you get up?" is usually a proposed action rather than a yes/no question, but can be either.

We examined the correspondence between syntactic frames and communicative acts by calculating their mutual information, which measures how much knowing one piece of information (an utterance's syntactic frame) reduces uncertainty about another (the utterance's communicative act) and vice versa. In this analysis, we included only syntactic frames that occurred more than 10 times in the entire corpus, to exclude extremely rare frames and reduce the sparsity of frame–act correspondence. We find that the mutual information between syntactic frames and communicative acts in children's speech increases over development, and in parents' speech stays roughly constant. To confirm this, we separately fit mixed effects linear models predicting mutual information from (log) age and a by-subject intercept for children and parents; (log) age significantly predicted mutual information for children ($\beta = 1.19, p < 0.001$) and a very slight change for parents ($\beta = 0.06, p < 0.001$).

Mutual information is bounded by the entropy of the lower-entropy distribution, in this case the communicative act distribution, which is also changing over development (as established in Part I). Thus, it is perhaps more meaningful to examine not just overall uncertainty reduction (mutual information) but the proportion of overall uncertainty that is reduced (mutual information divided by the entropy of communicative act use). Roughly, this expresses the proportion of diversity in communicative acts that is explained by syntactic frames: when you hear a specific syntactic frame, how much certainty do you gain about the communicative intent of your interlocutor?

For caregivers, on average, 83% of variability in communicative act use is explained by syntactic frame (Figure 4). For children, around 53% of variability in communicative act use is explained by syntactic frame at 14 months, reaching 76% by 34 months. Statistically confirmed as above, we find that the diversity in communicative acts explained by frames in children's speech increases over development ($\beta = 0.24, p < 0.001$) and in parents' speech does not change significantly over development ($\beta = 0.01, p = 0.07$). Thus, caregiver speech has a relatively tight correspondence between syntactic frames and communicative acts that is stable over development, providing a constrained problem for children learning to infer communicative intent from utterance form. Children approach an adult-like level of correspondence over the preschool years.

Discussion

Our work introduces a comprehensive large-scale longitudinal analysis of children's communicative acts. By examining children's communicative act development longitudinally, we were able to examine the individual communicative trajectories of children and child-caregiver dyads. We found that children's act repertoires expand quickly and that children's and parents' distributions of communicative acts become more similar over development. We also found that there are stable individual differences in communicative act repertoires, and that the diversity of communicative acts children use predicts their vocabulary size and diversity of syntactic frames. Further, we characterized the correspondence between utterance form and communicative act type in both parents' and children's speech, showing that parents have a tight form-function relationship that simplifies children's problem of inferring communicative intent from utterance form, and that children approach this level of correspondence over our age range. Taken together, we characterize communicative act use as a central language skill on which children stably differ, through which dyads echo one another, and which is integrally connected to other levels of language structure.

The findings presented here are constrained by current limitations of automated communicative act labeling. There may be certain communicative acts that the model struggles to capture because they are rare or the relationship between the utterance forms and act type is unreliable, and this may lead to higher correspondence between utterance form and communicative acts labeled by the model than those that would be labeled by a human. Further, this approach does not account for communicative acts children do not achieve through language: children often make communicative 'moves' through gesture, and there may be subtle prosodic markers of communicative act use that are not captured in these transcripts. Fully taking advantage of the nuanced and multimodal nature of children's communication would require training communicative act classifiers on large multimodal corpora, a future endeavor that will be central to advancing our understanding of children's communicative development.

References

Astington, J. W. (1988). Children's understanding of the speech act of promising. *Journal of Child Language*, *15*(1), 157–173.

Austin, J. L. (1975). *How to do things with words*. Oxford University Press.

Bates, E., Camaioni, L., & Volterra, V. (1975). The acquisition of performatives prior to speech. *Merrill-Palmer quarterly of behavior and development*, *21*(3), 205–226.

Bergey, C., Marshall, Z., DeDeo, S., & Yurovsky, D. (2022). Learning communicative acts in children's conversations: A hidden topic markov model analysis of the childes corpora. *Topics in Cognitive Science*, *14*(2), 388–399.

Bernicot, J., & Laval, V. (2004). Speech acts in children: The example of promises. In *Experimental pragmatics* (pp. 207–227). Springer.

Boyer, K., Ha, E. Y., Phillips, R., Wallis, M., Vouk, M., & Lester, J. (2010). Dialogue act modeling in a complex task-oriented domain. In *Proceedings of the sigdial 2010 conference* (pp. 297–305).

Bruner, J. S. (1975). The ontogenesis of speech acts. *Journal of child language*, *2*(1), 1–19.

Bruner, J. S. (1981). The social context of language acquisition. *Language & Communication*.

Cameron-Faulkner, T. (2014). The development of speech acts. *Pragmatic development in first language acquisition*, 37–52.

Cameron-Faulkner, T., & Hickey, T. (2011). Form and function in irish child directed speech.

Casillas, M., & Hilbrink, E. (2020). 3. communicative act development. *Developmental and clinical pragmatics*, *13*, 61.

Dunn, L. M., & Dunn, L. M. (1965). Peabody picture vocabulary test–.

Fenson, L., Dale, P. S., Reznick, J. S., Bates, E., Thal, D. J., Pethick, S. J., … Stiles, J. (1994). Variability in early communicative development. *Monographs of the society for research in child development*, i–185.

Frank, M. C., & Yurovsky, D. (2021). *Variability and consistency in early language learning*. MIT Press.

Goldin-Meadow, S., Levine, S. C., Hedges, L. V., Huttenlocher, J., Raudenbush, S. W., & Small, S. L. (2014). New evidence about language and cognitive development based on a longitudinal study: hypotheses for intervention. *American Psychologist*, *69*(6), 588.

Kidd, E., & Donnelly, S. (2020). Individual differences in first language acquisition. *Annual Review of Linguistics*, *6*, 319–340.

Kumar, H., Agarwal, A., Dasgupta, R., & Joshi, S. (2018). Dialogue act sequence labeling using hierarchical encoder with crf. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 32).

Nikolaus, M., Maes, E., Auguste, J., Prévot, L., & Fourtassi, A. (2022). Large-scale study of speech acts' development in early childhood. *Language Development Research*, *2*(1), 268–304.

Ninio, A. (1992). The relation of children's single word utterances to single word utterances in the input. *Journal of Child Language*, *19*(1), 87–110.

Ninio, A., Snow, C. E., Pan, B. A., & Rollins, P. R. (1994). Classifying communicative acts in children's interactions. *Journal of communication disorders*, *27*(2), 157–187.

Searle, J. R. (1969). *Speech acts: An essay in the philosophy of language*. Cambridge, UK: Cambridge University Press.

Shatz, M. (1979). How to do things by asking: Form-function pairings in mothers' questions and their relation to children's responses. *Child Development*, 1093–1099.

Snow, C. E., Pan, B. A., Imbens-Bailey, A., & Herman, J. (1996). Learning how to say what one means: A longitudinal study of children's speech act use. *Social Development*, *5*(1), 56–84.

Stolcke, A., Ries, K., Coccaro, N., Shriberg, E., Bates, R., Jurafsky, D., … Meteer, M. (2000). Dialogue act modeling for automatic tagging and recognition of conversational speech. *Computational linguistics*, *26*(3), 339–373.

Tomasello, M. (2010). *Origins of human communication*. MIT press.

Zufferey, S. (2014). *Acquiring pragmatics: Social and cognitive perspectives*. Routledge.