Development of Multi-level Linguistic Alignment in Child-Adult Conversations

Thomas Misiek Benoit Favre Abdellah Fourtassi

Aix-Marseille Univ, Universite de Toulon, CNRS, LIS, ILCB, Marseille, France
{thomas.misiek, benoit.favre, abdellah.fourtassi}@lis-lab.fr

Abstract

Interactive alignment is a major mechanism of linguistic coordination. Here we study the way this mechanism emerges in development across the lexical, syntactic, and conceptual levels. We leverage NLP tools to analyze a large-scale corpus of child-adult conversations between 2 and 5 years old. We found that, across development, children align consistently to adults above chance and that adults align consistently more to children than vice versa (even controlling for language production abilities). Besides these consistencies, we found a diversity of developmental trajectories across linguistic levels. These corpus-based findings provide strong support for an early onset of multi-level linguistic alignment in children and invite new experimental work.

1 Introduction

Linguistic alignment is the tendency that interlocutors have to change the way they talk to accommodate their conversational partners. This can happen through mirroring the partner’s linguistic behavior on many levels such as the choice of words, syntactic structures, and semantic topics. Linguistic alignment is considered an important mechanism for establishing common ground and rapport, fostering successful communicative interactions (Clark, 1996). In addition, understanding this coordination in its natural context is crucial for the design of conversational systems that interact with people in a natural and effective fashion (Zhao et al., 2016; Loth et al., 2015; Park et al., 2017).

While alignment has been largely studied with adults (Pickering and Garrod, 2004; Fusaroli et al., 2012; Dale et al., 2013; Doyle and Frank, 2016; Dideriksen et al., 2019), little has been done to investigate how it manifests in the context of child-adult early communication and how it evolves across development. This is a significant gap in

the literature. The child-adult early communication cannot be thought of as a simple extension of conversational dynamics between adults; it involves strong asymmetries in terms of cognitive abilities and social roles and, thus, requires more dedicated research (Clark, 2015). In addition, the study of child-caregiver linguistic interaction informs our theories of children’s cognitive development. On the one hand, children’s developing abilities in managing a conversation — through mechanisms such as interactive alignment — is a window into their emerging social-cognitive skills (Tomasello, 2009). On the other hand, the way caregivers use alignment across development allows us to understand whether and how adults tune their talk to children’s developing cognitive abilities. Such tuning has been suggested to play a pedagogical role, supporting linguistic and conceptual learning (Snow, 1972; Fourtassi et al., 2014, 2019).

Related Work and Novelty of the Current Study

Our study investigates children’s interactive alignment in natural conversations with adults. Previously, Dale and Spivey (2006) used recurrence analysis to investigate child-caregiver syntactic alignment (operationalized as sequences of parts of speech) and found evidence for syntactic coordination. Using a similar computational framework, Fernández and Grimm (2014) extended Dale and Spivey’s findings to the lexical and conceptual levels. Nevertheless, both studies were based on data from three children only. While such a small sample size allows for a detailed examination of development for specific children, it does not allow us to characterize general developmental patterns that could be shared by the majority of children. Indeed, both studies found large individual variability and, thus, no strong conclusions about development could be drawn.

In a more recent work, Yurovsky et al. (2016) studied a large-scale corpus of child-caregiver interactions containing two orders of magnitude more children than previous work. Using hierarchical Bayesian models, they found that both children and caregivers decreased their alignment over the first five years of development. Work by Yurovsky et al. (2016) thus provided a much more robust test of interactive alignment. However, it focused on the special case of function words. It is still an open question how development unfolds across the entire lexicon and along more abstract levels such as syntax and semantics. The current study is an effort to fill this gap in the literature. We leverage NLP tools to test interactive alignment at the lexical, syntactic, and conceptual levels, using a large-scale corpus of children’s natural language.

2 Data and Methods

2.1 Data

Following Yurovsky et al. (2016), and in order to maximize the statistical power and generalizability of our analysis, we selected all English-language data available in CHILDES (MacWhinney, 2000) for children between 2 to 5 years; a time of rapid growth in terms of expressive language and social-cognitive skills (Wellman, 2014). This resulted in 5,152 total transcripts across 725 unique children. The number of transcripts per child varied between 1 ($N = 381$) and 503 ($N = 1$), with an average of about 7 transcripts per child.

2.2 Levels of Linguistic Alignment

We study linguistic alignment, on a turn-by-turn basis. We measure alignment along the lexical, syntactic, and conceptual level, following largely operationalizations in Fernández and Grimm (2014).

Lexical Alignment

Lexical alignment characterizes the speaker’s re-use of words from the interlocutor’s previous utterance. We captured this phenomenon by counting the number of shared ngrams (unigrams, bigrams and trigrams) across pairs of turns, normalized by the number of all possible ngrams.

Syntactic Alignment

Syntactic alignment was approximated by measuring the extent to which the speaker re-uses Part of Speech (PoS) ngrams from the interlocutor’s previous utterance. In order to disentangle syntactic from lexical measures, we only took into account

the PoS ngrams where at least one word was not identical. We counted the number of shared PoS ngrams (bigrams and trigrams) across pairs of turns, normalized by the number of all possible ngrams. The PoS tags in CHILDES were automatically generated using the Morphological Analysis algorithm (MOR, MacWhinney 2000) which yields a high accuracy rate on CHILDES adult data (above 99%).

Conceptual Alignment

Conceptual alignment aims at quantifying the degree to which interlocutors talk about conceptually similar things, without necessarily using similar words or syntactic structures. We computed this similarity as follows. First, we represented the utterances with vectors in a high dimensional space. These vectors were constructed by adding up the vectors corresponding to the content words present in the utterance. One way to obtain word vectors is by training a distributional semantic model on CHILDES data. However, since we are not studying semantic similarity from the perspective of the child, there is a priori no reason to limit training only to the children’s surrounding language. Instead, we used vectors that were pre-trained on a much larger corpus than CHILDES, leading to a more robust similarity space. We used word2vec vectors (Mikolov et al., 2013) pre-trained on part of Google News dataset (about 100 billion words).¹ Second, we computed the standard cosine similarity between pairs of utterances’ vectors. In order to disentangle conceptual from lexical similarity, we only computed the similarity between turns with no lexical overlap.

2.3 Types of Interactive Alignment

All our measures of alignment were normalized, taking values between 0 (absence of alignment) and 1 (identical repetition). Investigation of data across all measures showed that only conceptual alignment approximated a normal distribution. In contrast, both lexical and syntactic alignment data had an exceeding number of zeros and (to a lesser extent) ones. If the zeros and ones are removed, the data approximate a normal distribution. A similar phenomenon (especially zero-inflation) has previously been documented with adult-adult alignment data as well (Dideriksen et al., 2019).

There is no simple parametric distribution that can capture simultaneously the three modes of the

¹We obtained these vectors from: <https://code.google.com/archive/p/word2vec/>

Figure 1: Changes in alignment across linguistic levels and alignment types for children between 2 and 5 years old (age was scaled). Data show the marginal effect of age by direction (Adults aligning to children or vice versa) as well as a baseline representing alignment between randomly selected pairs of child-adult turns from the same conversation/transcript. Ranges represent 95% confidence intervals. Here we show the results for the bigram-based measures only. The patterns were similar for the unigram- and trigram-based measures.

distribution. They are better analyzed in terms of three separate statistical processes (see [Ospina and Ferrari 2012](#)), which map on, conceptually, to three types of alignment. The first is a binary process that determines the presence or absence of alignment (i.e., 0 vs. non-0). This process maps on to the propensity to align at all. The second is a continuous process that determines the degree of alignment when this alignment is partial (i.e., with values in the open interval $]0,1[$). Such a measure allows us to characterize how alignment is used to manage a conversation: the speaker creates rapport by borrowing some words and structures from their interlocutor, while keeping the conversation going by adding new information. Finally, a binary process that determines total overlap as opposed to partial or zero overlap (i.e., 1 vs. non-1). This process maps on to the tendency to repeat the exact utterance (at least for the lexical measure). This type of alignment may play a different communicative role than partial alignment (e.g., echolalia in children with autism).

Critically, these phenomena may not follow the same developmental trajectory. For instance, it is possible that the propensity to alignment increases over development while the degree of this alignment decreases, or vice versa. Besides, both propensity and degree of alignment may be independent of the tendency to repeat the interlocutor's

utterances.

3 Analyses and Discussion

We model each type of alignment with a separate mixed-effects model, predicting alignment at a given linguistic level (lexical, syntactic, and conceptual) by age and alignment direction (Adults aligning to children or vice versa). To take into account possibly correlated data from the same child, we used the identity of the child as a random effect.² For the binary processes (i.e., 0 vs. non-0 and 1 vs. non-1), the model was fit using binomial regressions. For the continuous process (i.e., $]0,1[$), the model was fit using a linear regression. Figures 1 and 2 show the results.

Children Align Consistently

Children's alignment to adults was above the baseline (defined as alignment between randomly selected pairs of utterances in the same conversation) and this above-random alignment was consistent across development. This result was strikingly robust across all linguistic levels and alignment types, thus providing the strongest evidence to date about young children's abilities to engage in multi-level interactive alignment from as early as two years old.

²The model was specified as follows: $\text{Measure} \sim \text{Age} * \text{Direction} + (1 | \text{child})$

Figure 2: Changes in alignment at the conceptual level for children between 2 and 5 years old (age was scaled). Data show the marginal effect of age by direction. Ranges represent 95% confidence intervals.

Adults Align More

Another result — that was equally robust — is the fact that adults aligned more to children than children did to adults. Could this asymmetry in alignment be simply explained by the asymmetry in terms of language production abilities? Indeed, young children’s ability to align will naturally be hindered if caregivers utter words that these children do not yet know or produce. We tested this hypothesis as follows.

First, we counted, for every pair of turns, the number of words uttered by the adult that the child may not know, which we determined roughly based on whether or not this word occurs in the sum of all children’s production by that age/month in our corpus.³ Second, we ran different mixed effect models where, in an addition to age and direction as predictors (as in the original modeling), we add the proportion of unknown words by children. All models (7 in total) showed a significant effect of alignment direction, suggesting that child-caregiver alignment asymmetry reflects genuine differences in conversational strategies beyond mere language production abilities (at least as quantified by our rough measure of word knowledge).

Development Across Linguistic Levels

For the lexical measure, which reflects alignment at the surface form, we found a declining trend across the board. Both children and adults use less exact words from their interlocutors as children develop. This result validates the finding of

Yurovsky et al. (2016) and extends it beyond function words. As for the syntactic and conceptual measures, which reflect alignment at a more abstract level,⁴ we found evidence for an increase in alignment (but only in the propensity to alignment for syntax). Interestingly though, in both cases we observe a similar increase in the random baseline, suggesting that this development likely reflects general enrichment in children’s expressive language, rather than changes in their local conversational strategies. For example, if children start using more function words in their production, this would create more opportunity for syntactic alignment both locally (between adjacent turns) and globally (between random turns, i.e., the baseline). Similarly, if children start using more semantically specific words, relating to the topic of the discussion, this would increase conceptual similarity both between adjacent turns and (to a lesser extent) between more distant ones.

Development Across Alignment Types

The three alignment types did not follow a similar developmental trajectory. For example, we observed a clear dissociation in syntax between the general propensity to align (which increases) and the degree of alignment (which decreases). The tendency to repeat the structure of the entire utterance is the type of alignment where we found the largest difference in development between children and adults. While this tendency did not change much for children, it started relatively high with adults and almost converged with children’s level by 5 years old.

4 Conclusion and Future Work

This study offered the first large-scale test to multi-level alignment in the context of natural child language. The results confirm some previous findings and uncover new ones. One question for future experimental work is whether such patterns can be produced in controlled behavioral experiments. Finally, while we addressed English-speaking children’s interactions, we plan to generalize this work cross-culturally in future investigations.

5 Acknowledgement

Research supported by ANR16-CONV-0002

³The standard way to quantify children’s average age of word production is through the Communicative Development Inventory (CDI), but available CDI data only goes up to 3 years old.

⁴Remember that in both measures, we only used instances where there was no overlap with lexical alignment

References

- Eve V. Clark. 2015. Common ground. In *The Handbook of Language Emergence*, pages 328–353.
- Herbert H Clark. 1996. *Using language*. Cambridge university press.
- Rick Dale, Riccardo Fusaroli, Nicholas D Duran, and Daniel C Richardson. 2013. The self-organization of human interaction. In *Psychology of learning and motivation*, volume 59, pages 43–95. Elsevier.
- Rick Dale and Michael J Spivey. 2006. Unraveling the dyad: Using recurrence analysis to explore patterns of syntactic coordination between children and caregivers in conversation. *Language Learning*, 56(3):391–430.
- Christina Dideriksen, Riccardo Fusaroli, Kristian Tylén, Mark Dingemanse, and Morten H Christiansen. 2019. Contextualizing conversational strategies: backchannel, repair and linguistic alignment in spontaneous and task-oriented conversations. In *CogSci’19*, pages 261–267.
- Gabriel Doyle and Michael C Frank. 2016. Investigating the sources of linguistic alignment in conversation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 526–536.
- Raquel Fernández and Robert Grimm. 2014. Quantifying categorical and conceptual convergence in child-adult dialogue. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 36.
- Abdellah Fourtassi, Ewan Dunbar, and Emmanuel Dupoux. 2014. Self-consistency as an inductive bias in early language acquisition. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 36.
- Abdellah Fourtassi, Isaac Scheinfeld, and Michael C Frank. 2019. The development of abstract concepts in children’s early lexical networks. In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*, pages 129–133.
- Riccardo Fusaroli, Bahador Bahrami, Karsten Olsen, Andreas Roepstorff, Geraint Rees, Chris Frith, and Kristian Tylén. 2012. Coming to terms: Quantifying the benefits of linguistic coordination. *Psychological Science*, 23(8):931–939.
- Sebastian Loth, Katharina Jettka, Manuel Giuliani, and Jan P De Ruiter. 2015. Ghost-in-the-machine reveals human social signals for human–robot interaction. *Frontiers in psychology*, 6:1641.
- Brian MacWhinney. 2000. *The CHILDES Project: Tools for analyzing talk. transcription format and programs*, volume 1. Psychology Press.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositional-ity. In *Advances in neural information processing systems*, pages 3111–3119.
- Raydonal Ospina and Silvia LP Ferrari. 2012. A general class of zero-or-one inflated beta regression models. *Computational Statistics & Data Analysis*, 56(6):1609–1623.
- Hae Won Park, Mirko Gelsomini, Jin Joo Lee, and Cynthia Breazeal. 2017. Telling stories to robots: The effect of backchanneling on a child’s storytelling. In *2017 12th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 100–108.
- Martin J Pickering and Simon Garrod. 2004. Toward a mechanistic psychology of dialogue. *Behavioral and brain sciences*, 27(2):169–190.
- Catherine E Snow. 1972. Mothers’ speech to children learning language. *Child development*, pages 549–565.
- Michael Tomasello. 2009. *The cultural origins of human cognition*. Harvard university press.
- Henry M Wellman. 2014. *Making minds: How theory of mind develops*. Oxford University Press.
- Daniel Yurovsky, Gabriel Doyle, and Michael C Frank. 2016. Linguistic input is tuned to children’s developmental level. In *CogSci*.
- Ran Zhao, Tanmay Sinha, Alan W Black, and Justine Cassell. 2016. Automatic recognition of conversational strategies in the service of a socially-aware dialog system. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 381–392.