

The perception of stroke-to-stroke turn boundaries in signed conversation

Marisa Casillas^a, Connie de Vos^a, Onno Crasborn^b, & Stephen C. Levinson^a

a. Max Planck Institute for Psycholinguistics
Wundtlaan 1, 6525XD Nijmegen, Netherlands
{Marisa.Casillas, Connie.deVos, Stephen.Levinson}@mpi.nl

b. Centre for Language Studies, Radboud University
Erasmusplein 1, 6525HT Nijmegen, Netherlands
o.crasborn@let.ru.nl

Abstract

Speaker transitions in conversation are often brief, with minimal vocal overlap. Signed languages appear to defy this pattern with frequent, long spans of simultaneous signing. But recent evidence suggests that turn boundaries in signed language may only include the content-bearing parts of the turn (from the first stroke to the last), and *not* all turn-related movement (from first preparation to final retraction). We tested whether signers were able to anticipate “stroke-to-stroke” turn boundaries with only minimal conversational context. We found that, indeed, signers anticipated turn boundaries at the ends of turn-final strokes. Signers often responded early, especially when the turn was long or contained multiple possible end points. Early responses for long turns were especially apparent for interrogatives—long interrogative turns showed much greater anticipation compared to short ones.

Keywords: Turn taking; sign language; online prediction; questions

Introduction

Everyday conversation is built up from turns at talk; first one person talks, then another, forming sequences of contingent communicative acts. Human communication heavily relies on these spontaneous face-to-face interactions, yet we are only recently beginning to understand how interactional pressures might bear on the format of conversation and the ways in which we process and represent language.

Turn taking is a basic conversational behavior that shows strong cross-linguistic consistency in its implementation. For example, the timing of speaker transitions (from the previous speaker to the next), the timing of different response types, and the effect of visual cues on transition timing are implemented similarly across linguistic communities that are otherwise very different from one another (Stivers et al., 2009). For this reason, turn taking has garnered attention as a potential source for communicative universals that, in turn, could have consequences for linguistic processing. The pressure to temporally and semantically coordinate turns in conversation affects how speakers dedicate resources to comprehension, monitoring, and response-planning (Levinson, 2013). If these pressures naturally arise from conversational needs, they should affect linguistic processing universally. This hypothesis (‘Interaction Engine Hypothesis’; Levinson, 2006) has spurred several recent studies on language processing and conversation across a wide range of linguistic communities. But at least one major source of variation has remained relatively unexplored: signed languages.

Signed languages have the potential to challenge notions about turn taking that are based on spoken language. For example, the size (and thus quickness) of sign articulators is

dramatically different between signed and spoken languages (hands/arms/face vs. tongue/lips), and turn overlap (when more than one person is signing/talking at once) appears to occur far more frequently in signed than in spoken conversation (Coates & Sutton-Spence, 2001; de Vos, Torreira, & Levinson, 2015).

Recent corpus analyses of spontaneous conversation in Sign Language of the Netherlands (Nederlandse Gebarentaal; NGT) have revealed that, although turn overlap is more frequent in NGT than in many spoken languages, the additional overlap may come as a consequence of having larger and slower articulators (de Vos et al., 2015). For spoken languages, which rely on quick oral articulations, the start and end of an turn is clear: approximately when vocalization begins and ends. But for signed utterances, the beginnings and ends of utterances are bookended by preparatory and retractive movements—movements that don’t bear turn-related content (Arendsen, 2009; Kita, van Gijn, & ven der Hulst, 2006).

De Vos et al. (under review) hypothesized that, because of this, signers might perceive their turns as starting and ending with the content-bearing movements (stroke-to-stroke) and not with all turn-related movements (preparation-to-retraction). When they calculated the timing of turn transitions in NGT with stroke-to-stroke turn boundaries instead of preparation-to-retraction boundaries, they found that NGT transition timing and turn overlap were consistent with the documented averages for spoken turn taking. Stroke-to-stroke boundary perception is then a potentially critical mechanism for linking signed to spoken turn-taking behaviors. But there is currently no experimental evidence that supports the psychological reality of stroke-to-stroke turn boundaries for sign language users.

One way to test for the presence of turn boundaries is to find out when addressees *anticipate* upcoming turn boundaries during ongoing talk. Accurate anticipation of upcoming turn-boundaries is often necessary for addressees to respond at the right time; by monitoring an ongoing utterance and predicting when it will end, addressees can time their response planning (from concept to articulation) and coordinate perfectly with their interlocutor, yielding speaker transitions with minimal silence and minimal vocal overlap between turns (Levinson, 2013).

Several studies have now shown that, when asked to listen a conversational turn, participants can precisely indicate the moment just before that turn ends by pressing a button, though their accuracy depends on the available linguistic in-

formation (Bögels, Torreira, & Levinson, accepted; Magyari & De Ruiter, 2012; De Ruiter, Mitterer, & Enfield, 2006).

The present study uses this same experimental technique to explore whether sign language users can predict the upcoming end of ongoing turns when the “end” is defined as the end of the last stroke (the stroke-to-stroke hypothesis) rather than the end of the last movement (the preparation-to-retraction hypothesis). We report on the results of one experiment within a larger project about online turn boundary prediction in NGT (gebarentaalmpi.com).

Methods

Adapting the method used by de Ruiter and colleagues in their (2006) study, we measured participants’ ability to anticipate the end of an ongoing turn. Participants viewed short videoclips of spontaneous conversation between two signers and then, after receiving a cue to focus on only one signer, pressed a button when they anticipated that the signer’s turn was about to end. The final turns in the stimulus videos were clipped to their hypothetical turn-end stroke boundaries (de Vos et al., 2015) to test whether participants could anticipate turn-end boundaries with stroke-to-stroke conversational timing. Additionally, because there is significant variation in signing skill among members of the deaf community, we recruited a diverse sample of participants to test how factors like age of sign language acquisition and sign input source affect the ability to predict upcoming turn ends.

Participants

We recruited 52 deaf signers whose primary language of communication is NGT through advertisements and personal contact. NGT is used in the Netherlands by approximately 16,000 people, nearly one-third of whom use NGT as their native language (Crasborn, 2001). There are at least five dialects of NGT, originating from each of the deaf schools.

In the last 100 years, the philosophy for deaf education in the Netherlands (like many places in the world) has changed radically, shifting from a strict emphasis on oral language and sign-supported speech to more signed- and bilingually oriented education (Tijsseling, 2015). For that reason, the age at which people first gained access to NGT and their sources for linguistic input vary both historically and across individuals (e.g., individuals from deaf vs. hearing families, with earlier vs. later hearing loss, etc.)

Our sample is a diverse slice of the Dutch signing community, including signers from three different dialects, a wide range of sign acquisition backgrounds (first input: birth–32; input source: at home/primary school/adult education), and an even wider range of ages (10–77).

Materials

We recorded two 90-minute spontaneous dyadic conversations to create the videoclips used in the experiment. Participants sat in two different rooms for the conversation recording, communicating over a videochat set-up that was designed to capture high quality video of each interlocutor. In

the set-up, cameras were placed behind one-way glass onto which the image of the signer’s partner could be projected, thereby allowing each signer to look into the camera and at the image of their addressee simultaneously (Figure 1).

From each of the two recordings we extracted 80 conversation fragments (160 in total: 20 practice and 60 test fragments from each recording) whose last turn used non-overlapping stroke-to-stroke timing in the transition to the next turn.¹ Each fragment was split into two videoclips: “context” and “target-turn” (Figure 2). The context clips showed both signers, side-by-side, and provided a few seconds of context immediately preceding the target turn. The target-turn clips only showed one signer. Participants’ task was to focus on this single signer and press a button when they anticipated the signer’s would end.

Button-press findings with spoken language stimuli have shown that participants cannot always anticipate turn ends; anticipation depends on the predictability of the content within each turn (Magyari & De Ruiter, 2012) and the linguistic cues that are available (Bögels et al., accepted; De Ruiter et al., 2006). Prior work has handled this by adding two seconds of silence after the turn offset. When participants hear silence, they know the turn has ended and can press the button reactively (not in anticipation). Analogously, our target turn clips ended with two seconds of video that just showed the last frame of the turn (the signer appears to ‘freeze’). When participants see the frozen signer, they know the turn has ended and can press the button reactively.

Each target turn was coded for its duration, the number of potential turn ends (PTEs) it contained (single vs. multiple), and for its question status (question vs. non-question).² In signed language, as in spoken language, a single turn in conversation can contain multiple potential endpoints. Even though only one of the endpoints results in a transition to the next speaker, early opportunities can arise and then pass by (e.g., A: “I didn’t think it was that good. Did you?” — B: “I thought it was alright, actually.”) Participants who are focused on anticipating upcoming turn ends could recognize these early potential turn ends and respond to them before the actual turn end comes (De Ruiter et al., 2006). We therefore coded each utterance for whether it contained multiple potential turn ends or just one (the actual turn end). Prior studies of response anticipation have also demonstrated a strong effect of transition type: observers anticipate responses more quickly after hearing questions than non-questions. We therefore coded each item for its question status—question or non-question—in case participants anticipated turn ends more quickly for questions (Casillas & Frank, 2012, 2013).

¹This timing criterion enables us to compare the same conversational turns across two tasks: the button-press task reported here and also some anticipatory gaze data that we collected from the same participants.

²The present results rely initial coding from one native signer of NGT. We are currently collaborating to update these codes with two naive NGT signers and two non-native signing researchers.




Figure 1: Each signer sits in front of a camera as the image of his/her interlocutor is projected onto an angled one-way mirror.

Procedure

Participants were tested in a mobile lab at one of eight sites around the Netherlands. After filling out a short background information form, participants participated in three experimental tasks: (a) anticipatory gaze (not reported here), (b) the button-press measure, and (c) a reaction-time baseline task. At the start of the button-press task, participants saw a short instructional video in NGT, then consulted with a deaf research assistant to check their understanding, and then began the experiment with 20 practice trials. After practice, participants consulted with the assistant once more and, if they wanted, completed the same 20 practice trials again before proceeding with the 60 test trials. Many participants (48%; 25) opted for a second practice round before beginning the test trials. After 30 test trials, participants were given the opportunity to take a short break. Each participants saw video-clips from only one of the two conversation dyads; they were assigned randomly to one dyad or the other.

Each experimental item was presented as a sequence containing the context and target-turn video-clips (Figure 2). Participants were asked to watch the context and then, when one signer disappeared, to focus on the remaining signer and try to press the button at the moment they thought the last sign would end. They were told that, if they saw the screen freeze, they should press the button as quickly as possible because they had reached the end of the turn. This gave us a measure of their predictive and reactive button-presses across the stimuli. The task lasted approximately 20 minutes.

In a second, very short task, we measured participants' baseline button-press reaction times by having them press the button as quickly as possible when they saw a cross-hair on the monitor (50 trials, inter-stimulus-interval range between 500–5000 ms, randomized).

Data preparation

We excluded 9 participants prior to data analysis because: they did not complete the task (2), they did not follow instructions (1), their linguistic background was unclear (5), and because of significant motor problems (1).

We had expected that this task, which is challenging for cognitive control and requires metalinguistic skills, would be

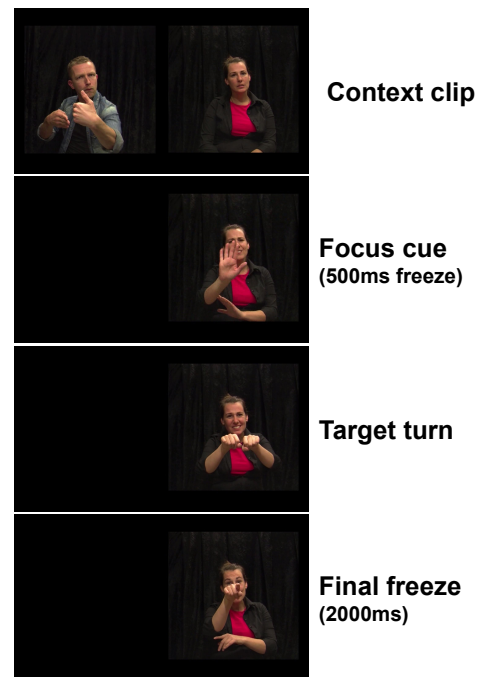


Figure 2: Example video sequence. View online at: <http://hdl.handle.net/1839/00-0000-0000-0020-6C0D-C@view>

difficult for a substantial portion of our sample—especially those who have had restricted access to language and formal instructional settings (like an experiment). We excluded 8 further participants who found the task too difficult to complete accurately; they either pressed the button during the context clips on more than 5% of trials (7) or pressed the button quite late (500+ msec after the turn-end freeze) on more than 10% of trials (1).³ We proceeded with the 35 remaining participants whose button presses we could confidently attribute to doing the task as intended. Despite the large number of exclusions, the remaining 35 participants still represented a diverse sample of the signing community (Table 1).

³The more generous criterion of 10% is because signed turns often end in a “hold” which can sometimes be ambiguous with the turn-end freeze.

We also excluded one item that was much longer than the others, plus any button presses that occurred within the first 720 msec of a turn or more than 500 msec after the turn-end freeze.⁴

Factor	Range (Mean; median)
Age	10–76 (46.57; 46)
Age of deafness	0–5 (0.69; 0)
Age of sign onset	0–32 (5.43; 3)

Factor	Distribution of participants
Input source	48% at home; 48% in primary school; 3% in adult coursework
Self-rated fluency	51% very good; 43% good; 6% reasonable
Education code	48% basic; 37% vocational; 14% post-secondary
Dialect fluency	23% North; 26% South; 89% West*

Table 1: Summary of sample characteristics after exclusions.*12 participants were fluent in more than one dialect.

Results

Participants responded with a button press to most trials ($M=55.6$; 93.3%), yielding a total of 1948 observations for the 35 participants. Because it takes a few hundred milliseconds to plan and execute a button press, we subtracted each participant’s baseline reaction time (their average in the cross-hair task) from their response latencies in the main task to estimate when, during each target turn, participants’ button presses were first triggered.

Participants often initiated their button presses before the stroke-to-stroke turn end; 1453 of the 1948 button presses (74.6%) were anticipatory (i.e., initiated before the end of the turn-final stroke), with the median overall response occurring 223 msec prior to the final stroke’s end ($M=-487$ msec; Figure 3).

Some turns contained multiple potential end points. In those turns, if participants responded to a non-final potential end point, their button press would be in extreme anticipation of the actual turn end. To remedy this, we focused exclusively on button presses targeted at the actual turn end by limiting our analyses to presses from the last 1500 msec of each turn (1215 of the 1453 anticipatory responses; 83.6%; Figure 3).

We modeled participants’ response latencies with mixed linear effects regression. Using an incremental model-building process, we added one predictor at a time, confirming for each that it improved the model’s goodness-of-fit; we used an ANOVA to compare pairs of models (one with and one without each added predictor) for significant improvement. Age of sign acquisition, age at test, and trial order did

⁴720 msec is the first point in the shortest stimulus where a turn end is imminent; Button presses more than 500 msec after the freeze are too late to reflect even reactive responding.




Figure 3: Button-press responses from all participants and items ($N=1948$; median latency is noted with a dashed vertical line). We analyzed anticipatory button presses targeted at the turn end (those in the shaded region; $N=1215$).

not contribute significantly to the model and were therefore not included in the analysis. The dependent variable—raw response latency—was negatively skewed, so we converted latency (negative; msec prior to the turn end) to an anticipation offset measure (positive; msec between the response and the turn end), log-transformed the new anticipation offset measure, and then removed outliers more than two standard deviations from the transformed mean.

The final model of log-transformed anticipation offset included fixed effects of target turn duration, question status (question/non-question), number of potential turn ends (PTE; single/multiple), the interactions between them, and, additionally, random effects of participant and item with maximal random effects structure.⁵

The largest predictor of response latency was the presence of a pre-final potential turn end: participants responded earlier when a potential turn end had occurred before the actual turn end ($\beta=4.34e-01$, $SE=1.67e-01$, $t=2.607$). This might indicate that the presence of pre-final potential turn ends helps participants respond earlier when the actual turn end arrives (e.g., because they have built up more certainty with the increased context). Alternatively, although we focused exclusively on responses within the last 1500 msec of the turn (i.e., targeted at the final possible turn end), the data set might still include responses to non-final potential turn ends, especially if they occurred shortly before the final turn end. To distinguish between these two explanations in future work, we will need to account for the timing of all potential turn-end boundaries within each item.

Relatedly, longer turns also resulted in earlier responses ($\beta=1.33e-04$, $SE=5.42e-05$, $t=2.451$), though this effect of turn duration was most clear for turns with only one potential turn end ($\beta=-1.47e-04$, $SE=5.35e-05$, $t=-2.754$). One possible explanation for this is, again, that increased context in longer turns allows participants to recognize the turn-final stroke earlier on and with more certainty. Alternatively, participants simply have more opportunities to respond early when turns are longer (e.g., compare the possibility for an early response

⁵Duration * Question * PTE + (1 + Question * PTE | Participant) + (1 | Item)




Figure 4: Anticipation offsets by turn duration for questions and non-questions in turns with a single potential end point (the equivalent graph for all turns is nearly identical).

in a 1- vs. 4-second turn). We would then expect a strong linear correlation between turn duration and average response anticipation, but the pattern is weak overall and non-linear, flattening out, or even decreasing for turns longer than two seconds (Figure 4).

We also saw a marginal interaction between turn duration and question status: the effect of turn duration was stronger for questions than non-questions, such that long questions showed consistently earlier responses than short questions (Figure 4; $\beta=9.05e-05$, $SE=5.31e-05$, $t=1.705$). This response pattern goes beyond a simple correlation of turn duration and early button presses. One explanation is that the timing of interrogative cues or the properties of interrogative speech acts push participants to give earlier responses while seeing longer turns (e.g., facilitated integration of late interrogative cue, like a *WH-* sign, because an early interrogative prosodic cue, like brow-raising).

To ensure that these findings were not driven by responses to pre-final potential turn ends (see above), we built a second model, restricting the data to target turns with a single potential endpoint. We used the same model-building process as before, resulting in model with fixed effects of turn duration, question status (question/non-question), and their interaction, and, additionally, random effects of participant and item with maximal random effects structure.⁶

The results of the second model mirror the findings from the first, showing significant effects for both turn duration ($\beta=2.77e-04$, $SE=6.08e-05$, $t=4.554$) and an interaction of turn duration and question status ($\beta=1.34e-04$, $SE=5.99e-05$, $t=2.241$). Notably, this second model—free from confounds of potential non-final end points—gives increased support for the interaction between turn duration and question status (Figure 4; correlation of anticipation offset and duration

for questions: $r^2 = 0.32$, $p < .001$; and non-questions: $r^2 = 0.024$, $p = 0.35$).

(a) Partial model output for all items
1161 observations, 35 participants, & 118 items

Factor	β	SE	t -value
PTE	4.34e-01	1.67e-01	2.607
Duration	1.33e-04	5.42e-05	2.451
Duration * PTE	-1.47e-04	5.35e-05	-2.754
Duration * Question	9.05e-05	5.31e-05	1.705

(b) Partial model output for items with only one potential turn end
786 observations, 35 participants, & 80 items

Factor	β	SE	t -value
Duration	2.77e-04	6.08e-05	4.554
Duration * Question	1.34e-04	5.99e-05	2.241

Table 2: Partial output from both statistical models reported. Potential turn end status (PTE: single/multiple); Question status (question/non-question); Duration (msec).

Discussion

In sum, we found that signers anticipated upcoming turn ends, even though the turn end was defined as the end of the last stroke (stroke-to-stroke) and not after turn-final retraction. Two primary factors affected the timing of participants' anticipations: the presence of early potential turn ends and longer turn durations were both associated with earlier average response latencies. Further, questions were more affected by turn duration than non-questions, suggesting that the effect of duration is not just an artifact of having more time to make early responses in longer turns.

Prior work using button-press measures of turn-end anticipation has also found that non-final potential turn ends can cause participants to press the button early (Bögels et al., accepted; De Ruiter et al., 2006), even when there is evidence (e.g., from prosody) that the speaker will continue speaking. In our first analyses, we tried to focus only on button presses targeted at the actual turn end, but we may have been unsuccessful in excluding all other types of responses. We can address this issue in the future by identifying the timing and linguistic content of pre-final potential ends. The potential turn ends in our signed stimuli may include prosodic or syntactic structures that very strongly suggest turn boundaries, despite other cues to continuation. By pursuing this line of inquiry, we could experimentally explore which linguistic cues are most likely to initiate early responses and which cues are more likely to signal continuation in signed conversation.

Earlier responses have been reported for longer turns in button-press measures of turn-end anticipation for spoken stimuli (De Ruiter et al., 2006), though they were cast as an artifact of having more time to make an early response. Our data attest to a more nuanced explanation because the effect of duration was (a) non-linear and (b) significantly more at play in interrogative turns. Both suggest that duration effects

⁶Duration * Question + (1 + Question | Participant) + (1 | Item)

interact with linguistic processing (i.e., in anticipating turn-end or interrogative content).

Prior work on turn anticipation has also found an effect of question status—participants watching videos of spoken conversation make earlier gaze switches to upcoming addresses after questions than non-questions (Casillas & Frank, 2012, 2013). Questions implicitly yield the floor to the addressee (and thereby forecast a turn ending). Question-marking cues might therefore be particularly salient for anticipating upcoming turn boundaries. Consider too that interrogative cues (e.g., an eyebrow raise, a palm facing upward, or a point to the addressee) can occur at different points in a turn; turns with early cues or turns with accumulated interrogative cues might result in earlier anticipated turn ends or even pre-final button presses. If participants focus in on early cues to question-hood, we might predict that early-initiated cues (e.g., brow raises) or early accumulations of cues (e.g., brow raises + palm up) would yield more anticipation than later cues alone. To test this with the current data we must first code each item for the presence and timing of interrogative markers.

Finally, because we choose to freeze the target turn video prior to retraction, it is possible that we reinforced the stroke-to-stroke responses we were trying to test. However, the lack of an order effect suggests that participants did not learn to become more accurate over the course of the experiment. That being said, a convincing follow-up would be to use stimuli in which the freeze occurs after the final retraction.

The findings presented here are the first to experimentally support the idea that signers use something like stroke-to-stroke turn boundaries to coordinate their turns in conversation. They also suggest that linguistic processing, here represented by question status, plays into the ability to use precisely-timed transitions in online conversation. In addition to digging further into item-based linguistic differences in the data, the next step is to directly compare the results from this task with more naturalistic measures of turn prediction (i.e., anticipatory gaze; not reported here) to determine which factors also spur *spontaneous* anticipation of upcoming turn structure. By combining multiple measures of turn prediction in NGT (i.e., button-press, spontaneous anticipation, and measures of spontaneous signing) we hope to extend our general knowledge about linguistic processing in conversation, increase our understanding about spontaneous conversation in signed language, and make concrete links between the implementation of turn-taking behaviors in signed and spoken languages.

Acknowledgments

We thank Richard Cokart, Mariko van der Garde, Said Jamal, Ellen Nauta, Tom Uittenbogert, Frouke van Winsum, Merel van Zuilen for their help in creating stimuli and collecting and coding data. We also thank our participants, especially those who hosted our mobile lab. This work was supported by ERC Advanced Grant 269484-INTERACT to SCL.

References

- Arendsen, J. (2009). *Seeing signs: On the appearance of manual movements in gestures*. Delft University of Technology, NL. (PhD thesis)
- Bögels, S., Torreira, F., & Levinson, S. C. (accepted). Intentional phrasing is necessary for turn-taking in spoken interaction. *Journal of Phonetics*.
- Casillas, M., & Frank, M. C. (2012). Cues to turn boundary prediction in adults and preschoolers. *Proceedings of the 16th Workshop on the Semantics and Pragmatics of Dialogue*.
- Casillas, M., & Frank, M. C. (2013). The development of predictive processes in children's discourse understanding. In *Proceedings of the 35th Annual Meeting of the Cognitive Science Society*.
- Coates, J., & Sutton-Spence, R. (2001). Turn-taking patterns in deaf conversation. *Journal of Sociolinguistics*, 5(4), 507–529.
- Crasborn, O. (2001). *Phonetic implementation of phonological categories in Sign Language of the Netherlands*. Utrecht, NL: LOT. (PhD thesis)
- de Vos, C., Torreira, F., & Levinson, S. C. (2015). Turn-timing in signed conversations: coordinating stroke-to-stroke turn boundaries. *Frontiers in Psychology*, 6.
- De Ruiter, J. P., Mitterer, H., & Enfield, N. J. (2006). Projecting the end of a speaker's turn: A cognitive cornerstone of conversation. *Language*, 82(3), 515–535.
- Kita, S., van Gijn, I., & van der Hulst, H. (2006). Movement phases in signs and co-speech gestures, and their transcription by human coders. In I. Wachsmuth & M. Fröhlich (Eds.), *Gesture and sign language in human-computer interaction* (pp. 23–35). Springer: Berlin, DE.
- Levinson, S. C. (2006). On the human “interaction engine”. In N. J. Enfield & S. C. Levinson (Eds.), *Roots of human sociality: Culture, cognition and interaction* (pp. 39–69). Oxford: Ber.
- Levinson, S. C. (2013). Action formation and ascriptions. In T. Stivers & J. Sidnell (Eds.), *The handbook of conversation analysis* (pp. 103–130). Wiley-Blackwell, Malden, MA.
- Magyari, L., & De Ruiter, J. P. (2012). Prediction of turn-ends based on anticipation of upcoming words. *Frontiers in Psychology*, 3:376, 1–9.
- Stivers, T., Enfield, N. J., Brown, P., Englert, C., Hayashi, M., Heinemann, T., ... others (2009). Universals and cultural variation in turn-taking in conversation. *Proceedings of the National Academy of Sciences*, 106(26), 10587–10592.
- Tijsseling, C. (2015). ‘School, waar?’ Een onderzoek naar de betekenis van het nederlandse dovenonderwijs voor de nederlandse dovegemeenschap, 1790–1990. Utrecht University, NL. (PhD thesis)