Automated mapping of competitive and collaborative overlapping talk in video meetings.

Eleni Margariti
Open Lab, Newcastle University, United Kingdom
e.k.margariti2@newcastle.ac.uk

Sean Rintel
Microsoft Research, Cambridge, United Kingdom
serintel@microsoft.com

Brendan Murphy
Microsoft Research, Cambridge, United Kingdom
bmurphy@microsoft.com

Abigail Sellen
Microsoft Research, Cambridge, United Kingdom
asellen@microsoft.com

ABSTRACT

Video meetings are notorious for difficulties with conversational turn-taking, which has impacts on inclusion and outcomes. We present a scalable automatic process to categorize turn-taking patterns in remote meetings based on eyes-off analysis of meeting transcripts. Drawing on a series of remote meetings (10 series, 34 total meetings) recorded in July-August 2021 by employees of a global technology company, we identified and parametrized patterns of cooperative and competitive overlaps of turns. The results show initial success characterizing people's behaviours as either likely to continue or cede turns based on either the amount of overlap that they produce during other's turns or the amount of overlap they experience in their own turns. With further development and validation, this method could be used to measure inclusion in remote and hybrid meetings.

CCS CONCEPTS

• **Human-centered computing**; • **Human computer interaction (HCI)**; • **Empirical studies in HCI**;

KEYWORDS

Audio/Video, Visualization, Conversational Analysis, Quantitative Methods

ACM Reference Format:
Eleni Margariti, Sean Rintel, Brendan Murphy, and Abigail Sellen. 2022. Automated mapping of competitive and collaborative overlapping talk in video meetings.. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI '22 Extended Abstracts), April 29–May 05, 2022, New Orleans, LA, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3491101.3519612

1 INTRODUCTION

Video meetings have longstanding challenges and opportunities for inclusion, many of which were exacerbated or highlighted during the COVID-19 pandemic [3]. One fundamental challenge is that cooperative competitive overlaps in conversational turn-taking, which are natural to dynamic conversations, are distorted by videoconferencing's asymmetrical perspectives [26] and aggravated by network jitter and latency [16, 17, 38]. Endless droning by one participant who refuses to yield the conversational floor, repeated false-starts, and missed opportunities to take the floor are all common spectres haunting video meetings [8, 9]. Meeting analytics may assist organisations improve meeting practices by providing insights on turn-taking behaviour. While technology will improve, at least a subset of problems arising turn-taking occur because people are socially blind to what is happening regardless of the technology. The challenge, then, is identifying these problems. Given that all widely used videoconferencing tools generate transcripts, we propose an automated method for using transcripts to detect cooperative competitive overlaps that is scalable and privacy compliant within any organisation that uses videoconferencing tools. As such, our research question was: How can cooperative competitive overlaps be automatically identified based on eyes-off processing of meeting transcripts? The findings suggest that the identification of cooperative competitive overlaps might be leveraged to identify behavioural profiles of meeting participants, taking steps towards addressing inclusivity in remote meetings.

We begin our report with a background on conversational turn-taking and cooperative/competitive overlaps, and prior work on quantitative analysis and meeting analytics. We then explain our methods, including recruitment and privacy, and detailing the capture, identification, parametrization, and visualization of overlap patterns. We report an initial set of results from a series of remote meetings (N series~10; N meetings ~34) recorded during July-August 2021 in a global technology company. The results provide a mapping of the impact of repeated competitive overlaps (attempts to take a turn) that are rejected, and competitive overtakes (changes of speakers that happen with an overlap). We discuss the outcomes of the results with respect to potential metrics for inclusive and dominating behaviour in meetings, and how this may enable group and individual behavioural profiles for meetings.

2 PRIOR WORK

2.1 Overlaps in conversational turn-taking

Speech overlaps are central to dynamic conversational turn-taking (Schegloff, E.A., 2000; Jefferson, G., 1986) [16, 36]. Simultaneous

CHI '22 Extended Abstracts, April 29–May 05, 2022, New Orleans, LA, USA

ACM ISBN 978-1-4503-9156-6/22/04…$15.00
https://doi.org/10.1145/3491101.3519612

speech is frequently observed in multi-party conversations and meetings in particular; with studies reporting that 30–50% of all turn exchanges in multi-party meetings contain some overlap [7, 11, 15]. Some instances of overlaps are cooperative while others are competitive [36].

Cooperative overlaps such as the 'continuers' (Schegloff, E.A., 2000) [36] or 'backchannels' e.g. 'uh huh', 'wow', "mm", "yeah" - provide attention, understanding, or agreement, rather than substantive content (Yngve, V., 1970 ; Sidnell, J., 2006) [36, 39, 45]. Other types of cooperative overlaps such as acknowledgement tokens e.g. 'yeah', 'right' or 'okay' (Jefferson, G., 1986) [16] can enhance turn stability, e.g. endorsing the current speaker by agreeing with them. Competitive overlaps are those in which listeners use various ways of taking the floor from current speakers [20]. There are different degrees to which non-speakers may infringe upon the rights of the current speaker to be heard e.g. rapport interruptions, process control interruptions; content control interruptions (Schegloff, E.A., 2000) [36]. Power interruptions are hostile attempts to take over the floor (where 'hostile' ranges from blithely unaware of disruption to deliberate disruption). These competitive overlaps on may inhibit conversational equity e.g. when a new speaker's overlapping leads to a current speaker cutting off their turn in progress [42]. Repeated overlaps during a turn may be cooperative, e.g. expressing encouragement, or competitive, e.g. insistently bidding for a turn. This relationship makes it crucial to model cooperative and competitive overlap together rather that separately.

So, overlapping speech is not necessarily the sign of a dysfunctional conversation—this very much depends on what function the overlaps serve, and both the sequential turn context and the social context in which they occur [2]. A crucial aspect of context for video meetings is that meeting function and attendee roles change over the many meetings in which people are involved. Some behaviours may be common to an individual in all meetings, but others may only occur in certain contexts. Videoconferencing's added contextual wrinkle is that latency disrupts both collaborative and competitive overlaps. Seuren et. al (2021) [38], building on previous work by [31], show that participants are generally not aware of latency itself, but treat its effects as meaningful, and frustration is common.

2.2 Quantitative Analysis of Turn-Taking

There have been many attempts to quantify turn-taking in order to investigate norms of distribution [6, 10, 22, 23]. In terms of raw quantification, there has been significant effort to operationalise the many sociolinguistic theories and definitions to clarify what is and isn't a turn by turning to acoustic measures such as 500ms or more uninterrupted talk spurts [6, 9, 10, 18, 22], and later also classifying VSUs (Very Short Utterances) [1], such as backchannels [13]. One step further, work has explored how to predict turn-taking. Levinson et al. (2015) [23] studied the timing of gaps and overlaps, the distribution of overlap durations, and where they start relative to interlocutors' turns. Their findings support that a participant in any given conversation predicts the end of the current speaker's turn to prepare their response in advance. Laskowsky's (2010) [22] model of turn-taking in multi-party conversation builds upon the automatic processing of conversations with a framework for computing turn-taking perplexity. Koiso et al. (2010) [19] using a prerecorded dataset, identified cues in overlapped utterances that elicit early initiation of a next turn and proposed a model for predicting a turn's completion in an incremental manner. Lala et al. (2018) [21] compare and evaluate deep learning models to predict turn-taking behaviour, working towards a generalised model that predicts turn-taking trained on speech and acoustic features. The final step after understanding turn distribution and prediction is to associate these with behaviours. Research on this includes Lindley et al [24, 25], who identified measures for the experience of participating in a multi-party dialog. These measures include conversational equality and freedom, conversational fluency, turn overlap, mean turn duration and turn synchronisation [24, 25]. Post-videoconference meeting inclusion dashboards that use such measures are currently trending in HCI research [34, 35] and commercially, such as Read.ai [31].

As noted above, most prior work focuses on audio conversations, builds upon speech and audio datasets [19, 21] or reflects on controlled experiments in a limited number of meetings [24, 25]. Recent work touches upon latency and the potentials of aggregated data to study turn-taking behaviour in remote meetings [16, 17, 37, 38, 40]; but often do not consider privacy constraints in collecting and using this data [4, 43, 44]. We contribute to prior work by presenting a GDPR-compliant method to detect cooperative competitive overlaps by automatically processing meeting transcripts as the input dataset instead of speech and audio corpora.

3 METHODS

Microsoft Teams videoconferencing was used for conducting the meetings, with automatic transcription provided via Microsoft's speech-to-text service [27]. Teams transcripts are generated real time with time-segmented and time-stamped speaker-attributed conversations. These transcripts are downloadable in VTT (time-aligned) format. A custom application was developed to process these transcripts and store extracted data in a secure GDPR-compliant database (see details in appendix 3.1). The application automatically de-identified participants' personal data, assigned unique IDs to meetings, and extracted selected participation metrics. It was possible to track a person's ID within a series of meetings but not identify them. This data was then exported and analysed for turn-taking patterns. Meeting metrics were then used for visual and statistical analysis: a) Seconds of talking per participant (each turn), b) Times talking (sum of turns), c) Start of talking per participant (milliseconds), d) Length of each talking interval (milliseconds), e) number of words spoken in each turn (word count).

Data was visualized as a timeline graph (see appendix 3.2), with each timestamped speech segment plotted along the X axis in milliseconds. Initially, overlap patterns were visually identified using these graphs, and validated using eyes-on transcripts using a few of the author team's internal meetings. Additional operationalisation parameters of the visual patterns were further derived from descriptive statistics of turn-length distributions. A subset of cooperative & competitive overlap patterns was selected for automatic identification in the sum of data.

3.1 Participants and meetings

We recruited n~52 participants for the study. This was a convenience sample of three project groups in a global technology company, with varied numbers of people in the regular meetings. Meeting attendees and owners all consented to their regular meetings being transcribed and used for research purposes. They met in regular project/status meetings across the organization. The researchers had no access to the transcript content, and all original files were deleted after their processing was complete, protecting the privacy and confidentiality of the groups involved. An IRB (ID: 10116) board authorized the processes and conduct of the study.

3.2 Latency

Measuring turn-taking in videoconferencing is difficult because of the idiosyncratic effects of network conditions. As noted above, latency and jitter [16, 17] are potentially quite disruptive to turn-taking in videoconferencing [38, 40]. Truly millisecond precision in the timing and duration of turns and overlaps between speakers would require audio recording at all endpoints and a complex formula for differences between endpoints and the cloud. The Microsoft Teams service produces a unified time-segmented transcript that reflects the cloud-mixed received audio rather than true endpoint audio, which makes measuring latency difficult. As such, we sought a way to reduce effort but enable consistent analysis. In a quick experiment, the authors clapped at successive times in rotation while recording themselves on phones and in Microsoft Teams. We compared the synchronized clips in Adobe Premiere and estimated a ~0.24s return delay. This made us reconsider the character of the overlaps smaller than 0.24s, because transcribed as an overlap might be perceived as a gap by some speakers, and vice versa. Therefore we excluded overlaps and gaps below 0.24s from our analysis; equal to 9% of all data. While acknowledging the constraints this imposes on accuracy, it saves much effort and time, and its consistency allows us to treat a certain size of overlaps as indetermined. Taking this error margin taken into consideration provides a common basis for analysis.

4 RESULTS

4.1 Visual Pattern Identification

Our first step was to analyse overlaps in some of own meetings, using Conversation Analytic methods from the work above (e.g. Schegloff, E.A., 2000) [36] and then visualize them in a manner similar to PauseCode [1], one example of which is shown in 1. Additionally, we documented patterns of conversational continuity, such as the repeated interchange of turns observed in the beginning or the end of the meeting, to characterize meeting types, and potential participant roles. We then conducted a quantitative analysis to assist with parametrizing the visually observed pattern. We plotted the duration distribution of all continuous speech utterances in the meeting to characterize the average size and distribution of each turn as transcribed. Median of turns was placed at ~2.4s (Median=2.38s, Mean=4.37s, Sd=5.146) indicating a fast-pasted conversation. Microsoft Teams software transcribed many short utterances (Mode=0.39s, 0.50s, 0.53s, 0.54s, 0.59s). We plotted the distribution of short utterances <2s to better understand their size and variation; this gave a Mean=0.92s and Median=0.77s, comprised by 1 (52%) and 2(17%) words (see appendix 4.1, 4.2).

Based on this analysis, Table 1 shows the identified overlap pattern categories. We highlighted the short overlaps (Mean=0.92s, Median=0.77s) occurring repeatedly within turns and in the beginning of an overtake (change of turn) as the most promising patterns for further investigation. Eyes-off analysis makes it hard to tell which patterns are purely competitive or cooperative, but some assumptions can be made. Pattern 6, a repeated overlap, may be a competitive attempt to take a turn or a cooperative backchannel. Similarly Pattern 7 may be a cooperative backchannel near the anticipated end of turn, or an intentional competitive interruption that results in a short pause that makes time for an overtake. Pattern 2 is likely cooperative due to the pauses (either a short pause at transition relevance place or the end of a turn) which provide time for an overtake. Patterns 3 and 8 are likely cooperative given the short length of the speech and overlap, their position in the beginning of a turn, and given that they are not repeated within short intervals. Patterns 4 and 10, on the other hand, are competitive, because a long overlap in the beginning of a turn is a strong indicator of competition for a turn. Patterns 1 and 9 present turn-taking patterns in different phases of the meeting. Pattern 1 shows short turns and long pauses between the turns with no overlaps, often observed at the beginning or the end of the meeting, and often the first speaker is the owner or moderator of the meeting. Pattern 9 also presents short alternating turns, but with overlaps and shorter pauses between turns. This faster pace and potentially competitive talk was observed during the middle or towards the end of the meeting.

We saw some of latency effects in our analysis. Often speech segments appeared with a very minor overlap (~0.2s) slightly delayed from the expected transition relevance place. Latency impacts the reception of the other's turn and increases the likelihood of speaking with overlap due to delayed signal. A likely example can be found in Figure 1, Pattern 7, where there is a slight delay in the response of P0. Still, we don't know if that is a delay generated by the equalization of response times by the transcription service, a delayed response due to latency that was then transcribed as such, or a competitive overlap that was transcribed as such. This provided a further reason for treating gaps/overlaps smaller than 0.24s as indeterminate in the current method.

We then expanded our visual analysis to eyes-off examples from full data set. Figure 2 shows a visualization from one meeting in which we identified the occurrence of Pattern 5: isolated overlap, Pattern 6: repeated overlap (cooperative or potentially competitive), 7: discussion with overlaps (competitive), 4: overtake with overlap (competitive) and 10: continuous overlap (competitive). We noted observations in each of pattern instances, making further interpretation of the turn-taking behaviour. In this meeting, we observed that P3 was talking for long time, and P3 also started the meeting. The assumption is that this person is giving a talk or a presentation, which was confirmed afterwards with the meeting participants. We observed that in meetings presentations, continuous uninterrupted turns (>3s) are often observed throughout the whole meeting, with no overlapping segments, occasional discussion segments (Pattern 9), and a fewer overtakes with overlap (Pattern 4). We observed P3 being interrupted by P1 (Figure10) with small overlaps – but it

00:07:08.830 -> 00:07:15.640	0		Is saying to you keep going? I'm listening to what you're saying or maybe even keep going and **[I prove what]** you're saying. And so we are collaboratively	**5**	Cooperative overlap, with a delay (0.2s)- transition point is [keep going and] but overlap happens a bit later
00:07:11.330 -> 00:07:12.100	1	**0.7s**	**[Exactly.]**	**5**	
00:07:16.200 -> 00:07:24.910	0		You know, engaged and you are. You are sort of your proceeding. Be **[cau]**se I'm asking you to proceed. Whereas **[if I'm]** doing more		
00:07:19.340 -> 00:07:19.610	1	**0.27s**	**[So]**	**6**	Attempt to take turn – delayed overlap (not in transition point) due to latency
00:07:23.030 -> 00:07:23.750	1	**0.62s**	**[Yes.]**	**6**	Another attempt to take turn (delayed)
00:07:26.580 -> 00:07:28.680	0		Uh, as you're **[talking?]**		Late notice
00:07:27.720 -> 00:07:37.280	1	**0.96s**	**[So in the]** yeah in that sense, even backchannel is a contribution to the discussion **[because or]** can be because it can	**4**	Competitive turn-taking – due to latency the person did not realize he/she was given space.
00:07:34.660 -> 00:07:35.420	0	**0.76s**	**[Yeah, what]**	**7**	attempt to take back turn, mapped in slight offset from the transition relevance point (latency)
		(3.2s)			
00:07:38.440 -> 00:07:39.300	1		Yeah	**7**	Gap, backchannel & abandoning the turn after the interruption.
00:07:39.260 -> 00:07:49.870	0		The contract well, what it contributes to it. It is not contributing to direct knew information. What is contributing to his conversational continuity?	**7**	

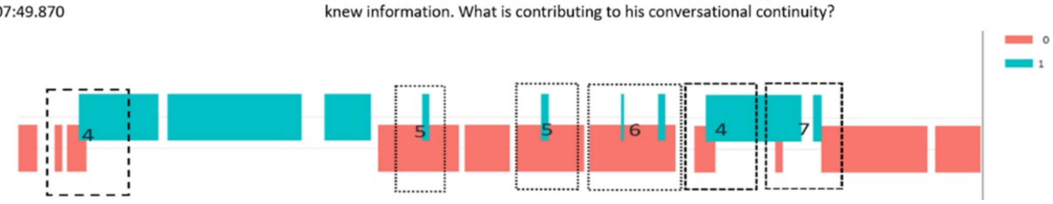

Figure 1: Patterns 5: Isolated overlap; 4: Overtake with overlap; 6: repeated overlap; 7: overlap at the end of turn

was unclear if these were competitive or not. As we examined the whole meeting, we observed that P3 was giving more room to P1 an overtake than to P4; with an extreme competitive overlap pattern (Pattern 10) observed between them. This example illustrates the potential of this type of analysis to show recurring patterns of competitive & cooperative overlaps, which may later allow profiling turn-taking behaviours of meeting participants and relationships between them (see appendix 4.4).

4.2 Automatic pattern identification

Excluding 9% of indeterminate data that fell within the 0.24s overlap/gap threshold, we conducted a descriptive statistical analysis of turn duration, overlap and overtake (change of turn) duration and distribution the remaining 91% of the full data set. The results of the distribution of all speech utterances resemble the outcome values of the initial experiment (Mean=5.14s, Median=2.57s, Sd=6.538s) and speech utterances <2.4s (Mean=0.89s, Median=0.63s; 53% one word, 16% two words). These results allowed us to further parametrize turn-taking analysis, using 3s (Median=2.57s) as the basis for the turn size and establishing a sample size ~0.9 (Mean=0.89s) for investigating overlaps. We further plotted the distribution of determined overlaps and gaps (>0.24s) for an example meeting series. Gaps were represented with a negative value, while overlaps with a positive value. The Mean = -1.45 suggests that there are more gaps than overlaps between turns in the meeting series. We observed some large negative values (e.g. long silences), which might have a wide range of potential technical or social causes and be positive or negative within the meeting context. The causes and results cannot be disambiguated by turn-taking patterns alone and would require some other contextual details. Additionally, we plotted the distribution of overtakes (e.g. change of interlocutor), and overtakes with overlap (>0.24s). Using 3s as the average turn size, the percentage of overtakes was 18%; overtakes with a determined overlap was 15% (see appendix 4.3).

We further investigated a subset of the identified patterns (from Table 1); selecting the ones that were more potentially relevant in terms of assessing non-inclusive behaviour. These are Pattern 5 (single overlap within a turn), Pattern 6 (repeated overlap within a turn),

Table 1: Patterns

Overlaps (Mean=0.92s, Median=0.77s) & no overtake			
Pattern 2: Gap – Overlap – Gap No Overtake Cooperative overlap	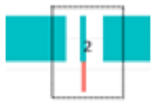	Pattern 3: Overlap before turn No Overtake Cooperative or Competitive	
Pattern 5: Single overlap within a turn No Overtake Cooperative or Competitive	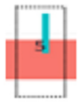	Pattern 8: Overlap in the beginning of a turn No Overtake Cooperative or Competitive	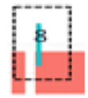
Pattern 6: Repeated overlap within a turn No Overtake Cooperative or Competitive	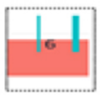	Pattern 7: Overlap at the end of a turn No Overtake Cooperative or Competitive	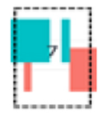
Overlaps in the beginning of a turn with an overtake			
Pattern 4: Overtake with overlap Overtake Competitive	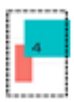	Pattern 10: Continuous overlap Overtake Competitive	
Conversational Continuity			
Pattern 1: Short interchanging turns (mean 2s) with no overlaps. Observed in the beginning and end of the meeting. Cooperative	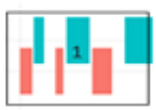	Pattern 9: Short interchanging overlapping turns across the duration of the meeting. Cooperative or Competitive	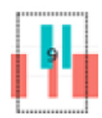

and Pattern 4 (overtakes with an overlap). We were primarily interested in mapping the turn-taking behaviour when these patterns occur. As the visual identification of the selected patterns in the sum of data would be time-inefficient, we developed an automatic process. We parametrised turns as any speech segment greater than 3s (when a person talks for equal or more than three seconds). Overlaps were parametrized based on duration (any speech segment smaller than 3s) and their relative position within a turn (e.g. excluded those occurring at the end of a turn and continuing for longer than the turn itself). Figure 3 visualises the turn-taking behaviour of each meeting participant in the following cases:

- Participant Initiating Turns (Figure 3 A, B, C):
- How often they maintain or abandon their turn if there is no overlap during their turn (Figure 3 A).
- How often they maintain or abandon their turn if there is a single overlap during their turn (Figure 3 B).
- How often they maintain or abandon their turn if there are repeated overlaps during their turn, initiated either by one or more other participants (Figure 3 C). The person that initiated most of the repeated overlaps is mentioned as the main backchannel.
- Participant Initiating Overlaps (Figure 3 D, E, F):
- How often they overtake if they initiate one overlap during a turn (Figure 3 D).
- How often they overtake if they initiate repeated overlaps during a turn (Figure 3 E).
- How often overtake if they are the main backchannel amongst others e.g. initiating more repeated overlaps during a turn amongst others initiating overlaps (Figure 3 F).

As illustrated above, some participants are much more likely to continue their turn when there is an overlap than others (comparing P3 and P4 in Figure 3 B). In the case of multiple overlaps occurring (Figure 3 C), the same participants presented the opposite behaviour, with P3 always abandoning turn, allowing main backchannel or other backchannel to continue (2/3 of the cases); whereas P4 tended to continue the turn (3/5) allowing 2/3 of the time the main backchannel to take over. On the other side, when participants initiate overlaps, we observe P3 initiating the greater

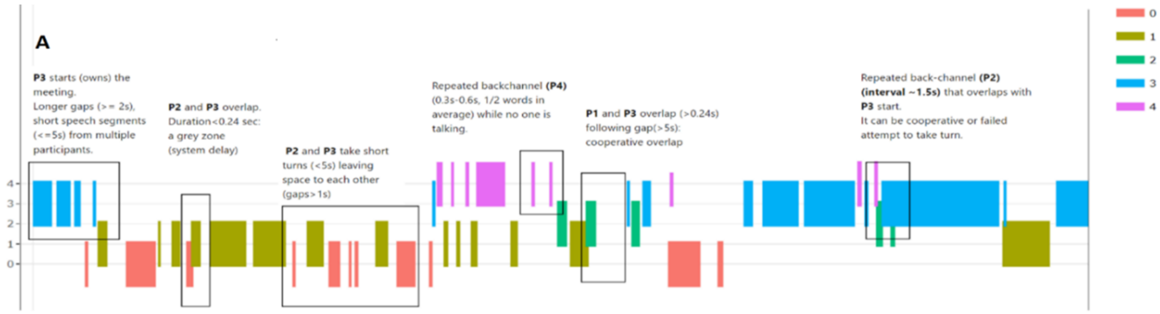

Figure 2: Identifying patterns in a Meeting of a Meeting Series (ID=4)

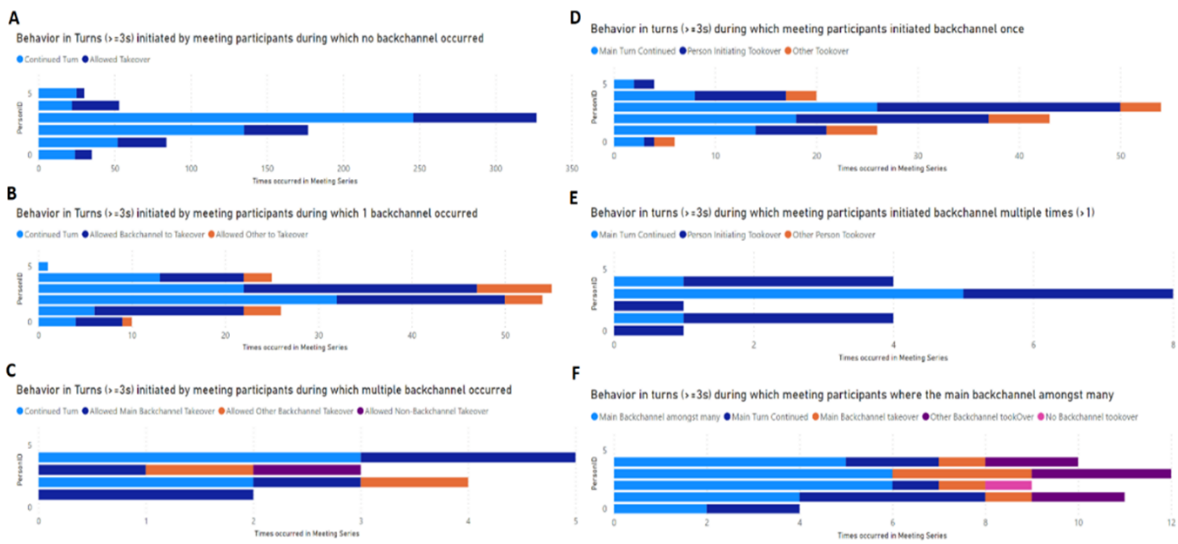

Figure 3: Turn-taking behaviour in a Meeting Series (ID=4)

amount of both isolated and repeated overlaps during turns. P3 approximately took over half of the times when initiating one overlap, and a bit less when initiating multiple overlaps. Lastly, P3 was able to take over half of the times when they were the main backchannel amongst many.

This type of automated and visual analysis can be meaningful to compare different behaviours of meeting participants within a meeting series, and may be the basis of developing profiles of turn-taking behaviours. We calculated the probability distributions (in percentages) in the sum of data examining the same cases and patterns as above. Calculating the probability of continuing an initiated turn (Mean= 63.51%) versus allowing someone to take-over when no overlap occurs (Mean= 31.2%), implies that *for any person in a multiparty conversation, initiating a turn is twice as likely to continue their turn than to allow someone else to take over when no overlap occurs during this turn*. Similarly, if a single overlap occurs during an initiated turn, the probability of continuing the turn or giving it to the person initiating the overlap are almost similar, with the second being slightly less likely (Mean = 34.98% and Mean= 31.30%). This illustrates that *the impact of a single overlap is significant*. On the other hand, the probability of a person who did not initiate any overlap to taking over is limited to Mean=7.4%. The probability distributions are described in more detail in the Supplemental Materials (see appendix 4.6).

We further compared the probability that a person initiated a single overlap during a turn will overtake. In most cases the current turn continues (Mean= 36.81%). This is 12.76% higher than the probability of the person who initiated this overlap to overtake (Mean= 24.05%). *The chance that a participant who does not generate backchannel will overtake someone else's turn is much lower than the other two* (Mean= 7.55%). We also examined the probability that the person initiating repeated overlaps during a turn would then overtake. This case includes a potential attempt to take turn, and whether an overtake is achieved or not. The results illustrate that *when a person overlaps someone else multiple times, it is more likely that an overtake will occur* (Mean =19.57%) than the current turn will continue (Mean=12.00%). At only Mean= 2.63% another speaker takes over.

5 FUTURE DIRECTIONS

The contribution of this work is to present an automated, scalable, privacy preserving and specifically, GDPR compliant way of identifying cooperative and competitive overlaps and their impact on turn-taking behaviour in remote meetings. This work provides the basis for assessing meeting inclusion based on a combination of metrics such as (a) the type, frequency and the amount of overtakes with an overlap that occur, and (b) the frequency and characteristics of overlaps from the same person during someone else's turn that did not result in an overtake (as these are potential attempts to take turn that are rejected). A key element for future work is the parametrization of turn-taking patterns and validation of the cooperative or competitive character of these patterns with a much larger data set. For instance, the repeated overlap pattern may have a cooperative or competitive character depending on the amount of repetition over a given turn, the size of the overlap, and the size of intervals between these repetitive segments.

The above analysis points to the potential for developing profiles of both inclusive and non-inclusive meeting participation [5]. For instance, participants that repeatedly do not cede turn after multiple overlaps occur and repeatedly overtake with overlap (Pattern 4 from 1) are enacting competitive overlap behaviours. Such competition may be factor in non-inclusive practices. However, this should not be judged prematurely without further insight into the meeting's context (function, role of attendees) and potentially a range of cultural indicators such as those proposed by Hofstede and Hofstede (2005) [14]. A participant with the recurring pattern above might have an acceptable social warrant for such behaviour in a particular type of meeting. As such, rich metrics must draw on and be validated against a range of contextual signals and extensive real-world testing. Interesting extensions to profiling would be mapping relationships between participants based on their turn-taking behaviour towards other participants; and study how turn-taking behaviour of the same person is shaped within and between different meeting series. Further, since conversation is a group process, profiles should also extend to groups. It may well be that group contextual features have a significant impact on turn-taking, perhaps more than individual profiles may suggest.

Of course, all individual or group profiling must be developed with precautions and deep considerations about the ethical and privacy implications it might have [4, 5, 24, 29, 30, 42]. Labelling behaviours can prove to be harmful for individuals, and linking behaviours such competitiveness or domination with race or gender should be avoided [4]. Domination and competitiveness are often integral parts of a multi-party communication, so these behaviours should be highlighted when they repetitively cause harm, exclude, or undermine others [12, 44, 45] or undermine group wellbeing [4]. There is therefore much discussion required around how these profiles will be developed in a manner that is both not biased and helpful, and how the behaviour tracking information will be used at an individual and organisational level [24, 30, 44].

The method presented was limited to remote meetings, but turn-taking in hybrid meetings has significant additional complexities [33, 34]. With technologies enabling speaker attribution in the room, it is possible to extend the above analysis to address turn-taking in hybrid meetings. Through characterizing the contribution of each person in a fully remote meeting series, it may be possible to measure the impact of being remote or local in a hybrid meeting, providing scope for understanding the effect of particular forms of remoteness and configurations of hybrid meetings.

6 CONCLUSIONS

This work contributes a novel GDPR compliant process of automatically identifying and categorizing cooperative & competitive overlaps from meeting transcripts and identifying the impact of these overlaps in turn-taking behaviour of participants. Results showed the impact of overlaps on maintaining or abandoning a turn is significant, and it varies depending on the number and repetition of these overlaps. This work also sets the basis for developing participation profiles and potentially assessing both inclusive and non-inclusive behaviours based on meeting metrics. Future work includes predicting more turn-taking behaviour based on transcripts, evaluating meeting inclusion from combined metrics, and exploring individual and group participation profiles to assess turn-taking behaviour in different types of meetings, and within both remote and hybrid meetings.

7 LIMITATIONS

The method's constraints have to do with the speech-to-text technology and the timing of the transcripts. We acknowledged many of these constraints in the Methods section above. We also noted above that this work shows promising initial results using a small dataset, which need to be validated with a larger data sample. In addition, though, it is also likely that application telemetry such as mute/unmute, reactions, raise hands, and, in particular, the use of the meeting's text chat, which has been shown to have effects on inclusion and accessibility [45], could provide a broader understanding of the range of turn-taking behaviour.

ACKNOWLEDGMENTS

Thank you to Neeltje Berger for the support; to the meeting organizers who helped collecting the transcripts, and the meeting participants who allowed us to monitor them.

REFERENCES

[1] Daniel Angus, Yeyang Yu, Paul Vrbik, Andrew Back, and Janet Wiles. 2018. PauseCode: Computational Conversation Timing Analysis. In Proceedings of the 4th International Workshop on Multimodal Analyses Enabling Artificial Agents in Human-Machine Interaction (MA3HMI'18). Association for Computing Machinery, New York, NY, USA, 43–47. https://doi.org/10.1145/3279972.3279975

[2] Dan Bohus and Eric Horvitz. 2011. Decisions about Turns in Multiparty Conversation: From Perception to Action. In Proceedings of the 13th International Conference on Multimodal Interfaces (ICMI '11). Association for Computing Machinery, New York, NY, USA, 153–160. https: //doi.org/10.1145/2070481.2070507

[3] Hancheng Cao, Chia-Jung Lee, Shamsi Iqbal, Mary Czerwinski, Priscilla N Y Wong, Sean Rintel, Brent Hecht, Jaime Teevan, and Longqi Yang. 2021. Large Scale Analysis of Multitasking Behavior During Remote Meetings. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21). Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3411764.3445243

[4] Nicole Chi, Emma Lurie, and Deirdre K Mulligan. 2021. Reconfiguring Diversity and Inclusion for AI Ethics. Association for Computing Machinery, New York, NY, USA, 447–457. https://doi.org/10.1145/3461702.3462622

[5] Amit K Chopra and Munindar P SIngh. 2018. Sociotechnical Systems and Ethics in the Large. In Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society (AIES '18). Association for Computing Machinery, New York, NY, USA, 48–53. https://doi.org/10.1145/3278721.3278740

[6] Steven E Clayman. 2012. Turn-Constructional Units and the Transition-Relevance Place. , 151–166 pages. https://doi.org/10.1002/9781118325001.ch8
[7] Stephen J Cowley. 1998. Of Timing, Turn-Taking, and Conversations. Journal of Psycholinguistic Research 27, 5 (1998), 541–571. https://doi.org/10.1023/A:1024948912805
[8] Ross Cutler, Yasaman Hosseinkashi, Jamie Pool, Senja Filipi, Robert Aichner, Yuan Tu, and Johannes Gehrke. 2021. Meeting Effectiveness and Inclusiveness in Remote Collaboration. Proc. ACM Hum.-Comput. Interact. 5, CSCW1 (apr 2021). https://doi.org/10.1145/3449247
[9] Paul Drew. 2009. Quit talking while I'm interrupting: a comparison between positions of overlap onset in conversation. In Talk in Interaction: Comparative Dimensions, Markku Haakana, Minna Laakso, and Jan Lindström (Eds.). Finnish Literature Society, 70–93.
[10] Paul Van Eecke and Raquel Fernández. 2016. On the Influence of Gender on Interruptions in Multiparty Dialogue. In Proc. Interspeech 2016. 2070–2074. https://doi.org/10.21437/Interspeech.2016-951
[11] J. Ganesh. 2020. Zoom and the lost art of interruption. https://www.ft.com/content/5745fc60-b0db-4958-bdf4-3bb6307e190d
[12] Kathryn Heath and Brenda F. Wensil. 2019. To Build an Inclusive Culture, Start with Inclusive Meetings. https://hbr.org/2019/09/to-build-an- inclusive-culture-start-with-inclusive-meetings
[13] Mattias Heldner and Jens Edlund. 2010. Pauses, gaps and overlaps in conversations. Journal of Phonetics 38, 4 (2010), 555–568. https://doi.org/10.1016/j.wocn.2010.08.002
[14] G. Hofstede, G. J. Hofstede, and Michael Minkov. 2010. Cultures and Organizations: Software of the Mind (3rd editio ed.). McGraw Hill. 576 pages.
[15] S.C. Holler, J., Kendrick, K.H., Casillas, M., Levinson. 2016. Turn-Taking in Human Communicative Interaction. Frontiers Media SA. https: //doi.org/10.3389/978-2-88919-825-2
[16] Gail Jefferson. 1986. Notes on 'latency' in overlap onset. Human Studies 9, 2 (1986), 153–183. https://doi.org/10.1007/BF00148125
[17] Hiroaki Kawashima, Takeshi Nishikawa, and Takashi Matsuyama. 2008. Visual Filler: Facilitating Smooth Turn-Taking in Video Conferencing with Transmission Delay. In CHI '08 Extended Abstracts on Human Factors in Computing Systems (CHI EA '08). Association for Computing Machinery, New York, NY, USA, 3585–3590. https://doi.org/10.1145/1358628.1358895
[18] Kobin H Kendrick. 2015. The intersection of turn-taking and repair: the timing of other-initiations of repair in conversation. Frontiers in Psychology 6 (2015). https://doi.org/10.3389/fpsyg.2015.00250
[19] Hanae Koiso and Yasuharu Den. 2010. Towards a precise model of turn-taking for conversation: A quantitative analysis of overlapped utterances.
[20] Emina Kurtić, Guy J Brown, and Bill Wells. 2013. Resources for turn competition in overlapping talk. Speech Communication 55, 5 (2013), 721–743. https://doi.org/10.1016/j.specom.2012.10.002
[21] Divesh Lala, Koji Inoue, and Tatsuya Kawahara. 2018. Evaluation of Real-Time Deep Learning Turn-Taking Models for Multiple Dialogue Scenarios. In Proceedings of the 20th ACM International Conference on Multimodal Interaction (ICMI '18). Association for Computing Machinery, New York, NY, USA, 78–86. https://doi.org/10.1145/3242969.3242994
[22] Kornel Laskowski. 2010. Modeling Norms of Turn-Taking in Multi-Party Conversation. In Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics (ACL '10). Association for Computational Linguistics, USA, 999–1008.
[23] Stephen C Levinson and Francisco Torreira. 2015. Timing in turn-taking and its implications for processing models of language . https://www.frontiersin.org/article/10.3389/fpsyg.2015.00731
[24] Rui Li, Chi Wang, and Kevin Chen-Chuan Chang. 2014. User Profiling in an Ego Network: Co-Profiling Attributes and Relationships. In Proceedings of the 23rd International Conference on World Wide Web (WWW '14). Association for Computing Machinery, New York, NY, USA, 819–830. https://doi.org/10.1145/2566486.2568045
[25] Siân E Lindley and Andrew F Monk. 2008. Social enjoyment with electronic photograph displays: Awareness and control. International Journal of Human-Computer Studies 66, 8 (2008), 587–604. https://doi.org/10.1016/j.ijhcs.2008.04.002
[26] Siân E Lindley and Andrew F Monk. 2013. Measuring social behaviour as an indicator of experience. Behaviour Information Technology 32, 10 (oct 2013), 968–985. https://doi.org/10.1080/0144929X.2011.582148
[27] Paul Luff, Christian Heath, Hideaki Kuzuoka, Jon Hindmarsh, Keiichi Yamazaki, and Shinya Oyama. 2003. Fractured Ecologies: Creating Environments for Collaboration. Human–Computer Interaction 18, 1-2 (jun 2003), 51–84. https://doi.org/10.1207/S15327051HCI1812_3
[28] Microsoft. 2022. What is Speech-to-Text. https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/speech-to-text
[29] Joon Sung Park, Michael S Bernstein, Robin N Brewer, Ece Kamar, and Meredith Ringel Morris. 2021. Understanding the Representation and Representativeness of Age in AI Data Sets. Association for Computing Machinery, New York, NY, USA, 834–842. https://doi.org/10.1145/3461702. 3462590
[30] Alicja Piotrkowicz, Vania Dimitrova, Antonija Mitrovic, and Lydia Lau. 2018. Using the Explicit User Profile to Predict User Engagement in Active Video Watching. In Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization (UMAP '18). Association for Computing Machinery, New York, NY, USA, 365–366. https://doi.org/10.1145/3209219.3209262
[31] Read.ai Read.ai. [n. d.]. Read - Real-Time In-Meeting Analytics for Zoom, Webex - Read Dashboard. https://www.read.ai
[32] Karen Ruhleder and Brigitte Jordan. 2001. Co-Constructing Non-Mutual Realities: Delay-Generated Trouble in Distributed Interaction. Comput. Supported Coop. Work 10, 1 (jan 2001), 113–138. https://doi.org/10.1023/A:1011243905593
[33] Banu Saatçi, Kaya Akyüz, Sean Rintel, and Clemens Nylandsted Klokmose. 2020. (Re)Configuring Hybrid Meetings: Moving from User-Centered Design to Meeting-Centered Design. Computer Supported Cooperative Work (CSCW) 29, 6 (2020), 769–794. https://doi.org/10.1007/s10606-020-09385-x
[34] Banu Saatçi, Roman Rädle, Sean Rintel, Kenton O'Hara, and Clemens Nylandsted Klokmose. 2019. Hybrid Meetings in the Modern Workplace: Stories of Success and Failure. In Collaboration Technologies and Social Computing, Hideyuki Nakanishi, Hironori Egi, Irene-Angelica Chounta, Hideyuki Takada, Satoshi Ichimura, and Ulrich Hoppe (Eds.). Springer International Publishing, Cham, 45–61.
[35] Samiha Samrose, Daniel McDuff, Robert Sim, Jina Suh, Kael Rowan, Javier Hernandez, Sean Rintel, Kevin Moynihan, and Mary Czerwinski. 2021. MeetingCoach: An Intelligent Dashboard for Supporting Effective Inclusive Meetings. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21). Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3411764.3445615
[36] Samiha Samrose, Ru Zhao, Jeffery White, Vivian Li, Luis Nova, Yichen Lu, Mohammad Rafayet Ali, and Mohammed Ehsan Hoque. 2018. CoCo: Collaboration Coach for Understanding Team Dynamics during Video Conferencing. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 1, 4 (jan 2018). https://doi.org/10.1145/3161186
[37] Emanuel A Schegloff. 2000. Overlapping talk and the organization of turn-taking for conversation. Language in Society 29, 1 (2000), 1–63. https://doi.org/DOI:10.1017/S0047404500001019
[38] Marwin Schmitt, Simon Gunkel, Pablo Cesar, and Dick Bulterman. 2014. Mitigating Problems in Video-Mediated Group Discussions: Towards Conversation Aware Video-Conferencing Systems. In Proceedings of the 2014 Workshop on Understanding and Modeling Multiparty, Multimodal Interactions (UM3I '14). Association for Computing Machinery, New York, NY, USA, 39–44. https://doi.org/10.1145/2666242.2666247
[39] Lucas M Seuren, Joseph Wherton, Trisha Greenhalgh, and Sara E Shaw. 2021. Whose turn is it anyway? Latency and the organization of turn-taking in video-mediated interaction. Journal of pragmatics 172 (jan 2021), 63–78. https://doi.org/10.1016/j.pragma.2020.11.005
[40] Jack Sidnell. 2006. Gene Lerner (ed.), Conversation analysis: Studies from the first generation. Language in Society 35, 5 (2006), 779–784. https://doi.org/DOI:10.1017/S0047404506350348
[41] Shinji Takao. 1999. The Effects of Narrow-Band Width Multipoint Videoconferencing on Group Decision Making and Turn Distribution. SIGSOFT Softw. Eng. Notes 24, 2 (mar 1999), 109–116. https://doi.org/10.1145/295666.295678
[42] Thomas Theodoridis, Symeon Papadopoulos, and Yiannis Kompatsiaris. 2015. Assessing the Reliability of Facebook User Profiling. In Proceedings of the 24th International Conference on World Wide Web (WWW '15 Companion). Association for Computing Machinery, New York, NY, USA, 129–130. https://doi.org/10.1145/2740908.2742728
[43] Khiet P Truong. 2013. Classification of cooperative and competitive overlaps in speech using cues from the context, overlapper, and overlappee. In Proc. Interspeech 2013. 1404–1408. https://doi.org/10.21437/Interspeech.2013-368
[44] Jess Whittlestone, Rune Nyrup, Anna Alexandrova, and Stephen Cave. 2019. The Role and Limits of Principles in AI Ethics: Towards a Focus on Tensions. In Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (AIES '19). Association for Computing Machinery, New York, NY, USA, 195–200. https://doi.org/10.1145/3306618.3314289
[45] Sarah Wille and Daphne Sajous-Brady. 2018. The Inclusive and Accessible Workplace. Commun. ACM 61, 2 (jan 2018), 24–26. https://doi.org/10.1145/3176410
[46] V H Yngve. 1970. On getting a word in edgewise. Chicago Linguistics Society, 6th Meeting (1970), 567–578. https://ci.nii.ac.jp/naid/10009705656/en/