Corresponding author: Ke Zhou, ke.zhou@nokia-bell-labs.com

The Role of Different Types of Conversations for Meeting Success

Ke Zhou
Nokia Bell Labs, Cambridge (UK); University of Nottingham (UK)

Marios Constantinides
Nokia Bell Labs, Cambridge (UK)

Luca Maria Aiello
IT University of Copenhagen (DK)

Sagar Joglekar
Nokia Bell Labs, Cambridge (UK)

Daniele Quercia
Nokia Bell Labs, Cambridge (UK); King's College, London (UK)

***Abstract*—While current meeting tools are able to capture key analytics from both text and voice (e.g., meeting summarization), they do not often capture important types of conversations (e.g., a heated discussion resulting in a conflict being resolved). We developed a framework that not only analyzes text and voice, but also quantifies fundamental types of conversations. Upon analyzing 72 hours of conversations from 85 real-world virtual meetings together with their 256 self-reported meeting success scores, we found that our quantification of types of conversations (e.g., social support, conflict resolution) was more predictive of meeting success than traditional voice and text analytics. These new techniques will be essential to uncover patterns in online meetings that might otherwise go unnoticed.**

THE INTRODUCTION Data analytics might help participants run their meetings more efficiently through accurate and real-time feedback on logistics, attendees, and environment [2, 7]. Simply put, meeting analytics partly suggest what is working, and what is not.

Current meeting tools offer useful insights through audiovisual or textual support. For example, a speech-based agent marks important 'items' in the spoken dialogue, while tools like ChAT identify topics or persons of interests within multi-party conversations [9]. While such tools often rely on textual or audiovisual analyses,

Published by the IEEE Computer Society

they do not capture all the aspects characterizing successful meetings. Consider, for example, a virtual meeting during which the host only enables the camera feed. In such a setting, while audio may convey, to a great extent, the sentiment and the prosody of the spoken words, the lack of physical presence and interaction makes it difficult to capture types of social interactions (e.g., a conversation full of support); a situation that many might have experienced during the COVID-19 pandemic. As more meetings are held in virtual rooms whose content can be recorded, we are faced with an unprecedented opportunity to automatically analyze their characteristics, and understand their language's nuances. To this end, we made three main contributions:

- We collected 72 hours of meeting conversations from 85 real-world virtual meetings, held on Cisco's WebEx in a corporate setting. Additionally, we collected 256 self-reported meeting success scores, which we used as the ground truth in our predictive models. (§DATASET).
- Using our dataset, we developed metrics based on the literature (§METHODOLOGY), which capture textual and verbal analytics, and types of social interactions expressed in the spoken dialogue. We built a model that predicts a meeting's success upon these metrics, and found that the quantification of types of social interaction was more predictive than verbal and textual analytics (§RESULTS).
- We discuss potential uses of this new set of analytics in current and future tools for monitoring and improving productivity in any organizations, such as predicting and enhancing online meeting experience (§DISCUSSION AND CONCLUSION).

RELATED WORK

Meeting analytics is key to productivity; in a sense, it provides a way to reflect and, ultimately, is a tool for improving meetings and running them more effectively. Previous research on meeting analytics focused on audiovisual analyses; examples include: generating speaker-annotated meeting transcripts [21]; identifying dominance and monitoring meeting participants' interactions [5]; detecting action items [15]; and, analyzing sentiment (or arguments) [19]. More recently, sensors that capture body postures and gestures were also used for quantifying human dynamics in meetings [4].

Most of this analytics automatically quantify human dynamics through metrics that range from visual (e.g., head movements) to physical (e.g., heart rate) to verbal (e.g., speech speed). However, previous studies fall short in understanding the relationship of those metrics with participants' meetings experience; contrary to previous research, our study correlates meeting metrics with participants' self-reported experience. Additionally, most of prior work relies on audio transcripts, and often overlooks communication nuances, which might reveal subtle social relationships [10], or key emotional expressions [8] that might be associated with meeting success.

DATASET

Using a Cisco's WebEx companion platform [2], we collected data from 85 virtual corporate meetings, approved by the HR department of authors' institutions. In total, these meetings lasted 4373 minutes with a median of 3 people participating in each meeting (). The dataset is comprised of a diverse range of meetings with varying duration (), hours of day (), days of week (), and days of month (). Meetings lasted for about 49 minutes on average, and all of them were conducted during business hours (8am to 6pm, Mon-Fri). The companion platform allowed participants to earmark key moments with a mobile app. These moments were then converted into one minute long audio chunks, which the meeting participants could playback in retrospect to get a quick audio summary of the meeting. The companion platform also allowed us to obtain self-reported measures that refer to the participants' meeting experience. More specifically, at the end of each meeting, the participants were prompted to answer two questions: one captured $Q_{psychological}$, which is the extent to which *[a participant]* felt listened or motivated to be involved, and the other captured $Q_{execution}$, which is the extent to which *[a participant]* felt that the meeting had a clear purpose and structure. The two questions were answered on a 1-7 Likert-scale, with 7 indicating greater extent. These two questions resulted from an extensive large-scale crowdsourcing study that determined the key predictors of a meeting's psychological experience [7], and are generalizable and independent from the specific analytics under study here.

Each meeting in the dataset was stored as a set of one-minute audio chunks of the earmarked moments, and the participant's self reported answers. We transcribed the earmarked audio chunks using the state-of-the-art Google's API Speech-to-Text service[1]; each meeting's transcript was used in our textual analyses, while a meeting's audio was used in our audio analyses. In total, all the 85 meetings contained 1007 earmarked moments, and 256 answers to the two questions.

METHODOLOGY

Using our collected dataset, we designed five metrics based on the literature that capture both verbal analytics (state-of-the-art) and types of social interactions (our proposal). Verbal analytics metrics are denoted with (V). To allow for experimental comparison, we developed two additional state-of-the-art metrics based on textual analyses, which are denoted with (T).

(A) State-of-the-art meeting analytics

1. Content (T). Following prior work [16], we consider a bag-of-words model that quantifies the frequencies of the most frequent uni-grams and bi-grams used in the meeting transcripts. To reduce sparsity, we counted the uni-grams and bi-grams that occur 5 times or more in the training set.

2. Sentiment (T). We applied sentiment analysis to capture the spectrum of sentiment expressed throughout the meeting. We applied both VADER (rule-based) [13] and FLARE (based on deep-learning) [1] to the meeting transcripts.

3. Sentiment (V). Verbal sentiment has been linked with people's perception of a meeting's experience [16]. We used a deep-learning speech-based sentiment classifier [12] to extract verbal sentiment for each meeting. The classifier was trained on an audio dataset annotated with eight emotions: neutral, calm, happy, sad, angry, fearful, surprise, and disgust.

[1]Speech-to-Text API: https://cloud.google.com/speech-to-text. It has been found that Google has superior performance on speech recognition compared to other platforms and tools [14].

4. Emotions from Pitch and Energy (V). In verbal communication, the pitch expresses emotional and paralinguistic information; it conveys emphasis, contrast, and intonation. Prior studies [8] have shown that prosodic features (e.g., pitch and energy) provide a reliable indication of the emotional status of conversational interactions. For example, the arousal state of a speaker (high activation versus low activation) affects the overall energy, and the energy distribution across the frequency spectrum [20]. To capture pitch and energy intensity patterns, for each meeting, we extracted the mean, the median, the standard deviation, the maximum, the minimum, and the range (max–min) of both the fundamental frequency and the energy. We also calculated the ratio of the up-slope to that of the down-slope of the pitch contour, which captures the fraction of high pitched voice regions.

5. Emotions from Speech rate (V). The arousal state of a speaker has been found to affect the frequency and duration of pauses. For example, an unusually high speaking rate has been linked to altered emotional states [17]. To capture speech rate, we used: *i)* the number of syllables per duration, *ii)* the number of syllables per phonation time, and *iii)* the ratio of duration of voiced and unvoiced regions.

6. Emotions from Prosody (V). In addition to time-dependent acoustic features (e.g., pitch, energy, and speech rate), spectral features are often selected as a short-time representation for speech signal. It is known that, during meetings, happy utterances have higher energy at high frequency range, while sad utterances have lower energy at the same frequency range [18]. For each meeting we computed the mel-frequency cepstrum (MFC) as it is a widely used representation of such short-term sound power spectrums [18].

(B) Our proposal: types of social interaction.

Prior studies [6, 10] showed that there are ten dimensions that capture, to a great extent, the type of social interactions in a wide variety of communication types in the workplace (e.g., email exchange). These dimensions were identified to be widely used ways of categorizing relationships. They were found so based on an extensive review of decades' worth of findings in sociology and social psychology [10].

These dimensions [6] include: *knowledge* (exchange of ideas or information; learning, teaching), *power* (having power over the behavior and outcomes of another), *status* (conferring status, appreciation, gratitude, or admiration upon another), *trust* (will of relying on the actions or judgments of another), *support* (giving emotional or practical aid and companionship), *romance* (intimacy among people with a sentimental or sexual relationship), *similarity* (shared interests, motivations or outlooks), *identity* (shared sense of belonging to the same community or group), *fun* (experiencing leisure, laughter, and joy), and *conflict* (diverging views, and conflict resolution).

Although these categories are not meant to cover exhaustively all possible social experiences, Deri et al. [10] provided empirical evidence that most people are able to characterize the nature of their relationships using these ten concepts only. Through a crowdsourcing experiment, they asked people to spell out keywords that described their social connections, and found that all of them fitted into the ten dimensions. We developed deep-learning classifiers to derive the ten social relationships from the conversational exchanges [6], and adopted such classifiers to quantify the types of social interactions within each meeting. We excluded the social dimension of *romance* as one expects, it was not substantially present in our meeting data.

While our metrics are all grounded in past work, they might not be exhaustive as they could be influenced by the diversity of meetings, or even cultures. To then evaluate our metrics, we set out to test the extent to which they are predictive of self-reported meeting success.

Self-reported meeting success score

To that end, we defined a "success" score from our participants' self-reports. We used this score as the outcome variable in our analyses, which has been previously validated in a large-scale crowdsourcing study [7]. In that previous study, a 28-item questionnaire was administered to 363 individuals whose answers were statistically analyzed through Principal Component Analysis (PCA). The analysis showed that two

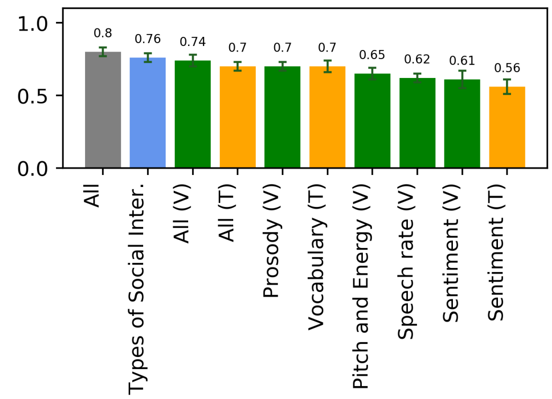

Figure 1: Evaluation (AUC) of our models trained on textual analytics (orange), verbal analytics (green), and types of social interaction (blue).

factors were sufficient to mostly capture whether a meeting is successful or not: (a) the extent to which participants felt listened during the meeting or motivated to be involved ($Q_{psychological}$), and (b) the extent to which the meeting had a clear purpose and structure ($Q_{execution}$). To this end, we obtained the loading factors of the first two components related to the two questions. Using these loading factors and the self-reports, we then computed an aggregated score of each attendee as: $success = (0.759 \cdot Q_{psychological}) + (0.673 \cdot Q_{execution})$. We binarized each $success$ score using the median computed across all meetings' scores, and assigned them to positive and negative classes (i.e., categorizing all meetings into "successful" and "unsuccessful").

RESULTS

To test the predictive power of our metrics, we developed classifiers to predict a meeting's success. We deployed a Logistic Regression (LR), a Support Vector Machine (SVM), a Random Forest (RF), a XGBoost, and an AdaBoost classifier. We chose these classifiers as they represent a wide range of linear and non-linear classification algorithms. These algorithms are also proven to be robust and perform well across datasets and applications. The best performing model was AdaBoost, which is an ensemble learning method (also known as "meta-learning"). AdaBoost uses an iterative approach to learn from the misclassifications of

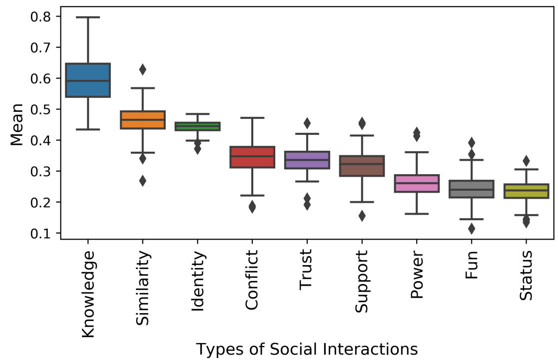

Figure 2: Distributions of types of social interaction expressed in our meetings.

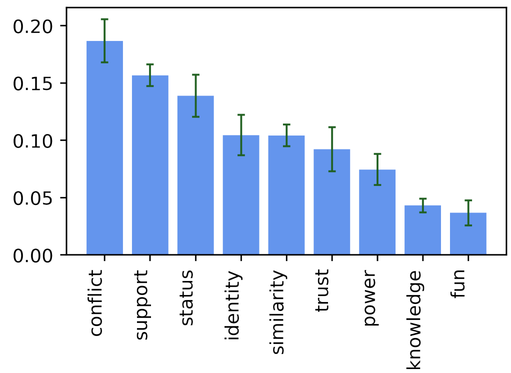

Figure 3: Feature importance (absolute value) of the AdaBoost model trained based on each type of social interaction.

weak classifiers, and builds a strong classifier by combining multiple weak classifiers; for brevity, we report only its results. We measured performance using a standard classification metric, that is, the area under curve (AUC), and employed a leave-one-out cross-validation. We report the averaged AUC across all folds.

General Evaluation. Figure 1 reports the AUC of our models trained on different combinations of our metrics.

- By inspecting each individual metric independently, we found that "types of social interaction" achieved the highest AUC of 76%, whereas the textual sentiment (T) metric, yielded the lowest AUC score of 56%. This is largely because most of the meeting transcripts do not contain explicit expressions of emotions.
- When comparing across three types of analytics (i.e., textual, verbal, and social), we found that the model trained on types of social interactions performed the best, achieving an AUC of 76%; followed by the model trained on verbal analytics, and the model trained on textual analytics. The model trained on verbal analytics (*all (V)* in Figure 1 that incorporates prosody, speech rate, pitch and energy, and verbal sentiment) achieved a relatively good result, obtaining an AUC of 74%. The model based on all the textual analytics (*all (T)* in Figure 1) performed the worst over all, yet it achieved an AUC as high as 70%.
- By combining all textual, verbal and social analytics (*All* in Figure 1), the best performing model achieved an AUC of 80%, demonstrating that these analytics are complementary to each other.

Analysis of Types of Social Interaction. As types of social interaction were found to be most predictive of meeting success, we then set out to determine which dimensions tend to be more predictive.

First, we inspected the distributions of those types of social interaction across all meetings (Figure 2). As one expects, meeting participants mostly exchange *knowledge*, with expressions of shared interests (*similarity*), and a sense of belonging to the same group (*identity*).

Secondly, we inspected the feature importance of the best performing AdaBoost model trained on types of social interactions (Figure 3), and this allowed us to understand which types (positively or negatively) contribute the most to the prediction accuracy. The classifier might capture non-linear relationships among the types of interaction. We found that conflict, support, and status were the most predictive types. *Conflict* reveals contrast or diverging views within the meeting, and, eventually, conflict resolution; *support* provides emotional or practical aid and companionship; and *status* confers status, appreciation, gratitude, or admiration upon another. Examples of such social interactions that led to meeting success can be found in Table 1. For example, it is not surprising that meetings that provide *support* and *status* updates would lead to a better experience (e.g., "*[...]* very interested

Table 1: Examples excerpts extracted from meeting transcripts, illustrating the use of language in the types of social interaction. Names appearing in the original dialogues were paraphrased, and quotes in boldface indicate language markers concerning corresponding interaction types.

Dimension	Examples
Conflict (contrast, diverging views, conflict resolution)	"Yeah, it's a problem. I think that's like *[...]* I don't know. **But, let me explain**. If you go on Instagram, you push that little heart button, and it starts floating up stuff. That's equivalent when you first get your phone. It's annoying, you know the vibration that most people turn off over time, right? *[...]* they [users] didn't have faith in the buttons being pushed, and over time, you know, maybe then you see the visual feedback. **You are right, it must be that way.**"
Support (giving emotional or practical aid and companionship)	"Human decisions about driver-less or autonomous cars is a very depressing topic. *[...]* The team, though, is **very interested to hear about these experiments**. Welcome everyone." "Thank you for the intro, and thank you for inviting me here. Can I start, right? Yeah, I'm going to talk about the moral machine experiment and a couple of *[...]*" "What about the experiments, and the **follow-up work with these amazing collaborators**?."
Status (appreciation , gratitude, admiration)	"If you have any interesting projects, topics, or ideas that you want to present, please don't hesitate *[...]*, just get in touch." "Exactly. Thank you Daniel, and thank you everyone for participating. See you soon to our next seminar series. Thank you. Thank you for joining us today. Thank you. Thanks for participating"
Identity (shared sense of belonging)	"I mean, but it's certainly one thing that people find interesting to talk about, and **they feel that they have something to share**. I think it's an important thing. Is that supposed to provide people with the opportunity to **give an opinion no matter how crazy or biased they are** *[...]*, but if they could give that opinion."
Similarity (shared interests, motivations or outlooks)	"**From your anecdotal example, it shows one amazing way** that artificial intelligence is used in healthcare. So I think this example is interesting for a couple of reasons. Firstly I think it really well illustrates the potential benefits we could have for medical AI. It also illustrates some of the high-stakes ethical decision-making that these kind of systems would end up being involved in."
Trust (will of relying on the actions of another)	"I don't know, if you guys have any comment on that." "Just one suggestion. Also, I would like to have some project updates. There is a lot of fuzziness around what the project entails, and we have not registered yet. But we don't know what the customer wants to. **We will figure it out, though.**"
Power (having power over the behavior of another)	"These people tend to have the **grace of God**." "*[...]* because of the way that our economy declines, we can now identify specific cultural dimensions."
Knowledge (exchange of ideas learning or teaching)	"**Let me start with an example**. *[...]* you have a typical prediction problem, and that is going to be used in a life-changing decision. *[...]*" "In the final part of the system, we will demonstrate how analog processing works *[...]*. **Any thoughts?**"
Fun (experiening leisure, laughter or joy)	" So, you guys here these sounds? *[...]* Not really in my side. *[...]* Oh, yeah. **It's like bird sounds**. I cannot here hear you on the bridge [referring to WebEx], but I hear voices like birds. **It's so funny**."

to hear about these experiments"). Contrary to conventional wisdom, we observed that *conflict* contributed positively to meeting success. This is partly explained by language exchanges that were mostly constructive, resulting in the definition of common goals or concrete action points (e.g., "*[...] But, let me explain.* [...] You are right, it must be that way".

Overall, these results show that these types of social interaction are instrumental elements in a meeting, and their presence or absence greatly matters for meeting success.

DISCUSSION AND CONCLUSION

While meeting tools translate, to a great extent, key aspects into analytics, we showed that there exist types of interactions in conversations that host important information linked to a meeting's success. If captured, these types of social interaction could potentially enrich meeting analytics, both in real-time and post-meeting. Our results reaffirmed previous findings [18] according to which verbal features (e.g., prosody and pitch) were found to complement textual sentiment and vocabulary ones. Interestingly, we found that types of interaction (e.g., conflict and support) were more predictive of meeting success than verbal features. In addition, both features were complementary to each other, and a combination of both was more predictive than what they were individually.

Our work has both theoretical and practical contributions. From a theoretical point of view, the types of social interactions used in our work could be theorized in the context of meetings, and widely adopted in Organizational and Management research. For example, these types of interaction could be linked to the concept of psychological safety. As Edmondson stated, psychological safety refers to "the absence of interpersonal fear that allows people to speak up with work-relevant content" [11]. As these types of social interaction greatly matter in meetings, if captured, it could help teams create safe environments for sharing and contribution. From a practical standpoint, our models could be deployed and integrated with any communication tool that provides voice recordings. MeetCues [2] is an example of such a tool: it allows participants to engage during a meeting, and reflect on their experience through visual and interactive features.

This work has limitations that call for future research efforts. First, our dataset draws mainly from business meetings, thus our findings might not generalize to other types of meetings. Secondly, we adopted audio as our main source to extract analytics. However, other aspects derived from facial expressions or body languages might be able to capture more nuanced emotions from meeting participants and about meeting structures (e.g., key turning points in a meeting). Finally, our types of social interaction capture the most frequent dynamics of interpersonal exchange in general settings, which are not specific to meetings. Tailoring those social interactions to the meeting context might boost the model performance, pushing it even further beyond our model's fairly high AUC of 76%.

Our work shed light on the importance of quantifying social interaction at scale. By monitoring the types of social interaction in any communication channel within an organization (e.g., company or university), we can measure the organizational productivity, and proactively take actions for improvements. For example, our analytics can be integrated as a plug-in for monitoring and improving online conference/meeting applications (e.g., Zoom). While this approach promises to improve organizational productivity, it also raises questions relating to workplace surveillance. It is often regarded that organizations and surveillance go hand in hand [3]. On a very pragmatic level, there is a handful of reasons as to why organizations opt in for employees' surveillance (e.g. maintaining productivity, monitoring resources used, protecting the organization from legal liabilities). The critics, however, rightly argue that there is a fine line between what organizations could be monitoring and what they should be monitoring. If crossed, it will have consequences on employees, affecting their well-being, work culture, and productivity. If future meetings tools incorporate any kind of employees' monitoring, they need to ensure that is done in a way that preserves an individual's rights, including that of privacy.

REFERENCES

1. A. Akbik, T. Bergmann, D. Blythe, K. Rasul, S. Schweter, and R. Vollgraf. Flair: An easy-to-use framework for state-of-the-art nlp. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 54–59, 2019.
2. B. A. Aseniero, M. Constantinides, S. Joglekar, K. Zhou, and D. Quercia. Meetcues: Supporting online meetings experience. In *Proceedings of IEEE Visualization Conference*, 2020.
3. K. Ball. Workplace surveillance: An overview. *Labor History*, 51(1):87–106, 2010.
4. I. Bhattacharya, M. Foley, N. Zhang, T. Zhang, C. Ku, C. Mine, H. Ji, C. Riedl, B. F. Welles, and R. J. Radke. A multimodal-sensor-enabled room for unobtrusive group meeting analysis. In *Proceedings of the 20th ACM International Conference on Multimodal Interaction*, pages 347–355, 2018.
5. C. Busso, P. G. Georgiou, and S. S. Narayanan. Real-time monitoring of participants' interaction in a meeting using audio-visual sensors. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP'07*, volume 2, pages II–685. IEEE, 2007.
6. M. Choi, L. M. Aiello, K. Z. Varga, and D. Quercia. Ten social dimensions of conversations and relationships. In *Proceedings of The Web Conference 2020*, pages 1514–1525, 2020.
7. M. Constantinides, S. Šćepanović, D. Quercia, H. Li, U. Sassi, and M. Eggleston. Comfeel: Productivity is a matter of the senses too. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(4):1–21, 2020.
8. E. Coutinho and N. Dibben. Psychoacoustic cues to emotion in speech prosody and music. *Cognition & emotion*, 27(4):658–684, 2013.
9. A. J. Cowell, M. L. Gregory, J. Bruce, J. Haack, D. Love, S. Rose, and A. H. Andrew. Understanding the dynamics of collaborative multi-party discourse. *Information Visualization*, 5(4):250–259, 2006.
10. S. Deri, J. Rappaz, L. M. Aiello, and D. Quercia. Coloring in the links: Capturing social ties as they are perceived. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW):1–18, 2018.
11. A. Edmondson. Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44(2):350–383, 1999.
12. C. Févotte, E. Vincent, and A. Ozerov. Single-channel audio source separation with nmf: divergences, constraints and algorithms. In *Audio Source Separation*, pages 1–24. Springer, 2018.
13. C. Hutto and E. Gilbert. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 8, 2014.
14. V. Këpuska and G. Bohouta. Comparing speech

recognition systems (microsoft api, google api and cmu sphinx). *Int. J. Eng. Res. Appl*, 7(03):20–24, 2017.

15. W. Morgan, P.-C. Chang, S. Gupta, and J. Brenier. Automatically detecting action items in audio meeting recordings. In *Proceedings of the 7th SIGdial Workshop on Discourse and Dialogue*, pages 96–103, 2006.
16. G. Murray. Uncovering hidden sentiment in meetings. In *Canadian Conference on Artificial Intelligence*, pages 64–72. Springer, 2016.
17. I. R. Murray and J. L. Arnott. Toward the simulation of emotion in synthetic speech: A review of the literature on human vocal emotion. *The Journal of the Acoustical Society of America*, 93(2):1097–1108, 1993.
18. T. L. Nwe, S. W. Foo, and L. C. De Silva. Speech emotion recognition using hidden markov models. *Speech communication*, 41(4):603–623, 2003.
19. S. Somasundaran, J. Ruppenhofer, and J. Wiebe. Detecting arguing and sentiment in meetings. In *Proceedings of the 8th SIGdial Workshop on Discourse and Dialogue*, pages 26–34, 2007.
20. C. E. Williams and K. N. Stevens. Vocal correlates of emotional states. *Speech evaluation in psychiatry*, pages 221–240, 1981.
21. T. Yoshioka, I. Abramovski, C. Aksoylar, Z. Chen, M. David, D. Dimitriadis, Y. Gong, I. Gurvich, X. Huang, Y. Huang, et al. Advances in online audio-visual meeting transcription. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 276–283. IEEE, 2019.

Ke Zhou is a Senior Research Scientist in the Social Dynamics team at Nokia Bell Labs Cambridge (UK) and Assistant Professor of Computer Science at University of Nottingham. He works in the areas of information retrieval and user modeling. Contact him at ke.zhou@nokia-bell-labs.com.

Marios Constantinides is a Research Scientist in the Social Dynamics team at Nokia Bell Labs Cambridge (UK). He is interested in human-computer interaction, mobile-sensing, and user modeling. Contact him at marios.constantinides@nokia-bell-labs.com.

Luca Maria Aiello is an Associate Professor at the IT University of Copenhagen (DK). He conducts interdisciplinary computational social science research. Contact him at luai@itu.dk

Sagar Joglekar is a Research Scientist in the Social Dynamics team at Nokia Bell Labs Cambridge (UK). He is interested in representation learning and its practical applications in the fields of social computing and urban informatics. Contact him at sagar.joglekar@nokia-bell-labs.com.

Daniele Quercia is the Department Head at Nokia Bell Labs in Cambridge (UK) and Professor of Urban Informatics at King's College London. He works in the areas of computational social science and urban informatics. Contact him at quercia@cantab.net.