

Detection of Listener Uncertainty in Robot-Led Second Language Conversation Practice

Ronald Cumbal

ronaldcg@kth.se

KTH - Royal Institute of Technology
Stockholm, Sweden

José Lopes

jd.lopes@hw.ac.uk

Heriot-Watt University
Edinburgh, United Kingdom

Olov Engwall

engwall@kth.se

KTH - Royal Institute of Technology
Stockholm, Sweden

ABSTRACT

Uncertainty is a frequently occurring affective state that learners experience during the acquisition of a second language. This state can constitute both a learning opportunity and a source of learner frustration. An appropriate detection could therefore benefit the learning process by reducing cognitive instability. In this study, we use a dyadic practice conversation between an adult second-language learner and a social robot to elicit events of uncertainty through the manipulation of the robot's spoken utterances (increased lexical complexity or prosody modifications). The characteristics of these events are then used to analyze multi-party practice conversations between a robot and two learners. Classification models are trained with multimodal features from annotated events of listener (un)certainty. We report the performance of our models on different settings, (sub)turn segments and multimodal inputs.

CCS CONCEPTS

• **Human-centered computing** → *Empirical studies in collaborative and social computing.*

KEYWORDS

Robot assisted language learning, conversation, social robotics

ACM Reference Format:

Ronald Cumbal, José Lopes, and Olov Engwall. 2020. Detection of Listener Uncertainty in Robot-Led Second Language Conversation Practice. In *Proceedings of the 2020 International Conference on Multimodal Interaction (ICMI '20)*, October 25–29, 2020, Virtual event, Netherlands. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3382507.3418873>

1 INTRODUCTION

Language learning has seen advantages from incorporating technology in training applications. These alternatives are great practice tools, but generally lack the personal interaction required in second language (L2) acquisition. Robot assisted language learning (RALL) aims to mitigate this drawback by including behavioral norms that resemble characteristics of human interaction. As teachers, assistants or peer-tutors [16], social robots can engage in activities designed to develop reading, grammar or speaking skills [24]. In the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICMI '20, October 25–29, 2020, Virtual event, Netherlands

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7581-8/20/10...\$15.00

<https://doi.org/10.1145/3382507.3418873>



Figure 1: Experiment settings. Left: (New) dyad conversation. Right: Three-party conversation. [8].

latter, it's common to apply explicit methods to improve this ability (e.g. pronunciation), while implicit approaches are less frequent. Conversation practice, therefore, stands as a novel path in RALL.

In a conversation exercise, as in other educational process, a learner is destined to experience moments of uncertainty when the complexity of the conversation is increased to reach the zone of proximal development. These events are generally examined in conversation analysis as communication problems, i.e. speaking, hearing, or understanding difficulties [20]. We can safely assume that in a language learning environment the occurrence of these problems tends to increase, and, therefore, an appropriate reaction from the (robotic) partner who guides the conversation is expected. In this investigation we focus primarily on **listener uncertainty**, i.e. *the case of a student failing to understand the information spoken by the conversational partner*. We designed an experiment to elicit listener uncertainty in a dyadic conversation through increased lexical complexity and modified prosody in the robot's spoken utterances. We then extended the analysis to include previously recorded robot-led practice conversations with pairs of language learners. Using multimodal features extracted from these recordings, we explore methods to automatically detect listener uncertainty from spoken output, physical expressions and time events.

1.1 Related Work

In educational settings, uncertainty has been evaluated in the form of *confusion*¹, as an affective state related to a student's learning performance [5, 13]. Various authors have build classification models by analyzing facial expressions [3], gaze exploration [17], body postures [7] and speech signals [12] with data extracted from academic tutoring systems. **However, these studies focus on the uncertainty of the student's answer, related to a fact-oriented question from the system, rather than the uncertainty of the student's semantic understanding of the question's content.**

¹uncertainty about what is happening, intended, or required.

Furthermore, research on uncertainty in speech production is in general focused on the *Feeling of Knowing* (FOK) paradigm [9], where the speaker struggles to recall or produce an idea, even if its concept is well understood. Previous studies have used prosody [19], lexical or linguistic features [6, 21], and audiovisual data [23] to detect speaker uncertainty. Similar to confusion detection, these efforts also focus on the certainty level of the students' response, which is related to, but not equivalent to, the participant's understanding of the tutor's utterance.

On the other hand, research on listener understanding or uncertainty aims to analyze the behavior of a conversation participant when communication errors occur. These phenomena are generally divided into misunderstandings and non-understandings [10], where the former implies an incorrect interpretation of the speaker's intention, while the latter indicates a complete absence or minimal confidence in any interpretation [22]. Normally, these problems are resolved in successive turns thorough different repair mechanism or common-knowledge validation (i.e., grounding), but L2 learners with low proficiency levels might be unable to handle problems in this manner, and may instead remain silent.

Close to our research, Kontogiorgos *et al.* [11] explored uncertainty in task-oriented interactions where the task-related environment enabled the grounding of the conversation (notably pointing or gazing). Our setting, on the other hand, evolves as a social conversation with minimum environmental effect.

Consequently, in this study we focus on listener uncertainty derived from non-understandings. Although misunderstandings are expected as well, the potential disruption of those errors might not be as drastic for a practice conversation. On the other hand, non-understandings reflect a stronger case of a miscommunication that could hinder the practice. Such events could have a bigger negative effect in the confidence of the learner and augment levels of stress or frustration. In order to build a system that can handle these events, we report our efforts to automatically detect listener uncertainty from video, audio and time signals recorded in dyadic and three-party practice conversations.

2 EXPERIMENT

We designed the uncertainty-eliciting experiment as a practice conversation between a Swedish second language learner and the anthropomorphic robot-head *Furhat* [1], with the setup shown in the left part of Figure 1. The social robot was controlled through a semi-automated Wizard-of-Oz interface by a Swedish native speaker. During the practice conversation, the wizard selected an option from a dynamically-updated list of up to 10 utterances (e.g. "Where do you come from?", "Could you elaborate?", "hmm", "yes") on topics that a learner could find relatable (e.g. "home country", "food preferences"). In a previous experiment where the system was employed, the participants reported difficulties of understanding related to the speed of the synthesizer and the complexity of utterances used by the robot. In the new experiment, in order to elicit uncertainty, we emphasized these observations by manipulating the robot utterances at the semantic and acoustic levels. In the former, we use samples from the TISUS² exam to define the threshold for high lexical complexity using the method proposed by Lu *et al.* [14]. We

then introduced utterances that surpassed this level in the dialogue system. In the latter, we accelerated the pronunciation of a set of low complexity utterances by modifying the speech rate of the synthesizer to "extra-fast". These manipulations were distributed equally along the dialogue. The experiment included a USB camera and microphone head-set to record the participants.

During the experiment, the participants were informed that they would interact with a robot in a short practice conversation. In an adjacent room, the wizard controlled the robot's interaction by listening to the participants' response and choosing the next robot utterance. The manipulated utterances were displayed in different colors in the computer interface and the wizards were instructed to prioritize their selection. After 30 dialogue transitions, the system presented utterances that could finish the conversation.

2.1 Preliminary Results

A group of 27 learners (16 females and 11 males) with ages between 21 and 51 years ($M=30$, $SD=8.9$), mainly recruited from a *Swedish for Immigrants* school and a university-organized Swedish Language Café, interacted with the robot. All participants had a Swedish proficiency level between A1 (basic) and B2 (intermediate).

Two thirds of the manipulated utterances resulted in a clear breakdown in the flow of the conversation. These events were characterized by the participant requesting a form of clarification (32 times), but more importantly, instances when the participants did not reply at all (4 cases), as noted in our preliminary work [4].

3 CHARACTERIZATION OF UNCERTAINTY

We consider four different response states to define different levels of listening uncertainty, as defined in Table 1. The manifestations of listener Uncertainty (Clarification request and No Response) differ considerably from a well understood Certain response (Thinking or Direct), where the learner follows the conversation with minimal interruption. As seen in Table 2, one particular difference between all states could be the length of the Silence Gaps. However, even if uncertain responses appear to have a longer silence gap between the question-answer turn, the temporal features are not sufficient to distinguish between a contemplative –certain– response, where the answer is first being considered, and an –uncertain– clarification response, where the repair is slowly requested. Human listeners tend use situational information about the participant to prepare for and handle these phenomena, as a robotic partner should. This means that it is not an alternative to simply let the dialogue manager wait a specific amounts of time before handling a delayed response.

	Category	Criterion
Certain	Thinking R.	Contemplation or meditation of the full answer.
	Direct R.	No hesitation, immediate response
Uncertain	Clarification	Request to repeat, slow down or rephrase previous utterance.
	No Response	No reply to the robot.

Table 1: Criteria for each category of (un)certainity.

²Test in Swedish for University Studies

Turns	Count (Dyad/Triad)	Length M (SD)	Silence Gap M (SD)
Certain			
Thinking R.	248 (150/98)	20.76 (13.56)	1.84 (1.08)
Direct R.	360 (190/170)	10.13 (7.28)	1.17 (0.60)
Uncertain			
Clarification	165 (69/96)+[165]	9.68 (5.96)	2.36 (1.55)
No Response	79 (28/51)+[79]	6.58 (2.52)	4.50 (2.18)

Table 2: Dataset characteristics. Augmented samples shown in brackets. Length and Silence Gap measured in seconds.

3.1 Annotation

We analyzed all dyadic sessions, as well as previously recorded conversations in the triad configuration, following the formerly addressed criteria. Both settings share the same interactive characteristics (dialogue content and style, robot behavior and wizard system), but only the new setup included manipulated utterances. In total, 42 sessions were annotated, 20 dyad and 22 triad recordings. Each annotated turn starts and ends at the beginning of consecutive robot utterances, as shown in Figure 2. Furthermore, we manually annotated 2 timing events between each turn: the end of the robot’s utterance and the beginning of the participant’s response. In the latter case, fillers (or hesitation marks) were not considered as the start of a participant response, since their precise function in a conversation can be varied. Two non-expert annotators reviewed the video recordings, with one of them working only on the dyad settings. From the videos that were shared between annotators, we computed a kappa inter-annotator agreement of 0.73. We confirmed this value by distributing 20 random video turns to two additional annotators, resulting in a kappa coefficient of 0.67. A total of 852 turns were labeled with class distribution shown in Table 2, along with a description of the timing events.

4 FEATURE EXTRACTION

The videos of all annotated turns were processed with OpenFace [2] to obtain the presence and intensity of Facial Action Units (FAUs), gaze direction angles, head pose coordinates and head rotation angles from every image frame. Except for FAUs, the time derivative (Δ) of these features was also included as an additional feature. Moreover, from the audio segments we extracted spectral features and voice activity through python’s *Librosa* [15] and *Webrtcvad* packages, respectively. The spectral features included RMS, (13) MFCCs, and Mel-Bank Spectrogram (13 components). From the MFCCs, we also computed the time derivatives (Δ and $\Delta\Delta$). Finally, we included the manually annotated time duration of the silence gap and the complete turn as part of temporal features. Except for the latter two, the selection of features, and the corresponding processing software, had the objective of implementing a real-time system that could detect listener uncertainty.

5 TRAINING AND EVALUATION

We employed a Machine Learning approach that enables the assessment of the input features when training the classification model.




Figure 2: Turn description.

Among different tested methods (e.g. K-NN, SVM, Logistic Reg.), the Random Forest classifier from Scikit-learn [18], set with 100 estimators (grid searched), produced the best performance. Since these models do not account on sequential inputs, we processed the features through an additional step. We calculated the mean of the feature sequence and merged these statistical values in one vector that represents a complete (sub)turn. This step is not applied to the temporal features as these values are scalars (time in seconds).

In order to examine the sequential characteristic of the input features, we also tested a Bidirectional LSTM with a fix sequence length input. Since our data segments had varied lengths, we reduced each sequence by sampling data points on evenly spaced intervals. In this method, we did not include the Silence Gap value, as this would ideally be contained in the Voice Activation vector, and replaced the value of the segment length with an increasing time vector to represent the time duration of the (sub)turn.

Since the dataset was clearly imbalanced, we augmented the sample size of the less frequent labels (Clarification and No Response) by sub-sampling feature arrays of those categories, as shown in Table 2. Additionally, the models were trained on weighted-class inputs to decrease the effect of the imbalanced classes. We used a Leave One Out Cross Validation approach to accurately assess the model’s performance when data from a completely new participant is used for testing. We trained with this method until all participant were tested and averaged all performance results. In the case of the Bi-LSTM, we randomly choose two additional participants’ data as validation set and used an Early Stop approach to avoid over-fitting.

Finally, we employed two segments of data from each turn to train our classifiers: data acquired before the beginning of the participants’ speech (BPS: sub-turn) and data from the complete turn (CT: turn). The former refrains from using full speech data from the learner, while the latter uses the complete input data to classify the turn, as shown in Figure 2.



(a) Random F. - BPS. (b) Random F. - CT. (c) Bi-LSTM - CT.

Figure 3: Normalized Confusion Matrix for Audio + Visual + Temporal features, trained on All settings data.

Modality	Setting	BPS		CT	
		Acc.	F1	Acc.	F1
Visual	Dyad	0.499	0.454	0.581	0.572
	Triad	0.473	0.445	0.547	0.533
	All	0.484	0.460	0.565	0.545
Audio	Dyad	0.389	0.344	0.488	0.453
	Triad	0.360	0.319	0.440	0.400
	All	0.379	0.333	0.476	0.440
Audio + Visual	Dyad	0.526	0.489	0.640	0.613
	Triad	0.493	0.463	0.596	0.583
	All	0.501	0.472	0.612	0.616
Audio + Visual + Temporal	Dyad	0.548	0.517	0.658	0.640
	Triad	0.516	0.506	0.696	0.687
	All	0.572	0.544	0.669	0.653

Table 3: Model performance on Accuracy and F1-score. Training data divided in different types and conversation settings. *Majority Baseline Accuracy at 0.33.

6 RESULTS

When we look at the different settings in Table 3, the models with one participant (dyadic) appear to have a better performance than the models trained on multi-party data. This difference might be explained by the skewed positioning of the camera and the expected additional noise from the recordings with a conversational partner. Nevertheless, when we merge the data from all settings, the results tend to fall between the values of the individual settings' score. Hence, to simplify our analysis, we used the joined settings data ("All") for our further examination.

In terms of multimodal data, it's clear that the combination of modalities improves the performance of the models. Nonetheless, it's relevant to point that models trained on audio data alone exhibit the worst values for both segments of input data (BPS and CT). Naturally, the contribution of audio is slightly higher on the CT data, as these segments include speech from the participant, but the visual components remain more informative. This characteristic is visible in Figure 4, where the cumulative importance of FAUs can easily outweigh other type of features (e.g. CT: FAU pres. + int. = 0.28). Furthermore, for both segments of data (BPS and CT) the inclusion of temporal data does indeed improve the performance of the models, but this contribution is not as important when we compare them to the collective importance of other types of features. For example, FAUs consist of 34 values –representing the presence and intensity for 17 units–, which individually are insignificant, but united form an essential information source, i.e. facial expressions are substantial for this analysis.

The most notorious difference is found between the two evaluated segments of data. As expected, the models trained on the complete turn (CT) have a higher score (Acc/F1: 0.67/0.65) compared to models trained on segments of data before the participants speak (BPS) (Acc/F1: 0.57/0.54). From Figure 3a and 3b, we notice that the biggest improvement is found in the classification of Thinking Response. Nevertheless, even these best results are far from optimal, since the models appear to have a high difficulty differentiating Thinking and Clarifications Response from a Direct Response. On




Figure 4: Feature Importance (Audio + Visual + Temporal).

the other hand, No Responses appear to be correctly classified most of the times, even when we only consider the BPS segments.

The Bi-LSTM model performed below expectations even after performing regularization and model reduction techniques. For consistency, we report the highest score corresponding to the full set of data types and complete turn (CT) segments: Acc/F1: 0.46/0.44. As seen in Figure 3c, this model struggled to classify most classes.

7 DISCUSSION

The task of detecting listening uncertainty in L2 (social) practice conversations is a complex one. In our effort to understand this phenomena, we proposed a novel interpretation of (un)certainly that expands the simple binary approach. With this annotation scheme, we found good classification results for No Responses and Direct Responses, but the intermediate classes were modest at best. The classification of these –certain– Thinking and –uncertain– Clarification responses became one of the harder tasks for the models. This distinction can be problematic for human partners as both responses could often be confused by a shared "thinking face". In these cases, people tend to rely on additional lexical data, which could be explored in our future work. Even if our results are not optimal, it does indicate that the correct detection of different response states is a favorable path. We acknowledge that averaging sequential data had a negative effect in our approach, as important information might have been removed, but we expected the Bi-LSTM model to account on these characteristics and improve on the results. Unfortunately, this method requires a large quantity of data to encode these properties. Hence, in our future work we plan to increase the size of our data samples. Nevertheless, it's evident that the integration of various modalities generates the best results, but the particular contribution of facial expressions (i.e. FAUs) and commonly utilized audio spectral features (i.e. MFCCs and Mel Spectrogram) is meaningful. In particular, the detection of No Response is important for RALL systems, as this quality would enable a dialogue manager to prepare an adequate response when the system detects a this state of listening uncertainty, that might not be repaired by the learner alone.

ACKNOWLEDGMENTS

We thank Per Fallgren, Mattias Bystedt and Marine Bastidas for their important assistance to this study.

REFERENCES

- [1] Samer Al Moubayed, Jonas Beskow, Gabriel Skantze, and Björn Granström. 2012. Furhat : A Back-projected Human-like Robot Head for Multiparty Human-Machine Interaction. In *Cognitive Behavioural Systems : COST 2102 International Training School, Dresden, Germany, February 21-26, 2011, Revised Selected Papers (Lecture Notes in Computer Science)*. 114–130. https://doi.org/10.1007/978-3-642-34584-5_9 QC 20121123.
- [2] Tadas Baltrušaitis, Amir Zadeh, Yao Chong Lim, and Louis-Philippe Morency. 2018. OpenFace 2.0: Facial Behavior Analysis Toolkit. In *2018 13th IEEE International Conference on Automatic Face Gesture Recognition (FG 2018)*. 59–66. <https://doi.org/10.1109/FG.2018.00019>
- [3] Nigel Bosch, Yuxuan Chen, and Sidney D'Mello. 2014. It's Written on Your Face: Detecting Affective States from Facial Expressions while Learning Computer Programming. In *Intelligent Tutoring Systems*, Stefan Trausan-Matu, Kristy Elizabeth Boyer, Martha Crosby, and Kitty Panourgia (Eds.). Springer International Publishing.
- [4] Ronald Cumbal, José Lopes, and Olov Engwall. 2020. Uncertainty in Robot Assisted Second Language Conversation Practice. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20)*. Association for Computing Machinery, New York, NY, USA, 171–173. <https://doi.org/10.1145/3371382.3378306>
- [5] Sidney D'Mello, Blair Lehman, Reinhard Pekrun, and Art Graesser. 2014. Confusion can be beneficial for learning. *Learning and Instruction* 29 (2014), 153 – 170. <https://doi.org/10.1016/j.learninstruc.2012.05.003>
- [6] Jeroen Dral, Dirk Heylen, and Riëks Akker. 2011. *Detecting Uncertainty in Spoken Dialogues: An Exploratory Research for the Automatic Detection of Speaker Uncertainty by Using Prosodic Markers*. 67–77. https://doi.org/10.1007/978-94-007-1757-2_6
- [7] Sidney D'mello and Arthur Graesser. 2007. Mind and Body: Dialogue and Posture for Affect Detection in Learning Environments. In *Proceedings of the 2007 Conference on Artificial Intelligence in Education: Building Technology Rich Learning Contexts That Work*. IOS Press, NLD, 161–168.
- [8] Olov Engwall, José Lopes, and Ann Åhlund. 2020. Robot Interaction Styles for Conversation Practice in Second Language Learning. *International Journal of Social Robotics* 11 (03 2020). <https://doi.org/10.1007/s12369-020-00635-y>
- [9] J. T. Hart. 1965. Memory and the feeling-of-knowing experience. *Journal of Educational Psychology* 56 (1965), 208–216.
- [10] Graeme Hirst, Susan McRoy, Peter Heeman, Philip Edmonds, and Diane Horton. 1994. Repairing conversational misunderstandings and non-understandings. *Speech Communication* 15, 3 (1994), 213 – 229. [https://doi.org/10.1016/0167-6393\(94\)90073-6](https://doi.org/10.1016/0167-6393(94)90073-6) Special issue on Spoken dialogue.
- [11] Dimosthenis Kontogiorgos, Andre Pereira, and Joakim Gustafson. 2019. Estimating Uncertainty in Task-Oriented Dialogue. In *2019 International Conference on Multimodal Interaction (ICMI '19)*. Association for Computing Machinery, New York, NY, USA, 414–418. <https://doi.org/10.1145/3340555.3353722>
- [12] Rohit Kumar, Carolyn P Rosé, and Diane J Litman. 2006. Identification of confusion and surprise in spoken dialog using prosodic features. In *Ninth International Conference on Spoken Language Processing*.
- [13] Zhongxiu Liu, Visit Pataranutaporn, Jaclyn Ocumpaugh, and Ryan Shaun Joazeiro de Baker. 2013. Sequences of Frustration and Confusion, and Learning. In *EDM*.
- [14] Xiaofei Lu. 2012. The Relationship of Lexical Richness to the Quality of ESL Learners' Oral Narratives. *The Modern Language Journal* 96, 2 (2012), 190–208. https://doi.org/10.1111/j.1540-4781.2011.01232_1.x
- [15] Brian McFee, Colin Raffel, Dawen Liang, Daniel Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. 2015. librosa: Audio and Music Signal Analysis in Python. 18–24. <https://doi.org/10.25080/Majora-7b98e3ed-003>
- [16] Omar Mubin, Catherine Stevens, Suleman Shahid, Abdullah Mahmud, and Jian-Jie Dong. 2013. A review of the applicability of robots in education. *Technology for Education and Learning* 1 (06 2013). <https://doi.org/10.2316/Journal.209.2013.1.209-0015>
- [17] Mariya Pachman, Amael Arguel, L. Lockyer, Gregor Kennedy, and Jason Lodge. 2016. Eye tracking and early detection of confusion in digital learning environments: Proof of concept. *Australasian Journal of Educational Technology* 32 (12 2016). <https://doi.org/10.14742/ajet.3060>
- [18] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *Journal of machine learning research* 12, Oct (2011), 2825–2830.
- [19] Heather Pon-Barry and Stuart Shieber. 2011. Recognizing Uncertainty in Speech. *EURASIP journal on advances in signal processing* 2011 (03 2011). <https://doi.org/10.1155/2011/251753>
- [20] Emanuel Schegloff, Gail Jefferson, and Harvey Sacks. 1977. The Preference for Self-Correction in the Organization of Repair in Conversation. *Language* 53 (06 1977), 361–382. <https://doi.org/10.2307/413107>
- [21] Tobias Schrank and Barbara Schuppler. 2015. Automatic detection of uncertainty in spontaneous German dialogue. In *INTERSPEECH*.
- [22] Gabriel Skantze. 2007. *Error Handling in Spoken Dialogue Systems : Managing Uncertainty, Grounding and Miscommunication*. Ph.D. Dissertation. KTH, Speech, Music and Hearing, TMH. QC 20100812.
- [23] Marc Swerts and Emiel Krahmer. 2005. Audiovisual prosody and feeling of knowing. *Journal of Memory and Language* 53, 1 (2005), 81 – 94. <https://doi.org/10.1016/j.jml.2005.02.003>
- [24] Rianne van den Berghe, Josje Verhagen, Ora Oudgenoeg-Paz, Sanne Van der Ven, and Paul Leseman. 2018. Social Robots for Language Learning: A Review. *Review of Educational Research* (12 2018). <https://doi.org/10.3102/0034654318821286>