



HAL
open science

How Confident are You? Exploring the Role of Fillers in the Automatic Prediction of a Speaker's Confidence

Tanvi Dinkar, Ioana Vasilescu, Catherine Pelachaud, Chloé Clavel

► To cite this version:

Tanvi Dinkar, Ioana Vasilescu, Catherine Pelachaud, Chloé Clavel. How Confident are You? Exploring the Role of Fillers in the Automatic Prediction of a Speaker's Confidence. Workshop sur les Affects, Compagnons artificiels et Interactions, CNRS, Université Toulouse Jean Jaurès, Université de Bordeaux, Jun 2020, Saint Pierre d'Oléron, France. hal-02933476

HAL Id: hal-02933476

<https://inria.hal.science/hal-02933476>

Submitted on 8 Sep 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

How Confident are You? Exploring the Role of Fillers in the Automatic Prediction of a Speaker’s Confidence

Tanvi Dinkar

tanvi.dinkar@telecom-paris.fr

Institut Mines-Telecom, Telecom Paris, CNRS-LTCL, Paris, France

Catherine Pelachaud

catherine.pelachaud@upmc.fr

CNRS-ISIR, UPMC, Paris, France

Ioana Vasilescu

Ioana.Vasilescu@limsi.fr

LIMSI, CNRS, Université Paris-Saclay, Orsay, France

Chloé Clavel

chloe.clavel@telecom-paris.fr

Institut Mines-Telecom, Telecom Paris, CNRS-LTCL, Paris, France

ABSTRACT

“Fillers”, example “um” in English, have been linked to the “Feeling of Another’s Knowing (FOAK)” or the listener’s perception of a speaker’s expressed confidence. Yet, in Spoken Language Processing (SLP) they remain unexplored, or overlooked as noise. We introduce a new challenging task for educational applications, that is the prediction of FOAK. We design a set of filler features based on linguistic literature, and investigate their potential in FOAK prediction. We show that the integration of information related to implicature meanings allows an improvement in the FOAK model and that the different functions of fillers are differently correlated with confidence.

KEYWORDS

Spoken Language Processing, Disfluency, Fillers, Confidence, Educational Applications

1 INTRODUCTION

This paper will be part of the proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP) 2020. This paper presents a study which is a part of ANIMATAS (Advancing intuitive human-machine interaction with human-like social capabilities for education in schools), an H2020 Marie Skłodowska Curie European Training Network that aims to enhance the pedagogical environment through human-agent interaction. In this project, we investigate the role of spoken language in this educational context [5].

An essential part of spoken language are disfluencies; which can be defined as breaks, irregularities or non-lexical vocables that occur within the flow of otherwise fluent speech. They are frequent in spoken language, as spoken language is rarely fluent. In the past few years, there has been a widespread interest in SLP. However, methodologies in SLP to study the disfluencies of the speakers and the information they can provide, are often overlooked.

Fillers, are a type of disfluency that can be a sound (*um* or *uh* in English) or word/phrase (*well*, *you know*) filling a pause in an utterance or conversation. Research on fillers conclude that they are informative in understanding spoken language [1–3, 31]. There are several roles that fillers play within the verbal message. The speaker can use a filler to indicate a pause in speech [2] or hesitation [18]. A speaker can use fillers to inform about the linguistic

structure of their utterance, such as in their (difficulties of) selection of appropriate vocabulary while maintaining their turn (in dialogue). Importantly, fillers are linked to the metacognitive state of the speaker. It was observed that fillers and prosodic cues are linked to a speaker’s *Feeling of Knowing (FOK)* or *expressed confidence*, that is, a speaker’s certainty or commitment to a statement [26].

However, the meanings of fillers are contextual, and dependent on the perception of the listener [2, 10]. The speaker encodes the meaning into their speech, while the listener decodes (or perceives) this meaning depending on the context. Hence, studies have also looked at the *comprehension* of fillers, that is, by taking into account the listener’s understanding of speech uttered by the speaker [4]. [1] found that listeners can *perceive* a speaker’s metacognitive state, by using fillers and prosody as cues to study the *Feeling of Another’s Knowing (FOAK)*; or the listener’s perception of a speaker’s expressed confidence/ commitment to their speech.

Despite the rich literature regarding fillers, from an SLP perspective, fillers remain mostly unexplored, or overlooked as noise. Given their link to phenomena such as hesitation and FOAK, their analysis could have an important role to play in educational applications [5]. The aim of this paper is to do a feature study of fillers, to see whether the functions of fillers, can contribute to the prediction of perception of expressed confidence, or FOAK. We think that this task of FOAK prediction, is important and has widespread applicability, given the increasing popularity of automatic processing of educational and job interviews, reviews and speeches. We approach this task from the listener’s perspective, as ultimately, the listener’s perspective in these contexts determines the outcome of the speech.

The research questions are as follows: **RQ1**. Are the different functions of fillers correlated differently to the listener’s FOAK? **RQ2**. Can fillers be informative in the prediction of the listener’s FOAK?

The research on the FOK has focused on the link to prosody [1, 19, 26], fillers [1, 26], facial expressions and gestures [27], and overt lexical cues [12]. By overt lexical cues, we mean words that explicitly mark uncertainty/certainty, such as (*I’m unsure, definitely...*). Similar to our task, are studies that automatically predict points of uncertainty in speech [6, 22].

In affective computing, fillers (such as filler count) are commonly used as an attribute to study persuasiveness [17] and big 5 personality traits [15]. The Computational Paralinguistics Challenge 2013,

focused on sensing fillers, which they considered to be a *social signal* [23]. Research in personality computing has the most consistent correlation, with observations made from speech; including paralinguistic, such as fillers [7, 30]. To our knowledge, only a few studies use fillers as the *focal* feature in machine learning tasks (unless in disfluency detection); fillers are found to be successful in *stance* prediction (stance referring to the subjective spoken attitudes towards something [11]) [13], *turn-taking* prediction [21], and to enhance the *naturalness* of the text-to-speech systems [14].

FOK studies are usually statistical; there has not been a vast interest in the automatic prediction of confidence. These studies are limited to a narrow range of question-answering (QA) tasks [22]. The speaker is typically asked to give an answer (usually the length of a sentence) to a question with a filler inserted, based on a script. The listener is then asked to form a perception based on this answer, and may or may not notice the use of filler due to the short answer length. However, in natural conversation, listener’s are not aware of the use of fillers, unless overused or used in the wrong context [28]. We propose to study free speech in a monologue, which is significantly longer, and has naturally occurring fillers. The annotations that the listener gives are impressions created after hearing the entire monologue (details in Section 3).

This highlights why our task differs from detecting points of uncertainty in speech [6, 22]; as uncertainty detection typically obtains fine grained uncertainty labels based on overt lexical cues, and not the listener’s *overall* impression. Points of uncertainty in speech, could still lead to a general impression of confidence of the speaker. The applications of filler features in affective computing are common, but there aren’t studies that do a focused analysis of them. Certain studies also predetermine based on the context (such as in job interviews, where fluency is desirable) whether the impact of fillers is positive or negative [20], which may not always be the case. When fillers are used as features, typically their count is taken [20], or n-gram features [13], but richer representations of their roles are required. Based on state of the art linguistic literature, we propose a novel way to *represent* our filler features, as outlined in Section 2.

The rest of the paper is organised as follows: In Section 2, we describe how we design our filler features based on linguistic literature. In Section 3, we discuss the dataset used for our experiments. Section 4 outlines the methodology, the experiments and the results of the research questions. Section 5 discusses the conclusion of the experiments.

2 DESIGNING FILLER FEATURES

In the *filler-as-word hypothesis*, a filler functions as an interjection, therefore its meaning is highly dependent on the context. Fillers are hence distinguished by their *basic meaning* and their implied meaning, or *implicature* [2]. We use this hypothesis of *basic and implicature* interjections as a basis for our feature representation. The **basic meaning** of a filler, is to announce the initiation, at t (*filler*) by the speaker, of what is expected to be a delay in speech [2]. Although fillers have one basic meaning; they have several **implicatures**. A filler can have an *unable to proceed* implicature, indicating hesitation, nervousness, or uncertainty of the speaker

[18]. A filler can be used by the speaker in a *syntactic marking* implicature, to indicate pausing to (re)formulate thoughts at discourse boundaries. Fillers can have an implicature linked to a speaker’s *stance*, and *stance polarity* [13]. Fillers are associated with the cognitive load of the speaker [24]; disfluency rates increase for longer sentences, suggesting an increase in cognitive load.

We thus design two sets of filler features; a *basic set* and an *implicature set*. An example of the extracted features is given in Figure 1. The basic set includes three features: f_num , f_uh , f_um , corresponding to the total number of fillers, the total number of filler *uh* and the total number of filler *um*, respectively. The implicature set includes the following features:

- $f_uncertain$: the number of fillers present in sentences that contain stutter markings (fillers to denote an *unable to proceed implicature*).
- f_stance : the number of fillers occurring in sentences that contain tokens (*token* refers to both fillers and words) that are very positive, positive, negative and very negative (to denote fillers associated with *stance*). Token-level stance and polarity labels for the dataset are taken from [8].
- f_start , f_mid : the number of fillers that occur at the start of the sentence and within the sentence respectively (fillers to denote *syntactic marking*).
- f_cog , sen_len : the length of the sentence in tokens (sen_len) and the position of the f_mid in the sentence (as an indication of *Cognitive load*).
- $rate$: the rate of speaking that is computed by dividing the number of tokens by the duration of the video (as a way to measure deletions vs. repetitions : acoustic analysis [25] shows that speakers that use deletions “John likes, *uh* loves Mary”, have a faster speaking rate than speakers that use repetitions “it’s *stutter um* it gets hilarious”).

3 DATASET

We use the Persuasive Opinion Mining (POM) dataset [17], an English multimedia corpus of 1000 movie review videos. Speakers record videos of themselves giving a movie review. After watching the review, annotators are asked to label the video for high-level attributes, such as confidence. Annotators (3 per video) were asked “How confident was the reviewer?”, and had to give a label based on a Likert scale; from 1(Not Confident) to 7(Very Confident), at the granularity of the video level. The inter-annotator score for confidence is high, (Krippendorff’s alpha : 0.73) [17]. Further details can be found in [17].

The corpus has been manually transcribed to include the two fillers *uh* and *um*. *stutter* markings have been annotated, although the annotation context varies, and is dependent on the perception of the transcriber. We highlight relevant information about the dataset in Table 1. The occurrences of the fillers are not sparse, as seen in Table 1, where $\approx 4\%$ of the speech consists of fillers. This is relatively high; the Switchboard [9] dataset of human-human free speech for example, has $\approx 1.6\%$ of tokens that are fillers [24]. The count of each *uh* and *um* filler is roughly the same. Out of the videos used, only 100 of them do not contain fillers.

We take the Root Mean Squared (RMS) value for the 3 annotations per video as the final label, to reflect higher annotation scores.

Video₁

Today I'll be reviewing Saving Private Ryan DVD. (umm) Saving Private Ryan is a very good movie starring Tom Hanks. (f_start) It's about four brothers all with the last name Ryan who are in World War Two, and three of them die and the last one has to be saved because of (uhh) U.S. law. (f_cog)

...

And (umm) The very good movie teaches you a lot about history and about the soldiers and what they went through during World War Two. It's (uhh) one of my favorite movies and (umm) it's won five Academy Awards... (f_mid)

Video₂

I think there's been great reviews (stutter) about this movie and it's very (umm) fun to watch with friends and I (f_uncertain) have always had a fun time watching it.

Figure 1: Example annotation scheme for the fillers. Video transcripts are taken from the corpus.

Description	Value
Videos that contain fillers	792
Total number of videos used	892
Total <i>um</i> fillers in the corpus	4969
Total <i>uh</i> fillers in the corpus	4967
Total fillers in the corpus	9936
Number of tokens in the corpus	230462
% of tokens that are fillers	4.31
Average length (in tokens) of a video	255.9

Table 1: Details about the POM dataset.

We remove the labels in the 1-2 range, due to sparsity of these labels.

4 EXPERIMENTS AND RESULTS

The contextual information provided by fillers, may contain information that is essential for predicting the listener’s FOAK of the speaker. We incorporate such information into our experiments through feature representation that is based on state of the art linguistic literature [1, 2, 13, 18, 24, 26], statistical analysis and linear machine learning models.

We use the transcripts of the dataset, and duration information as our input. We decided not to utilise certain audio features already provided by the CMU-Multimodal SDK, due to the poor results of the forced alignment [16, 32] algorithms. We are not able to pinpoint specific audio regions for fillers.

We compute textual features as a representation of the whole video to be used in statistical analysis and as input to our linear models. We compute f_num , f_um , f_uh , f_start , f_mid , $f_uncertain$ and f_stance by counting each respectively, and then normalising by the total number of tokens in each video. Both sen_len and f_cog are normalised by the average sentence length of the video.

RQ1: Are the different functions of fillers correlated differently to the FOAK?

We use a Pearson’s correlation coefficient between all the filler features and the label of confidence, as shown in Figure 2. We take the pValues for correlation between Confidence and other features, where $p < .001$.

Given Table 2, many of the features show strong correlation with Confidence, which is consistent with previous studies as discussed in Section 1. In figure 2, and we see that there exists a negative correlation between the label of Confidence and the f_num , f_mid . The negative correlation with f_start is smaller than f_mid , meaning




Figure 2: Feature correlation of all features and Confidence using a Pearson’s correlation coefficient. Confidence is given in the last column.

Interjection	Feature name	pValue
Basic	f_num, f_uh, f_um	$P_s < .001$
Implicature	$f_uncertain, f_mid, f_start, rate$	$P_s < .001$
	f_stance	.066
	f_cog	.05
	sen_len	.26

Table 2: pValues for correlation between Confidence and filler features. Highlighted values indicate statistical significance at $p < .001$.

Model	Features	MSE
RV baseline		1.95
RF	Basic	1.11
	Imp	0.85
	Basic + imp	0.84
Ridge	Basic	0.89
	Imp	0.84
	Basic + imp	0.83

Table 3: Results of the models described in Section 4. *Imp.* stands for implicature features.

that there is a distinction between the placement of the fillers. These negative correlations could reflect that listener’s are typically not aware that fillers have been used, unless overused, or used in the wrong context [29]. The higher the filler count is, the more conscious the listener becomes of them, and this will impact the confidence score. Interestingly, $f_uncertain$ is negatively correlated with confidence, indicating that listeners do perhaps perceive these filler-stutter occurrences as speech disturbances [17]. Hence, the functions of fillers do correlate differently with the FOAK.

RQ2: Can fillers be informative in the prediction of the listener’s FOAK?

We use the original standard training, testing and validation folds provided in the CMU-Multimodal SDK [33]. Our baseline is a *Random Vote* (RV); where 100 random draws respecting the train dataset balance were made. We use a mean squared error (MSE) to evaluate our models. For RV, the MSE is averaged over these 100 samples. We take 2 classic Machine learning algorithms, that is *Random Forest* (RF) and *Ridge regression* (RR), with respective hyper-parameter searches on the validation set. We choose RR as we have multi-collinear features, and both RF and RR have easy interpretability for feature importance.

The main experiments and their results are listed in Table 3. Please refer to Section 2 for the description of the features. In order to get further insight into which features seem relevant to our task, we utilise a RF feature importance, as shown in Table 4. In Table 3, the MSE for the models that use basic filler features, is higher than our baseline in predicting the FOAK. We can see that the impact of adding the implicature features decreases the MSE.

Looking at Table 4, we see that the rate of speech (*rate*) is the most important feature. In [25], it was observed that faster speakers “get ahead of themselves”; using more deletions in their speech and beginning anew, whereas slower speakers take more time to plan, increasing hesitations such as repetitions. We interpret that slower speakers may have more moments of hesitations in the videos, and *rate* may be a better indicator of this than $f_uncertain$. However, since rate is also shown as an important general speech feature in personality computing [30], it is not surprising that *rate* is ranked highest.

We would like to focus on the two sets of features that are highlighted in the Table 4, which are the syntactic feature f_mid , and the cognitive load measured by f_cog and sen_len . These are both features that relate to planning and structure of the speech. Given the top 5 features, we interpret that the functions of fillers that are given importance by the models are the *structural functions* of fillers, and not the *metacognitive functions*, such as (f_stance , $f_uncertain$). FOAK could be related to how well formulated the speaker is, and this could be captured by our model. We interpret that the impression the listener has of the speaker could be based on how well structured their argument was.

5 CONCLUSIONS

In this paper we presented a feature study of the role of fillers in the prediction of the FOAK. The contributions of this paper are: 1. To introduce a new challenging task for educational application (FOAK/Confidence prediction) 2. To design filler features and observe their role in the FOAK, in the context of free speech, and 3. To

Rank	Feature	Importance
1	rate	0.210
2	f_num	0.110
3	sen_len	0.104
4	f_mid	0.102
5	f_cog	0.088

Table 4: Top 5 features calculated for the RF model.

highlight that deeper representations of fillers could be beneficial to SLP/affective computing tasks. Future work will be dedicated to expand our features set by mixing filler features with text embedding representations and acoustic features to develop a FOAK prediction system. Given our results, we would also like to investigate the role of discourse features in FOAK prediction. For the acoustic filler features, this will require a corpus with precise speech and text alignment. It would be interesting to see if other SLP tasks benefit from a more fine-grained representation of fillers at a syntactic level, especially if it is able to capture structural elements of a speaker’s argument, given that spoken language often lacks structure.

ACKNOWLEDGEMENTS

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 765955.

REFERENCES

- [1] Susan E Brennan and Maurice Williams. 1995. The feeling of another’s knowing: Prosody and filled pauses as cues to listeners about the metacognitive states of speakers. *Journal of memory and language* 34, 3 (1995), 383–398.
- [2] Herbert H Clark and Jean E Fox Tree. 2002. Using uh and um in spontaneous speaking. *Cognition* 84, 1 (2002), 73–111.
- [3] Martin Corley, Lucy J MacGregor, and David I Donaldson. 2007. It’s the way that you, er, say it: Hesitations in speech affect language comprehension. *Cognition* 105, 3 (2007), 658–668.
- [4] Martin Corley and Oliver W Stewart. 2008. Hesitation disfluencies in spontaneous speech: The meaning of um. *Language and Linguistics Compass* 2, 4 (2008), 589–602.
- [5] Tanvi Dinkar, Ioana Vasilescu, Catherine Pelachaud, and Chloé Clavel. 2018. Disfluencies and Teaching Strategies in Social Interactions Between a Pedagogical Agent and a Student: Background and Challenges.
- [6] Jeroen Dral, Dirk Heylen, and Riëks op den Akker. 2011. Detecting uncertainty in spoken dialogues: an exploratory research for the automatic detection of speaker uncertainty by using prosodic markers. In *Affective Computing and Sentiment Analysis*. Springer, 67–77.
- [7] Paul Ekman, Wallace V Friesen, Maureen O’Sullivan, and Klaus Scherer. 1980. Relative importance of face, body, and speech in judgments of personality and affect. *Journal of personality and social psychology* 38, 2 (1980), 270.
- [8] Alexandre Garcia, Slim Essid, Florence d’Alché-Buc, and Chloé Clavel. 2019. A multimodal movie review corpus for fine-grained opinion mining. *CoRR* abs/1902.10102 (2019). arXiv:1902.10102 <http://arxiv.org/abs/1902.10102>
- [9] John J Godfrey, Edward C Holliman, and Jane McDaniel. 1992. SWITCHBOARD: Telephone speech corpus for research and development. In *[Proceedings] ICASSP-92: 1992 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Vol. 1. IEEE, 517–520.
- [10] H Paul Grice, Peter Cole, Jerry Morgan, et al. 1975. Logic and conversation. *1975* (1975), 41–58.
- [11] Pentti Haddington. 2004. Stance taking in news interviews. *SKY Journal of Linguistics* 17 (2004), 101–142.
- [12] Xiaoming Jiang and Marc D Pell. 2017. The sound of confidence and doubt. *Speech Communication* 88 (2017), 106–126.
- [13] Esther Le Grezause. 2017. *Um and Uh, and the expression of stance in conversational speech*. Ph.D. Dissertation.
- [14] Yaniv Leviathan and Yossi Matias. 2018. Google Duplex: An AI System for Accomplishing Real World Tasks Over the Phone. *Google AI Blog* (2018).

- [15] François Mairesse, Marilyn A Walker, Matthias R Mehl, and Roger K Moore. 2007. Using linguistic cues for the automatic recognition of personality in conversation and text. *Journal of artificial intelligence research* 30 (2007), 457–500.
- [16] R. M. Ochshorn and M. Hawkins. [n.d.]. Gentle forced aligner [computer program]. ([n.d.]).
- [17] Sunghyun Park, Han Suk Shim, Moitreya Chatterjee, Kenji Sagae, and Louis-Philippe Morency. 2014. Computational analysis of persuasiveness in social multimedia: A novel dataset and multimodal prediction approach. In *Proceedings of the 16th International Conference on Multimodal Interaction*. ACM, 50–57.
- [18] Joseph P Pickett. 2018. *The American heritage dictionary of the English language*. Houghton Mifflin Harcourt.
- [19] Heather Pon-Barry. 2008. Prosodic manifestations of confidence and uncertainty in spoken language. In *Ninth Annual Conference of the International Speech Communication Association*.
- [20] Sowmya Rasipuram and Dinesh Babu Jayagopi. 2016. Automatic assessment of communication skill in interface-based employment interviews using audio-visual cues. In *2016 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*. IEEE, 1–6.
- [21] Divya Saini. 2017. The Effect of Speech Disfluencies on Turn-Taking.
- [22] Tobias Schrank and Barbara Schuppler. 2015. Automatic detection of uncertainty in spontaneous German dialogue. In *Sixteenth Annual Conference of the International Speech Communication Association*.
- [23] Björn Schuller, Felix Weninger, Yue Zhang, Fabien Ringeval, Anton Batliner, Stefan Steidl, Florian Eyben, Erik Marchi, Alessandro Vinciarelli, Klaus Scherer, et al. 2019. Affective and behavioural computing: Lessons learnt from the First Computational Paralinguistics Challenge. *Computer Speech & Language* 53 (2019), 156–180.
- [24] Elizabeth Shriberg. 2001. To ‘errrr’ is human: ecology and acoustics of speech disfluencies. *Journal of the International Phonetic Association* 31, 1 (2001), 153–169.
- [25] Elizabeth E Shriberg. 1999. *Phonetic consequences of speech disfluency*. Technical Report. SRI INTERNATIONAL MENLO PARK CA.
- [26] Vicki L Smith and Herbert H Clark. 1993. On the course of answering questions. *Journal of memory and language* 32, 1 (1993), 25–38.
- [27] Marc Swerts and Emiel Krahmer. 2005. Audiovisual prosody and feeling of knowing. *Journal of Memory and Language* 53, 1 (2005), 81–94.
- [28] Gunnel Tottie. 2014. On the use of uh and um in American English. *Functions of Language* 21, 1 (2014), 6–29.
- [29] Gunnel Tottie. 2014. On the use of uh and um in American English. *Functions of Language* 21, 1 (2014), 6–29.
- [30] Alessandro Vinciarelli and Gelareh Mohammadi. 2014. A survey of personality computing. *IEEE Transactions on Affective Computing* 5, 3 (2014), 273–291.
- [31] Etsuko Yoshida and Robin J Lickley. 2010. Disfluency patterns in dialogue processing. In *DISS-LPSS Joint Workshop 2010*.
- [32] Jiahong Yuan and Mark Liberman. 2008. Speaker identification on the SCOTUS corpus. *Journal of the Acoustical Society of America* 123, 5 (2008), 3878.
- [33] Amir Zadeh. [n.d.]. CMU-Multimodal SDK. https://github.com/A2Zadeh/CMU-MultimodalSDK/blob/master/mmsdk/mmdatask/dataset/standard_datasets/POM/pom_std_folds.py.