

WHAT MAKES A GOOD PAUSE? INVESTIGATING THE TURN-HOLDING EFFECTS OF FILLERS

Bing'er Jiang, Erik Ekstedt, Gabriel Skantze

KTH Speech, Music and Hearing, Stockholm, Sweden
{binger, erikekst, skantze}@kth.se

ABSTRACT

Filled pauses (or *fillers*), such as *uh* and *um*, are frequent in spontaneous speech and can serve as a turn-holding cue for the listener, indicating that the current speaker is not done yet. In this paper, we use the recently proposed Voice Activity Projection (VAP) model, which is a deep learning model trained to predict the dynamics of conversation, to analyse the effects of filled pauses on the expected turn-hold probability. The results show that, while filled pauses do indeed have a turn-holding effect, it is perhaps not as strong as could be expected, probably due to the redundancy of other cues. We also find that the prosodic properties and position of the filler has a significant effect on the turn-hold probability. However, contrary to what has been suggested in previous work, there is no difference between *uh* and *um* in this regard.

Keywords: Hesitation, fillers, turn-taking, spoken dialog, computational modelling

1. INTRODUCTION

In spontaneous speech, filled pauses (or *fillers*), such as *uh* and *um*, are frequent and are typically associated with a hesitation on part of the speaker [1]. Some studies have found an association between the use of fillers and higher cognitive load in both monolog and dialog [2]. From the perspective of coordinating turn-taking in conversation [3, 4, 5], it has been suggested that fillers work as a *turn-holding cue* for the listener, indicating that the current speaker is not done yet [6, 7]. To what extent such fillers are produced intentionally for this purpose, and how much information they carry, is not yet clear [8].

According to [7], *uh* and *um* are to be considered as words in the speaker's vocabulary. While these words do not have any propositional content, it is argued that fillers conform to the "phonology, prosody, syntax, semantics, and pragmatics of English words". In this view, they should be seen as interjections (similar to *ah* and *oh*), indicating that

the speaker wants to keep the floor. In addition, [7] claim that *uh* and *um* have different meanings, where the latter signals a longer upcoming delay. This claim is based on an analysis of the pause length following these fillers in several spontaneous speech corpora. Later studies of other corpora have also found such differences [2], while others have not [9]. Another possibility is that it is not the lexical form of the filler that is important for their function in signalling delay, but rather their prosodic characteristics. If so, it would perhaps be wrong to regard them as carrying distinct symbolic meaning. However, we are not aware of any studies that have investigated this.

A general problem with corpus-based analyses is that we can only find correlations, which do not necessarily imply causation. Thus, if we compare places with filled pauses with places without them, we do not know whether it is the filled pauses that actually cause the observed differences, or whether they were due to other factors that happened to correlate with the usage of fillers. One alternative is to perform perception experiments where stimuli are systematically manipulated, while all other factors are kept constant (e.g. [10, 6, 11, 12]). However, such experiments are costly to perform on a larger scale and they are often done with a small set of made-up examples that do not necessarily reflect general distributions. In this paper, we take a third route: To train a model that makes predictions in conversation and then use that model to investigate the effects of large-scale systematic stimuli manipulation.

Using the recently proposed **Voice Activity Projection** model [13] (VAP), which models the dynamics of conversation and can be used to predict the **turn-hold probability (THP)** (i.e., the probability that the turn will not shift to the other speaker) at any given point, we systematically remove and insert fillers in order to investigate their effects. We want to answer the following questions:

1. How is the THP affected when fillers are removed from their natural occurrences?
2. How is the THP affected when fillers are inserted at places which should be clearly turn-

yielding?

- How do the properties of the filler (position, lexical form, pitch, intensity, duration) affect the THP?

2. VOICE ACTIVITY PROJECTION

Voice Activity Projection [13], VAP, is a predictive model of conversational dynamics developed for turn-taking research. The model objective is to continuously predict the upcoming voice activity of both speakers in a dialog. The voice activity is defined in binary terms (i.e., whether they will be speaking or not at a certain point in time), and the two speakers' future activities are jointly encoded into a discrete label that represents the next 2s of dialog.

Similar to a language model (text-to-text), the VAP model (speech-to-activity) is optimized through MLE (cross-entropy). The model processes raw spoken dialog audio and outputs a likelihood distribution over each of the discrete activity labels. This likelihood can then be used for controlling the turn-taking in a spoken dialog system, or as a tool to analyze the likelihood of turn-taking events in recorded dialogs. In previous work, it has been shown that the model is sensitive to prosodic cues [14]. In this work, we utilize the VAP model to investigate the effect of a filler has on how long the current turn holds until the other speaker takes the turn.

The activity projection is defined by eight time window bins of a 2s window, four for each speaker (0.2s, 0.4s, 0.6s, 0.8s respectively), which indicates if the speaker is active or not in that time window. The various combinations of these bins result in $2^8 = 256$ discrete activity labels. While these labels represent various turn-taking-related events (such as turn-shifts and backchannels), we are here primarily interested in calculating the probability of who will be the next speaker. To do this, we take a weighted average over the probability assigned to each label, focus on the bins that cover our region of interest (0-600ms), and compare the contributions for each speaker. The result is a value between 0 and 1 that represents the probability that the current speaker will continue, defined as **turn-hold probability (THP)**. In our study, we apply the model on dialogs and get incremental predictions of THP.

The model used in this experiment is a stereo version of the original VAP model that operates on two separate waveforms (one for each speaker) and is trained on subsets of the Fisher part 1 and the Switchboard corpus [15, 16]. We define a held-




Figure 1: The THP following an utterance ending with *uh*. Top: the spectrogram of the spoken utterance; Middle: the THP with the filler; Bottom: the THP without the filler. The blue/yellow areas corresponds to A/B being the most likely next speaker. The black number bars represent the duration of silence before a turn-shift prediction (THP falls under 0.5).

out test set containing 577 sessions from the dialog act annotated subset [17], swda, of the switchboard corpus to be used in the subsequent analysis.

3. EXPERIMENT 1: EXCLUDING FILLERS

We evaluate the effect of fillers through examining how 'soon' the VAP model predicts a turn shift after the end of a speech segment. In Experiment 1, we investigate the effect of excluding a valid filler present in the data. As illustrated in Figure 1, we apply the model to 20s of dialog (a pair of data samples including and excluding the filler), followed by 10s of artificial silence, and track the THP over the silence segment. While a pause of 10s is very unlikely in real life conversations, we use a longer length anyways to fully examine the prediction of the turn shift. The time t where the THP drops below 50% is considered the position of turn shift.

We select filler candidates *uh* and *um*, based on word-level transcriptions, and define three criteria in order to consider them valid fillers: 1) the duration of the candidates must be longer than 0.2s; 2) the fillers should be followed by a pause of at least 0.2s of the same speaker and 'isolated' from activity of the listener of at least 1s, before and after the filler; 3) there should be at least 20s of past dialog context (for the model to make its predictions). This results in 5316 valid fillers from the 577 test sessions.

We use *Survival Analysis* to examine the results, as some cases may not reach the turn shift predicted by the model even at the end of the 10s silence. The *Survival Analysis* includes a *censoring* status to




Figure 2: Kaplan-Meier curve of turn hold predictions through 10s of silence for Experiment 1 (excluding fillers), with or without fillers.

identify whether the expected event has happened or not (*uncensored* vs. *censored*). In our case, dialogs that are not predicted to reach turn shift by the end of the silence are marked as *censored*.

In order to examine whether the filler makes a difference, we plot the Kaplan-Meier curve of the two conditions. Figure 2 shows the proportion of data samples that keep the hold (y-axis) as a function of time (x-axis). Note that the y-axis does *not* reflect the THP; instead, it refers to percentages of cases where the THP are above 50%, i.e., percentages of data that are likely to keep the hold at time t .

Figure 2 shows that the filler indeed makes a difference in holding the turn, confirmed by the log-rank test ($p < 0.001$). Moreover, the effect changes throughout time. In particular, the filler does not help holding the turn at the very beginning of the silence, as shown by the drastic drop of the green line. At the later stage of the silence (after 1.2s), the effect of keeping the hold by the filler becomes more prominent, reaching the greatest between 4s and 5s, and then decreases until the end of the silence.

While there is a clear effect of the filler, it is interesting to note that the effect is perhaps not as big as expected; the turn hold effect is quite strong even with the filler removed. This is likely due to other redundant turn-holding cues present before the filler (prosody, syntax, etc).

4. EXPERIMENT 2: INSERTING FILLERS

While Experiment 1 indeed shows a turn-holding effect of existing fillers, it is unclear how strong the effect is, given the confounded turn-holding context. In Experiment 2, we further investigate the effects of fillers in contexts that are clearly turn-yielding. Specifically, we insert fillers at the end of yes/no questions and examine the difference in THP between the original data samples and the ones with

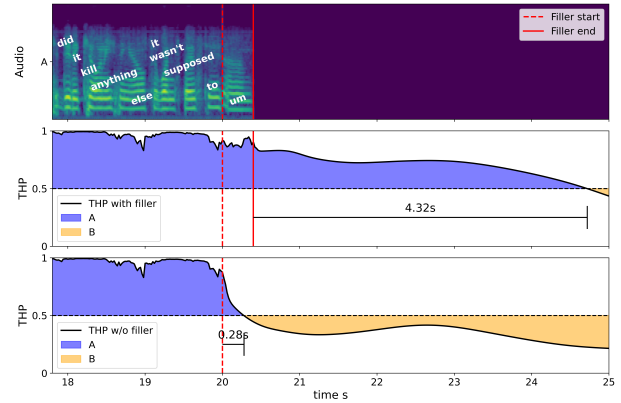


Figure 3: The THP following a yes/no question. Top: the spectrogram of the spoken utterance; Middle: the THP with the insertion of the filler *um*; Bottom: the THP of the original utterance.

inserted filler at the end.

We define valid utterances by the following conditions: we first select all utterances that are labeled as a yes/no-question dialog act according to the *swda* annotation [17]. The original utterance annotations may include multiple dialog act segments in a single utterance and we omit those that do not end with a yes/no-question dialog act. Valid utterances must be followed by a pause longer than 0.5s (from the same speaker) and ‘isolated’ from activity of the listener of at least 0.5s, before and after the end of the utterance. The candidates with a context duration of less than 20s are omitted. The final condition is that the model assigns a turn-shift in the first 5s of silence when 20s of context up and including the utterances are processed. This approach result in 245 valid turn-shift utterances.

For each valid utterance we select all the valid fillers, from the same speaker and session, as candidates to add to the utterance (1436 total filler candidates). Similarly to Experiment 1, we create paired data samples constituting 20s of dialog context, up to and including the utterances (with or without the added filler), followed by 10s of silence. This is illustrated in Figure 3.

Following the same preprocessing as described in Experiment 1, we plot the Kaplan-Meier curve, as shown in Figure 4. The figure shows that inserting a filler consistently lengthens the hold, and the two curves differ significantly from the log-rank test ($p < 0.001$). As can be seen, for the original yes/no questions, the model clearly predicts a fairly quick turn shift in most cases. While adding a filler does not flip all of them to turn-holds, the fillers seem to have a stronger effect here than in Experiment 1, and so the model clearly is sensitive to them.




Figure 4: Kaplan-Meier curve of turn hold predictions through 10s of silence for Experiment 2 (inserting fillers after yes/no questions), with or without fillers.

5. EXPERIMENT 3: PROSODIC AND LEXICAL EFFECTS

Experiment 3 examines how prosodic and lexical properties of the filler affect the turn hold. For capturing prosodic information, we include by-speaker standardized mean **pitch**, by-speaker standardized mean **intensity** and filler **duration** on the log scale. We also examine whether the **lexical form** (*uh* or *um*), as well as the **position** of the filler in the current dialogue act (*mid* or *start*) has any influence on the THP. All continuous variables (prosody) are standardized by subtracting the mean and divided by two standard deviations. Note that the pitch and intensity variables are normalized with regard to each speaker, while duration is not.

We fit a Cox proportional hazards model on data with fillers from Experiment 1 to examine the joint effects from the five factors. The independent variables are the five factors and an interaction term between F0 and the lexical form, and the dependent variables are time reaching the turn shift derived from the VAP model and the censor status.

Table 1, summarizing the Cox model, shows that except for the filler type, all four other variables are statistically significant. The coefficients reflect the change of hazards function (here, the probability of the turn shift): a negative coefficient estimate corresponds to less ‘hazard’ of turn shift (higher probability of turn hold), while a positive coefficient estimate indicates an increased probability of turn shift. The exponentiated coefficients are termed as *hazards ratio*, which represents the difference in the hazard of an event at any given time.

The results from Table 1 reveals that the following factors help keep the hold (decrease the ‘hazard’): higher pitch, start of the dialog act, higher intensity and longer duration. Moreover, while the lexical

	coef	coef(exp)	SE	Pr(> z)
F0	-0.725	0.484	0.246	0.003
Intensity	-0.127	0.879	0.035	0.0003
Lex _{um}	-0.077	0.925	0.050	0.12
Duration	-0.118	0.888	0.037	0.001
Pos _{mid}	-0.305	0.736	0.065	<0.0001
F0:Lex _{um}	-1.237	0.237	0.290	0.007

Table 1: Model summary of the Cox regression model. Bold *p* values are significant.

form of the filler itself does not significantly affect the hold, it has an indirect influence through pitch: *uh* has higher pitch and thus results in longer holds.

Specifically, if a speaker raises their voice by 0.5 standard deviation (SD) of their pitch range, the probability of holding the turn determined by the model will increase by 51.6% (here, the change of hazards of turn shift from *hazards ratio* is $0.484 - 1 = -0.516$). Increasing 0.5 SD of the intensity, however, only raises the probability by 12.1%. The duration also affects the turn hold: if the filler is lengthened by 0.5 SD, the probability is raised by 11.2%. In terms of the position, if a filler is at the beginning of the dialog act, it tends to raise the probability by 26.4% compared to one in the middle.

6. DISCUSSION AND CONCLUSIONS

In this study, we use computational modelling to evaluate the effects of fillers on holding a turn, and examine a variety of prosodic, positional and lexical properties of the fillers. Experiment 1 shows that excluding existing fillers indeed results in shorter holds, while the effect is perhaps not as strong as expected. This is likely due to the redundancy of other turn-holding cues. Experiment 2 further reveals that fillers have a strong turn-holding effect when added to clearly turn-yielding contexts (yes/no questions). Experiment 3 shows that a number of properties of the filler contribute to the turn-hold effect: higher pitch, higher intensity, and longer duration of the filler are associated with longer turn hold. Moreover, utterance-initial fillers have a stronger turn-hold effect than mid-utterance fillers. However, there is no significant difference between the lexical forms *uh* and *um*. Nevertheless, the two types of fillers still indirectly affects the THP through increasing the pitch of the filler. Thus, instead of the lexical forms, it is their phonetic realization that makes a difference, as *um* includes a nasal which associates with lower pitch compared to the oral *uh*. This speaks against the interpretation given by [7], where *uh* and *um* are seen as different words with different meanings.

7. ACKNOWLEDGEMENTS

This work was supported by Riksbankens Jubileumsfond (RJ), through the project *Understanding predictive models of turn-taking in spoken interaction* (P20-0484), as well as the Swedish Research Council, through the project *Prediction and Coordination for Conversational AI* (2020-03812).

8. REFERENCES

- [1] J. E. Fox Tree, "Interpreting Pauses and Ums at Turn Exchanges," *Discourse Processes*, vol. 34, no. 1, pp. 37–55, Jul. 2002.
- [2] R. L. Rose, "Um and Uh as Differential Delay Markers: The Role of Contextual Factors," in *Disfluency in Spontaneous Speech (DiSS)*, 2015.
- [3] H. Sacks, E. Schegloff, and G. Jefferson, "A simplest systematics for the organization of turn-taking for conversation," *Language*, vol. 50, pp. 696–735, 12 1974.
- [4] S. Duncan, "Some Signals and Rules for Taking Speaking Turns in Conversations," *Journal of Personality and Social Psychology*, vol. 23, no. 2, pp. 283–292, 1972.
- [5] G. Skantze, "Turn-taking in Conversational Systems and Human-Robot Interaction : A Review," *Computer Speech & Language*, vol. 67, 2021.
- [6] P. Ball, "Listeners' responses to filled pauses in relation to floor apportionment," *British Journal of Social and Clinical Psychology*, vol. 14, no. 4, pp. 423–424, 1975.
- [7] H. Clark and J. E. Fox Tree, "Using uh and um in spontaneous speaking," *Cognition*, vol. 84, no. 1, pp. 73–111, 2002.
- [8] M. Corley and O. W. Stewart, "Hesitation Disfluencies in Spontaneous Speech: The Meaning of um," *Language and Linguistics Compass*, vol. 2, no. 4, pp. 589–602, 2008.
- [9] D. C. O'Connell and S. Kowal, "Uh and um revisited: Are they interjections for signaling delay?" *Journal of Psycholinguistic Research*, vol. 34, no. 6, pp. 555–576, 2005.
- [10] M. Cook and M. G. Lalljee, "The interpretation of pauses by the listener," *British Journal of Social and Clinical Psychology*, vol. 9, pp. 375–376, 1970.
- [11] J.-P. de Ruyter, H. Mitterer, and N. J. Enfield, "Projecting the end of a speaker's turn: A cognitive cornerstone of conversation," *Language*, vol. 82, no. 3, pp. 515–535, 2006.
- [12] S. Bögels and F. Torreira, "Listeners use intonational phrase boundaries to project turn ends in spoken interaction," *Journal of Phonetics*, vol. 52, pp. 46–57, Sep. 2015.
- [13] E. Ekstedt and G. Skantze, "Voice Activity Projection: Self-supervised Learning of Turn-taking Events," in *Proc. Interspeech 2022*, 2022, pp. 5190–5194.
- [14] —, "How Much Does Prosody Help Turn-taking? Investigations using Voice Activity Projection Models," in *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*. Edinburgh, UK: Association for Computational Linguistics, Sep. 2022, pp. 541–551.
- [15] C. Cieri, D. Miller, and K. Walker, "The fisher corpus: a resource for the next generations of speech-to-text," in *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*. Lisbon, Portugal: European Language Resources Association (ELRA), May 2004.
- [16] J. J. Godfrey, E. C. Holliman, and J. McDaniel, "Switchboard: Telephone speech corpus for research and development," in *Proceedings of the 1992 IEEE International Conference on Acoustics, Speech and Signal Processing*. USA: IEEE Computer Society, 1992, p. 517–520.
- [17] D. Jurafsky, E. Shriberg, and D. Biasca, "Switchboard SWBD-DAMSL shallow-discourse-function annotation coders manual, draft 13," University of Colorado, Boulder Institute of Cognitive Science, Boulder, CO, Tech. Rep. 97-02, 1997.