

---

# Fillers in Spoken Language Understanding: Computational and Psycholinguistic Perspectives

Tanvi Dinkar\* — Chloé Clavel\*\* — Ioana Vasilescu\*\*\*

\* Heriot Watt University, *T.Dinkar [at] hw.ac.uk*

\*\* Télécom Paris, *chloe.clavel [at] telecom-paris.fr*

\*\*\* University Paris-Saclay, *ioana [at] limsi.fr*

---

*ABSTRACT. Disfluencies are ubiquitous to spoken discourse. Fillers (“uh”, “um”, ...) occur the most frequently compared to other kinds of disfluencies. Yet, to the best of our knowledge, there isn't a resource that brings together the research perspectives influencing Spoken Language Understanding (SLU) on these speech events. The aim of this article is to synthesise a breadth of perspectives in a holistic way; i.e. from underlying (psycho)linguistic theory on fillers, to their annotation and consideration in Automatic Speech Recognition (ASR) and SLU systems, to lastly, their study from a generation and Text-to-Speech (TTS) standpoint. The article aims to present the perspectives in an approachable way to the SLU and Conversational AI community, and discuss what we believe are the trends and challenges in each area.*

*KEYWORDS: Disfluencies, Fillers, Spoken Language Understanding.*

*RÉSUMÉ. Les disfluences sont omniprésentes dans le discours, et les fillers ("euh", ...) sont le type de disfluence le plus fréquent. Pourtant, il n'existe aucune ressource qui rassemble les perspectives de recherche sur ces événements discursifs dans le cadre de la compréhension de la langue parlée (CLP). L'objectif de cet article est de synthétiser un large éventail de perspectives de manière holistique, comprenant la théorie (psycho-)linguistique des fillers, leur annotation et leur prise en compte dans les systèmes de reconnaissance automatique de la parole et de CLP, ainsi que leur étude dans le cadre de la génération. L'article a pour but de présenter ces perspectives de manière accessible à la communauté de la CLP et des systèmes conversationnels, et de discuter de ce que nous voyons comme les tendances et défis de chaque domaine.*

*MOTS-CLÉS: Disfluences, Fillers, Compréhension de la Langue Parlée*

---

## 1. Introduction

Speech production is a complex process; i.e. “respiratory, phonatory, and articulatory gestures are timed in such a way as to produce an acoustic signal that adequately conveys the intended message both quickly and smoothly” (Lickley, 2015). *Disfluencies* can be thought of as elements that break this fluent flow of speech during speech production. Formally, they are interruptions in the regular flow of speech, such as pausing silently, repeating words, or interrupting oneself to correct something said previously (Fraundorf *et al.*, 2018). They occur between intentional signals (a gesture used to point at an object, ...) and *unintentional* signals (slips of the tongue, ...) of communication (Corley and Stewart, 2008). With the increasing popularity of voice assistant technologies, an open challenge that remains is the ability to design systems that can comprehend the different signals of speech communication. However, when taking stock of a machine’s capability to understand language, there is an emphasis on the learning of *forms* (such as in Language Modelling (LM), where the task is string prediction), but not on *meaning*; or the relationship between linguistic form and communicative intent (Bender and Koller, 2020). This issue is now widely acknowledged, for instance, the UnImplicit workshop<sup>1</sup> state as motivation “... an important question that remains open is whether such methods are actually capable of modelling how linguistic meaning is shaped and influenced by *context*, or if they simply learn superficial patterns that reflect only *explicitly* stated aspects of meaning ...”. From a research standpoint, this is especially challenging when considering disfluencies, which can have *implicit* and *contextual* meanings – such as when a speaker says “uh”, “er” and so on. Hence, this work is motivated by the present observations:

**People rarely speak in the same manner with which they write.** As Bailey and Ferreira (2003) state, “The processes involved in speaking and in writing differ substantially from each other, and so the products of the two systems are not the same”. The following is an example of a transcription taken from a corpus of conversational speech:

**A:** For a while there I, I, I, uh, subscribed to *New York Times*, a-, actually a couple of newspapers because, uh, you know, my fiance, well, she was unemployed for a while ...

**B:** Uh-huh.

**A:** ...so she, you know, really needed to look at the, the, want-, help wanted ads.

The above transcript is not *fluent*, and not easily *readable*. Already, some modifications have been done to the raw transcript keeping in mind the goal of readability. For instance, punctuation markers have been added to introduce sentence structure and capture the prosodic cues used by the speaker. In the transcripts, we can see that people tend to repeat themselves (“I, I, I”), interrupt each other (“Uh-huh”), rapidly shift the focus of the topic in a conversation (from newspaper subscriptions to unemployment), and are in general, *disfluent*. One of the departures from written text

1. <https://unimplicit.github.io>.

is the presence of disfluencies. Disfluencies are frequent in speech, as fluent speech is rarely the norm; with estimates that natural human-human conversations comprise of  $\approx 5 - 10\%$  of disfluencies (Shriberg, 1994). *Thus disfluencies are ubiquitous to spontaneous speech.*

**Despite this, there are varying attitudes in the treatment of disfluencies depending on the field,** that influence research in Spoken Language Understanding (SLU). Generally, the field of psycholinguistics (dealing with communicative and cognitive aspects of language) focuses on the role disfluencies play in the production and comprehension process of speech, with many works to show their importance in speech communication. For instance, they inform us about the linguistic structure of an utterance: such as in the (difficulties of) selection of appropriate vocabulary while circumventing interruption. Other perspectives, such as a *computational* one, focus on recognising disfluencies in order to remove them – for the improvement of automatic speech recognition (ASR) systems.

**Furthermore, “disfluencies” is an umbrella term used to describe a wide variety of communicative phenomena.** It is difficult to find an overarching definition of disfluency due to several confusions in terminology. Lickley (2015) categorises all taxonomies of disfluencies as based on *form* or *function*. Terminology to describe disfluencies may be based on form when the objective is to describe the “patterns of words and syntactic units that disfluencies display”. On the other hand, function descriptions may be used when disfluencies are described with respect to the planning processes involved in speech production, deliberating over the reasons for a departure from fluency (though overlap exists between the two categories). For instance, “uh” may be described as an interjection (word or phrase independent of grammatical connections to other words or phrases), or a “filled pause”/ “filler” (sound filling a pause in flow of speech) when considering form, or a “hesitation” when considering function. When “uh” is used as a (meta-) “discourse marker” (a unit of talk that brackets speech (Schiffrin, 1987)), it considers both form and function – i.e. functions to mark discourse structure, and occurs at discourse boundaries (form). Other examples of functional disfluencies can include “self”/ “other-repair” (“repair” specifically used to indicate that something has gone wrong in the planning process, and needs to be corrected). For more details of annotation used for disfluency detection tasks (which is typically form-based), please refer to Sec. 3.1 Annotation for Disfluency Detection.

**This work is concerned with fillers** (“uh”, “um” ...), compared to other disfluencies. These speech events may not contribute to the final message when considering purely a lexical level (Vasilescu *et al.*, 2010), and thus do not have *explicit* meaning (Meteer *et al.*, 1995). Yet, they occur with *high frequency* in speech datasets, compared to other structures of disfluencies. Shriberg (2001) shows that the number of fillers per word across corpora exceed other kinds of disfluencies. Additionally, they occur in an *intersection between speech and text*. Fillers are a common property of spontaneous speech and are shown to have *distinct acoustic characteristics/ paralinguistic properties* (Shriberg, 1999; Vasilescu *et al.*, 2010), yet they can be transcribed in text, without the requirement of more detailed annotation schemes. Thus, with the

growing demand of voice assistant technologies, there is an increase in research that deals with spontaneous speech corpora – where naturally, *these tokens will frequently occur*.

**Yet so far, there isn't a body of work that brings together the research relevant to SLU on these specific disfluencies.** To the best of our knowledge, there is no *introductory work* that brings together several perspectives surrounding fillers, i.e. from psycholinguistic to computational. While a linguistic versus computational perspective is discussed in other works (Lickley, 1994), there is a narrow focus within this perspective on disfluency detection for the purposes of removal. This ignores the broader applications of SLU that deal with *social communication*. Additionally, while other linguistic works give rigorous details about the annotation of disfluencies (Nicholson, 2007; Grosman, 2018), including the history behind these schemes (Lickley, 1994), they are not accessible for the SLU community. This work is more targeted towards researchers in SLU and dialogue (or indeed, linguists interested in the work on fillers and disfluencies in SLU). Given the recent benchmark release of a filler detection task and dataset (Zhu *et al.*, 2022), there is a growing interest in the community on these unique disfluencies.

Thus, **the aim of this paper** is to bring together a *breadth* of perspectives on these spontaneous speech phenomena in a holistic way, considering fillers at different stages in a pipeline; i.e. from ASR to generation. Furthermore, we draw parallels in the fields of psycholinguistics and SLU; discussing research in both fields that considers fillers informative, noise . . . Throughout the work, we discuss the challenges in each field; from issues in safety and robustness, to a lack of contextual analysis and feasibility in in-the-wild scenarios. To do so, we contrast computational approaches to disfluencies distinguished from psycholinguistic approaches, loosely adopted from Lickley (1994). Broadly, psycholinguistic theories deal with the communicative and cognitive aspects of language. For instance, Levelt (1983) studies how speakers monitor and correct their speech, and in turn, how listeners are able to integrate new material correcting the previous material. Thus in Sec. 2 we focus on psycholinguistic research that studies the role disfluencies in the *planning/production* of speech, and the *comprehension* of speech. As discussed, we focus on the research surrounding fillers; highlighting works that study them as informative signals of communication. However, we also present research on other disfluencies when relevant – this is to discuss general concepts and findings from the field. Then, in Sec. 3, we discuss the computational perspectives on disfluencies; i.e. approaches more concerned with the *recognition/processing* of disfluencies. Here, we discuss how the treatment of disfluencies can vary depending on the *type of task* in SLU. Sec. 4 then gives the conclusion of the article.

**This paper can be read non-linearly by referring to the legend given in Table 1.** It is difficult to create a one-to-one mapping in the fields, given that the methodologies and approaches to each field are different. However, the intention behind some works are similar, and we believe that they can be grouped together as complementary perspectives.

	<b>Psycholinguistic</b>	<b>Computational</b>
<i>Production</i>	2.1.1 Cognitive load 2.1.2 Communicative Function	3.2.1 Broader SLU
<i>Comprehension</i>	2.2.1 Informative cues 2.2.3 Results on incremental processing 2.2.4 Time-buying measures	3.2.2 Generation
<i>Noisy channel</i>	2.2.2 Filtered out noise	S3 (intro) Computational Perspectives S3.1 SLU for SDS
<u>Note on annotation</u> S1 Introduction	X	3.1.3 Annotation for Disfluency Detection

**Table 1.** Legend to show the sections of each perspective that can be read as complementary to the other.

## 2. Psycholinguistic Perspectives: from Production to Comprehension

In this section, we discuss works that study the planning/production of speech by the speaker, and the comprehension of speech by the listener<sup>2</sup>. Please note, we consistently utilise the term “fillers” in this work, including when citing previous research<sup>3</sup> that may use other terms to describe the same phenomena<sup>4</sup>.

### 2.1. Production

#### 2.1.1. Cognitive load

A common theme in the planning process is the idea of *cognitive load*, i.e. the amount of cognitive *effort* required in the planning process. The production of disfluencies resulting from the planning process has been considered at different linguistic levels. For example, a prosodic/acoustic analysis found that speakers tend to main-

2. We utilise both terms “planning” and “production”, because it is not clear how intentional or unintentional disfluencies are as signals uttered by the speaker (Nicholson, 2007).

3. Please note, there are several other linguistic perspectives not discussed in this paper. Since disfluencies are ubiquitous to spontaneous speech, the literature on disfluencies is vast. For example, a socio-linguistic focus; such as studying the effect that gender, regional background, etc. have on the the production of disfluencies in Shriberg (2001). Some of these perspectives, including works from psycholinguistics, *do consider disfluencies as noise*. We refer the reader to the works of Nicholson (2007) and Lickley (1994), which give a comprehensive overview of these perspectives.

4. Clark and Fox Tree (2002) introduced the term “fillers”; as the term “filled pauses” seemed to indicate that there is a pause in speech *filled* by some sort of (meaningless) sound.

tain a fixed speaking rate during most utterances, but often adopt a faster or slower rate, depending on the cognitive load (O’Shaughnessy, 1995). From this, it was found that disfluent speech may be a result of cognitive load; i.e. speakers slow down their speech when having to make unanticipated choices (for example, using more fillers), and accelerate their speech when repeating some words. This shows that the *types of disfluencies produced* may give further insight to this planning process. The link between the rate of speech and the different disfluencies produced was also found in Shriberg (2001). Two groups of speakers were identified – *repeaters*, who produce more repetitions (when what was said is exactly repeated) than deletions (when previous material that was uttered is abandoned), and *deleters*, who produce more deletions than repetitions. The *repeater-deleter* difference was not only due to stylistic variation in speakers; deleters have a faster speaking rate than repeaters in terms of words per unit time. The interpretation suggested here (contrary to O’Shaughnessy (1995)) is that speakers with a slower speaking rate (repeaters) “take more time to plan”, leading to an increase in repetitions, while faster speakers (deleters) “get ahead of themselves”, and recant what was said to begin again. The different conclusions could be due to the *difference in dataset size, context and domain*<sup>5</sup>. For instance, speakers have been shown to be more disfluent in dialogues compared to monologues (Oviatt, 1995), in human-human conversations than human-machine conversations (Oviatt, 1995), and disfluencies are affected by dialogue role and domain (Colman and Healey, 2011).

The analysis of disfluencies at other levels also reveal *characteristics of the planning process*. In an acoustic-syntactic analysis, it was found that high-frequency monosyllabic function words (such as “the” or “I”) are more likely to be prolonged or have a fuller form when there are neighbouring fillers (“uh” and “um”), indicating that the speaker was encountering problems in planning the utterance (Bell *et al.*, 2003). At an utterance level, Shriberg (2001) found that the longer the utterance, the more disfluencies they contain, also suggesting an increase in cognitive load of the speaker. Disfluencies were also found to occur at the start of an utterance, due to higher cognitive load in planning an utterance (Maclay and Osgood, 1959). At a discourse level, Swerts (1998) found that fillers can be used by the speaker to indicate pausing to (re)formulate thoughts, particularly at discourse boundaries.

### 2.1.2. *Communicative function*

Beattie and Butterworth (1979) suggested that instead of cognitive load, there is an *element of speaker choice* in the planning process. They first establish that generally, speakers are disfluent both when producing low frequency words and improbable words in the context, focusing on fillers “ah”, “er”, and “um”. However, even when frequency of a word was maintained by the speaker, they were still disfluent when producing words with low contextual probability. Like this, there are two main positions behind the speaker’s production of disfluencies. One is that disfluencies are acciden-

5. Indeed, Shriberg (2001)’s findings were based on conversational style dialogues (human-human) in addition to task-oriented dialogues, while O’Shaughnessy (1995) focused on task-oriented dialogues (human-machine).

tally caused in speech due to *cognitive burden* of the speaker (such as Bard *et al.* (2001)). Other works study disfluencies as an important *communicative function* used in dialogue. This view is based on the *strategic modelling view*, where the speaker strategically updates the listener, by using disfluencies as cues (Nicholson, 2007). The distinction between the two views is that the former is an unconscious by-product of speech produced by cognitive (over)load, while the latter is an intentional and strategic production by the speaker. Often studies will look at both of these positions, by analysing the individual disfluencies of a speaker as well as the collective disfluencies produced by interlocutors. The results for the production of disfluencies are often mixed, such as in Nicholson (2007) and Yoshida and Lickley (2010), with evidence suggesting that they occur in both cases; i.e. speakers may both strategically and unconsciously produce disfluencies depending on the context.

## 2.2. Comprehension

Research also focuses on the *comprehension* of disfluent speech, i.e. taking into account the listener’s understanding of the speaker’s disfluencies (Corley and Stewart, 2008), and not on why the disfluency itself was produced (Nicholson, 2007). As Corley and Stewart (2008) state, “it is hard to determine the reason that a speaker is disfluent, especially if the investigation is carried out after the fact from a corpus of recorded speech”. Works that study the effect of disfluencies on listener comprehension state that listeners *must* have developed a comprehension system to process disfluencies, given how frequent they are in spoken language.

### 2.2.1. Informative cues

Research shows that listeners can use disfluencies as signals to *understand and resolve incoming information* in the flow of speech, regardless of whether the speaker intentionally used disfluencies in that way. Consider the following example taken from Brennan and Williams (1995):

**A:** Can I borrow that book?

**B:** ... {F um} ... all right.

Here, speaker **B** used a filler {F...}, which causes **A** to note that **B** might have had a different intention compared to if **B** answered “all right” immediately. While this example is on a pragmatic level (i.e. the listener notes the discrepancy between what was said in essence, and how it was said), works also suggest that listener’s can use disfluencies as *communicative cues* in other ways.

For instance, it was found that fillers helped in the faster recognition of a target word for listeners, indicating that they cause listeners to pay more attention to the upcoming flow of speech (Fox Tree, 1995). The use of fillers also *biases* listeners towards new referring expressions rather than ones already introduced into the discourse (Arnold *et al.*, 2004). Arnold *et al.* (2007) additionally showed that listeners have expectations on the upcoming material to contain *difficult to describe/unconventional*

referring expressions when preceded by the filler “uh”. Listeners expect the speaker to refer to something new following the filler “um”, compared to noise of the same duration (such as a cough or sniffle) (Barr and Seyfeddinipur, 2010). This result was found to be *speaker specific* for the listener. This means that the listener was able to consider what was an already introduced referring expression versus a new referring expression for the current speaker; not just what was old or new for themselves (the listener). Barr and Seyfeddinipur (2010) suggest that this is evidence for the *perspective taking account of language comprehension*. This account suggests that listeners are able to interpret fillers as delay signals, and then infer on plausible reasons for this delay in speech by considering the speaker’s perspective. This implies that fillers can have a *metacognitive (i.e. assessment of knowledge state) effect*, with the listener using fillers as cues to interpret the speaker’s *metacognitive state*. Fillers and prosodic cues were also found to impact listener’s attributions of a speaker’s metacognitive state, specifically the estimation of a speaker’s level of certainty on a topic (Brennan and Williams, 1995).

### 2.2.2. *Filtered out noise*

As a counterpoint to the works discussed, research also suggests that disfluencies may be perceived as noise and thus filtered out by the human listener. In experiments, Fox Tree (1995) noted that removing repetitions from the utterance digitally did not affect the rating of perceived naturalness of speech. However, this could be explained from later work by Shriberg (1999), who found that repetitions have similar pitch contours, but just stretched out over time. Lickley *et al.* (1991) found that listeners could not pinpoint the exact interruption point of a disfluency, and tended to point it out later (up to one word) in the flow of speech. Lickley and Bard (1998) studied a listener’s ability to identify (several types of) disfluencies to find that they are not reliably predictable unless a noticeable pause or abandoned word is apparent. Note, the works discussed here on disfluencies as cues to integrate information are mainly targeted towards fillers, compared to these works that study more complex disfluencies<sup>6</sup>.

### 2.2.3. *Results on incremental processing*

Bailey and Ferreira (2003) point out that the idea of “filtering” out disfluencies (despite their prosody remaining intact, location not easily remembered by listeners) does not account for the *incremental* nature of speech processing; i.e. the processing starts before the input is complete. Filtering would imply that processing needs to occur *after* the removal of disfluencies, and that the listener then, would need to wait for the entire utterance to be completed by the speaker before processing. However, the processes involved in comprehension are continuous and *incremental*, as humans process utterances incrementally. Listeners *must* have developed a comprehension system in order to process disfluencies, given how frequently they occur in spoken language (Bailey and Ferreira, 2003). Disfluencies are part of the incremental processing of the

6. Confusions in terminology arise from research itself, as “disfluencies” is an all-encompassing term used for many works, including when the works are *only* concerned with fillers.



flow of speech; for instance Bailey and Ferreira (2003) show that disfluencies can affect the internal syntactic parser of the listener. The disfluencies considered here was the filler “uh”, but also any kind of interruption (noise) was also deemed to be a disfluency. Brennan and Schober (2001) show that listeners may use disfluencies as cues to avoid integrating what they deem to be incorrect material in an online processing task. Thus, there is evidence to support that fillers are not filtered out by the listener, and indeed, included in the *online processing of the utterance*.

Fillers used as cues to understand new information has even been shown *neurologically*, and in an online processing task. Corley *et al.* (2007) studied the effect of filler (“um”, called a “hesitation” in the work) on the listener’s comprehension using the *N400* function of an Event-related potential (ERP). The *N400* effect can be observed during language comprehension, typically occurring 400 ms after the word onset; it is a negative charge recorded at the scalp consequent to hearing an unpredictable word. The *N400* effect was first established in listeners who heard unpredictable words compared to predictable words. Then, when the filler preceded the unpredictable word, the *N400* effect in listeners was *visibly reduced*. In a subsequent memory test on the listener, words preceded by this filler were more likely to be remembered.

#### 2.2.4. *Time-buying measures and effects on memory*

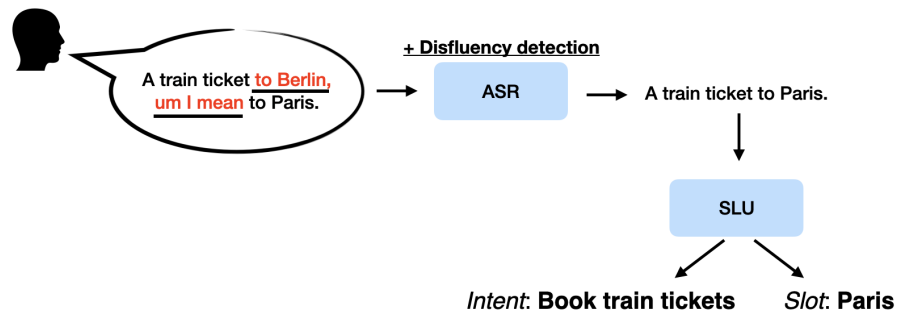
Works also hypothesised whether all these effects of integration can be explained by the *processing time hypothesis*. This means, considering whether the disfluent speech is more memorable/noticeable simply because disfluencies add more time to the speech utterance (referring to research that studies the role of fillers in giving pause to the discourse). Thus, does this effect – i.e. a pause in the utterance – cause the listener to simply give more attention to the utterance? Fraundorf and Watson (2011) examined this in a study on how fillers affect the memory of the listeners. They exposed the listeners to fluent speech containing fillers versus coughs of equal duration artificially spliced into the speech. They found that while fillers facilitated recall, therefore being beneficial on memory, coughs negatively hampered recall accuracy. Disfluent speech (fillers) is hence more likely to be remembered by the listener, and this is *not solely based on the additional time added to the utterance*. Fraundorf and Watson (2011) also manipulated the location of the fillers in speech to study the effect of *position* of fillers on comprehension. This was based on the findings of Swerts (1998) (discussed previously), who found that following fillers, listeners may expect a speaker to shift topics as they carry information about larger topical units – therefore, acting as cues for discourse structure. However, Fraundorf and Watson (2011) found that fillers benefit listener’s recall accuracy regardless of its typical or atypical location. Tottie (2014) found that fillers are noticed more when overused or used in the wrong context, so while they may facilitate recall regardless of location, they still may be more noticeable in atypical locations.

Thus the speaker produces disfluencies, and we have illustrated works that show that they may be strategically used, but also that they can indicate problems in the planning process. When comparing production versus comprehension, regardless of

whether the speaker intentionally uses disfluencies (Barr and Seyfeddinipur, 2010), the listener still *integrates* them as cues for upcoming information. However, it seems that the type of disfluency may also matter in the process of comprehension and whether the listener integrates these disfluencies or filters them out.

### 3. Computational Perspectives: Removal versus Integration

We briefly defined the computational perspective on disfluencies; i.e. works concerned with the computational processing of disfluencies in order to remove them; rendering the utterance more “fluent” and text like. In this distinction, we also consider linguistic work motivated by this goal. For example, phonetic works that characterise the acoustic environments that disfluencies occur in, with the ultimate goal of aiding Automatic Speech Recognition (ASR) systems. While it is a general trend to consider *all* disfluencies as noise in this perspective, we also distinguish research that has found disfluencies informative in SLU. SLU is broad and not one defined task, it “combines speech processing and natural language processing (NLP) by leveraging technologies from machine learning (ML) and artificial intelligence (AI)” (Tur and De Mori, 2011). Thus we distinguish between what is generally considered SLU (SLU for Spoken Dialogue Systems), and broader SLU (SLU tasks concerned with social communication), loosely adopting the distinction made in Tur and De Mori (2011). We make this distinction, because the treatment of disfluencies varies greatly – from *removal to integration* – depending on the SLU task.



**Figure 1.** Example of SLU for SDS, where the disfluent part of the utterance is highlighted in red and underlined. As shown, the disfluencies are removed in post-processing using disfluency detection systems, after ASR systems transcribe the input speech. Then, the input is further collapsed into a semantic frame, consisting of intent and slot.

#### 3.1. SLU for Spoken Dialogue Systems (SDS)

In SLU for SDS, the objective is to collapse the input utterance into a *semantic frame*, consisting of an *intent* and a *slot* (Louvan and Magnini, 2020).

Consider the disfluent utterance “A train ticket to Berlin uh I mean to Paris”. If “to Berlin uh I mean” is not removed (marked in red in Fig. 1 to get the sanitised utterance; “A train ticket to Paris”), it leads to confusion in the subsequent (semantic) processing of the utterance – the correct intent would be identified (i.e. book train tickets) and but not the correct slot (i.e “Paris” not “Berlin”). Disfluencies in this regard are discarded as noise. Thus in a standard SDS pipeline, the output transcripts of ASR systems are cleaned of disfluencies in post-processing as shown in Fig. 1, using disfluency detection systems. From a computational perspective, the intent to process of disfluencies in present day is due to automatic Natural Language Understanding (NLU) systems not being *robust* to them (Ginzburg *et al.*, 2014). Representations of this semantic frame focus on the lexical aspect of the utterance but neglect the non-lexical channel. Thus removing disfluencies such as fillers could remove important information about the *social* aspects of communication (subsequently discussed in Sec. 3.2).

There has been an interest in studying the characteristics of disfluencies for recognition purposes, such as how they are distributed in corpora (for instance, Shriberg (1994) and Shriberg (2001)), and also, at different linguistic levels. The motivation of these works was to identify the properties of disfluent speech to *enable a comparison with its fluent counterparts*. For instance, at a phonetic level, Shriberg (1999) found that the vowels in fillers gave much longer durations than the same vowels in fluent contexts, and Shriberg and Lickley (1993) found that the intonation of fillers is not independent of prior prosodic context. These studies were done with the aim of ideally *leveraging these features for ASR to be more robust to disfluencies*, but with mixed results (such as in Shriberg *et al.* (1997)). Other linguistic levels have also been studied, such as morpho-syntactic features (Goryainova *et al.*, 2014), contextual occurrence (Vasilescu *et al.*, 2010), pragmatic levels (Shriberg *et al.*, 1998) and so on.

### 3.1.1. Disfluency detection

Johnson and Charniak (2004) were the first to introduce the task of disfluency detection, framing it as a noisy channel problem; i.e. noise (disfluencies) has been added to the fluent source utterance. Presently, the task of detection is framed in many ways. For instance, as a *sequence-tagging* task, where typically each token in the input utterance needs to be classified as either in the beginning, inside, or outside (BIO) of a disfluent region. This has been adapted to neural architectures, for instance Zayats and Ostendorf (2019), with a bidirectional Long Short Term Memory (biLSTM) network, or a Transformer model from Lou and Johnson (2020b). Encoder-decoder frameworks allow for detection to be framed as a *sequence-to-sequence neural translation task*, where the encoder learns a representation of the disfluent utterance, and the decoder generates a fluent version of that utterance (see Wang *et al.* (2018)). Similar to this, the objective of end-to-end *speech translation* is to output fluent text from an input of disfluent *audio*. This frames the detection problem as an interim step between ASR and a downstream task such as machine translation (see Salesky *et al.* (2019)). Lou and Johnson (2020a) propose an end-to-end speech recognition and disfluency removal system, so that the fluent output may be used for any downstream

task. Hough and Schlangen (2017) propose a *joint task* of disfluency detection and utterance segmentation. They find that the task is better approached jointly, with the combined system outperforming results on individual tasks. Detection may be combined with other tasks in a multi-task learning (MTL) scenario, showing improvements when there is a secondary task of language modelling (Shalyminov *et al.*, 2018), or even simultaneous tasks of detection, POS-tagging, language modelling and utterance segmentation (Rohanian and Hough, 2020).

A drawback on state-of-the-art (SOTA) systems is they are not robust to transcribing disfluencies correctly in-the-wild, with criticisms that they could not work in an online processing task (Chen *et al.*, 2022) nor with raw transcripts from ASR (Rohanian and Hough, 2021). Specifically for *disfluency detection*, often the *textual input* used is conditioned on already appropriate segmentation and accurate transcription – which is unrealistic in a real-world scenario. From Rohanian and Hough (2021), works on disfluency detection “are almost *exclusively* conducted on pre-segmented utterances of the Switchboard (SWBD) (Godfrey *et al.*, 1992) corpus of telephone conversations”; contradicting the generalisation capabilities of such systems. Works are exploring the feasibility of detection with online processing (Chen *et al.*, 2022; Rohanian and Hough, 2021; Shalyminov *et al.*, 2018; Hough and Schlangen, 2017), with transcripts arising from ASR (Rohanian and Hough, 2021). With hybrid conversations, where user interactions switch between task-oriented and open domain requests (Kim *et al.*, 2021), this topic of detecting disfluencies is starting to come to the forefront once more. The improvement of these systems would be beneficial for other types of research, i.e. better automatic annotation of disfluencies in corpora in order to study disfluent phenomena – particularly for areas such as clinical SLU, to help determine a patient’s cognitive state.

### 3.1.2. *Filler detection*

Systems may be better adapted to recognising specific kinds of disfluencies depending on their properties; fillers for instance, are acoustically distinct. Off-the-shelf open-source speech recognisers such as CMU Sphinx (Lamere *et al.*, 2003) provide functionality to define a filler phone dictionary of *paralinguistic sounds* that are not present in the Language Model (LM). Das *et al.* (2019) leverage the acoustic properties of fillers by training a convolutional recurrent neural network (CRNN) applied directly to *audio recordings* to do filler segmentation, i.e. segmenting fillers in order to output a more fluent utterance. Fillers here are considered to be “unprofessional speech”, stating that the speaker sounds more “fluent, confident and practised compared to the original recorded speech” evaluated by automatic measures<sup>7</sup>, rather than human annotation judgements. Regardless, it would be beneficial to many branches of research to automatically label fillers reliably. However, as pointed out in Das *et al.* (2019), these systems can suffer from false positives (for instance the “um” in

7. Proposed in Kormos and Dénes (2004) to study the perception of fluency in second language learners. Indeed they state in the work that the number of fillers *did not affect the perception of fluency the raters had* (see further discussion in Sec. 3.2, Broader SLU).

“umbrella”). Recently, a new filler detection benchmark and dataset was introduced, where the objective was to label filler candidates given an incoming speech input (Zhu *et al.*, 2022). The system leverages voice activity detection (VAD) and ASR to detect possible filler candidates from other tokens, and then a classifier to categorise them. They evaluate the ASR approach to find it outperforms a keyword spotting approach. The publicly available dataset is based on podcast episodes and contains a diverse range of speaking styles and topics – which could facilitate many future research directions. Thus, there is attention now given to individually detecting fillers from raw audio as a task, without gold-standard transcripts.

Interestingly, while Google Cloud Speech-to-Text services map paralinguistic sounds such as fillers to silence, it has introduced a functionality in its medical ASR model<sup>8</sup> to include fillers when transcribing medical documents. To contrast, in the Google Dialogflow API<sup>9</sup>, there are guidelines given in order to *avoid* fillers as giving unnecessary variety to utterances, as they are ignored by subsequent NLU models (that collapse the input into a semantic frame as shown in Fig. 1). What is interesting here is the *exclusion* of fillers in an SLU for SDS task (i.e. Google Dialogflow), and the *inclusion* of fillers in an ASR task with a broader SLU context – i.e. to pick up social cues in medical transcription. Thus the treatment of fillers and disfluencies is evolving in SLU research to consider specifically the nature of the task; compared to the previous view that all disfluencies must be removed as noise. Hence detection may focus on *removal* if the subsequent task is to collapse the input into a semantic frame, or *integration* if the social aspects of communication may be further considered.

### 3.1.3. Annotation for disfluency detection

Annotation schemes for disfluency detection are almost always form based, as the intent is to sanitise the utterance of disfluencies. For this, there needs to be precise, form-based characterisations of the environment in which disfluencies occur, with no interest in deliberating as to why there was a departure from fluency. We briefly outline the most commonly used annotation scheme in disfluency detection, consistent with the Switchboard repair mark-up (Meteer *et al.*, 1995). Consider the following example:

$$\text{Archie } \underbrace{[\text{likes}]_{\text{RM}}} + \underbrace{\{\text{F uh}\}_{\text{IM}}} \underbrace{\text{loves}]_{\text{RP}}} \text{ Veronica.} \quad [1]$$

At a high level, there is the notion of *some kind of erroneous speech that is to be replaced by corrected speech* – here, “loves” to replace “likes”. This is used as a starting point for annotation (originally proposed by Levelt (1983)). While variations of this scheme exist, the underlying reasoning *specifically* for detection tasks remains the same. Formally, there is: i) the *reparandum phase* (*RM*), i.e. or the entire region to be

8. <https://cloud.google.com/speech-to-text/docs/medical-models>.

9. <https://cloud.google.com/dialogflow/cx/docs/concept/agent-design>.

deleted, and ii) the *repair phase (RP)*, i.e. what replaces the *RM*. This was adopted by Shriberg (1994), who also proposes the term iii) *interregnum phase IM*, which is an optional interruption point, where the speaker may realise that a correction needs to be made. *Non-sentence elements* (such as fillers) can occur within this *IM* structure or outside, also called *isolated edit terms/ single edit tokens*. From this structure, several types of disfluencies can be formally defined. For instance *repetitions* – i.e. when the *RM* phase is repeated exactly in the *RP* phase, and *restarts* – i.e. when the *RM* phase is discarded. This annotation structure is useful for disfluency detection and SLU because the reasoning is to only keep the corrected speech in the utterance (i.e. the *RP*) for downstream processing (see Fig. 1), and discard the rest as noise (*RM+IM*).

However, a drawback of this scheme for a detection task is that it spans only one speaker turn and can only be initiated by the speaker producing the disfluencies. In many contexts, repair could span different speaker turns and different interlocutors (see a detailed discussion in Purver *et al.* (2018)). Disfluency detection thus can fail on longer disfluencies (Zayats *et al.*, 2019), and disfluencies spanning multiple turns (Purver *et al.*, 2018). Additionally, many other annotation schemes exist when the objective is not constrained to detection alone. For instance, Christodoulides *et al.* (2014) proposed a tool to holistically do a form-based annotation of several phenomena characteristic of spontaneous speech: with disfluency detection and annotation, and multi-word unit recognition (including POS, syllable, phone tagging ...). Rather than an annotation specific to one task, these layers of annotation could be beneficial for several downstream SLU tasks. The annotation scheme is more realistically designed for spoken language – such as using prosodic cues for segmentation in the absence of punctuation. There are also language specific schemes proposed (such as Benzitoun *et al.* (2012), Kahane and Gerdes (2020) and Eshkol *et al.* (2010) for French), as **a limitation of this survey is that we constrain ourselves to English**<sup>10</sup>.

Problems may also arise when transcribing (disfluent) audio itself. Le Grezause (2017) found that transcribers who had transcribed a few conversations tended to have *substantially higher transcription errors*, compared to transcribers that had transcribed a large number of conversations. Zayats *et al.* (2019) found that tokens generally related to spontaneous speech phenomena (fillers ...) have a high frequency of transcription errors, which they *discuss could be due to these tokens being non-standard*, i.e. unaccounted for in annotation instructions. Fillers in particular, are among some of the most likely tokens to be mis-transcribed (whether inserted, deleted or substituted), and transcriber experience has a noticeable effect on the accuracy of transcribing fillers specifically (Le Grezause, 2017).

10. Note also that the issue of annotation has been extensively studied from a non-computational linguistic perspective (Grosman, 2018), with extensive summaries of terminology from different linguistic works (Lickley *et al.*, 1991; Nicholson, 2007; Lickley, 2015), and new and emergent annotation schemes proposed (Crible *et al.*, 2015) which are not discussed here.

### 3.1.4. Text representations of disfluencies

With the increasing popularity of voice assistant technologies, a trend has emerged in Natural Language Understanding (NLU) research to overlap with SLU; i.e. considering the *textual processing of speech transcripts* (as discussed in Ruder (2020) as “Speech first-NLP”). However, discrepancies may arise in the processing of speech transcripts compared to processing grammatically written text, if utilising the same (NLP) systems. For example, Tran *et al.* (2017) present an attention-based encoder-decoder model for parsing conversational sentences arising from speech transcripts. An important finding from this work was that the integration of acoustic-prosodic features showed *the most gains over disfluent and longer sentences* compared to fluent ones. This empirically *shows a discrepancy* between transcripts that are more “speech-like” (i.e. are disfluent and have acoustic variation) compared to transcripts that are more grammatical and resemble written text. It is worth noting that for this task of parsing, there is the assumption of *already structured* spontaneous speech transcripts to be used as input (i.e. annotated punctuation, input being pre-segmented sentences . . .) – so indeed, there is a further gap between perfectly fluent, grammatical text versus raw ASR/verbatim transcripts. Hervé *et al.* (2022) investigate how deep contextualised embeddings could be pre-trained on massive amounts of ASR generated transcripts for more realistic samples in spoken language modelling. The downstream tasks including parsing conversational utterances show improvements of the model pre-trained on raw ASR data (which they call FlauBERT-Oral) compared to the original model trained on written data (i.e. FlauBERT, a French LM (Le *et al.*, 2020)). They point out that the ASR generated only lowercase transcripts, and that adding capitalisation and punctuation to the transcripts could benefit the model i.e. *re-introducing some degree of sentence structure*.

Tran *et al.* (2019) showed that deep contextualised embeddings pre-trained on large written corpora can be fine-tuned on smaller spontaneous speech datasets to improve parsing on conversational speech transcripts. While this indicates that some *general* characteristics of spontaneous speech may be learnt in the fine-tuning stage (and as shown by Hervé *et al.* (2022) using models pre-trained on raw ASR), results are mixed when specifically considering the text representation of fillers. Barriere *et al.* (2017) for instance showed that pre-trained word embeddings such as Word2vec (Mikolov *et al.*, 2013) have poor representation of spontaneous speech phenomena such as “uh”, as they are trained on written text and do not carry the same meaning as when used in speech. Interestingly, this discrepancy may not always negatively impact the *robustness* of the system when purely considering fillers. For instance, Dinkar *et al.* (2020) investigated the representations of fillers using deep contextualised word embeddings. They found that Bi-directional Encoder Representations (BERT) (Devlin *et al.*, 2019) has existing representations of fillers, despite being pre-trained on massive amounts of written text. They found (similar to Stolcke and Shriberg (1996)’s results on *ngrams*) that the inclusion of fillers *reduces* the uncertainty of a language model in a spoken language modelling task. This is despite research that shows that overall, speech disfluencies occur at higher perplexities (Sen, 2020). Thus, in addition to psycholinguistic work, there is computational research to show that indeed, fillers

may be able to provide information regarding the neighbouring words to the right in a language modelling task. However, Dinkar *et al.* (2020) found that BERT is unable to distinguish between the two fillers “uh” and “um”, despite research to show that they occupy different functions in discourse (Le Grezause, 2017; Dinkar, 2022). To conclude, there is now the general awareness that spoken speech is not like written text, and increasing studies work on how to learn representations of spontaneous speech given that architectures are usually developed on large amounts of written data.

### 3.2. Broader SLU and Generation

#### 3.2.1. Broader Spoken Language Understanding (SLU)

We consider broader SLU as the analysis of the flow of speech by leveraging NLP and ML techniques<sup>11</sup> – particularly for higher order tasks that are concerned with social/affective computing. Though the methodologies differ from the psycholinguistic approach to study the production contexts of disfluencies, the underlying intent is similar. That is, to study the context of disfluencies produced (and indeed, many other lexical and non-lexical features) and their link to a variety of socio-communicative and cognitive phenomena. The findings are useful for feature engineering in broader SLU tasks. For example, disfluencies in several works have been found to be an informative *social signal* (Ekman *et al.*, 1980; Mairesse *et al.*, 2007; Vinciarelli and Mohammadi, 2014; Schuller *et al.*, 2019). Fillers specifically are commonly used as an attribute to study big 5 personality traits (Mairesse *et al.*, 2007). The Computational Paralinguistics Challenge focused on detecting fillers, which they considered to be a “social signal” (Schuller *et al.*, 2019), acknowledging the importance of fillers in *tasks concerned with social communication*. Along these lines, research in personality computing has the most consistent correlation, with observations made from speech; including *paralanguage* such as fillers (Ekman *et al.*, 1980; Vinciarelli and Mohammadi, 2014). We cannot exhaustively account for these works, as they are numerous. Some examples are research on the role of disfluencies (mainly fillers) in the prediction of a speaker’s *emotions* (Moore *et al.*, 2014; Tian *et al.*, 2015), *stance* (Le Grezause, 2017), perceived *persuasiveness* (Park *et al.*, 2014), perceived *confidence* (Dinkar *et al.*, 2020), etc. Dufour *et al.* (2014) point out that spontaneous speech is not the same as prepared speech, where the utterances are well formed and closer to written documents. They focus on classifying the *degree of spontaneity* of speech in order to do SLU tasks, such as characterising speaker roles (example, an interviewer versus an interviewee) using fillers and other acoustic-linguistic features. Recent work argues that disfluencies will be useful in dialogue based computer-assisted language learning; i.e. detecting and analysing a learner’s disfluencies (including silences and laughter) could potentially help a system determine appropriate pedagogical interventions (Skidmore and Moore, 2022). Disfluencies may play a crucial role in *clinical SLU* tasks, such as in dementia recognition (Rohanian *et al.*, 2021). As stated, Google

11. Which we loosely adopt from Tur and De Mori (2011).



Cloud Speech-to-Text has included in its medical ASR model options to include fillers, for the purpose of transcribing medical documents.

However, an open challenge that remains in analysis is the inclusion of context. Research may predetermine based on the scenario whether disfluencies are positive or negative. For instance, in the automatic processing of job interview data (Rasipuram *et al.*, 2016), fluency is assumed to be desirable. While indeed, the speaker’s production of disfluencies may have an effect on the outcome of a job interview; there are nuances of how the speaker utilises such spontaneous speech phenomena; i.e. both production and perception will vary based on socio-linguistic background, context, domain and so on. Analysis may be based on the idea that the frequency of words (including, non-lexical tokens such as fillers present in transcripts) can represent underlying affective traits (Boyd and Schwartz, 2021). In this aspect, the link between fillers and a wide variety of phenomena is to be expected – from linguistic levels to higher affective levels; with Barr (2001) even describing fillers as *vocal gestures*. In a recent survey, Boyd and Schwartz (2021) discuss this drawback regarding research in the intersection between psychology and language analysis – but also considering interdisciplinary fields such as social computing; stating that it may not always be a case of “paying attention to  $X$  is correlated with  $Y$ ”. Thus we should cautiously assume a linear relationship between the fillers produced and the task under consideration. Dinkar *et al.* (2021) for instance found that speakers tend to stylistically use fillers before introducing new information in the discourse, but that listeners may not associate this specific use of fillers with their estimation of the speaker’s confidence.

An additional challenge is the varied terminology used to describe fillers in the context of broader SLU. While standardised terminology may be used in a benchmark task such as detection, it is not the case for less standardised tasks. This is problematic, because it can lead to a lack of availability and transparency of the findings. For instance, the *Linguistic Inquiry and Word Count (LIWC)* (Pennebaker *et al.*, 2001), a commonly used text analysis software in personality computing (Mehta *et al.*, 2020), gives guidelines to annotate “nonfluencies” (fillers) – “uh”, “um”, “er”, and “stuttering”. Here, “stuttering” broadly denotes the general phenomena of being disfluent<sup>12</sup>.

### 3.2.2. Generation

Similar to psycholinguistic approaches, it is to be noted that there are works that study the *perception* of disfluencies from a generation standpoint, i.e. using artificially synthesised disfluencies in speech. While this is not under the umbrella of SLU as such, we briefly describe some research specific to fillers. For instance, in a work directly motivated by psycholinguistic perspectives of comprehension, Wollermann *et al.* (2013) explore the listener’s perception of disfluencies using Text-to-Speech (TTS). This was based on the work of Brennan and Williams (1995), that discusses

12. Please note, the term is not to be confused with the clinical sense of “stuttering”. In clinical literature, the term commonly used is “stutter-like disfluencies” (SLDs). This use seems borrowed from Mahl (1956), who used the term to refer to repetition of partial words (Lickley, 2015).

the role of fillers and prosodic cues in a listener’s evaluation of how uncertain they think the speaker is regarding a topic. They had the system exhibit “uncertain” behaviour through disfluent TTS responses in a question-answering context. They found that disfluencies in combination with prosodic cues (i.e. delays + fillers) increased a listener’s perception of uncertainty towards the system’s answers. Similarly, Kirkland *et al.* (2022) found that when fillers were not present in synthesised speech, it leads to a perception of more confident sounding utterances, while utterance-medial fillers lead to a perception of the least-confident sounding utterances (both in addition to other prosodic features).

Some works may not have basis in psycholinguistic theory. For instance, there is research that considers how disfluencies enhance the *naturalness* of the synthesised speech. Pfeifer and Bickmore (2009) evaluate an agent that uses fillers “uh” and “um” in speech. The motivation behind this was to improve the naturalness of speech in an Embodied Conversational Agent (ECA), as ECAs often try to emulate humans in gestures and facial expressions, yet speak in fluent sentences. Results are mixed, with some participants saying that fillers enhanced the naturalness of the conversation, while others expected that an agent should speak fluently, and fillers were deemed inappropriate. Székely *et al.* (2019) discuss approaches for treating fillers in TTS tasks, i.e. suggesting methods that will result in them being synthesised naturally (both distributionally and perceptually) in the generated output. Disfluencies may also be utilised as a *communicative strategy* in generation. Skantze *et al.* (2015) studied how a system can use multi-modal turn-taking signals (including fillers) as a *time-buying measure*, i.e. to buy time for generating a response as the next move of the robot is decided.

While a common goal of AI is to work towards more human-like (anthropomorphic) agents, a challenge that remains is to consider the trade-off between the naturalness of a system and the safety of its deployment. Consider Google Duplex (Leviathan and Matias, 2018), a TTS system for accomplishing real-world tasks over the phone. The *inclusion of disfluencies* (such as fillers and repairs) led to highly natural sounding generated responses, showing how ubiquitous disfluencies are to everyday communication. However, these responses made human callers think that they were conversing with another human. This illusion of agency may have negative consequences when considering *safety* in conversational AI. For example, attributing anthropomorphic traits to an agent in an IMPOSTOR EFFECT (Dinan *et al.*, 2021; Abercrombie and Rieser, 2022) scenario, i.e. where a system provides inappropriate advice in safety critical situations (such as when a user seeks medical advice). Despite this, from a generation perspective, fillers and disfluencies may still be desirable if used as a mitigation strategy, i.e. to exhibit uncertainty and doubt in safety critical situations (e.g. “uh... I’m not an expert here but...”). This is important, as Mielke *et al.* (2022) pointed out that neural dialogue agents are not linguistically calibrated – i.e. despite the agent being *factually inaccurate* it may still (inappropriately) verbalise expressions of confidence in responses. Thus, *the trade-off between naturalness and safety* merits consideration in future research, particularly when using disfluencies in generation tasks.

#### 4. Conclusion

The preliminary goal of this article was to bring together several different perspectives on fillers that influence research in SLU. The article was motivated by the increasing popularity of voice assistant technologies, where often, the corpora used will invariably contain these spontaneous speech events. Yet, to the best of our knowledge, there was not an introductory source on these specific phenomena, that offered a *breadth* of different perspectives and approaches – particularly targeted towards the SLU community. We were motivated to introduce these different perspectives, without exhaustively going into details that may overwhelm the reader. Our aim was to discuss these perspectives in a holistic way, that is considering them at various stages of the SLU pipeline; from underlying (psycho)linguistic theory, to annotation, ASR and SLU perspectives, and finally, looking briefly at research on the generation of these phenomena. We have tried to the best of our ability to synthesise the main perspectives in an approachable way, and suggest further reading and other sources when appropriate. Additionally in each section we pinpoint (what we believe) are the main challenges and trends of each area. To do so, we identified two areas of research, i.e. psycholinguistic and computational perspectives. In the former, we focused on the research available on the production and comprehension of disfluencies, particularly fillers. In the computational approaches, we discussed the treatment of disfluencies by distinguishing between SLU as it is typically considered for SDS, and broader SLU tasks more concerned with social communication. We try to highlight the commonalities regarding both perspectives as well, i.e. discussing works in both areas that find disfluencies informative, consider disfluencies as noise, consider the listener’s perception . . . . Going forward, some open challenges are to consider the balance between removing disfluencies with regard to the robustness of a system versus including them as speech events that could offer social context, the trade-off between how natural sounding the system is versus how safe it may be in deployment, the integration of fine-grained context, and lastly, the taxonomy of such phenomena in order to make the research on them more accessible.

#### Acknowledgements

Tanvi Dinkar is supported by ‘AISEC: AI Secure and Explainable by Construction’ (EP/T026952/1), and Tanvi Dinkar and Chloé Clavel were supported by ‘ANIMATAS: Advancing intuitive human-machine interaction with human-like social capabilities for education in schools’, European Union’s Horizon 2020 research and innovation programme under grant agreement No. 765955.

#### 5. References

Abercrombie G., Rieser V., “Risk-graded Safety for Handling Medical Queries in Conversational AI”, *arXiv preprint arXiv:2210.00572*, 2022.

- Arnold J. E., Kam C. L. H., Tanenhaus M. K., “If you say thee uh you are describing something hard: the on-line attribution of disfluency during reference comprehension.”, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 33, n° 5, p. 914, 2007.
- Arnold J. E., Tanenhaus M. K., Altmann R. J., Fagnano M., “The old and thee, uh, new: Disfluency and reference resolution”, *Psychological science*, vol. 15, n° 9, p. 578-582, 2004.
- Bailey K. G., Ferreira F., “Disfluencies affect the parsing of garden-path sentences”, *Journal of Memory and Language*, vol. 49, n° 2, p. 183-200, 2003.
- Bard E. G., Lickley R. J., Aylett M. P., “Is disfluency just difficulty?”, *ISCA Tutorial and Research Workshop (ITRW) on Disfluency in Spontaneous Speech*, 2001.
- Barr D. J., “Trouble in mind: Paralinguistic indices of effort and uncertainty in communication”, *Oralité et gestualité: Interactions et comportements multimodaux dans la communication*, p. 597-600, 2001.
- Barr D. J., Seyfeddinipur M., “The role of fillers in listener attributions for speaker disfluency”, *Language and Cognitive Processes*, vol. 25, n° 4, p. 441-455, 2010.
- Barriere V., Clavel C., Essid S., “Opinion Dynamics Modeling for Movie Review Transcripts Classification with Hidden Conditional Random Fields”, *Proceedings of Interspeech 2017*, Stockholm, Sweden, August, 2017.
- Beattie G. W., Butterworth B. L., “Contextual probability and word frequency as determinants of pauses and errors in spontaneous speech”, *Language and speech*, vol. 22, n° 3, p. 201-211, 1979.
- Bell A., Jurafsky D., Fosler-Lussier E., Girand C., Gregory M., Gildea D., “Effects of disfluencies, predictability, and utterance position on word form variation in English conversation”, *The Journal of the Acoustical Society of America*, vol. 113, n° 2, p. 1001-1024, 2003.
- Bender E. M., Koller A., “Climbing towards NLU: On meaning, form, and understanding in the age of data”, *Proc. of ACL*, 2020.
- Benzitoun C., Fort K., Sagot B., “TCOF-POS : un corpus libre de français parlé annoté en morphosyntaxe (TCOF-POS : A Freely Available POS-Tagged Corpus of Spoken French) [in French]”, *Proceedings of the Joint Conference JEP-TALN-RECITAL 2012, volume 2: TALN, ATALA/AFCP*, Grenoble, France, p. 99-112, June, 2012.
- Boyd R. L., Schwartz H. A., “Natural language analysis and the psychology of verbal behavior: The past, present, and future states of the field”, *Journal of Language and Social Psychology*, vol. 40, n° 1, p. 21-41, 2021.
- Brennan S. E., Schober M. F., “How listeners compensate for disfluencies in spontaneous speech”, *Journal of Memory and Language*, vol. 44, n° 2, p. 274-296, 2001.
- Brennan S. E., Williams M., “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*, vol. 34, n° 3, p. 383-398, 1995.
- Chen A., Zayats V., Walker D. D., Padfield D., “Teaching BERT to Wait: Balancing Accuracy and Latency for Streaming Disfluency Detection”, 2022.
- Christodoulides G., Avanzi M., Goldman J.-P., “DisMo: A Morphosyntactic, Disfluency and Multi-Word Unit Annotator. An Evaluation on a Corpus of French Spontaneous and Read Speech”, *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, European Language Resources Association (ELRA), Reykjavik, Iceland, May, 2014.

- Clark H. H., Fox Tree J. E., “Using uh and um in spontaneous speaking”, *Cognition*, vol. 84, n<sup>o</sup> 1, p. 73-111, 2002.
- Colman M., Healey P., “The distribution of repair in dialogue”, *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 33, 2011.
- Corley M., MacGregor L. J., Donaldson D. I., “It’s the way that you, er, say it: Hesitations in speech affect language comprehension”, *Cognition*, vol. 105, n<sup>o</sup> 3, p. 658-668, 2007.
- Corley M., Stewart O. W., “Hesitation disfluencies in spontaneous speech: The meaning of um”, *Language and Linguistics Compass*, vol. 2, n<sup>o</sup> 4, p. 589-602, 2008.
- Crible L., Dumont A., Grosman I., Notarrigo I., “Annotation des marqueurs de fluence et disfluence dans des corpus multilingues et multimodaux, natifs et non natifs v. 1.0”, 2015.
- Das S., Gandhi N., Naik T., Shilkrot R., “Increase Apparent Public Speaking Fluency by Speech Augmentation”, *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, p. 6890-6894, 2019.
- Devlin J., Chang M.-W., Lee K., Toutanova K., “BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding”, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Association for Computational Linguistics, Minneapolis, Minnesota, p. 4171-4186, June, 2019.
- Dinan E., Abercrombie G., Bergman A. S., Spruit S., Hovy D., Boureau Y.-L., Rieser V., “Anticipating Safety Issues in E2E Conversational AI: Framework and Tooling”, 2021.
- Dinkar T., Computational models of disfluencies: fillers and discourse markers in spoken language understanding, PhD thesis, Institut Polytechnique de Paris, 2022.
- Dinkar T., Colombo P., Labeau M., Clavel C., “The importance of fillers for text representations of speech transcripts”, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Association for Computational Linguistics, Online, p. 7985-7993, November, 2020.
- Dinkar T., LTCI T. P., Biancardi B., Clavel C., “From local hesitations to global impressions of a speaker’s feeling of knowing”, *Proceedings of the 25th Workshop on the Semantics and Pragmatics of Dialogue*, 2021.
- Dufour R., Esteve Y., Deléglise P., “Characterizing and detecting spontaneous speech: Application to speaker role recognition”, *Speech communication*, vol. 56, p. 1-18, 2014.
- Ekman P., Friesen W. V., O’Sullivan M., Scherer K., “Relative importance of face, body, and speech in judgments of personality and affect”, *Journal of personality and social psychology*, vol. 38, n<sup>o</sup> 2, p. 270, 1980.
- Eshkol I., Maurel D., Friburger N., “Eslo: From Transcription to Speakers’ Personal Information Annotation”, *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC’10)*, European Language Resources Association (ELRA), Valletta, Malta, May, 2010.
- Fox Tree J. E., “The effects of false starts and repetitions on the processing of subsequent words in spontaneous speech”, *Journal of memory and language*, vol. 34, n<sup>o</sup> 6, p. 709-738, 1995.
- Fraundorf S. H., Arnold J., Langlois V. J., “Disfluency”, 2018.
- Fraundorf S. H., Watson D. G., “The disfluent discourse: Effects of filled pauses on recall”, *Journal of memory and language*, vol. 65, n<sup>o</sup> 2, p. 161-175, 2011.

- Ginzburg J., Paris-diderot U., Fernández R., Schlangen D., “Disfluencies as intra-utterance dialogue moves”, *Semantics and Pragmatics*, 2014.
- Godfrey J. J., Holliman E. C., McDaniel J., “SWITCHBOARD: Telephone Speech Corpus for Research and Development”, *Proceedings of the 1992 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, vol. 1, IEEE, p. 517-520, 1992.
- Goryainova M., Grouin C., Rosset S., Vasilescu I., “Morpho-Syntactic Study of Errors from Speech Recognition System”, *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, European Language Resources Association (ELRA), Reykjavik, Iceland, p. 3045-3049, May, 2014.
- Grosman I., Évaluation contextuelle de la (dis)fluence en production et perception: pratiques communicatives et formes prosodico-syntaxiques en français, PhD thesis, UCL-Université Catholique de Louvain, 2018.
- Hervé N., Pelloin V., Favre B., Dary F., Laurent A., Meignier S., Besacier L., “Using ASR-Generated Text for Spoken Language Modeling”, *Proceedings of BigScience Episode #5 – Workshop on Challenges & Perspectives in Creating Large Language Models*, Association for Computational Linguistics, virtual+Dublin, p. 17-25, May, 2022.
- Hough J., Schlangen D., “Joint, incremental disfluency detection and utterance segmentation from speech”, *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, p. 326-336, 2017.
- Johnson M., Charniak E., “A TAG-based noisy-channel model of speech repairs”, *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*, p. 33-39, 2004.
- Kahane S., Gerdes K., “Annotation syntaxique du français parlé: Les choix d’ORFÉO”, *Langages*, n° 3, p. 69-86, 2020.
- Kim S., Liu Y., Jin D., Papangelis A., Gopalakrishnan K., Hedayatnia B., Hakkani-Tur D., “How Robust RU?: Evaluating Task-Oriented Dialogue Systems on Spoken Conversations”, *arXiv e-prints*. arXiv-2109, 2021.
- Kirkland A., Lameris H., Székely E., Gustafson J., “Where’s the uh, hesitation? The interplay between filled pause location, speech rate and fundamental frequency in perception of confidence”, *Proc. Interspeech 2022*p. 4990-4994, 2022.
- Kormos J., Dénes M., “Exploring measures and perceptions of fluency in the speech of second language learners”, *System*, vol. 32, n° 2, p. 145-164, 2004.
- Lamere P., Kwok P., Gouvea E., Raj B., Singh R., Walker W., Warmuth M., Wolf P., “The CMU SPHINX-4 speech recognition system”, *Ieee intl. conf. on acoustics, speech and signal processing (icassp 2003), hong kong*, vol. 1, p. 2-5, 2003.
- Le Grezause E., Um and Uh, and the expression of stance in conversational speech, PhD thesis, 2017.
- Le H., Vial L., Frej J., Segonne V., Coavoux M., Lecouteux B., Allauzen A., Crabbé B., Besacier L., Schwab D., “FlauBERT: Unsupervised Language Model Pre-training for French”, *Proceedings of the Twelfth Language Resources and Evaluation Conference*, European Language Resources Association, Marseille, France, p. 2479-2490, May, 2020.
- Levelt W. J., “Monitoring and self-repair in speech”, *Cognition*, vol. 14, n° 1, p. 41-104, 1983.
- Leviathan Y., Matias Y., “Google Duplex: An AI System for Accomplishing Real World Tasks Over the Phone.”, *Google AI Blog*, 2018.

- Lickley R. J., Detecting disfluency in spontaneous speech, PhD thesis, University of Edinburgh, 1994.
- Lickley R. J., “Fluency and Disfluency”, *The handbook of speech production* p. 445, 2015.
- Lickley R. J., Bard E. G., “When can listeners detect disfluency in spontaneous speech?”, *Language and speech*, vol. 41, n° 2, p. 203-226, 1998.
- Lickley R. J., Shillcock R. C., Bard E. G., “Processing Disfluent Speech: How and when are disfluencies found?”, *Second European Conference on Speech Communication and Technology*, 1991.
- Lou P. J., Johnson M., “End-to-End Speech Recognition and Disfluency Removal”, *Findings of the Association for Computational Linguistics: EMNLP 2020*, Association for Computational Linguistics, Online, p. 2051-2061, November, 2020a.
- Lou P. J., Johnson M., “Improving Disfluency Detection by Self-Training a Self-Attentive Model”, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Online, p. 3754-3763, July, 2020b.
- Louvan S., Magnini B., “Recent Neural Methods on Slot Filling and Intent Classification for Task-Oriented Dialogue Systems: A Survey”, *arXiv preprint arXiv:2011.00564*, 2020.
- Maclay H., Osgood C. E., “Hesitation phenomena in spontaneous English speech”, *Word*, vol. 15, n° 1, p. 19-44, 1959.
- Mahl G. F., “Disturbances and silences in the patient’s speech in psychotherapy.”, *The Journal of Abnormal and Social Psychology*, vol. 53, n° 1, p. 1, 1956.
- Mairesse F., Walker M. A., Mehl M. R., Moore R. K., “Using linguistic cues for the automatic recognition of personality in conversation and text”, *Journal of artificial intelligence research*, vol. 30, p. 457-500, 2007.
- Mehta Y., Majumder N., Gelbukh A., Cambria E., “Recent trends in deep learning based personality detection”, *Artificial Intelligence Review*, vol. 53, n° 4, p. 2313-2339, 2020.
- Meteer M. W., Taylor A. A., MacIntyre R., Iyer R., *Dysfluency annotation stylebook for the switchboard corpus*, University of Pennsylvania Philadelphia, PA, 1995.
- Mielke S. J., Szlam A., Dinan E., Boureau Y.-L., “Reducing conversational agents’ overconfidence through linguistic calibration”, *Transactions of the Association for Computational Linguistics*, vol. 10, p. 857-872, 2022.
- Mikolov T., Chen K., Corrado G., Dean J., “Efficient Estimation of Word Representations in Vector Space”, *arXiv preprint arXiv:1301.3781*, 2013.
- Moore J. D., Tian L., Lai C., “Word-level emotion recognition using high-level features”, *International Conference on Intelligent Text Processing and Computational Linguistics*, Springer, p. 17-31, 2014.
- Nicholson H. B. M., “Disfluency in dialogue: attention, structure and function”, 2007.
- O’Shaughnessy D., “Timing patterns in fluent and disfluent spontaneous speech”, *1995 International Conference on Acoustics, Speech, and Signal Processing*, vol. 1, IEEE, p. 600-603, 1995.
- Oviatt S., “Predicting spoken disfluencies during human-computer interaction”, *Computer Speech and Language*, vol. 9, n° 1, p. 19-36, 1995.
- Park S., Shim H. S., Chatterjee M., Sagae K., Morency L.-P., “Computational Analysis of Persuasiveness in Social Multimedia: A Novel Dataset and Multimodal Prediction Approach”,

- Proceedings of the 16th International Conference on Multimodal Interaction, ICMI 2014*, Association for Computing Machinery, New York, NY, USA, 2014.
- Pennebaker J. W., Francis M. E., Booth R. J., “Linguistic inquiry and word count: LIWC 2001”, *Mahway: Lawrence Erlbaum Associates*, vol. 71, n° 2001, p. 2001, 2001.
- Pfeifer L. M., Bickmore T., “Should agents speak like, um, humans? The use of conversational fillers by virtual agents”, *International Workshop on Intelligent Virtual Agents*, Springer, p. 460-466, 2009.
- Purver M., Hough J., Howes C., “Computational Models of Miscommunication Phenomena”, *Topics in Cognitive Science*, vol. 10, n° 2, p. 425-451, 2018.
- Rasipuram S., Rao S. P., Jayagopi D. B., “Automatic prediction of fluency in interface-based interviews”, *2016 IEEE Annual India Conference (INDICON)*, p. 1-6, 2016.
- Rohanian M., Hough J., “Re-framing Incremental Deep Language Models for Dialogue Processing with Multi-task Learning”, *CoRR*, 2020.
- Rohanian M., Hough J., “Best of Both Worlds: Making High Accuracy Non-incremental Transformer-based Disfluency Detection Incremental”, *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, p. 3693-3703, 2021.
- Rohanian M., Hough J., Purver M., “Alzheimer’s Dementia Recognition Using Acoustic, Lexical, Disfluency and Speech Pause Features Robust to Noisy Inputs”, *arXiv preprint arXiv:2106.15684*, 2021.
- Ruder S., “ICLR 2021 Outstanding Papers, Char Wars, Speech-first NLP, Virtual conference ideas”, Apr, 2020.
- Salesky E., Sperber M., Waibel A., “Fluent Translations from Disfluent Speech in End-to-End Speech Translation”, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Association for Computational Linguistics, Minneapolis, Minnesota, p. 2786-2792, June, 2019.
- Schiffrin D., *Discourse markers*, n° 5, Cambridge University Press, 1987.
- Schuller B., Weninger F., Zhang Y., Ringeval F., Batliner A., Steidl S., Eyben F., Marchi E., Vinciarelli A., Scherer K. *et al.*, “Affective and behavioural computing: Lessons learnt from the First Computational Paralinguistics Challenge”, *Computer Speech & Language*, vol. 53, p. 156-180, 2019.
- Sen P., “Speech Disfluencies occur at Higher Perplexities”, *Proceedings of the Workshop on the Cognitive Aspects of the Lexicon*, Association for Computational Linguistics, Online, p. 92-97, December, 2020.
- Shalyminov I., Eshghi A., Lemon O., “Multi-Task Learning for Domain-General Spoken Disfluency Detection in Dialogue Systems”, *CoRR*, 2018.
- Shriberg E., “To ‘errrr’ is human: ecology and acoustics of speech disfluencies”, *Journal of the International Phonetic Association*, vol. 31, n° 1, p. 153-169, 2001.
- Shriberg E., Bates R. A., Stolcke A., “A prosody only decision-tree model for disfluency detection.”, *Eurospeech*, vol. 97, Citeseer, p. 23832386, 1997.
- Shriberg E. E., Preliminaries to a theory of speech disfluencies, PhD thesis, Citeseer, 1994.
- Shriberg E. E., Phonetic consequences of speech disfluency, Technical report, SRI INTERNATIONAL MENLO PARK CA, 1999.



- Shriberg E. E., Lickley R. J., “Intonation of clause-internal filled pauses”, *Phonetica*, vol. 50, n<sup>o</sup> 3, p. 172-179, 1993.
- Shriberg E., Stolcke A., Jurafsky D., Coccaro N., Meteer M., Bates R., Taylor P., Ries K., Martin R., Van Ess-Dykema C., “Can prosody aid the automatic classification of dialog acts in conversational speech?”, *Language and speech*, vol. 41, n<sup>o</sup> 3-4, p. 443-492, 1998.
- Skantze G., Johansson M., Beskow J., “Exploring turn-taking cues in multi-party human-robot discussions about objects”, *Proceedings of the 2015 ACM on international conference on multimodal interaction*, p. 67-74, 2015.
- Skidmore L., Moore R., “Incremental Disfluency Detection for Spoken Learner English”, *Proceedings of the 17th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2022)*, Association for Computational Linguistics, Seattle, Washington, p. 272-278, July, 2022.
- Stolcke A., Shriberg E., “Statistical Language Modeling for Speech Disfluencies”, *Proceedings of the 1996 IEEE International Conference on Acoustics, Speech, and Signal Processing Conference Proceedings (ICASSP)*, vol. 1, IEEE, p. 405-408, 1996.
- Swerts M., “Filled Pauses as Markers of Discourse Structure”, *Journal of Pragmatics*, vol. 30, n<sup>o</sup> 4, p. 485 - 496, 1998.
- Székely É., Henter G. E., Beskow J., Gustafson J., “How to train your fillers: uh and um in spontaneous speech synthesis”, *The 10th ISCA Speech Synthesis Workshop*, 2019.
- Tian L., Moore J. D., Lai C., “Emotion recognition in spontaneous and acted dialogues”, *2015 International Conference on Affective Computing and Intelligent Interaction (ACII)*, IEEE, p. 698-704, 2015.
- Tottie G., “On the use of uh and um in American English”, *Functions of Language*, vol. 21, n<sup>o</sup> 1, p. 6-29, 2014.
- Tran T., Toshniwal S., Bansal M., Gimpel K., Livescu K., Ostendorf M., “Joint modeling of text and acoustic-prosodic cues for neural parsing”, *arXiv preprint arXiv:1704.07287*, 2017.
- Tran T., Yuan J., Liu Y., Ostendorf M., “On the Role of Style in Parsing Speech with Neural Models”, *Proceedings of Interspeech 2019*, p. 4190-4194, 2019.
- Tur G., De Mori R., *Spoken language understanding: Systems for extracting semantic information from speech*, John Wiley & Sons, 2011.
- Vasilescu I., Rosset S., Adda-Decker M., “On the Role of Discourse Markers in Interactive Spoken Question Answering Systems”, *LREC*, 2010.
- Vinciarelli A., Mohammadi G., “A survey of personality computing”, *IEEE Transactions on Affective Computing*, vol. 5, n<sup>o</sup> 3, p. 273-291, 2014.
- Wang F., Chen W., Yang Z., Dong Q., Xu S., Xu B., “Semi-Supervised Disfluency Detection”, *Proceedings of the 27th International Conference on Computational Linguistics*, Association for Computational Linguistics, Santa Fe, New Mexico, USA, p. 3529-3538, August, 2018.
- Wollermann C., Lasarczyk E., Schade U., Schröder B., “Disfluencies and uncertainty perception—evidence from a human–machine scenario”, *Sixth Workshop on Disfluency in Spontaneous Speech*, 2013.
- Yoshida E., Lickley R. J., “Disfluency patterns in dialogue processing”, *DiSS-LPSS Joint Workshop 2010*, 2010.

Zayats V., Ostendorf M., “Giving attention to the unexpected: using prosody innovations in disfluency detection”, *arXiv preprint arXiv:1904.04388*, 2019.

Zayats V., Tran T., Wright R., Mansfield C., Ostendorf M., “Disfluencies and Human Speech Transcription Errors”, *arXiv preprint arXiv:1904.04398*, 2019.

Zhu G., Caceres J.-P., Salamon J., “Filler Word Detection and Classification: A Dataset and Benchmark”, *arXiv preprint arXiv:2203.15135*, 2022.