

THE EFFECT OF HEARING LOSS ON THE INTELLIGIBILITY OF SYNTHETIC SPEECH

XXX

XXX

XXX

ABSTRACT

Many factors affect the intelligibility of synthetic speech. One aspect that has been severely neglected in past work is hearing loss. In this study, we investigate whether pure-tone audiometry thresholds across a wide range of frequencies (0.25–20kHz) are correlated with participants' performance on a simple task that involves accurately recalling and processing reminders. Participants' scores correlate not only with thresholds in the frequency ranges commonly associated with speech, but also with extended high-frequency thresholds.

Keywords: ageing, speech synthesis, intelligibility, audiometry, memory

1. INTRODUCTION

In this paper, we present the first systematic investigation of the effects of age-related changes in hearing thresholds on the intelligibility of unit selection speech synthesis. Older people are a key user group for a wide range of voice interfaces, including applications such as smart home and home care systems [11], automatic reminder systems [13] and systems for delivering health care interventions [2].

1.1. Influence of Ageing on Intelligibility

Older listeners report a wide range of hearing problems [18]. The anatomical and physiological changes that underlie these problems generally reflect the greater vulnerability of the basal portion of the cochlea. Effects of auditory ageing need to be controlled for in any intelligibility study of older listeners. This is illustrated clearly by the results of two studies, one using diphone synthesis [14] and one using unit-selection synthesis [6]. Roring, Hines, and Charness [14] report that subjects performed consistently worse than younger subjects when listening to synthetic speech produced by a diphone synthesiser, even in the presence of context. However, this age effect vanished when hearing loss was taken into account (binaural audiogram, frequencies: 0.25, 0.5, 1, 2, 3, 4, 6, and 8 kHz). In Langner and Black's study, [6], participants asked to

transcribe sentences which were produced by a human recorded in silence (natural speech), a human who was recorded while listening to a multi-speaker babble (natural speech in noise), a synthetic voice (synthetic speech), and the synthetic voice modified to sound like the natural speech in noise. While older listeners (60–90+) performed best in the natural speech in noise and synthetic speech conditions, older participants with self-reported hearing difficulties performed significantly worse than participants with no self-reported hearing problems.

In addition, older people tend to have a lower working memory capacity [15], which affects their ability to briefly store information needed for cognitive processing. Even though synthetic speech produced by formant synthesis can be just as intelligible as natural speech on the segmental level [5], it has been shown that this type of synthetic speech can stretch limited cognitive resources even further [8]. This could be due to a number of factors, including lack of acoustic information in the signal [4] and missing or wrong prosodic cues [10]. Auditory stimuli that are difficult to process lead to increased cognitive load, leaving fewer resources for working memory [12]. Although later studies of concatenative speech synthesis have failed to replicate this effect [16, 14], cognitive factors clearly need to be taken into account as potential confounders.

1.2. Hypotheses

In this pilot study, we investigate the effect of hearing loss on participants' ability to understand synthetic speech as produced by a unit selection speech synthesis system, Cerevoice [1]. Although unit selection technology is state-of-the-art in most commercial synthesis systems, there is next to no research on its intelligibility for older listeners, with one main exception [6]. Hearing loss is assessed using pure-tone audiometry. The frequency range we used, 0.25–20 kHz, covers both more apical (lower frequencies) and more basal regions (higher frequencies) of the cochlea. Conventional pure tone audiometric thresholds are routinely determined for 0.25–8.0 kHz only. We decided to include higher frequencies for two reasons. First, the more basal

region of the cochlea is often first and most severely affected by the effects of ageing [7]. Secondly, it has been hypothesised that hearing loss in extended high frequencies may be an indicator of early, subclinical damage in more apical regions of the cochlea [9]. To the best of our knowledge, the effect of extended high frequency hearing loss on the intelligibility of synthetic speech has never been investigated before. The hypotheses we are testing here are as follows:

Intelligibility Difference: Synthetic stimuli are less intelligible than natural speech stimuli.

Effect of Hearing Loss: Hearing loss affects the intelligibility of synthetic and natural speech.

Effect of Working Memory: Working memory capacity will explain some of the variation in scores not covered by hearing loss.

2. METHOD

2.1. Participants

This paper reports results from 20 participants recruited for a larger study of auditory ageing. 8 participants were younger (age 28 ± 5 years), 12 participants were aged between 50 and 60 (age 54 ± 4 years). The 50–60 age group provides a good baseline for subsequent work: Even though there is often already significant hearing loss in the high frequencies in this age group, clinical hearing problems are still relatively rare.

2.2. Audiological and Cognitive Tests

In this section, we present only that part of the full assessment battery which is relevant to the results discussed here. All participants completed a working memory test [17] that was presented visually and scored from an answer sheet. Visual presentation was chosen because auditory presentation might affect scores. Pure-tone (PTA) and extended high-frequency (EHF) audiometry was measured on a recently calibrated audiometer (Grason-Stadler, Milford, NH; model GSI 61) in a double-walled sound-proofed room (Industrial Acoustics Corporation, Staines, Middlesex, UK). Air-conduction thresholds were measured for each ear at 0.25, 0.5, 1, 2, 3, 4, 6, and 8 kHz following the procedure recommended by the British Society of Audiology [3]. EHF thresholds were established at 9, 10, 11.2, 12.5, 14, 16, 18, and 20 kHz. Testing always began with the better ear in all subjects. Since there are significant differences between the two ears, data from the right and the left ear will be reported separately in this analysis.

For each ear, we computed average hearing thresholds for four frequency groups:

Trad: 0.5, 1, 2, and 4 kHz, the frequencies conventionally used for screening participants in speech

synthesis experiments

F1: 0.25, 0.5, and 1 kHz, the frequency range of F1

F2: 1, 2, and 3 kHz, the frequency range of F2

EHF: 9, 10, and 11.2 kHz, the lower range of extended high frequencies that were measured. These EHF's were chosen because they could be measured in all subjects.

2.3. Synthesis Experiment

The intelligibility of synthesis systems has traditionally been evaluated using highly artificial stimuli. For this study, we used stimuli that are closely modelled on a real-life application—task reminders. 32 reminders were generated, 16 reminders to meet a person at a given time, and 16 reminders to take medication at a given time. In each group, time preceded person or medication in eight sentences, with the order reversed in the other eight. This allowed us to systematically vary the difficulty of target stimuli, with times being the easiest, person names of medium difficulty, and medication names the most difficult. Person names were monosyllabic CVC words with both consonants being nasal or oral stops, the only exception being the name “Rick”. Names had been chosen to ensure the existence of at least two other names with the same rhyme (-VC sequence). Medication names were constructed using morphemes taken from actual medication names to yield words of 3-4 syllables that did not resemble any existing or commonly used medication. While person names were phonologically simple, but easily confoundable, medication names were both unfamiliar and phonologically complex, making them very difficult to remember. This was intended as a safeguard against ceiling effects.

All 32 reminders were synthesised using Cerevoice. Medication names were transcribed by hand, with the transcriptions adjusted to render them maximally intelligible. The reminders were also read by the same speaker, a young Scottish female, who provided the source material for the synthetic voice. The natural speech was then postprocessed to match the procedures used for creating synthetic speech. They were first high-pass filtered with a cut-off frequency of 70 kHz, then downsampled to 16kHz, and finally encoded and decoded with the tools `speexenc` and `speexdec`.

Participants were asked to recall one aspect of the reminder, either the time or the person/medication. In each list, participants recalled 16 times, 8 persons, and 8 medications. 8 times, and 4 person and 4 medication names were presented using natural speech, the other 16 were presented using synthetic speech. If participants' responses were a valid pronunciation of the orthographic form of the target word,

a score of 1 was assigned, otherwise, a score of 0 was assigned. This procedure takes into account differences in accent between the participants and the Scottish English voice that produced the reminders, such as rhoticity.

Four stimulus lists were created. Each reminder was presented using the synthetic voice in two lists, and using natural speech in the remaining two. Reminders were followed by a short question, recorded using the same natural voice as that used for the reminders. In two lists (one synthetic, one natural), participants were asked for the first item of a given reminder, while in the other two conditions, participants were asked for the second item. This was done in order to control for recency effects. The sequence of reminders was randomised once and then kept constant for all lists.

3. RESULTS

3.1. General Results

Four scores were computed for each participant: total score for all reminders (**Total**), total score for reminders presented by a human voice (**Natural**), total score for reminders presented by the synthetic voice (**Synthetic**), and difference between the scores for natural and synthetic speech (Δ (**Syn,Nat**)).

As Table 1 shows, synthetic speech is more difficult to understand than natural speech ($p < 0.005$, $V = 15.5$, Mann-Whitney test). The main source of the difference are medication names: some medication names are consistently more difficult to understand in the synthetic version than in the natural speech. On average, participants identified 93% of person names correctly, 98% of times, but only 59% of medication names, with the percentage of accurately identified medication names dropping to 45% for synthesised stimuli. The differences in mean scores between the two age groups are not significant (t-test).

Table 1: Mean scores for each group

Group	Total	Synthetic	Natural	Δ
Younger	28.63 (1.06)	13.63 (1.06)	15.00 (0.53)	-1.38 (1.30)
50–60	27.42 (1.68)	13.00 (1.65)	14.42 (1.31)	-1.42 (2.47)

3.2. Influence of Hearing Loss

If there is no effect of age, is there is an effect of hearing loss? To answer this question, we first need to establish whether there are indications for hearing loss in our sample. Table 2 shows that even though

all participants would have passed traditional initial screening tests, with threshold **Trad** above 20kHz for the better ear, thresholds are significantly higher for the older group ($p < 0.0001$ or better, t-test, for all thresholds).

The effect of hearing loss on participants' scores was assessed using ANOVAs with the four frequency thresholds plus working memory span as independent, additive variables. Power was insufficient for testing interactions. Separate ANOVAs were computed for thresholds from each ear and for each of the four target scores. The results are summarised in Table 3 (*: $p < 0.05$, **: $p < 0.01$).

The results show that there are clear *effects of hearing loss*. Hearing thresholds on the non-dominant left ear correlate somewhat better with participants' scores than hearing thresholds for the dominant right ear. This effect is strongest for synthetic speech. Of all four thresholds, **EHF**, corresponding to the lowest three extended high frequencies, and **LowF1**, corresponding to the frequency range of F1, explain the largest variation in scores. Judging solely from the spectrum of the speech that was presented, **EHF** should be completely irrelevant, since the highest frequency present in a signal sampled with 16 kHz is 8 kHz. The difference between scores on human and synthetic speech is best predicted by **LowF1** and **LowF2**, but again this holds mainly for the left ear.

Our hypothesis about the *effects of working memory*, on the other hand, needs to be rejected. Even though there were significant differences in working memory span scores between the two participant groups (younger: 37 ± 7.4 , 50–60: 26.83 ± 10.58 , maximum score: 42), there was no effect of Working Memory Span on any of the four scores.

Table 2: Average hearing Thresholds per Frequency Group, right ear, in dB (standard deviation)

Group	Trad	F1	F2	EHF
Younger	2.81 (5.08)	0.63 (3.88)	3.33 (4.71)	9.38 (16.38)
50–60	10.42 (6.69)	8.89 (9.93)	10.14 (5.97)	38.19 (14.20)

4. CONCLUSION

To our knowledge, this is the first study to assess the influence of conventional and extended high frequency audiometry on the understanding of synthetic speech. Although our sample is small, our data lead us to conclude that extended high

Table 3: Percentage of Variation Explained by Hearing Thresholds

Score	EHF	F1	F2	Trad
Left Ear				
Total	38.0%**	6.8%	1.0%	0.0%
Natural	6.8%	5.8%	12.0%	0.3%
Synthetic	21.6%*	21.3%*	13.4%	0.3%
$\Delta(\text{Syn,Nat})$	3.6%	20.9%*	19.9%*	0.5%
Right Ear				
Total	25.4%*	0.3%	1.1%	4.1%
Natural	2.6%	16.9%	3.4%	7.9%
Synthetic	17.7%	13.9%	6.4%	0.0%
$\Delta(\text{Syn,Nat})$	4.5%	23.5%*	7.8%	2.1%

frequency hearing loss (> 8 kHz) clearly predicts the intelligibility of synthetic speech, even in an age group that is not normally regarded as “older”, namely 50–60 year-olds. The main effects come from unfamiliar, phonologically complex target words, medication names, since there were ceiling effects for the other two target types. EHF thresholds, which are not usually measured, emerge as one of the main predictors of participant performance, in addition to thresholds below 5 kHz which cover the regions of the first two formants. The correlation between EHF thresholds and participant scores clearly needs to be investigated further. Furthermore, both this correlation and the large amount of unexplained variation indicate that aspects of hearing loss other than pure-tone thresholds need to be investigated to obtain a clearer picture of the mechanisms at work. We are currently gathering data from a larger group of participants which allows better control for potential confounders. We also plan to correlate findings with tests of auditory processing that are sensitive to potential sub-clinical damage of the cochlea. In future experiments, we will compare different types of synthesis systems in order to relate our results to previous work and vary familiarity and phonological complexity of our targets more systematically.

5. REFERENCES

[1] M. A. Aylett, C. J. Pidcock, and M. E. Fraser. The cerevoice blizzard entry 2006: A prototype database unit selection engine. In *Proceedings of Blizzard Challenge Workshop, Pittsburgh, PA, 2006*.

[2] T. Bickmore and T. Giorgino. Health dialog systems for patients and consumers. *J. Biomed. Inform.*, 39(5):556–571, 2006.

[3] British Society of Audiology. Pure tone air and bone conduction threshold audiometry with and

without masking and determination of uncomfortable loudness levels, 2004.

[4] S. Duffy and D. Pisoni. Comprehension of synthetic speech produced by rule: A review and theoretical interpretation. *Language and Speech*, 35:351–389, 1992.

[5] L. E. Humes, K. J. Nelson, and D. B. Pisoni. Recognition of synthetic speech by hearing-impaired elderly listeners. *Journal of Speech and Hearing Research*, 34:1180–1184, 1991.

[6] B. Langner and A. W. Black. Using Speech In Noise to Improve Understandability for Elderly Listeners. In *Proceedings of ASRU, San Juan, Puerto Rico, 2005*.

[7] F. S. Lee, L. J. Matthews, J. R. Dubno, and J. H. Mills. Longitudinal study of pure-tone thresholds in older persons. *Ear Hear*, 26:1–11, 2005.

[8] P. Luce, T. Feustel, and D. Pisoni. Capacity demands in short-term memory for synthetic and natural speech. *Human Factors*, 25:17–32, 1983.

[9] O. D. Murnane and J. K. Kelly. The effects of high-frequency hearing loss on low-frequency components of the click-evoked otoacoustic emission. *J Am Acad Audiol*, 14:525–33, 2003.

[10] C. R. Paris, M. H. Thomas, R. D. Gilson, and J. P. Kincaid. Linguistic cues and memory for synthetic and natural speech. *Human Factors*, 42:421–431, 2000.

[11] M. Perry, A. Dowdall, L. Lines, and K. Hone. Multimodal and ubiquitous computing systems: Supporting independent-living older users. *IEEE Transactions on Information Technology in Biomedicine*, 8:258–270, 2004.

[12] M. K. Pichora-Fuller, B. A. Schneider, and M. Daneman. How young and old adults listen to and remember speech in noise. *J.Acoust.Soc.Am.*, 97:593–608, 1995.

[13] J. Pineau, M. Montemerlo, M. Pollack, N. Roy, and S. Thrun. Towards robotic assistants in nursing homes: challenges and results. *Robotics and Autonomous Systems*, 42:271–281, 2003.

[14] R. W. Roring, F. G. Hines, and N. Charness. Age differences in identifying words in synthetic speech. *Hum Factors*, 49:25–31, 2007.

[15] T. A. Salthouse, R. L. Babcock, and R. J. Shaw. Effects of adult age on structural and operational capacities in working memory. *Psychol.Aging*, 6:118–127, 1991.

[16] G. Sonntag, T. Portele, and F. Haas. Comparing the comprehensibility of different synthetic voices in a dual task experiment. In *Proc. Third ESCA Workshop on Speech Synthesis, Jenolan Caves*, pages 5–10, 1998.

[17] N. Unsworth and R. Engle. Simple and complex memory spans and their relation to fluid abilities: Evidence from list-length effects. *Journal of Memory and Language*, 54:68–80, 2006.

[18] J. F. Willott. *Ageing and the Auditory System*. Singular, San Diego, CA, 1991.