Effects of the availability of visual cues during adaptation to noise-vocoded speech

Anthony S Trotter ^{a)}, Briony Banks ^{b)}, and Patti Adank ^{a)}

^{a)} Speech, Hearing and Phonetic Sciences, Chandler House, Wakefield Street, University College
London, London, UK

^{b)} Department of Psychology, Fylde College, Lancaster University, Lancaster, LA1 4YF, UK.

Abstract

The ability to quickly adapt to distorted speech signals, such as noise-vocoding, is one of the mechanisms listeners employ to understand one another in challenging listening conditions. In addition, listeners have the ability to exploit information offered by visual aspects of speech, and being able to see the speaker's face while perceiving distorted speech improves perception of and adaptation to these distorted speech signals. However, it is unclear how important viewing specific parts of the speaker's face is to the successful use of visual speech information – particularly, does looking at the speaker's mouth specifically improve recognition of noise-vocoded speech, or is it equally effective to view the speaker's entire face? This study aimed to establish whether viewing specific parts of the speaker's face (eyes or mouth), compared to viewing the whole face, affected perception of and adaptation to distorted sentences. In a secondary aim, we wanted to establish whether it was possible to replicate results on processing of noise-vocoded speech from lab-based experiments in an online setting. We monitored speech recognition accuracy online while participants were listening to noise-vocoded sentences in a between-subjects design with five groups. We first established if participants were able to reliably perceive and adapt to audiovisual noise-vocoded sentences when the speaker's whole face was visible (AV Full). Four further groups were tested: a group in which participants could only view the moving lower part of the speaker's face – i.e., the mouth (AV Mouth), only see the moving upper part of the face (AV Eyes), a group in which participants could not see the speaker's moving lower or upper face (AV Blocked), and a group in which they were presented with an image of a still face (AV Still). Participants repeated around 40% of key words correctly for the noise-vocoded sentences and adapted over the course of the experiment but only when the moving mouth was visible (AV Full and AV mouth). In contrast, performance was at floor level and no adaptation took place in conditions when the moving mouth was not visible (AV Blocked, AV Eyes, and AV Still). Our results show the importance of being able to observe relevant visual speech information from the speaker's mouth region, but not the eyes/upper face region when listening and adapting to speech under challenging conditions online. Second, our results also demonstrated that it is feasible to run speech perception and adaptation studies online, but that not all findings reported for lab studies necessarily replicate.

Key words: Noise-vocoded speech; adaptation; speech perception, audiovisual speech.

1. Introduction

We often interact with others in listening situations that are less than ideal, e.g., in a crowded cafeteria, at a busy railway station, when interacting online, or over a poor audio and/or video connection. Indeed, most of us can deal with these distortions, although our speech recognition tends to be poorer compared to in clear listening conditions (Mattys, Davis, Bradlow, & Scott, 2012). However, one of the strategies that listeners use to deal with challenging listening situations is fast adaptation to some distortions of the speech signal. Such fast adaptation can occur in a relatively short time frame; listeners can improve their response speed and accuracy after exposure to fewer than 30 distorted sentences. Such rapid adaptation has, for example, been reported for noise-vocoded speech (Davis, Johnsrude, Hervais-Adelman, Taylor, & McGettigan, 2005; Hervais-Adelman, Davis, Johnsrude, & Carlyon, 2008) (Davis, Johnsrude, Hervais-Adelman, Taylor, & McGettigan, 2005; Hervais-Adelman, Davis, Johnsrude, & Carlyon, 2008), accented (Adank & Janse, 2010; Banks, Gowen, Munro, & Adank, 2015a, 2015b) and time-compressed speech (Pelle & Wingfield, 2005; Sebastián-Gallés, Dupoux, Costa, & Mehler, 2000).

Much of the research into rapid speech adaptation has used noise-vocoded speech, which is an artificial distortion of the speech signal in which harmonic components are replaced with bands of noise. The distorted signal has lost much of the fine spectral and harmonic detail, but amplitude modulation information is largely preserved (Shannon, Zeng, Kamath, Wygonski, & Ekelid, 1995). The speech signal is first divided into separate frequency bands (generally between 4 and 32). Next, the amplitude envelope is extracted, which is subsequently used to manipulate a broadband carrier signal. This type of distortion has been used as a simulation of how speech and other sounds are transmitted in people with a cochlear implant (Faulkner, Rosen, & Smith, 2000; Rosen, Faulkner, & Wilkinson, 1999), which is an implanted device that restores hearing in those who have severe or complete hearing loss. When normal-hearing listeners are exposed to noise-vocoded speech they generally show adaptation (i.e., improvement in speech perception performance over time). The distorted speech generated by noise-vocoding it is presumed to sound similar to what is perceived by someone with a cochlear implant. It is generally more difficult to understand noise-vocoded speech with a lower number of frequency bands (4 or 6) than a higher number of bands (Dorman, Loizou, & Rainey, 1997; Faulkner et al., 2000; Sohoglu, Pelle, Carlyon, & Davis, 2014).

The studies discussed so far focused on perception of and adaptation to noise-vocoded speech used auditory-only stimuli. However, being able to see as well as hear the speaker can considerably improve perception of different types of distorted or unclear speech (e.g., in background noise), a phenomenon referred to as the *audiovisual benefit* (Erber, 1975; MacLeod & Summerfield, 1987; Sommers, Tye-Murray, & Spehar, 2005; Sumby & Pollack, 1954). Listeners benefit from the availability

of visual cues and are thought to integrate them with auditory speech cues, which then in turn improves speech perception performance. The audiovisual benefit has also been studied for perceptual adaptation to noise-vocoded speech (Banks, Gowen, Munro, & Adank, 2020; Bernstein, Auer, Eberhardt, & Jiang, 2013; Kawase et al., 2009; Pilling & Thomas, 2011). For instance, Pilling & Thomas (2011) and Banks et al (2020) compared adaptation to noise-vocoded sentences with and without audiovisual speech cues. When visual speech cues were made available, listeners adapted more than for auditory-only conditions, although the audiovisual benefit was smaller and earlier in Banks et al., occurring after exposure to only 75 sentences, compared to 228 in Pilling & Thomas. Similar results were reported by Bernstein et al., who report that the presence of visual speech cues leads to more adaptation to noise-vocoded syllables. Visual speech cues can thus benefit listeners during rapid adaptation to distorted speech. However, it remains unclear whether it is only visual cues from the mouth that benefit listeners, or whether cues from other parts of the face (e.g., the eyes) are also able to help listeners adapt.

Several speech perception studies using eye-tracking have shown that listeners look more at a speaker's mouth during perception of speech in noise (Buchan, Paré, & Munhall, 2007, 2008; Lansing & McConkie, 2003) and noise-vocoded speech (Banks et al., 2020). Notably, fixations on the mouth increase for poorer signal-to-noise ratios (Vatikiotis-Bateson, Eigsti, Yano, & Munhall, 1998). These findings suggest that cues from a speaker's mouth are more important than other potential cues from a speaker's face – for example, movements from the eyebrows or forehead. In addition, it may also be the case that directing visual attention specifically to the speaker's mouth can benefit adaptation. Indeed, Banks et al. (2020) observed a relationship between the duration of fixations on a speaker's mouth and speech perception accuracy for noise-vocoded sentences, whereby longer fixations were related to more accurate perception, but the evidence for this relationship was relatively weak. Furthermore, when a listener directs their foveal vision towards (i.e., fixates or looks directly at) a speaker's mouth, other cues from the speaker's face are still accessible in peripheral vision and may contribute to overall improvements in speech perception. Although foveal vision provides the greatest visual acuity, Munhall, Kroos, Josan & Vatikiotis-Bateson (2004) have shown that high spatial frequency is unnecessary for visual speech cues to benefit perception of speech in noise. Similarly, Paré, Richler, ten Hove & Munhall (2003) demonstrated that direct fixation of a speaker's mouth is neither required nor related to the presence of a McGurk effect. Thus, the importance of specifically viewing a speaker's mouth in difficult listening conditions is still not fully clear. Listeners may benefit from viewing a speaker's face as a whole, as they can integrate multiple visual cues from a speaker's face with auditory cues. Conversely, it might be more beneficial if observers can *only* look at the speaker's mouth during adaptation to noise-vocoded speech, as their visual attention would be fully

directed to the most salient visual speech cues. That is, listeners might be able to benefit more from focusing solely on the mouth if the eyes region is inaccessible to them.

A recent study tested to what extent listeners relied on information from the mouth region while listening to noise-vocoded sentences (Drijvers & Özyürek, 2017), although they did not measure adaptation over time. Drijvers and Özyürek's primary aim was to establish how co-speech gestures contribute to information from visible speech to enhance noise-vocoded speech perception, but their design also included conditions in which the speaker's mouth region was obscured. They presented 20 normal hearing native speakers of Dutch with videos of a female speaker producing an action verb with different levels of noise-vocoding, in a free-recall task. The videos were presented in three multimodal arrangements: audio-only speech with an video with the mouth area blurred, audiovisual speech with the mouth area blurred, audiovisual speech with the mouth area blurred plus an iconic gesture. Participants were tested in a within-group design and completed all nine conditions. We will focus on the results omitting the effects of the presence of the iconic gesture. Results showed ceiling effects for the clear speech conditions (100% correct responses), floor effects for the audio-only and audiovisual two-band conditions with the blurred mouth region (<10% correct). Compared to the conditions in which the mouth region was blurred, participants performed on average 10-20% better for the two vocoding conditions. However, as Drijvers and Özyürek did not test whether and how availability of visual speech information displayed by the mouth affected adaptation to noise-vocoded speech, this question remains unaddressed.

The current study aimed to establish to what extent the audiovisual benefit during perception of and adaptation to audiovisual noise-vocoded sentences relies on viewing articulatory visual cues from different parts of the face. We hypothesised first that participants who are able to see the mouth region will show better speech perception performance, and greater adaptation, than those who cannot see the mouth region (Hypothesis 1). This hypothesis would be supported if both perception (the main effect of condition) and the adaptation (the two-way interaction between condition and block or trial) fail to show statistical differences between the AV Full and AV Mouth conditions, and these two conditions differ with AV Eyes, AV Block and AV Still conditions. Second, we hypothesized that those who are required to focus on the mouth region (i.e., with the upper part of the face not visible) will show better speech perception performance, and greater adaptation, than those who can see the moving mouth and the eyes regions (Hypothesis 2). Support for this hypothesis would require significant differences in the main effect of condition and the two-way interaction between condition and block or trial between the AV Full and AV Mouth conditions. Third, if visual cues from the eyes region are useful, then being able to see the eyes should show a different pattern of adaptation compared to when no movement is visible or when the mouth movements are not visible (Hypothesis

3). In this case, we would expect to see significantly better perception and adaptation for the AV Eyes condition than in the AV Still, and AV Block conditions. Additionally, there should be significant differences present for adaptation and perception between the AV Eyes and AV Mouth conditions, with similar performance in the AV Full and AV Eyes Conditions.

First, we tested perception of 4-band audiovisual noise-vocoded sentences also used in Banks et al. (2020), for a group of participants who were exposed to audiovisual stimuli with the whole face of the speaker visible (AV Full). We then tested four other conditions: first, a group of listeners who heard the noise-vocoded sentences in the presence of a still image of the speaker (AV Still) – namely, with no useful visual cues at all. Second, we aimed to establish the relative relevance of different parts of the face for adaptation to and perception of audiovisual noise-vocoded sentences, so tested a group of participants who could not see the eyes region, but who could see the mouth region, and a group had access to the eyes region, but not the mouth region. Another group of participants was exposed to a video of the speaker with the mouth and eyes obscured from view, and a final group was shown a still image of the speaker while being tested. Second, we aimed to replicate the behavioural results reported in Banks et al. (2020) in an online experimental paradigm to demonstrate that participants were able to adapt outside the lab. We also asked participants to give us an indication of their perceived effort as different circumstances in which people process distorted speech have been shown to affect performance in similar ways yet be associated with different levels of perceived effort (McGarrigle, Dawes, Stewart, Kuchinsky, & Munro, 2017; McGarrigle et al., 2014; Pichora-Fuller et al., 2016). Finally, to ensure participants attended to the speaker's face in the experiment, we queried them afterwards about how much attention they paid to the speaker's face and how much they thought being able to see the speaker's face helped them during the task.

2. Method

Participants

We tested 150 participants (18-30 years of age, 125F and 25M), who all declared to be native monolingual speakers of British English and be resident in the UK at the time of the experiment. All declared to have good hearing and vision, and to not have any neurological or psychiatric disorders (including dyslexia). All participants were recruited through the online platform Prolific.co, and the experiment was hosted on Gorilla.sc. We tested 30 participants per condition. Participants were randomly allocated to each condition. Our minimal sample size, per group as well as the ratio of female and male participants was based on Banks et al. (2020), and we tested 30 participants with a ratio of 25F:5M participants (see the Analysis section for further justification of the selected sample size). We replaced one male participant in condition AV Blocked and one male participant in condition Eyes,

both for not engaging with the task (i.e., not giving a single response). The demographics were as follows across the five conditions: AV Full 25F|5M, 24.6M, 3.8SD, AV Mouth 25F|5M, 23.9M 3.6SD, AV Blocked 25F|5M, 22.8M, 3.2SD, AV Still 25F|5M, 24.5M, 3.8SD, AV Eyes 25F|5M, 24.3M, 4.0SD. Participants all consented to take part and were paid upon completion of the experiment at a rate corresponding to £7.50 per hour. The experiment was approved by UCL's Research Ethics Committee (UREC, #0599.001).

Materials

We used the same materials as in Banks et al. (2020), and adapted them to create the stimuli for the specific conditions in the present study. Banks et al. originally used 91 randomly selected Institute of Electrical and Electronics Engineers Harvard sentences (IEEE, 1969). Stimuli were recorded in a soundproofed laboratory using a Shure SM58 microphone and a High Definition Canon HV30 camera. A 26-year-old female native British English speaker recited the sentences, and was asked to look directly at the camera, to remain still and to maintain a neutral facial expression throughout the recordings to minimise head movement (see Figure 1). Video recordings were subsequently imported into iMovie 11, running on an Apple MacBook Pro, as large (960 x 540) high definition digital video (.dv) files. Video recordings were then edited to create a video clip per sentence. The audio tracks for each clip were extracted as audio (.wav) files, then normalised by equating the root mean square amplitude. Next, they were resampled at 22 kHz in stereo, cropped at the nearest zero crossings at voice onset and offset, and vocoded using Praat speech processing software (Boersma & Weenink, 2012) and custom scripts. Speech recordings were noise-vocoded (Shannon et al., 1995) using four frequency bands (cut-offs: 50 Hz → 369 Hz → 1160 Hz → 3124 Hz → 8000 Hz), selected to represent equal spacing along the basilar membrane (Greenwood, 1990). Of the 91 sentences that were originally recorded, we randomly selected a subset of 60 for inclusion in the online experiment. To ensure that timing of the audio and video was synchronous, we attached the noise-vocoded audio stimuli as an audio track to the video stimuli using Final Cut Pro as a mono track to be played over both channels of a participant's headphones. We repeated the same procedure for an additional single sentence in quiet, to be used in the practice trial presented prior to the main experiment. Audiovisual stimuli were saved as MPeg-4 movie (MP4) files with a resolution of 1920x1080. We also created white rectangular shapes that were used to cover (parts of) the speaker's face in four conditions. The rectangle used to cover the eyes or the mouth in the conditions AV Mouth, AV Eyes, and AV Blocked was a width of 1920 pixels and a height of 720 pixels and a resolution of 300 pixels per inch. For the condition AV Still, we used a screenshot of the speaker's face in PNG format with a width of 1907 pixels and a height of 1074 pixels and a resolution of 300 pixels per inch.

Procedure

The experiment was conducted online, via the Gorilla Experiment Builder ([Gorilla.sc](https://gorilla.sc)) (Anwyl-Irvine1, Massonnié, Flitton, Kirkham, & Evershed, 2020) and participants were recruited via Prolific ([Prolific.co](https://prolific.co)). Upon receiving an email invitation from Prolific, participants entered the online study and were linked through to the experiment hosted in Gorilla. They were then given information on the study, before providing consent. Participants who did not provide consent were rejected from the study. Next, they were asked to enable auto play of video and audio on their internet browser, maximise their screen, and plug in their headphones (Bluetooth headphones were excluded). They were subsequently routed to a page where they could check their sound levels where they were played a short sound consisting of one second of white noise. They were asked to replay this sound over their headphones and adjust their volume to a comfortable level before progressing to the headphone check.

The next check was previously developed to allow for more control over sound presentation in online experiments by providing a test to establish whether participants are wearing headphones (Woods, Siegel, Traer, & McDermott, 2017). This test was designed to be difficult to complete if the participant is not wearing headphones, through the manipulation of anti-phase attenuation rather than differences in intensity between the tones. The headphone check is designed as a 3AFC task in which six sets of three sine wave tone stimuli are played. After participants clicked at the start button, a new page appeared where three 200Hz tones were played with a duration of 1000ms, with 100ms on- and off-ramps, two at -14dB (in-phase) and one at -20dB (180° out of phase). The stimulus duration per triad of tones was four seconds (tone duration: 900ms, interstimulus interval: 600ms, time before first stimulus onset: 100ms, time after the last stimulus offset: 100ms). Participants listened to six trials in total. The participants were to decide which tone they perceived as having the lowest intensity by selecting one of three buttons labelled “FIRST sound is SOFTEST”, “SECOND sound is SOFTEST”, and “THIRD sound is SOFTEST”. They had to select the correct stimulus for five of the six trials (accuracy level of 83.3%), or they were rejected.

Participants who successfully completed the headphone check were subsequently routed through to the instructions and a single undistorted and visually unobstructed practice sentence. For the conditions AV Full, AV Mouth, AV Eyes, and AV Blocked, they saw an the same MP4 video and heard the corresponding undistorted sentence. For the condition AV Still, they were presented with the same still PNG image as used in the main experiment. Participants were asked attend to the video and spoken sentence and to type into a response text box any words they thought they had heard. After the single practice trial, they were shown a screen explaining what they should have typed in the response box. Next, they were told that the main experiment would start next and that all trials would

move on to the next trial automatically, so they would not be able to take a break until the main task finished.

In the main task, participants transcribed 60 noise-vocoded sentences. Participants triggered the start of the experiment and each subsequent trial by pressing the “Next” Button at the bottom of each screen. In each trial, the audiovisual noise-vocoded sentence and corresponding visual stimulus was presented. The visual part of this stimulus was different for the five conditions (see Figure 1). Participants in the Full condition saw the unobstructed video. Participants in the AV Mouth and AV Eyes conditions were shown the video with a white rectangle covering the eyes or mouth of the speaker, respectively. As can be seen in Figure 1, the block covered either the upper or lower part of the face. The tip of the nose and chin were mostly visible in the AV Eyes and AV Mouth videos, but sometimes not visible due to the speaker moving while speaking. Participants in the AV Blocked condition were shown the video with a white block covering the mouth, nose, and eyes of the speaker. Here, the chin and forehead of the speaker were visible, but the space in between was covered, so that only small head movements were visible. Finally, participants in the AV Still condition were shown a still image of the speaker, where the entire face was visible, but no movement.

After the main task, the programme moved to a final response screen, where participants typed in their response. They were asked to type in “/” if they could not decipher any words in the sentence. After finishing typing, they could move to the next stimulus by pressing the “Next” button. If they did not press this button, the experiment moved to the next trial automatically after 23 seconds. After the main task was completed, participants were shown a screen with three response sliders and a response text box. Participants were asked to provide ratings of their perceived effort as follows: *“Question 1: Please indicate using the slider below how effortful you found it to understand the sentences (0 = Not effortful, 100 very effortful).”*. They were also asked to rate what proportion of the time the video was presented they looked at the speaker’s face as follows: *“Question 2: Please indicate using the slider below what proportion of the time you spent looking at the speaker’s face when the video was presented (0% of the time - 100% of the time).”*. A final question queried whether being able to see the speaker’s face helped their speech performance: *“Question 3: How much do you think looking at the speaker’s face helped you understand the sentences? (0 Not at all - 100 Very much).”*. At the bottom of the page was a response box where they were invited to type in any comments. After they clicked next, they were returned to Prolific to be evaluated for payment.

All data were collected in a single session lasting approximately 20 minutes. However, as the experiment was in part self-paced, durations differed across participants, although the main transcription part of the study lasted maximally 30min if participants did not manually progress each trial. A single participant (in the condition AV Blocked) took 30 minutes for the main task, but as their

responses were within the ranges specified for accuracy, we included their data in the final analysis. Average durations for the entire session was 21min and 4s across all 150 participants (SD 7min and 14s). The session timed out automatically after 90 minutes. The average duration for the main transcription part of the session task was 12min and 30s (SD 3min and 18s) and as follows for individual conditions: AV Full 13min and 18s (SD 4min and 19s), AV Mouth 13mins and 27s (SD 2min and 18s), AV Blocked 13min and 31s (SD 2min and 18s), AV Still 10min and 28s (SD 3min and 19s), and AV Eyes took 13 minutes and 31 seconds (SD 2 minutes and 18 seconds). Data from all participants in the Condition AV full was collected first, followed by the AV Mouth condition, the AV Blocked condition, the AV Still condition, and the AV Eyes condition. Online testing took place in May-June 2020. Participants were paid £2.50 for their time.

Figure 1. Still images from conditions AV Full, AV Mouth, AV Eyes, AV Blocked, and AV Still.

Design and Analysis

The experiment measured speech perception performance as the percentage of words accurately entered as the dependent variable. The independent variables were Block and Condition. Block represented the percentage of correctly entered key words across 15 sentences and this factor therefore has four levels: Block 1 represented this percentage for sentences 1-15, Block 2 for sentences 16-30, Block 3 for sentence 31-45, and Block 4 for sentences 46-60. The factor Condition had five levels: AV Full, AV Mouth, AV Eyes, AV Blocked, AV Still. The AV Mouth condition was included to establish to which extent forcing participants to focus on the speaker's mouth affects perception of and adaptation to noise-vocoded sentences (Hypothesis 1). The AV Eyes condition was included to test if and how being able to see the eyes supports perception of and adaptation to noise-vocoded speech. The AV Blocked condition was included to determine if and how removing information conveyed by the speaker's mouth and eyes affected perception/adaptation. The AV Still condition was included to test if the presence of moving visual information affected perception/adaptation. If Hypothesis 1 (being able to see the moving mouth improves adaptation and perception) is correct, then participants in condition AV mouth and AV Full should show better speech perception performance and greater adaptation, than participants in conditions AV Block, AV Eyes, and AV Still. If Hypothesis 2 (having to focus on the mouth region improves adaptation/perception) is correct, then participants in condition AV Mouth should show better perception/adaptation, than participants in

conditions AV Full, AV Block, AV Eyes, and AV Still. Finally, if Hypothesis 3 (being able to see the moving eyes while the moving mouth is not visible improves perception/adaptation) is correct, then participants in condition AV Eyes should show better perception/adaptation than participants in conditions AV Block and AV Still.

We retrospectively scored participants' responses according to how many key words (content or function words) they correctly repeated out of a maximum of four. We included/excluded (typed) responses as follows. Responses were scored as correct despite incorrect suffixes (such as -s, -ed, -ing) or verb endings; however if only part of a word (including compound words) was repeated this response was scored as incorrect following Banks et al. (2015, 2020). It should be noted that Banks et al. audio-recorded participants' verbal responses, and these responses were subsequently judged by an experimenter. In contrast, as we asked participants to type in their responses, we therefore also included homonyms (e.g., "weak" instead of "week"), compound words separated by a space (e.g., "door knob" instead of "doorknob", as well obvious typos (e.g., "whire" instead of "wire"). Moreover, we excluded participants as follows: participants who had an average % error rate greater than three standard deviations (3SD) away from the group mean were excluded from further analysis and replaced. Participants were excluded if they failed to provide responses to a number of trials >2SD from the group mean.

As condition AV Full was intended to closely replicate the design of the audiovisual condition in Banks et al. (2020), we initially chose to collect 30 participants as a minimum sample and then used sequential hypothesis testing with Bayes Factors to determine our final sample size (Schönbrodt, Wagenmakers, Zehetleitner, & Perugini, 2017). After collecting the initial 30 participants, we calculated BF_{10} to assess whether we reached a pre-defined level of evidence ($BF_{10} > 3$ in favour of the alternative hypothesis, and $BF_{10} < 0.2$ in favour of the null hypotheses). BF_{10} indicates how likely the data are to occur under the alternative hypothesis. If $BF_{10} > 0.2$ and < 3.0 , we aimed to collect additional participants. After collecting an additional participant for each group, we would calculate BF_{10} until we met the conditions noted above. If more participants were required, we planned to minimise the danger of type 1 and type 2 errors by graphing BF_{10} after running each additional participant to assess whether any changes in the BF were stable. When the BF was stable for four consecutive participants, we would cease data collection. However, BF_{10} exceeded the criterion value of 3.0 after collecting 30 participants for each condition. As such, additional data collection was not necessary.

To calculate BF_{10} , we utilised Bayes Information Criterion (BIC) values obtained during model comparison of linear mixed effects (LME) hierarchical regression models. Specifically, we used the *lme4* (Bates & Sarkar, 2005) package, while the *lmerTest* (Kuznetsova, Brockhoff, & Christensen, 2017)

package was used to find the best-fitting models. *lmerTest* utilises the backward model-selection strategy, which takes as input the most complex model (including random effects) that should a priori fit the data and then iteratively removes model terms, performing a test of model fit after the removal of each term until the optimal model fit is determined. After obtaining the best-fitting model (H1), we compared it against a model excluding the effect of interest (H0), used the *anova* function to obtain BIC values for each. Following Jarosz and Wiley (2014), we used the difference in BIC to compute the Bayes Factor (BF_{10}) using the following equation:

$$BF_{01} = e^{\Delta BIC/2}$$

$$BF_{01} = e^{BIC_{H0} - BIC_{H1}/2}$$

$$BF_{10} = 1/BF_{01}$$

We initially collected and analysed the data from condition AV Full only, as all hypotheses relied on whether it is possible to measure perceptual adaptation to distorted speech in an online paradigm. In this first stage, we tested whether accuracy increases over the course of the experiment, as measured over the course of the four testing blocks, or over the 60 trials. In this case, the H1 BIC value corresponded to a model predicting accuracy including main effect of trial (see section 3.1.1) or block (see supplementary II), whilst H0 BIC was for a model only including random by-participants and by-items slopes. In the second stage, we analysed data for all five conditions and tested main effects of testing block, trial (as a linear and polynomial) and condition, plus the interaction of both, using LMEs as described above. The H1 BIC will therefore be for a model including the critical two-way interaction between trial (see section 3.1.2) or block (see supplementary materials on: <https://osf.io/2w6j4/>) and condition, while the H0 BIC will be for a model including on the main effects.

We also analysed the data collected in the questionnaire presented to participants in the As this dataset was comprised of one observation per participant, we utilised simple linear models to analyse the data after the main task. As this dataset was comprised of one observation per participant, we utilised simple linear models to analyse this data.

The design of conditions AV Full, AV Mouth, and AV Blocked was preregistered on www.AsPredicted.org under number #41527 “Transcribing distorted audiovisual speech.” The inclusion and design of conditions AV Still and AV Eyes was preregistered on www.AsPredicted.org under number #42910 “Transcribing distorted audiovisual speech, a follow-up study.” In all analyses we discuss results for the five conditions in the order they were collected. All raw data plus analysis scripts can be found on the Open Science Framework: <https://osf.io/2w6j4/>.

3. Results

3.1. Accuracy

For 12 trials, stimulus materials could not be loaded by Gorilla across all 150 participants. In seven cases for the Full condition, two for the Mouth condition, two for the Eyes condition, the video mp4 file could not be loaded (which happened to a different stimulus every time and appeared a random occurrence or glitch in Gorilla). In a single case for the Still condition, the audio file could not be loaded. These 12 cases were therefore removed from the data set.

3.1.1 AV Full

Participants in the AV Full condition reported a mean of 1.7 (SD = 1.4) key words correct across the 60 sentences. The mean accuracy (taken as percentage of correct key words) in Block 1 was 33.7% (SD = 4.0%), 42.6% (SD = 36.4%) in Block 2, 45.9% (SD = 35.6%) in Block 3, and 50% (SD = 36.3%), indicating that participants' performance differed across blocks, with an increase in performance of 16.2% between Blocks 1 and 4 (see figure 2). We first examined the effect of Block and Trial (as both a linear and quadratic effect) to test our hypothesis that participants could adapt to noise-vocoded sentences using the *lme4* (Bates & Sarkar, 2005) and *lmerTest* (Kuznetsova et al., 2017) packages. In this analysis, inclusion of the by-participants ($p = .296$) and by-items ($p = .767$) slopes did not significantly improve model fit. The regression also indicated that when trial was included as a main effect, that block did not improve model fit ($p = .345$). The analysis of the effect of Block is reported in the supplementary materials. The best fitting model therefore included only by-participants and by-items random intercepts, and the main effect of trial (see Table I below). In this case, the alternative hypothesis states that participant performance would increase over trials. Therefore, we compared the best fitting model (BIC = 17197) against a model including only the random effects (BIC = 17275). BF_{10} was > 150 , indicating that the evidence in favour of the alternative hypothesis – that adaptation will occur across trials – was very strong, following Raferty (1995). The model outcomes for the linear effect of trial was significant ($t = 10.387, p < .0001$), indicating that participants in the AV Full condition adapted to the masked speech over trials. Whilst the quadratic effect of trial also reached significance ($t = -2.329, p = .02$), this indicates that the effect of trial was better modelled as a linear effect. This effect is illustrated in figure 2 below - generated using the effects package to extract model estimates at five moments of the distribution for the linear function (Fox et al., 2019) - which displays the model estimates of performance by Trial. To conclude, participants showed an increase in performance across the trials for the AV Full condition. Following these results, we decided to collect and analyse data for the four follow-up conditions.

Figure 1. Model estimates of percentage correctly reported key words across trials in condition AV Full, error bars represent one standard error.

3.1.2 Analysis of all five conditions

The mean accuracy in the AV Full condition was 43.069 (SD = 36.03), 43.806 (SD = 36.13) in the AV Mouth condition, 5.486 (SD = 13.625) in the AV Eyes condition, 5.333 (SD = 14.122) in the AV Block condition, and 3.861 (SD = 11.323) in the AV still condition. Figure 3 displays the per block accuracy across conditions. Taken together, the descriptive statistics suggest that performance was almost identical when participants were able to see the speaker's mouth and speech movements.

Figure 2. Mean accuracy across the four blocks in all five conditions, error bars represent one standard error.

To analyse all five conditions, we followed the same procedure as the analysis for the AV Full condition, while including testing condition (factor-coded) as an additional main effect, plus the interactions between condition and block, and condition and trial. The maximal model upon which we conducted the backward step-wise model comparison therefore included by-items, by-participants random intercepts, a random intercept for participant nested within condition, the main effects of block, trial (linear and quadratic) and condition, and the two-way interactions between block and condition, and trial and condition. Random slopes were excluded from the analysis, as their inclusion resulted in issues of singular model fit. The backwards stepwise model selection indicated that the inclusion of the main effect of block ($p = .183$), the interaction between block and condition ($p = .766$) and the simple by-participants random effect did not improve model fit ($p = 1$). As a result, the final model included a random by-items random intercept, a random-intercept for participant nested within condition, the main effects of trial and condition, and the two-way interaction between condition and trial (see supplementary Table 4 for full model syntax and model summary). The analysis including the effect of block can be found in the supplementary materials on the Open Science Framework: <https://osf.io/2w6j4/>.

To assess the likelihood of the alternative hypothesis – that different conditions would elicit different levels of adaptation – we performed model comparison between a model including the interaction (H_1 , BIC = 80876) against a null model only including the main effects (H_0 , BIC = 80976). BF10 was therefore $5.185 \times 10^{+21}$, vastly exceeding Raftery's (1995) threshold for strong evidence (> 150) in favour of the alternative hypothesis. This reflects the floor performance seen in the AV Blocked, AV Eyes, and AV Still conditions relative to the AV Full, and AV Mouth conditions.

The outcomes of the linear model that adaptation occurred in the AV Full ($t = 12.659$, $p < .0001$) and AV Mouth conditions ($t = 12.657$, $p < .0001$). However, participants in the AV Block ($t = 0.876$, $p = .381$), AV Still ($t = 0.664$, $p = .507$), and the AV Blocked ($t = 0.875$, $p = .381$) conditions did not show adaptation with increased exposure. Adaptation, did not differ significantly between the AV Full and AV Mouth conditions ($t = -0.008$, $p = .994$). For the AV Full condition, the data was better described by a linear function of trial, indicating that participants adapted consistently over the course of training. Both the linear and quadratic per-condition estimates are illustrated in figure 3 below. Two of the quadratic estimates (between trials 0 to 20, and 50 to 60) for the AV Mouth condition differ from the linear estimates, suggesting that the largest increase in performance occurred in the first twenty trials, and tailed off slightly in the last ten trials. However, the model estimates demonstrate the fit was better for the linear ($t = 12.657$, $p < .0001$) relative to the quadratic ($t = -5.104$, $p < .0001$) term. This demonstrates that adaptation largely proceeded linearly across trials, with only minor deviations from this trend over training.

In summary, participants in the AV Full and AV Mouth condition showed increased speech perception performance and demonstrated adaptation (i.e., better accuracy for later trials). In contrast, when participants were unable to see the speaker's mouth (AV Block, AV Still, AV Eyes) speech perception was impaired, and participants were unable to adapt to the vocoded speech. In comparison to the AV Mouth condition, the only additional information was being able to see the speaker's eyes in AV Eyes. As perception and adaptation did not differ statistically between these conditions, the results suggest that being able to see the speaker's moving mouth primarily determines participants' perception and adaptation. Taken together, the results support Hypothesis 1 - being able to see the moving mouth improves adaptation and perception - as adaptation and perception did not differ in conditions where participants were able to see the speaker's moving mouth. Hypothesis 2 - that having to focus on the mouth region improves adaptation and perception - did not receive support, however, as participant performance in the AV Mouth and AV Full conditions did not differ, despite being able to see the eyes in the latter. Hypothesis 3 - being able to see the moving eyes while the moving mouth is not visible supports improves adaptation and perception - was also not supported, as there were no statistical differences in adaptation or perception between the AV Eyes, AV Still, or AV Blocked conditions.

Figure 3. Model accuracy estimates across trials in all conditions, error bars indicate one standard error.

3.2. Effort questionnaire

Ratings from two participants in the AV Mouth condition, and from one in the AV Blocked condition were removed as they were $>3SD$ separated from the average for that respective condition. (Participants in the condition AV Full provided on average an effort score of 91.2% (SD 9.7%) and estimated that they had looked at the speaker's face 91.7% (SD 14%) of the time and 65% (SD 28.7%) stated that being able to see the speaker's face helped speech perception. Participants in the condition AV Mouth provided on average an effort score of 85.6% (SD 10.3%) rated they looked at the face 94.6% (SD 8.8%) of the time and 61.3% (SD 31.9%) stated that seeing the speaker's face helped speech perception. Participants in the condition AV Blocked gave an average effort score of 97.2% (SD 11.2%), rated they looked at the face 66.2% (SD 28.8%) of the time and 17.2% (SD 22.2%) stated that seeing the speaker's face helped speech perception. Participants in the condition AV Still gave an average effort score of 98.8% (SD 4.2%), rated they looked at the face 41.9% (SD 30.6%) of the time and 2.1% (4.7%) stated that seeing the speaker's face helped speech perception. Participants in the condition AV Eyes gave an average effort score of 97.9% (SD 5.1%), rated they looked at the face 69.4% (SD 26.3%) of the time and 26.1% (SD 22.7%) stated that seeing the speaker's face helped speech perception.

Three separate models were conducted of the dependent variables Effort (perceived effort score), Face (estimated proportion of time spent looking at the face), and Face (estimation of how much being able to see the face was helpful) with condition (AV Full, AV Mouth, AV Block, AV Still, AV Eyes) as a factor. In each case, AV Full was taken as the reference level for the condition factor. The Effort model revealed that participants reported significantly lower effort in the AV Mouth compared to the AV Full condition ($t = -2.490$, $p = .014$), whilst the AV Block ($t = 2.716$, $p = .007$), AV Eyes ($t = 3.481$, $p = .003$) and AV Still ($t = 3.481$, $p = .0007$) conditions reported significantly higher effort, suggesting that participants found the AV Mouth condition the least effortful (see figure 4). The Face model indicated that participants spent a similar amount of time looking at the speaker's face in the AV Full and AV Mouth conditions ($t = 0.471$, $p = .638$), however participants in the AV Block ($t = -4.159$, $p < .0001$), AV Eyes ($t = -3.679$, $p = .0003$), and AV Still ($t = -7.941$, $p < .0001$) conditions spent significantly less time looking at the face. The Face Helped model suggested that participants in the AV Full and AV Mouth condition found a similar benefit from seeing the face ($t = -0.561$, $p = .575$), while participants found the face helped significantly less in the AV Block ($t = -7.678$, $p < .0001$), AV Eyes ($t = -6.292$, $p < .0001$) and AV Still conditions ($t = -10.211$, $p < .0001$).

Overall, the participant ratings on these three factors align well with the experimental results; participants who could see the space reported lower required effort. Participants reported lower effort for the AV Mouth relative to the AV Full condition, however, this effect was not reflected in the

accuracy data. This pattern offers some degree of support for Hypothesis 2, in that removing the information provided by the eyes in the AV Mouth condition was associated with reduced perceived effort. When participants could not see the speaker's mouth, effort was increased. Participants who were able to see the speaker's mouth (AV Full, AV Mouth) reported the highest benefit from being able to see the face, and that seeing the face assisted, in contrast to participants who could not (AV Block, AV Eyes, AV Still). As a result, it appeared that in this online testing environment, being able to see the speaker's moving mouth improved both objective (adaptation and perception) and subjective measures of performance (perceived effort).

Figure 4. Participant estimates of perceived effort, time spent looking at the speaker's face (Face), and whether the face being visible helped participants during the task (Face Helped). Error bars indicate one standard error.

4. Discussion

This study aimed to establish if viewing only the mouth or eyes of the speaker's face affected perception of and adaptation to noise-vocoded sentences when compared to viewing their whole face. We ran an online experiment with five listener groups, who could either see the full moving face of the speaker (AV Full), see the moving face with the eyes blocked (AV Mouth), see the moving face with the mouth blocked (AV Eyes), see the face with the eyes and mouth blocked (AV Blocked), or were presented with a still image of the speaker's face (AV Still). All groups transcribed 60 4-band noise-vocoded sentences. The results showed clear differences between the five conditions, with participants in the conditions AV Full and AV Mouth showing considerably better overall accuracy

scores than the participants in the other three groups, where performance was effectively at floor level. There was no difference in overall accuracy between conditions AV Full and AV Mouth, and no differences were found between AV Eyes, AV Block, and AV Still. Second, the results showed an interaction between condition and block, indicating that participants in the conditions AV Full and AV mouth improved their accuracy scores over the course of the experiment, while no such pattern was found for the other three conditions. Therefore, perceptual adaptation to noise-vocoded speech was only found when the moving mouth area of the speaker's face was visible.

Full and Still conditions

The results for the AV Full condition in part replicate the results from the audiovisual condition in Banks et al. (2020), as participants adapted to the noise-vocoded sentences over the four blocks. However, participants performed overall worse than in Banks et al.'s audiovisual condition, as our participants showed an accuracy of 43% correct, and accuracy improved from 33.7% to 50% over the four blocks. Participants in Banks et al.'s audiovisual condition repeated an average of 54% of key words correctly, and this accuracy percentage improved from 42% to 61% over their six blocks (participants improved between 42% to 59% over the first four blocks, i.e., over the first 60 sentences). In contrast, the results for the condition Still, which replicates the audio-only condition in Banks et al., show a very different picture. Banks et al. report an average performance of 35% of key words correctly repeated, with performance increasing from 24% to 43% over their six blocks (participants improved between 24% to 37% over the first four blocks, i.e., over the first 60 sentences). We found that average performance was 4% for our Still condition on average, with performance remaining largely stable. Therefore, while we mostly replicated the (patterns in) the results for the Full condition, such replication was clearly not found for the Still condition. Furthermore, baseline and overall accuracy was lower in the AV Full condition in the present study compared to the audiovisual condition in Banks et al (2020).

It is not clear what factors can account for the differences in results between our and Banks et al.'s results for the AV Full and AV Still conditions. It seems plausible that this difference might be accounted for by differences across both studies, the most prominent of which is the difference in testing platform. Banks et al. tested their participants in a sound-proofed, light-controlled lab, and participants were tested using the same stimulus delivery parameters (e.g., headphones, intensity and sound card, screen size and resolution) and in the absence of any other distractions. In contrast, participants in the current study were tested online. They all wore headphones, but these headphones varied in quality, and by our estimation only a very small number (two out of 150 participants) of headphones listed by our participants could match the audio quality delivered by the headphones used in Banks et al. (Sennheiser HD 25-SP II). In Banks et al. participants were tested in a more

controlled environment in terms of focusing their attention on the task, while in our experiment, we could not control their testing environment and whether they were refraining from engaging in other distractions (e.g., looking at their phone). Also, Banks et al. recorded participants' gaze using eye-tracking while participants adapted and could therefore closely monitor whether and where participants looked at the speaker while listening to the sentences (in both their conditions).

As participants were tested in their own environment and eye gaze was not monitored, we cannot be certain that participants' attention was focused on the task alone. However, all participants were asked in a final questionnaire whether and how much they looked at the speaker's face after the main task ended. On average, participants in the condition AV Full estimated that they had looked at the speaker's face 91.7% of the time and 65% stated that being able to see the speaker's face helped speech perception. For the condition AV Still, participants on average looked at the face 98.8% (even though it displayed no movement) of the time even though only 2% stated that being able to see the speaker's face helped speech perception. Second, we presented participants with 30 fewer stimuli than Banks et al., who exposed them to 90 sentences in total. However, it does seem likely that this issue can explain the observed difference in the results for the AV Full and AV Still conditions, as participants in Banks et al. did not improve significantly beyond their block 4, and baseline accuracy between the two studies was vastly different. It therefore seems plausible that participants in our study reached their maximal performance level after four blocks. A final reason might be due to differences in participant sample, particularly given that participants in Banks et al. were recruited from a University (and were therefore mostly undergraduate students), whereas the sample in the present study was drawn from the general population, or as far as participants on Prolific represent this population. However, as we included two conditions where both the mouth and eyes were blocked, using slightly different visual stimuli (a still image compared to the video of the speaker with eyes and mouth obscured), both of which had similarly poor overall accuracy (i.e., at floor), this explanation also seems unlikely. It is thus likely that differences between our results for condition AV Full and AV Still and Banks et al.'s were mostly related to the differences in testing conditions: online versus lab-based.

Mouth, Eyes, and Block conditions

We included the AV Mouth, Eyes, and Block conditions to test the three hypotheses of this study. Hypothesis 1 stated that being able to see the moving mouth region will show better speech perception performance, and greater adaptation, than those who cannot see the mouth region (AV Full and AV mouth). AV Mouth was also included to test whether participants would perform better and adapt more if their attention was focused on the mouth region per Hypothesis 2. The condition AV Eyes and AV Block were included to establish whether information from the eyes was useful per

hypothesis 3. The results from AV Mouth, in which participants could see the mouth moving but the eyes were blocked, were nearly identical to those reported for the AV Full condition, and therefore also replicate in part the results for the audiovisual condition in Banks et al. The results from the conditions AV Full and AV Mouth thus confirm Hypothesis 1. Next, we predicted that participants might show better overall speech perception performance, and greater adaptation, as their attention would be focused specifically on the speaker's articulatory movements. This prediction was not supported, as performance did not differ from that in the AV Full group, and there was also no difference in the rate and amount of perceptual adaptation. Thus, the finding that participants in the condition AV Mouth did not perform better than those in AV Full refutes Hypothesis 2 with respect to objective task performance. However, we found that being able to see the speaker's moving mouth improved subjective measures of performance (perceived effort) compared to the AV Full condition. Thus Hypothesis 2 could be refuted in part with respect to subjective task performance, but we decided to not do so, as we did not specifically account for subjective performance when we designed the experiment and submitted the preregistration. Yet, we are planning to explore the subjective performance differences between distortion conditions further in future online studies. In addition, it can be concluded that not being able to see the speaker's eyes and the upper part of the speaker's face (as these were blocked off in the AV Mouth condition), did not negatively affect speech perception nor adaptation. Furthermore, overall accuracy and adaptation were almost identical in the AV Full and AV mouth conditions, indicating that viewing the speaker's full face offered no additional benefits over and above mouth movements. It is therefore unhelpful for participants to be able to see the speaker's eyes during perception of noise-vocoded audiovisual speech, as listeners were not able to benefit from any visual cues that might be present in the upper part of a speaker's face, for example eyebrow or forehead movements, thus refuting Hypothesis 3.

The results for the conditions AV Eyes and AV Blocked mirrored the results for the condition AV Still. For all three conditions a floor effect was found, with participants on average reporting 4.9% of key words correctly. In addition, participants did not improve their performance over the course of the experiment, although there was a nonsignificant trend towards improvement over the four blocks in the condition AV Eyes. These results were unexpected as they do not replicate findings reported by other studies using noise-vocoded sentences. The majority of studies examining adaptation to noise-vocoded speech were conducted using audio-only stimuli, yet still manage to find evidence that participants adapt after short-term exposure to a low number of sentences or words (Davis et al., 2005; Huyck & Johnsrude, 2012; Kennedy-Higgins, Devlin, & Adank, 2020; Paulus, Hazan, & Adank, 2020). Studies using audiovisual stimuli report adaptation for their audio-only conditions. For instance, Pilling & Thomas (2011) presented two groups of participants with noise-vocoded sentences

in an audiovisual and an audio-only condition. Participants in both groups adapted readily to the noise-vocoded sentences, although participants in the audiovisual group adapted more than those in the auditory-only condition. Participants were exposed to three blocks of 76 sentences (pre-training, training, post-training) and reported key words. Nevertheless, Pilling & Thomas used an 8-band noise-vocoder with a pitch-shift that aimed to approximate a cochlear implant with a 6 mm place mismatch, while we used 4-band noise-vocoded speech without a pitch shift. The degradations are therefore different and it is likely that the degradation used in Pilling & Thomas resulted in overall more intelligible speech.

Our results for the AV Blocked, AV Still, and AV Eyes conditions are instead similar to those reported in (Drijvers & Özyürek, 2017, 2020), as they report close to floor-level performance for their conditions with 2- and 6-band noise-vocoded speech where the mouth area was not clearly visible. Our results are also in line with those reported by (Rosen et al., 1999), who examined perception of and adaptation to 4-band spectrally-shifted noise-vocoded speech in a live setup. Participants were trained to the distorted speech with a connected discourse tracking task. In this task they were to repeat words when communicating with a live speaker who could be seen through a glass partition and whose speech was vocoded and pitch-shifted in real-time. Participants were exposed to eight blocks of five minutes of speech and reported back what they could understand. After the first block, they could only understand around 1% of the key words, but after the training had finished, they could understand over 40%, showing a clear adaptation effect. It therefore appears that our participants' performance was comparable to that of the participants in Rosen et al. after the first training block. However, we do not know if our participants would subsequently also improve, as our experiment ended after 60 sentences had been presented.

It is unclear why participants were unable to understand or adapt to the noise-vocoded sentences in the AV Blocked, AV Still, and AV Eyes conditions. It seems likely that participants simply 'gave up' as they were not able to understand most of the sentence when they first heard them. It could be that participants in fact need to be able to understand at least part of the sentence upon the first encounter for perception to improve over time. However, this explanation cannot fully explain our results, as participants in Rosen et al. and Pilling & Thomas also started out at a similarly low performance level of <5% correct, and participants all improved over the course of both studies. It should be noted that their participants were provided with a significantly larger number of sentences/utterances than was the case in our study, Pilling & Thomas exposed participants to 228 noise-vocoded sentences, and participants in Rosen et al. listened to a speaker whose speech was noise-vocoded and spectrally-shifted for a total of 40 minutes in eight five-minute blocks. In follow-up studies, it could be considered to provide participants with substantially more training sentences

and to present these to them in separate sessions to further facilitate learning and avoid potential fatigue effects. In addition, it might be useful to examine if participants would show better performance and adaptation in similar online conditions to AV Block, AV Still, and AV Eyes for 6, 9, or 10-channel noise-vocoded sentences. Moreover, a final possibility could be that participants in the online conditions simply did not adapt because they were not made aware that they actually *could* adapt to this type of distortion. We did not mention in the instructions that we expected them to adapt and the title of the study on Prolific was “Transcribing distorted audiovisual speech”. Perhaps participants would adapt more if they were ‘primed’ to learn in the instructions or if learning or adaptation was mentioned in the name of the experiment. We asked participants to provide comments after the main task ended, and most comments could be summarised as they all found the task very difficult and the speech near-impossible to understand.

Conclusion

The results from our study thus demonstrate that it is essential to be able to see a speaker’s moving mouth while trying to understand noise-vocoded sentences in an online setup. When the mouth was not visible, participants could not understand the noise-vocoded sentences at all. Our prediction that being able to see the speaker’s mouth movements would benefit observers when listening to noise-vocoded speech, compared to situations when these movements were not visible was confirmed. However, our prediction that participants who were forced to focus more on the mouth rather than the whole face would perform better, was not confirmed as participants in both AV Full and AV Mouth conditions performed equally well. In addition, it also appears that being able to see the eyes region, but not the mouth region, does not support speech perception or adaptation. Finally, from our results it can also be concluded that, while we partially replicate the lab-based results from Banks et al., it should not be assumed that lab-based tasks will necessarily replicate in online designs, especially when these tasks are particularly difficult.

References

- Adank, P., & Janse, E. (2010). Comprehension of a novel accent by young and older listeners. *Psychology and Aging, 25*(3), 736-740. doi:10.1037/a0020054
- Anwyl-Irvine1, A. L., Massonnié, J., Flitton, A., Kirkham, N., & Evershed, J. K. (2020). Gorilla in our midst: An online behavioral experiment builder. *Behavior Research Methods, 52*, 388-407. doi:10.3758/s13428-019-01237-x
- Banks, B., Gowen, E., Munro, K., & Adank, P. (2015a). Audiovisual cues benefit recognition of accented speech in noise but not perceptual adaptation. *Frontiers in Human Neuroscience, 9*(422), 1-13. doi:10.3389/fnhum.2015.00422

- 701 Banks, B., Gowen, E., Munro, K., & Adank, P. (2015b). Cognitive predictors of perceptual adaptation
702 to accented speech. *Journal of the Acoustical Society of America*, 137(4), 2015-2024.
703 doi:10.1121/1.4916265
- 704 Banks, B., Gowen, E., Munro, K., & Adank, P. (2020). Individual differences in eye gaze during
705 perceptual adaptation to audiovisual noise-vocoded speech. *PsyArXiv*.
706 doi:10.31234/osf.io/65849
- 707 Bates, D., & Sarkar, D. (2005). lme4: Linear Mixed-Effects Models Using Eigen and R. [http://CRAN.R-](http://CRAN.R-project.org)
708 [project.org](http://CRAN.R-project.org).
- 709 Bernstein, L. E., Auer, E. T., Eberhardt, S. P., & Jiang, J. (2013). Auditory perceptual learning for speech
710 perception can be enhanced by audiovisual training. *Frontiers in Neuroscience*, 7(34).
711 doi:10.3389/fnins.2013.00034
- 712 Boersma, P., & Weenink, D. (Producer). (2012). Praat: doing phonetics by computer (Version 5.3.05).
713 Retrieved from <http://www.praat.org/>
- 714 Buchan, J. N., Paré, M., & Munhall, K. G. (2007). Spatial statistics of gaze fixations during dynamic face
715 processing. *Social Neuroscience*, 2(1), 1-13. doi:10.1080/17470910601043644
- 716 Buchan, J. N., Paré, M., & Munhall, K. G. (2008). The effect of varying talker identity and listening
717 conditions on gaze behavior during audiovisual speech perception. *Brain Research*, 25, 162-
718 171. doi:10.1016/j.brainres.2008.06.083
- 719 Davis, M. H., Johnsruide, I. S., Hervais-Adelman, A. G., Taylor, K., & McGettigan, C. (2005). Lexical
720 information drives perceptual learning of distorted speech: evidence from the comprehension
721 of noise-vocoded sentences. *Journal of Experimental Psychology: General*, 134(2), 222-241.
722 doi:10.1037/0096-3445.134.2.222
- 723 Dorman, M. F., Loizou, P. L., & Rainey, D. (1997). Speech intelligibility as a function of the number of
724 channels of stimulation for signal processors using sine-wave and noise-band outputs. *Journal*
725 *of the Acoustical Society of America*, 102(4), 2403-2411. doi:10.1121/1.419603
- 726 Drijvers, L., & Özyürek, A. (2017). Visual context enhanced: The joint contribution of iconic gestures
727 and visible speech to degraded speech comprehension. *Journal of Speech, Language, and*
728 *Hearing Research*, 60(1), 212-222. doi:10.1044/2016_JSLHR-H-16-0101
- 729 Drijvers, L., & Özyürek, A. (2020). Non-native Listeners Benefit Less from Gestures and Visible Speech
730 than Native Listeners During Degraded Speech Comprehension *Language and Speech*, 62(3),
731 209-220. doi:10.1177/0023830919831311
- 732 Erber, N. P. (1975). Auditory-visual perception of speech. *Journal of Speech and Hearing Disorders*,
733 40(4), 481-492. doi:10.1044/jshd.4004.481

- Faulkner, A., Rosen, S., & Smith, C. (2000). Effects of the salience of pitch and periodicity information on the intelligibility of four-channel vo- coded speech: Implications for cochlear implants. *Journal of the Acoustical Society of America*, 108, 1877-1887. doi:10.1121/1.1310667
- Fox, J., Weisberg, S., Friendly, M., Hong, J., Anderson, R., Firth, D., & Taylor, S. (2019). Effects Package: Effect Displays for Linear, Generalized Linear, and Other Models:. R package version 4.1-4. Retrieved from <https://cran.r-project.org/web/packages/effects/index.html>
- Hervais-Adelman, A. G., Davis, M. H., Johnsrude, I. S., & Carlyon, R. P. (2008). Perceptual learning of noise vocoded words: effects of feedback and lexicality. *J Exp Psychol Hum Percept Perform*, 34(2), 460-474. doi:10.1037/0096-1523.34.2.460
- Huyck, J. J., & Johnsrude, I. S. (2012). Rapid perceptual learning of noise-vocoded speech requires attention. *J Acoust Soc Am*, 131(3), EL236-242. doi:10.1121/1.3685511
- IEEE. (1969). IEEE recommended practice for speech quality measurements. *IEEE Transactions on Audio and Electroacoustics*, AU-17, 225-246.
- Jarosz, A. F., & Wiley, J. (2014). What are the odds? A practical guide to computing and reporting Bayes factors. *The Journal of Problem Solving*, 7(7), 2. doi:10.7771/1932-6246.1167
- Kawase, T., Sakamoto, S., Hori, Y., Maki, A., Suzuki, Y., & Kobayashi, T. (2009). Bimodal Audio-Visual Training Enhances Auditory Adaptation Process. *NeuroReport*, 20(14), 1231-1234. doi:10.1097/WNR.0b013e32832fbef8
- Kennedy-Higgins, D., Devlin, J. T., & Adank, P. (2020). Investigating the cognitive mechanisms underpinning successful perception of different speech distortions. *The Journal of the Acoustical Society of America*, 147(4), 2728-2740. doi:10.1121/10.0001160
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. (2017). lmerTest package: tests in linear mixed effects models. *Journal of statistical software*, 83(13), 1-26. doi:10.18637/jss.v082.i13
- Lansing, C. R., & McConkie, G. W. (2003). Word identification and eye fixation locations in visual and visual-plus-auditory presentations of spoken sentences. *Perception & Psychophysics*, 65(4), 536-552. doi:10.3758/BF03194581
- MacLeod, A., & Summerfield, A. Q. (1987). Quantifying the contribution of vision to speech perception in noise. *British Journal of Audiology*, 21(October), 131-141. doi:10.3109/03005368709077786
- Mattys, S. L., Davis, M. H., Bradlow, A. R., & Scott, S. K. (2012). Speech recognition in adverse conditions: A review. *Language and Cognitive Processes*, 27(7/8), 953-978. doi:10.1080/01690965.2012.705006

- 766 McGarrigle, R., Dawes, P., Stewart, A. J., Kuchinsky, S. E., & Munro, K. J. (2017). Pupillometry reveals
767 changes in physiological arousal during a sustained listening task. *54*(2), 193-203.
768 doi:doi.org/10.1111/psyp.12772
- 769 McGarrigle, R., Munro, K. J., Dawes, P., Stewart, A. J., Moore, D. R., Barry, J. G., & Amitay, S. (2014).
770 Listening effort and fatigue: what exactly are we measuring? A British Society of Audiology
771 Cognition in Hearing Special Interest Group 'white paper. *International Journal of Audiology*,
772 *53*(7), 433-440. doi:10.3109/14992027.2014.890296
- 773 Munhall, K. G., Kroos, C., Jozan, G., & Vatikiotis-Bateson, E. (2004). Spatial frequency requirements for
774 audiovisual speech perception. *Perception & Psychophysics*, *66*(4), 574-583.
775 doi:10.3758/BF03194902
- 776 Paré, M., Richler, R. C., ten Hove, M., & Munhall, K. G. (2003). Gaze behavior in audiovisual speech
777 perception: The influence of ocular fixations on the McGurk effect. *Perception &*
778 *Psychophysics*, *65*(4), 553-567. doi:10.3758/BF03194582
- 779 Paulus, M., Hazan, V., & Adank, P. (2020). The relationship between talker acoustics, intelligibility and
780 effort in degraded listening conditions. *Journal of the Acoustical Society of America*, *147*.
781 doi:10.1121/10.0001212
- 782 Peelle, J. E., & Wingfield, A. (2005). Dissociations in perceptual learning revealed by adult age
783 differences in adaptation to time-compressed speech. *Journal of Experimental Psychology:*
784 *Human Perception and Performance*, *31*(6), 1315-1330. doi:10.1037/0096-1523.31.6.1315
- 785 Pichora-Fuller, M. K., Kramer, S. E., Eckert, M. A., Edwards, B., Hornsby, B. W., Humes, L. E., . . .
786 Wingfield, A. (2016). Hearing Impairment and Cognitive Energy: The Framework for Under-
787 standing Effortful Listening (FUEL). *Ear and Hearing*, *37*, 5S-27S.
788 doi:10.1097/AUD.0000000000000312
- 789 Pilling, M., & Thomas, S. (2011). Audiovisual Cues and Perceptual Learning of Spectrally Distorted
790 Speech. *Language and Speech*, *54*(4), 487-497. doi:10.1177/0023830911404958
- 791 Raftery, A. E. (1995). Bayesian model selection in social research. *Sociological methodology*, *25*, 111-
792 163. doi:10.2307/271063
- 793 Rosen, S., Faulkner, A., & Wilkinson, L. (1999). Adaptation by normal listeners to upward spectral shifts
794 of speech: Implications for cochlear implants. *Journal of the Acoustical Society of America*,
795 *106*(6), 3629-3636. doi:10.1121/1.428215
- 796 Schönbrodt, F. D., Wagenmakers, E.-J., Zehetleitner, M., & Perugini, M. (2017). Sequential hypothesis
797 testing with Bayes factors: Efficiently testing mean differences. *Psychological Methods*, *22*(2),
798 322–339. doi:10.1037/met0000061

- Sebastián-Gallés, N., Dupoux, E., Costa, A., & Mehler, J. (2000). Adaptation to time-compressed speech: Phonological determinants. *Perception & Psychophysics*, 62, 834-842. doi:10.3758/BF03206926
- Shannon, R. V., Zeng, F. G., Kamath, V., Wygonski, J., & Ekelid, M. (1995). Speech recognition with primarily temporal cues. *Science*, 270, 303-304. doi:10.1126/science.270.5234.303
- Sohoglu, E., Pelle, J. E., Carlyon, R. P., & Davis, M. H. (2014). Top-Down Influences of Written Text on Perceived Clarity of Degraded Speech. *Journal of Experimental Psychology: Human Perception and Performance*, 40(1), 186-199. doi:10.1037/a0033206
- Sommers, M. S., Tye-Murray, N., & Spehar, B. (2005). Auditory-visual speech perception and auditory-visual enhancement in normal-hearing younger and older adults. *Ear and Hearing*, 26, 263-275.
- Sumby, W. H., & Pollack, I. (1954). Visual contribution of speech intelligibility in noise. *Journal of the Acoustical Society of America*, 26, 212-215. doi:10.1121/1.1907309
- Vatikiotis-Bateson, E., Eigsti, I. M., Yano, S., & Munhall, K. G. (1998). Eye movement of perceivers during audiovisual speech perception. *Perception & Psychophysics*, 60(6), 926-940. doi:10.3758/BF03211929
- Woods, K. J., Siegel, M. H., Traer, J., & McDermott, J. H. (2017). Headphone screening to facilitate web-based auditory experiments. *Attention, Perception & Psychophysics*, 79(7), 2064-2072. doi:10.3758/s13414-017-1361-2