

Towards Estimating the Upper Bound of Visual-Speech Recognition: The Visual Lip-Reading Feasibility Database

Adriana Fernandez-Lopez, Oriol Martinez and Federico M. Sukno

Department of Information and Communication Technologies, Pompeu Fabra University, Barcelona, Spain

Abstract—Speech is the most used communication method between humans and it involves the perception of auditory and visual channels. Automatic speech recognition focuses on interpreting the audio signals, although the video can provide information that is complementary to the audio. Exploiting the visual information, however, has proven challenging. On one hand, researchers have reported that the mapping between phonemes and visemes (visual units) is one-to-many because there are phonemes which are visually similar and indistinguishable between them. On the other hand, it is known that some people are very good lip-readers (e.g: deaf people). We study the limit of visual only speech recognition in controlled conditions. With this goal, we designed a new database in which the speakers are aware of being read and aim to facilitate lip-reading. In the literature, there are discrepancies on whether hearing-impaired people are better lip-readers than normal-hearing people. Then, we analyze if there are differences between the lip-reading abilities of 9 hearing-impaired and 15 normal-hearing people. Finally, human abilities are compared with the performance of a visual automatic speech recognition system. In our tests, hearing-impaired participants outperformed the normal-hearing participants but without reaching statistical significance. Human observers were able to decode 44% of the spoken message. In contrast, the visual only automatic system achieved 20% of word recognition rate. However, if we repeat the comparison in terms of phonemes both obtained very similar recognition rates, just above 50%. This suggests that the gap between human lip-reading and automatic speech-reading might be more related to the use of context than to the ability to interpret mouth appearance.

I. INTRODUCTION

Speech is the most used communication method between humans, and it is considered a multi-sensory process that involves perception of both acoustic and visual cues since McGurk demonstrated the influence of vision in speech perception. Many authors have subsequently demonstrated that the incorporation of visual information into speech recognition systems improves their robustness [1], [2].

Visual information usually involves position and movement of the visible articulators (the lips, the teeth and the tongue), speaker localization, articulation place and other signals not directly related to the speech (facial expression, head pose and body gestures) [3], [4], [5]. Even though the audio is in general much more informative than the video signal, speech perception relies on the visual information to help decoding spoken words as auditory conditions are degraded [3], [6], [7], [8]. Furthermore, for people with hearing impairments, the visual channel is the only source of information to understand spoken words if there is no sign language interpreter [2], [9], [10]. Therefore, visual speech recognition is implicated in our speech perception process

and is not only influenced by lip position and movement but it also depends on the speaker's face, as it has been shown that it can also transmit relevant information about the spoken message [4], [5]. Much of the research in Automatic Speech Recognition (ASR) systems has focused on audio speech recognition, or on the combination of both modalities using Audio-Visual Automatic Speech Recognition (AV-ASR) systems to improve the recognition rates, but Visual Automatic Speech Recognition (VASR) systems have been less frequently analyzed alone [11], [12], [13], [14], [15], [16], [17]. The performance of audio only ASR systems is very high if there is not much noise to degrade the signal. However, in noisy environments AV-ASR systems improves the recognition performance when compared to their audio-only equivalents [2], [11]. In contrast, in visual only ASR systems the recognition rates are rather low [18]. This can be partially explained by the higher difficulty associated to decoding speech through the visual channel, when compared to the audio channel.

One of the key limitations of VASR systems resides on the ambiguities that arise when trying to map visual information into the basic phonetic unit (phonemes), i.e. not all the phonemes that are heard can be distinguished by observing the lips. There are two types of ambiguities: *i*) there are phonemes that are easily confused because they look visually similar between them (e.g: /p/, /b/ and /m/). For example, the phones /p/ and /b/ are visually indistinguishable because voicing occurs at the glottis, which is not visible; *ii*) there are phonemes whose visual appearance can change (or even disappear) depending on the context. This is the case of the *velars*, consonants articulated with the back part of the tongue against the soft palate (e.g: /k/ or /g/), because they change their position in the palate depending on the previous or following phoneme. Specifically, velar consonants tolerate palatalization (the phoneme changes to palatal) when the previous or following phoneme is a vowel or a palatal [19]. Other drawbacks associated to lipreading have also been reported in the literature, such as the distance between the speakers, illumination conditions or visibility of the mouth [3], [20], [21]. However, the latter can be easily controlled, while the ambiguities explained above are limitations intrinsic to lip-reading and constitute an open problem.

On the other hand, it is known that some people are very good lip-readers. In general, visual information is the only source of reception and comprehension of oral speech for people with hearing impairments, which leads

to the common misconception that they must be good lip-readers. Indeed, while many authors have found evidence that people with hearing impairments outperform normal-hearing people in comprehending visual speech [22], [23], [24], [25], [26], there are also several studies where no differences were found in speech-reading performance between normal-hearing and hearing-impaired people [27], [28]. Such conflicting conclusions might be partially explained by the influence of other factors beyond hearing impairment. For example, it is well known that human lip-readers use the context of the conversation to decode the spoken information [3], [5], [20], thus it has been argued that people who are good lip-readers might be more intelligent, with more knowledge of the language, and with a more comprehensible oral speech for others [21], [27], [29], [30].

While the above complexities may provide some explanation to the rather low recognition rates of VASR systems, there seems to be a significant gap between these and human lip-reading abilities. More importantly, it is not clear what would be the upper bound of visual-speech recognition, especially for systems not using context information (it has been argued that humans can *read* only around 30% of the information from the lips, and the rest is filled-in from the context [21], [31]). Thus, it is not clear if the poor recognition rates obtained in VASR systems are due to the inappropriate or incomplete design or because there is an intrinsic limitation in visual information that causes the impossibility of perfect decoding of the spoken message.

Contributions: In this work we explore the feasibility of visual speech reading with the aim to estimate the recognition rates achievable by human observers under favorable conditions and compare them with those achieved by an automatic system. To this end, we focus on the design and acquisition of an appropriate database in which recorded speakers actively aim to facilitate lip-reading but conversation context is minimized. Specifically, we present a new database recorded with the explicit goal of being visually informative of the spoken message. Thus, data acquisition is especially designed with the aim that a human observer (or a system) can decode the message without the help of the audio signal. Concretely, lip-reading is applied to people that are aware of being read and has been instructed to make every effort so that they can be understood based exclusively on visual information. Then, the database deals with sentences that are uttered slowly, with repetitions, well pronounced and viewed under optimal conditions ensuring good illumination and mouth visibility (without occlusions and distractions).

In this database we divided the participants in two groups: 9 hearing-impaired subjects and 15 normal-hearing subjects. In our tests, hearing-impaired participants outperformed the normal-hearing participants but without reaching statistical significance. Human observers outperform markedly the VASR system in terms of word recognition rates, but in terms of phonemes, the automatic system achieves very similar accuracy to human observers.

II. AUDIO-VISUAL SPEECH DATABASES

Visual only speech recognition spans over more than thirty years, but even today is still an open problem in science. One of the limitations for the analysis of VASR systems is the accessible data corpora. Despite the abundance of audio speech databases, there exist a limited number of databases for audio-visual or visual only ASR research. That is explained in the literature because the field is relatively young, and also, because the audio-visual databases add some challenges such as database collection, storage and distribution, not found as a problem in audio corpora. Acquisition of visual data at high resolution, frame rate and image quality, with optimal conditions and synchronized with the audio signal requires expensive equipment. In addition, visual storage is at least one or two orders of magnitude to the audio signal, making his distribution more difficult [13], [32].

Most databases used in audio-visual ASR systems suffer from one or more weaknesses. For example, they contain low number of subjects ([33], [34]), small duration ([33], [34], [35], [36]), and are addressed to specific and simple recognition tasks. For instance, most corpora are centered in simple tasks such as isolated or connected letters ([33], [34], [35]), digits ([35], [36], [37], [38], [39]), short sentences ([36], [40], [41], [42], [43], [44]) and only recently continuous speech ([38], [45], [46], [47]). These restrictions make more difficult the generalization of methods and the construction of robust models because of the few samples of training. Additional difficulties are that some databases are not freely available.

As explained in Section I the aim of this project is to apply continuous lip-reading to people that are conscious of being read and is trying to be understood based exclusively on visual information. Thus, from the most common databases, only VIDTIMIT [40], AVICAR [35], Grid [41], MOBIO [42], OuluVS [43], OuluVS2 [44], AV@CAR [45], AV-TIMIT [46], LILiR [47] contain short sentences or continuous speech and could be useful to us. However, we rejected the use of them because the participants speak in normal conditions without previous knowledge of being lip-read. In addition, most of the databases have low technical aspects and limited number of subjects with restricted vocabularies centred in repetitions of short utterances. Subsequently, we decided to develop a new database designed specifically for recognizing continuous speech in controlled conditions.

III. VISUAL LIP-READING FEASIBILITY DATABASE

The Visual Lip-Reading Feasibility (VLRf) database is designed with the aim to contribute to research in visual only speech recognition. A key difference of the VLRf database with respect to existing corpora is that it has been designed from a novel point of view: instead of trying to lip-read from people who are speaking naturally (normal speed, normal intonation,...), we propose to lip-read from people who strive to be understood.

Therefore, the design objective was to create a public database visually informative of the spoken message in which

it is possible to directly compare human and automatic lip-reading performance. For this purpose, in each recording session there were two participants: one speaker and one lip-reader. The speaker was recorded by a camera while pronouncing a series of sentences that were provided to him/her; the lip-reader was located in a separate room, acoustically isolated from the room where the speaker was located. To make the human decoding as close as possible to the automatic decoding, the input to the lip-reader was exclusively the video stream recorded by the camera, which was displayed in real time by means of a 23" TV screen.

After each uttered sentence, the lip-reader gave feedback to the speaker (this was possible because it was possible to enable audio feedback from the lip-reading room to the recording room, but not conversely). Each sentence could be repeated up to 3 times, unless the lip-reader decoded it correctly in fewer repetitions. Both the speaker utterances and the lip-reader answers (in each repetition) were annotated.

Participants were informed about the objective of the project and the database. They were also instructed to make their best effort to be easily understood, but using their own criteria (e.g: speak naturally or slowly, emphasize separation between words, exaggerate vocalization,...).

Each recording session was divided in 4 levels of increasing difficulty: 3 levels with 6 sentences and 1 level with 7 sentences. We decided to divide the session in different levels to make it easier for participants to get accustomed to the lip-reading task (and perhaps also to the speaker). Specifically, in the first level the sentences are short with only few words, and as the level increases the difficulty increases in terms of number of words. The sentences are unrelated among them and only the context within the sentence is present. Thus, in the first sentences participants had to read fewer words but with very little context and in the last sentences the context was considerably more important and would certainly help decoding the sentence. To motivate participants and to ensure their concentration during all the session, at the end of each level both participants changed their roles.

Finally, because our objective was to determine the visual speech recognition rates that could be achievable, we also recruited volunteers which were hearing-impaired and accustomed to use lip-reading in their daily routine. Then, we will also compare the capability of lip-reading of normal-hearing and hearing-impaired people.

A. Participants

We recruited 24 adult volunteers (3 male and 21 female). Thirteen are University students, one is Teacher of Sign Language at UPF and the other 10 participants are members of the Catalan Federation of Associations of Parents and Deaf (ACCAPS) [48]. The 24 participants were divided in two groups: normal-hearing people and hearing-impaired people.

– *Normal-hearing participants.* Fifteen of the volunteers are normal-hearing participants (14 females and 1 male), who were selected from a similar educational range (e.g: same degree) because, as explained in Section I, lip-reading abilities have been related to intelligence and language



Fig. 1. Scheme of the recording setup and snapshots of the VLR database.

knowledge. Two of the participants were more than 50 years old and have a different education level while the other 13 subjects of this group shared educational level and age range.

– *Hearing-impaired participants.* There were nine hearing-impaired participants, all above 30 years old (7 female and 2 male). Eight of them have post-lingual deafness (the person loses hearing after acquiring spoken language) and one has pre-lingual deafness (the person loses hearing before the acquisition of spoken language). There were 4 participants with cochlear implants or hearing aids.

B. Utterances

Each participant was asked to read 25 different sentences, from a total pool of 500 sentences, proceeding similarly to [41]. The sentences were unrelated between them to avoid that lip-readers could benefit from conversation context. Sentences had different levels of difficulty, in terms of their number of words. There were 4 different levels, from 3-4 words, 5-6 words, 7-8 words and 8-12 words. We decided to divide the sentences in different levels for two reasons. Firstly, to allow lip-readers to get some *training* with the short sentences of the first level (i.e. to get acquainted and gain confidence with the setup, the task and the speaker). Secondly, to compare the effect of the context in the performance of human lip-readers. The utterances with fewer words have very little context, while longer sentences contained considerable context that should help the lip-reader when decoding the message.

Overall, there were 10200 words in total (1374 unique). The average duration per sentence was 7 seconds and the whole duration of the database was 180 minutes (540,162 frames). The sentences contained a balanced phonological distribution of the Spanish language, based on the phonological distribution of the balanced utterances used in the AV@CAR database [45].

C. Technical aspects

The database was recorded in two contiguous soundproof rooms (Fig. 1). The distribution of the recording equipment into the rooms is shown in Fig. 1. A Panasonic HPX 171 camera was located with a tripod PRO6-HDV in front of the chair of the speaker, to ensure an approximately frontal face shot, with a supplementary directional microphone mounted on the camera to ensure a directional coverage in

the direction of the speaker. The camera recorded a close up shot (Fig.1) at 50 fps with a resolution of 1280×720 pixels and audio at 48 kHz mono with 16-bit resolution. Two Lumatek ultralight 1000W Model 53-11 were used together with reflecting panels to obtain a uniform illumination and minimize shadows or other artifacts on the speaker's face. When performing the lip-reading task, the lip-reader was located in the control room. The position of the lip-reader was just in front of a 23" LG Flatron M2362D PZ TV. This screen was connected to the camera so that it reproduced in real time what the camera was recording. Only the visual channel of the camera was fed into the control room, although both audio and video channels are recorded for post processing of the database. The rooms were acoustically isolated between them except for the feedback channel composed by a microphone in the control room and a loudspeaker in the recording room. This channel was used after each utterance to let the speaker know what message was decoded by the lip-reader.

D. Data labeling

The ground-truth of the VLRF database consists in labelling each frame with the phoneme being pronounced at that time instant.

We used the EasyAlign plug-in from Praat [49], which allows to locate the phoneme in each time instant based on the audio stream. Specifically, the program locates the phonemes semi-automatically and there is usually the need for manual intervention to adapt the boundaries of each phoneme to more precise positions. The phonemes used are based on the phonetic alphabet SAMPA [50]. For the Spanish language, the SAMPA vocabulary is composed of the following 31 phonemes: /p/, /b/, /t/, /d/, /k/, /g/, /tS/, /j/, /f/, /B/, /T/, /D/, /s/, /z/, /x/, /G/, /m/, /n/, /N/, /J/, /l/, /L/, /r/, /4/, /j/, /w/, /a/, /e/, /i/, /o/, /u/.

IV. RESULTS

In this section we show the word- and phoneme-recognition rates obtained in our experiments. We start by analyzing the human lip-reading abilities and comparing the performance of hearing-impaired and normal-hearing participants. Then, we analyse the influence of training and context in human performance. Finally, we compare the performance of our automatic system to the results obtained by human observers.

The use of two separate measures (word and phoneme rates) is necessary to analyze different aspects of our results. On one hand, phonemes are the minimum distinguishable units of speech and directly constitute the output of our automatic system. However, the ultimate goal of lip-reading is to *understand* the spoken language, hence the need to focus (at least) on words. It is important to notice that acceptable phoneme recognition rates do not necessarily imply good word recognition rates, as will be shown later.

The word recognition rate was computed as the fraction of words correctly understood in a given sentence. The phoneme recognition rate was computed as the fraction of




Fig. 2. Top: Analysis of word recognition rate for each user (normal-hearing (H) and hearing-impaired people (H-Imp)) at each repetition; Bottom: Word accuracy for normal-hearing (H) and hearing-impaired people (H-Imp) at each repetition.

video frames in which the correct phoneme was assigned. Consequently, 25 accuracy measures were computed for each participant and each repetition. Recognition rates for the automatic system were computed in the same manner, except that there were no multiple repetitions.

A. Experimental setup

Our VASR system starts by detecting the face and performing an automatic location of the facial geometry (landmark location) using the Supervised Descend Method (SDM) [51]. Once the face is located, the estimated landmarks are used to fix a bounding box around the region (ROI) that is then normalized to a fixed size. Later on, local appearance features are extracted from the ROI based on early fusion of DCT and SIFT descriptors in both spatial and temporal domains. As explained in Section I there are phonemes that share the same visual appearance and should belong to the same class (visemes). Thus, we constructed a phoneme to viseme mapping that groups 32 phonemes into 20 visemes based on an iterative process that computes the confusion matrix and merges at each step the phonemes that show the highest ambiguity until the desired length is achieved. Then, the classification of the extracted features into phonemes is done in two steps. Firstly, multiple LDA classifiers are trained to convert the extracted features into visemes and secondly, at the final step, one-state-per-class HMMs are used to model the dynamic relations of the estimated visemes and produce the final phoneme sequences. This system was shown to produce near state-of-the-art performance for continuous visual speech-reading tasks (more details in [52]).

B. Human lip-reading

As explained in Section I, it is not clear if hearing-impaired people are better lip-readers than normal-hearing people. Fig. 2 (Top) shows the word recognition rates for both groups at each repetition and Fig. 2 (Bottom) shows the word recognition rates for each participant and repetition. Analyzing each participant individually, it is difficult to observe any group-differences between hearing-impaired and

TABLE I
STATISTICAL COMPARISON BETWEEN HEARING-IMPAIRED AND
NORMAL-HEARING PARTICIPANTS FOR EACH REPETITION.

Attempt	Wilcoxon signed rank	Unpaired two-sample
1	$p = 0.116$	$p = 0.094$
2	$p = 0.094$	$p = 0.088$
3	$p = 0.041$	$p = 0.037$

normal-hearing participants. However, we do observe large performance variations within each of the groups, i.e. there are very good and quite poor lip-readers regardless of their hearing condition.

On the other hand, looking at the results globally, split only by group (Fig. 2 (*Top*)), they suggest that hearing-impaired participants outperform normal-hearing participants in the lip-reading task for all three repetitions. However, the results differ about 20% in terms of word recognition rate and thus we need to study if this difference is statistically significant.

To do so, we estimated the word accuracy of each participant as the average accuracy across the 25 sentences that he/she had to lip-read. Then, we performed statistical tests to determine if there were significant differences between the 9 hearing-impaired samples and the 15 normal-hearing samples. Because we only want to test if the hearing-impaired participants were better than normal-hearing participants, we performed single-tailed tests where the null hypothesis was that the mean or median (depending on the test) performance of hearing-impaired participants was not higher than the performance of normal-hearing participants. We ran two tests (summarized in Table I) for each of the 3 repetitions: *Wilcoxon signed rank test* and *Unpaired two-sample t-test*. Taking the conventional significance threshold of $p < 0.05$ it could be argued that at the third repetition the performance of hearing-impaired participants was significantly better than that of normal-hearing participants. However, this was not observed in the first two repetitions. Moreover, the 9 hearing-impaired subjects did better than the 15 normal-hearing, but taking into account that the sample size is relatively small, current trends in statistical analysis suggest that the obtained p-values are not small enough to claim that this would extrapolate to the general population. On the other hand, looking at the p-values, with the current number of subjects we are not far from reaching significance [53].

In Fig. 2 we also show the influence of repetitions into the final performance: as the number of repetitions increases the recognition rate increases too. This effect can be seen split by group and analysing each participant separately.

C. Training and context influence on lip-reading

The context is one of the human resources more used in lip-reading to complete the spoken message. To analyse the influence of the context, the participants were asked to read four different types of sentences, in terms of number of words (explained in Section III). Thus, as the level increases, sentences are longer and the context increases too.

In Fig. 3 we can observe how the first level has the lowest word recognition rates for all repetitions, while the last




Fig. 3. Average word recognition rate for each participant at each level.




Fig. 4. Cumulative average per sentence of all participants in each repetition.

level has the highest rates. There are two factors that could contribute to this effect: 1) Context: humans use the relation between words to try decoding a meaningful message, and 2) Training: as the level increases the participants are more acquainted to the speaker and to the lip-reading task.

The results of Fig. 3 are not enough to determine whether the effect is due to context, training or both. Thus, in Fig. 4 we analyze the variation of performance per sentence (with a cumulative average) instead of per level, which should make clearer the effect of training. This is because training occurs continuously from one sentence to another while context only increases when we change from one level to the next one. Thus, the effect of training can be seen as the constant increase performance in each of the curves (up to 20%). As the users have lip-read more sentences they tend to become better lip-readers. On the other hand, the influence of context is better observed by comparing the different repetitions. In the first attempt, the sentence was completely unknown to the participants, but, in the second and third repetitions there was usually some context available because the message had been already partially decoded, hence constraining the possible words to complete the sentence.

D. Human observers and automatic system comparison

The results of the automatic system are only computed for the first attempt, since it was not designed to benefit from repetitions. The resulting word-recognition rates are shown in Fig. 5 (*Top*). Notice that now the participant number indicates the person that was pronouncing the sentences as the recognition is always performed by the system. Thus, this figure provides information about how well the system was able to lip-read each of the participants. The system produced the highest recognition rates for participants 1, 8, 17 and 21. Interestingly, these participants had good pronunciation and visibility of the tongue and teeth.

We are interested in comparing the performance of humans lip-reading and a VASR system. Focusing on Fig. 6 (*Top*)




Fig. 5. Top: System performance in terms of word recognition rate for each participant. Bottom: Phoneme accuracy recognized by the VASR system for each participant.




Fig. 6. Top: Average word recognition for each participant for human observers (Repetition 1) and for the VASR system; Bottom: Phoneme accuracy recognized by the VASR system for each participant.

we can observe how the word recognition rates are lower for the system in most of the cases. However, we have to take into account that the system does not use the context into the sentence. Indeed, the system is not even targeting words but phonemes, which are later merged to form words. In contrast, people directly search for correlated words with the lip movements of the speaker. Thus, it is reasonable to expect a considerable gap between human and automatic performance, which will be shown to reduce considerably if the comparison is done in terms of phonemes.

In the same figure (Fig. 6) we can observe a direct comparison of the mean recognition rates of each participant identified by humans and by the automatic system. The system gives an unbiased measure about the facility to lip-read participants because it evaluates each of them in the same manner. In contrast, human lip-reading was performed in couples (couples are organized in successive order, e.g. participants 1 and 2, 3 and 4, etc), hence each participant was only lip-read by its corresponding partner. Analyzing Fig. 6 we can identify which users were good lip-readers and also good speakers. For example, participant 7 was lip-read by participant 8 with high word recognition rate.




Fig. 7. Top: Number of wrong detected phonemes. The red columns represent the false negatives phonemes and the green ones the false positives.; Bottom: Precision and Recall of each phoneme.

Then, in the curve corresponding to human performance, we observe a high value for participant 8, meaning that he/she was very successful at lip-reading. When we look at system's performance, however, the value assigned to participant 8 corresponds to the rate obtained by the system and is therefore a measure related to how participant 8 spoke rather than how he/she lip-read. For this specific participant, the figure shows that system performance was also high, hence he/she is a candidate to be good lip-reader and speaker.

The word recognition rates reported by our system are rather low compared to those obtained by human observers. However, as stated earlier, our system is trying to recognize phonemes and convert them to words, so it is also interesting to analyze its performance in terms of phoneme recognition. The phoneme recognition rates obtained by the system are between 40% and 60%, as shown in Fig. 5 (Bottom) and Fig. 6 (Bottom). It is interesting to note that system performance was much more stable across participants than human performance. In addition, in terms of phoneme units, the global mean of the automatic system was 51.25%, very close to the global mean of 52.20% obtained by humans.

There are several factors that help understanding why the system achieves significantly higher rates in terms of phonemes than in terms of words: 1) Phoneme accuracy is computed at frame level because that is the output rate of the system. Thus, the temporal resolution used for phonemes is much higher than that of words and correctly recognizing a word implies the correct match of a rather long sequence of contiguous phonemes. Any phoneme mismatch, even if in a single frame, results in the whole word being wrong. 2) The automatic system finds it easier to recognize concrete phonemes (e.g: vowels) with high appearance rates in terms of frames (vowels are usually longer than consonants). This implies that a high phoneme recognition rate does not necessarily mean that the message is correctly decoded. To analyze this, system performance is displayed in Fig. 7. Specifically, in Fig. 7 (Top) we can observe the number of phonemes that were wrongly detected, distinguishing false negatives (in red color) and false positives (in green),

while Fig. 7 (*Bottom*) shows the corresponding values of precision and recall. Most of the consonants have very high precision, but many samples are not detected, deriving in a low recall. In contrast, vowels have an intermediate precision and recall because they are assigned more times than their actual occurrence. Close inspection of our data suggests that this effect is partially explained by the difficulty in correctly identifying the temporal limits of phonemes.

V. DISCUSSION AND CONCLUSIONS

In this work we explore visual speech reading with the aim to estimate the recognition rates achievable by human observers and by an automatic system under optimal and directly comparable conditions. To this end, we recorded the VLRF database, appropriately designed to be visually informative of the spoken message. For this purpose we recruited 9 hearing-impaired and 15 normal-hearing subjects. Overall, the word recognition rate achieved by the 24 human observers ranged from 44% (when the sentence was pronounced only once) to 73% (when allowing up to 3 repetitions). These results are compatible to those from Duchnowski et al. [31], who stated that even under the most favorable conditions (including repetitions) "speech-readers typically miss more than one third of the words spoken".

We also tested the performance of participants grouped by their hearing condition to compare their lip-reading abilities and verify if these are superior for hearing-impaired subjects, as suggested in some studies. Concretely, we found that hearing-impaired participants outperformed normal-hearing participants on the lip-reading task, but without statistical significance. The performance difference, which averaged 20%, was not sufficient to conclude significance with the current number of subjects. Hence, future work will address the extension of the VLRF database so that it includes sufficient subjects to reach a clearer conclusion.

The participation of hearing-impaired people was very important given their daily experience in lip-reading. During the recording sessions they explained that lip-reading in our database was a challenge because they did not know the context of the sentence beforehand. For them, it is easier to lip-read when they know the context of the conversation. The conversation topic constrains the vocabulary that can appear in the talk. Furthermore, we mentioned before that lip-reading is related to the intelligence and the language knowledge. During the recording sessions we noticed that sentences directly related to daily life were easier to understand than sentences with words not used in colloquial language.

Another important aspect to consider is how easy or difficult is to lip-read different speakers. As explained in Section III, participants were instructed to use their own criterion to facilitate lip-reading. It is difficult to objectively judge the effectiveness of the techniques that were used, but we observed some interesting tendencies during the recordings. Firstly, facial expressions help decoding the spoken message adding context to the sentence (e.g: sad expression if you are speaking about something unfortunate); hearing-impaired

participants used this technique more often than normal-hearing subjects. Secondly, it is more useful to separate clearly between words than to exaggerate pronunciation. That is because the human system is searching words that fit the lip movements. We noticed that when pronunciation was exaggerated the separation between words was not clear or even lost considerably increasing the difficulty of lip-reading.

The above is important when interpreting the results of human observers, as they are conditioned both by the lip-reading abilities of the lip-reader and by the pronunciation abilities of the speaker. Recall that, in our experiments, each participant only lip-read his/her corresponding partner. It would be interesting to separate these factors, which could be done by randomizing the combinations of speakers and lip-readers on a per-sentence basis. In particular, the most interesting aspect would be to estimate the level of difficulty to lip-read each of the speakers, which could be done by having several subjects lip-reading the same speaker. There would be several advantages in doing so: 1) it would allow a more direct comparison to the performance of the system, as speaker performance will not be conditioned to a single human reader; 2) speakers that are too difficult could be excluded from the analysis, at least when seeking for the theoretical limit of lip-reading in optimal conditions; 3) it would help understand which are the best speaking techniques to use to facilitate lip-reading understanding.

As just explained, in our experiments, human observers reached word accuracy of 44% in the first attempt while our visual-only automatic system achieved 20% of word recognition rate. However, if we repeat the comparison in terms of phonemes, the automatic system achieves recognition rates quite similar to human observers, just above 50%. These results are comparable with those reported by Almajai et al. [54] who tested in the RM corpus, using 12 speakers and 6 expert lip-readers. Concretely, their human lip-readers reached 52.63% viseme accuracy (in our case 52.20% phoneme accuracy) and their system obtained 46% viseme accuracy (our system 51.25% phoneme accuracy). Therefore, in terms of viseme/phoneme accuracy, both Almajais and our system reach near-human performance. But this does not happen in terms of word accuracy: Almajai et al. reported human word accuracy of 21% (ours 44%) and system word accuracy of 14% (ours 20%).

When trying to explain the above, we found that the low word recognition rates were related to: 1) the fact that it is quite easy to make mistakes at frame level and a mistake in a single frame results in the whole word being wrong; 2) the imbalance in the occurrence frequencies of phonemes. The latter is especially important because it highlights that the system, while achieving similar phoneme rates to those from humans, does not actually perform equally well. In other words, the phoneme sequences returned by humans always make some sense, which is not generally true for the system as it does not include higher-level constraints (e.g. at the word- or phrase-level). Hence, future directions should focus on introducing constraints related to bigger speech structures such as connected phonemes, syllables or words.

ACKNOWLEDGEMENTS

This work is partly supported by the Spanish Ministry of Economy and Competitiveness under the Ramon y Cajal fellowships and the Maria de Maeztu Units of Excellence Programme (MDM-2015-0502), and the Kristina project funded by the European Union Horizon 2020 research and innovation programme under grant agreement No 645012.

REFERENCES

- [1] H. McGurk and J. MacDonald, "Hearing lips and seeing voices," *Nature*, vol. 264, pp. 746–748, 1976.
- [2] G. Potamianos, C. Neti *et al.*, "Recent advances in the automatic recognition of audiovisual speech," *P IEEE*, vol. 91, pp. 1306–1326, 2003.
- [3] S. Hilder, R. Harvey, and B.-J. Theobald, "Comparison of human and machine-based lip-reading," in *AVSP*, 2009, pp. 86–89.
- [4] J. J. Williams, J. C. Rutledge *et al.*, "Frame rate and viseme analysis for multimedia applications to assist speechreading," *J VLSI Signal Process Syst Signal Image Video Technol*, vol. 20, pp. 7–23, 1998.
- [5] A. Chişu and L. J. Rothkrantz, "Automatic visual speech recognition," *Speech enhancement, modeling and recognition*, p. 95, 2012.
- [6] N. P. Erber, "Auditory-visual perception of speech," *J Speech Hear Disord*, vol. 40, pp. 481–492, 1975.
- [7] W. H. Sumby and I. Pollack, "Visual contribution to speech intelligibility in noise," *J Acoust Soc Am*, vol. 26, pp. 212–215, 1954.
- [8] R. E. Ronquest, S. V. Levi, and D. B. Pisoni, "Language identification from visual-only speech signals," *Atten Percept Psychophys*, vol. 72, pp. 1601–1613, 2010.
- [9] E. Antonakos, A. Roussos, and S. Zafeiriou, "A survey on mouth modeling and analysis for sign language recognition," in *FG*, 2015, pp. 1–7.
- [10] R. Seymour, D. Stewart, and J. Ming, "Comparison of image transform-based features for visual speech recognition in clean and corrupted videos," *Eurasip J Image Vide*, p. 14, 2008.
- [11] S. Dupont and J. Luetin, "Audio-visual speech modeling for continuous speech recognition," *IEEE T Multimedia*, vol. 2, pp. 141–151, 2000.
- [12] A. V. Nefian, L. Liang *et al.*, "A coupled hmm for audio-visual speech recognition," in *ICASSP*, 2002, pp. II–2013.
- [13] Z. Zhou, G. Zhao *et al.*, "A review of recent advances in visual speech decoding," *Image Vis Comput*, vol. 32, pp. 590–605, 2014.
- [14] W. C. Yau, D. K. Kumar, and H. Weghorn, "Visual speech recognition using motion features and hidden markov models," in *CAIP*, 2007, pp. 832–839.
- [15] C. Sui, M. Bennamoun, and R. Togneri, "Listening with your eyes: Towards a practical visual speech recognition system using deep boltzmann machines," in *ICCV*, 2015, pp. 154–162.
- [16] J. S. Chung, A. Senior *et al.*, "Lip reading sentences in the wild," *arXiv preprint arXiv:1611.05358*, 2016.
- [17] S. Petridis and M. Pantic, "Deep complementary bottleneck features for visual speech recognition," in *ICASSP*, 2016, pp. 2304–2308.
- [18] Z. Zhou, X. Hong *et al.*, "A compact representation of visual speech data using latent variables," *IEEE Trans Pattern Anal Mach Intell*, vol. 36, pp. 1–1, 2014.
- [19] K. L. Moll and R. G. Daniloff, "Investigation of the timing of velar movements during speech," *J Acoust Soc Am*, vol. 50, pp. 678–684, 1971.
- [20] J. N. Buchan, M. Paré, and K. G. Munhall, "Spatial statistics of gaze fixations during dynamic face processing," *Soc Neurosci*, vol. 2, pp. 1–13, 2007.
- [21] I. d. I. R. R. Ortiz, "Lipreading in the prelingually deaf: what makes a skilled speechreader?" *Span J Psychol*, vol. 11, pp. 488–502, 2008.
- [22] G. Potamianos and C. Neti, "Automatic speechreading of impaired speech," in *AVSP*, 2001.
- [23] L. E. Bernstein, M. E. Demorest, and P. E. Tucker, *What makes a good speechreader? First you have to find one*. Hove, United Kingdom: Psychology Press Ltd. Publishers, 1998.
- [24] C. M. Capek, M. MacSweeney *et al.*, "Cortical circuits for silent speechreading in deaf and hearing people," *Neuropsychologia*, vol. 46, pp. 1233–1241, 2008.
- [25] T. Ellis, M. MacSweeney *et al.*, "Tas: A new test of adult speechreading-deaf people really can be better speechreaders," in *AVSP*, 2001.
- [26] B. Lyxell and I. Holmberg, "Visual speechreading and cognitive performance in hearing-impaired and normal hearing children (11-14 years)," *Br J Educ Psychol*, vol. 70, pp. 505–518, 2000.
- [27] I. R. Rodríguez-Ortiz, D. Saldaña, and F. J. Moreno-Perez, "How speechreading contributes to reading in a transparent orthography: the case of spanish deaf people," *J Res Read*, 2015.
- [28] F. E. Kyle, R. Campbell *et al.*, "Speechreading development in deaf and hearing children: introducing the test of child speechreading," *J Speech Lang Hear Res*, vol. 56, pp. 416–426, 2013.
- [29] T. Mohammed, R. Campbell *et al.*, "Speechreading and its association with reading among deaf, hearing and dyslexic individuals," *Clin Linguist Phon*, vol. 20, pp. 621–630, 2006.
- [30] F. E. Kyle and M. Harris, "Concurrent correlates and predictors of reading and spelling achievement in deaf and hearing school children," *J Deaf Stud Deaf Educ*, vol. 11, pp. 273–288, 2006.
- [31] P. Duchnowski, D. S. Lum *et al.*, "Development of speechreading supplements based on automatic speech recognition," *IEEE T Bio-Med Eng*, vol. 47, pp. 487–496, 2000.
- [32] G. Potamianos, C. Neti *et al.*, "Audio-visual automatic speech recognition: An overview," *AVSP*, vol. 22, p. 23, 2004.
- [33] I. Matthews, T. F. Cootes *et al.*, "Extraction of visual features for lipreading," *IEEE Trans Pattern Anal Mach Intell*, vol. 24, pp. 198–213, 2002.
- [34] S. J. Cox, R. Harvey *et al.*, "The challenge of multispeaker lip-reading," in *AVSP*, 2008, pp. 179–184.
- [35] B. Lee, M. Hasegawa-Johnson *et al.*, "Avicar: audio-visual speech corpus in a car environment," in *Interspeech*, 2004.
- [36] K. Messer, J. Matas *et al.*, "Xm2vtsdb: The extended m2vts database," in *AVBPA*, 1999, pp. 965–966.
- [37] E. K. Patterson, S. Gurbuz *et al.*, "Cuave: A new audio-visual database for multimodal human-computer interface research," in *ICASSP*, 2002, pp. II–2017.
- [38] J. Huang, G. Potamianos *et al.*, "Audio-visual speech recognition using an infrared headset," *Speech Commun*, vol. 44, pp. 83–96, 2004.
- [39] P. J. Lucey, G. Potamianos, and S. Sridharan, "Patch-based analysis of visual speech from multiple views," *AVSP*, 2008.
- [40] C. Sanderson, "The vidtimit database," IDIAP, Tech. Rep., 2002.
- [41] M. Cooke, J. Barker *et al.*, "An audio-visual corpus for speech perception and automatic speech recognition," *J Acoust Soc Am*, vol. 120, pp. 2421–2424, 2006.
- [42] C. McCool, S. Marcel *et al.*, "Bi-modal person recognition on a mobile phone: using mobile phone data," in *ICMEW*, 2012, pp. 635–640.
- [43] G. Zhao, M. Barnard, and M. Pietikainen, "Lipreading with local spatiotemporal descriptors," *IEEE T Multimedia*, vol. 11, pp. 1254–1265, 2009.
- [44] I. Anina, Z. Zhou *et al.*, "Ouluvs2: A multi-view audiovisual database for non-rigid mouth motion analysis," in *FG*, 2015, pp. 1–5.
- [45] A. Ortega, F. Sukno *et al.*, "Av@car: A spanish multichannel multimodal corpus for in-vehicle automatic audio-visual speech recognition," in *LREC*, 2004.
- [46] T. J. Hazen, K. Saenko *et al.*, "A segment-based audio-visual speech recognizer: Data collection, development, and initial experiments," in *ICMI*. ACM, 2004, pp. 235–242.
- [47] R. Bowden, "LILiR language independent lip reading," <http://www.ee.surrey.ac.uk/Projects/LILiR/datasets.html>, 2010, accessed: 2016-08-16.
- [48] "ACCAPS federació d'associacions catalanes de pares i persones sordes," <http://www.acapps.org/web/>, accessed: 2016-08-16.
- [49] P. Boersma *et al.*, "Praat, a system for doing phonetics by computer," *Glott international*, vol. 5, pp. 341–345, 2002.
- [50] J. C. Wells *et al.*, "Sampa computer readable phonetic alphabet," *Handbook of standards and resources for spoken language systems*, vol. 4, 1997.
- [51] X. Xiong and F. De la Torre, "Supervised descent method and its applications to face alignment," in *CVPR*, 2013, pp. 532–539.
- [52] A. Fernandez-Lopez and F. M. Sukno, "Automatic viseme vocabulary construction to enhance continuous lip-reading," *VISAPP*, 2017.
- [53] D. Colquhoun, "An investigation of the false discovery rate and the misinterpretation of p-values," *Open Science*, vol. 1, p. 140216, 2014.
- [54] I. Almajai, S. Cox *et al.*, "Improved speaker independent lip reading using speaker adaptive training and deep neural networks," in *ICASSP*, 2016, pp. 2722–2726.