

Some observations on computer lip-reading: moving from the dream to the reality

Helen L. Bear ^a, Gari Owen^b, Richard Harvey^a, and Barry-John Theobald^a

^aUniversity of East Anglia, Norwich, NR4 7TJ, UK.

^bAnnwvyn Solutions, Bromley, Kent, BR1 3DW, UK.

ABSTRACT

In the quest for greater computer lip-reading performance there are a number of tacit assumptions which are either present in the datasets (high resolution for example) or in the methods (recognition of spoken visual units called “visemes” for example). Here we review these and other assumptions and show the surprising result that computer lip-reading is not heavily constrained by video resolution, pose, lighting and other practical factors. However, the working assumption that visemes, which are the visual equivalent of phonemes, are the best unit for recognition does need further examination. We conclude that visemes, which were defined over a century ago, are unlikely to be optimal for a modern computer lip-reading system.

Keywords: Lip-reading, speech recognition, pattern recognition

1. BACKGROUND

There has been consistent and sustained interest in building computer systems that can understand what humans are saying without hearing the audio channel.¹⁻⁴ There are obvious applications for such systems in security but also in noisy environments such as cockpits, battlefields and crowds where audio recognition is likely to be impossible or highly degraded. Early work consisted of very small vocabularies (often fewer than 10 words),⁵ single speakers, high-definition video (often the camera would be zoomed into the lip region or the frame rate would be greater than 60 fields per second)⁶ and, often, the talker would wear special lipstick to allow easy segmentation and analysis of the lips.⁷ Subsequently, our understanding of the problem has improved such that lip-reading in outdoor conditions (which requires very robust lip-tracking) and with 1000-voclabularies (which requires good machine learning) looks feasible. The problem of speaker dependence is still only partially solved.⁸ One surprising recent result was a characterisation of the effect of resolution on lip-reading. An informal understanding was that relatively high resolution was required (at least a couple of hundred pixels to span the lips). In practice, it was reported in⁹ that, provided the tracking was perfect, then fewer than 10 pixels can give acceptable results. A further observation¹⁰ was that off-axis lip-reading gave slightly better performance than full frontal (which is the default for most experiments). It seems, when it comes to lip-reading, one’s intuition might often be wrong – indeed experimenters in the field are often confounded by one of the most counter-intuitive illusions in the field – the McGurk effect.¹¹

Experimental recognition systems for audio are almost always built using phonemes. There appears to be good agreement as to which phonemes appear in the major languages and what their expected frequency might be. Once these phonetic units have been recognised then the sequence (together with their probabilities and next-most probable sequence and so on) is fed into a language model which generates hypotheses for words and sentences. In modern speech recognition language models are powerful and important and have been the subject of decades of work. There is clearly a huge advantage in a lip-reading system re-using the language model so many lip-reading systems recognise using the visual units, visemes, and then feed the sequence into an acoustic language model modified to cope with visemes. If visemes exist in the form postulated by linguists e.g.,^{12,13} then there are many choices of visemes. However there has been surprisingly few examinations of which visemes give the best performance or how fragile that performance is compared to phonetic recognition.

Further author information: (Send correspondence to RWH)
RWH: E-mail: r.w.harvey@uea.ac.uk

2. INTRODUCTION

A phoneme is generally regarded as the smallest sound which can be uttered.¹⁴ A viseme, which is often said to be the visual equivalent of a phoneme, is not so precisely defined^{15–17} so we use the working definition: ‘a viseme is a set of phonemes that have identical appearance on the lips’. Therefore any phoneme falls into one viseme class but a viseme may represent many phonemes: a many to one mapping.

A typical lip-reading system is a sequence of tasks as in Figure 1 and our work is focused within the recognition step.




Figure 1. Steps in a typical lip-reading system

Similar to a simplified version of audio recognition whereby we seek to identify a string of unique phonemes, each recognizer is based upon training data of the correctly labelled phoneme. In visual-only recognition we use the same concept of building recognizers based upon visual-only training samples correctly labelled according to a viseme mapping. There is still debate over what the *correct* phoneme-to-viseme mapping is and many have been suggested, e.g.^{16, 18–20} but our interest is in the contribution of each viseme to the recognition performance. We look for any particular visemes (or combinations of phonemes) that contribute more to the recognition accuracy.

We aim to measure the reduction of each unique visemic recogniser in contribution value to the whole task of accurate recognition in continuous speech. To demonstrate the influence of reduced recogniser classes in visual speech recognition we compare the outputs with those of audio recognition of the same data. For a fair comparison we use the same groupings of phonemes into faux ‘audio-viseme’ recognisers on the audio data. Audio recognition has a higher quantity of classifiers (phonemes) than proposed viseme classes, therefore we hypothesise visual classes have bigger variance in use/purpose towards the whole recognition task. We anticipate, fewer visemes will be used in visual speech recognition than ‘audio-visemes’ in audio recognition.

3. DATASET AND FEATURE EXTRACTION

For the first two steps in Figure 1 we use full face Active Appearance Models (AAMs)²¹ to track the faces through the videos, and lip-only AAMs (one for shape and another for appearance) and using the methods of²² we produce two sets of talker-dependent features; shape-only visual features and appearance-only visual features.

Table 1. Frame images from each video.

Video	Num. of AAM train images	Video Length (frames)	Duration
Talker1 - 1	10	21,658	00:06:01
Talker1 - 2	10	21,713	00:06:02
Talker2 - 1	11	31,868	00:08:52
Talker2 - 2	11	33,338	00:09:17

Shape features (1) are based solely upon the lip shape and positioning during the duration of the talker speaking e.g. the landmarks in Figure 2. The landmark positions can be compactly represented using a linear model of the form:

$$s = s_0 + \sum_{i=1}^m s_i p_i \quad (1)$$

where s_0 is the mean shape and s_i are the modes. The appearance features are computed over pixels, the original images having been warped to the mean shape. So $A_0(x)$ is the mean appearance and appearance is described as a sum over modal appearances:

$$A(x) = A_0(x) + \sum_{i=1}^l \lambda_i A_i(x) \quad \forall x \in S_0 \quad (2)$$




Figure 2. Showing full face shape landmarks for talker T1 (a) and a lip shape landmarks for talker T1 (b).

The Rosetta Raven data is four videos of recitations of Edgar Allen Poe’s poem ‘The Raven’. There are two talkers, one male, one female. Neither are trained actors and they do not recite the poem with the intended trochaic octameter.²³ The videos were recorded at 1440×1080 resolution (non-interlaced) at 60 frames per second. Table 1 summarises the video data.

A set of images are extracted from each video (one image per frame) via ffmpeg using image2 encoding at full high-definition resolution (1440×1080). To construct an initial AAM we select the first frame and nine or ten others randomly. These *training frames* are hand-labelled with a shape model of a face and lips to build a preliminary model for each talker. These models are then fitted, via inverse compositional fitting²² to the remaining frames (Table 1). Thus we get tracked and fitted full-face talker-dependent AAMs (Figure 2 left) on full resolution lossless PNG frame images (Figure 1 step 1).

Next we create a sub-model of only the lips for each talker by decomposing the two full face models (Figure 2 right). From the fitted landmarks, the shape and appearance parameters for each frame are extracted. For talker1 (T1), we retain 6 shape and 14 appearance parameters and for talker2 (T2), 7 shape and 14 appearance parameters. We restrict the feature parameters to retain 95% of variation from the mean AAM model produced using the whole tracked video data.²¹ (Figure 1 step 2.)

We did not implement $\Delta\Delta$ ’s into our extracted features to address co-articulation because we used a phonetic-alignment in the production of our ground-truth benchmark and forced-alignment within the training process of our HMM recognizers.

Table 2. Phone to viseme mapping.

vID	Phones	vID	Phones
v01	/p/ /b/ /m/	v10	/i/ /ih/
v02	/f/ /v/	v11	/eh/ /ae/ /ey/ /ay/
v03	/th/ /dh/	v12	/aa/ /ao/ /ah/
v04	/t/ /d/ /n/ /k/ /g/ /h/ /j/ /ng/ /y/	v13	/uh/ /er/ /ax/
v05	/s/ /z/	v14	/u/ /uw/
v06	/l/	v15	/oy/
v07	/r/	v16	/iy/ /hh/
v08	/sh/ /zh/ /ch/ /jh/	v17	/aw/ /ow/
v09	/w/	v18	silence




Figure 3. Viseme counts for both talker transcripts

To have a benchmark for measuring our recognition outputs we produce a ground-truth viseme transcription using the Carnegie Mellon University (CMU) North American pronunciation dictionary,²⁴ and a word transcription. We convert a phonetic transcript to a viseme transcript assuming 15 visemes, listed in Table 2 which is a combination of Montgomery *et al's* vowel mapping and Walden's consonant mapping.^{25,26}

The limited availability of large datasets is documented² so we work within the restrictions of short datasets. Here we note these may not provide adequate training examples of all visemes. Where this happens, we group the untrainable visemes into a single garbage viseme. In this case we select a 150 sample threshold so visemes /v08/, /v09/, /v14/ and /v15/ are grouped. Figure 3 shows the occurrence of visemes listed in Table 3 in our data and Table 4 shows our revised viseme mapping.

Table 3. Phone to viseme mapping modified to accomodate restrictions in dataset.

vID	Phones	vID	Phones
v01	/p/ /b/ /m/	v11	/eh/ /ae/ /ey/ /ay/
v02	/f/ /v/	v12	/aa/ /ao/ /ah/
v03	/th/ /dh/	v13	/uh/ /er/ /ax/
v04	/t/ /d/ /n/ /k/ /g/ /h/ /j/ /ng/ /y/	v16	/iy/ /hh/
v05	/s/ /z/	v17	/aw/ /ow/
v06	/l/	v18	silence
v07	/r/	garb	/u/ /uw/ /oy/ /w/ /sh/ /zh/ /ch/ /jh/
v10	/i/ /ih/		

For each talker, a test fold is randomly selected as 42 of the 108 lines in the poem with replacement. The remaining lines are used as training folds. Repeating this five times gives five-fold cross-validation. Note that visemes cannot be equally represented in all folds.

For recognition we use Hidden Markov Models (HMMs) implemented in the Hidden Markov Toolkit (HTK).²⁷ An HMM is initialised using the ‘flat start’ method using a prototype of five states and five mixture components and the information in the training samples. We choose five states and five mixture components based upon.²⁸ We define an HMM for each viseme plus silence and short-pause labels (Table 3) and re-estimate the parameters four times with no pruning.

Next, we use the HTK tool `HHed` to tie together the short-pause and silence models between states two and three before re-estimating the HMMs a further two times. Then `HVite` is used to force-align the data using the word transcript*. The HMMs are now re-estimated twice more, however now we use the force-aligned viseme transcript rather than the original viseme transcript used in the previous HMM re-estimations.

To complete recognition using our HMMs we require a word network as we have a continuous speech dataset. We use `HLStats` and `HBuild` to make a Bi-gram Word-level Network (BWN). Finally `HVite` is used with the network support for the recognition task and `HResults` gives us both correctness and accuracy viseme recognition values and a viseme confusion matrix for all folds. We have provided the reader with technical details to enable repeatability of our experiments. Please contact the author for original videos.

4. RESULTS

We have extracted figures from the `HResults` confusion matrices for analysis. For each viseme we have calculated the inverse probability of its recognition $\Pr\{v|\hat{v}\}$.

Figure 4 shows the probability of correct recognition using shape-only features (mean and ± 1 standard error) plotted against the probability of correct recognition using appearance-only features for each viseme. As usual some talkers are better recognised with shape and some with appearance^{1†}. Note that the top right-hand point is the visual silence phoneme. In general, visual silence can be quite variable compared to audio silence because talkers breathe and show emotion. However here, because the source text is a poem, there are well-defined visual silence periods at the start of each line.

Table 4. Ranked mean viseme recognition for Shape, Appearance, Talker 1 and Talker 2.

Feature	Viseme order
Shape	/v18/ {/v04/, /v12/} /v11/ /v01/ /v07/ /v05/ {/v02/ /v06/ /garb/} /v10/ {/v03/ /v13/} /v16/ /v17/
Appearance	/v18/ /v04/ /v12/ /v11/ /v01/ /v07/ {/v02/, /v05/} /v06/ /garb/ /v10/ /v03/ {/v13, /v16/} /v17/
Talker 1	/v18/ {/v04/, /v12/} /v11/ /v01/ /v07/ {/v02/, /v05/} /v06/ {v10/, /garb/} /v03/ /v13/ /v16/ /v17/
Talker 2	/v18/ {/v04/, /v12/} /v11/ /v01/ /v07/ {/v02/, /v05/} {/v06/, /garb/} /v10/ /v03/ /v13/ /v16/ /v17/
Overall	/v18/ /v04/ /v12/ /v01/ /v11/ /v07/ {/v02/, /v05/} /v19/ /v06/ /v10/ /v13/ /v03/ /v16/ /v17/

Figures 5 and 6 show, for the T1 and T2 shape and appearance models, the probability of correctly recognising the top ten visemes, $\Pr\{v|\hat{v}\}$. They also show, the audio performance measured on visemes. The x -axis varies by performance; the best performing viseme is on the left hand side which for visual shape and appearance features is silence for all features.

*We use the `-m` flag with `HVite` with the manual creation of a viseme version of the CMU dictionary for word to viseme mapping so that the force-alignment produced uses the break points of the words.

[†]The conventional wisdom is that appearance features give the best results but only in studio-type conditions with good tracking.




Figure 4. Relationship between Shape and Appearance model features for both talkers.

It has been observed in human lip reading there are few visual cues that are reliable and humans use these combined with rich contextual information to interpret or ‘fill in the gaps’ of what a talker is saying.^{29,30} Therefore our hypothesis is that robust audio recognition is based upon a large spread of recognised phones and the resilience in recognition is due to the number of phones contributing to the accuracy. Visually, as with human lip-readers, it is anticipated that fewer visemes would perform the equivalent recognition and, as such, the graph would demonstrate a steeper performance decline over the top performing visemes.

In Figure 5 we do see a greater decline from left to right over the top ten visemes for visual features than for audio for both talkers. We also note that the error bars after the 5th position viseme increase, which is consistent with our hypothesis that audio recognition is spread over more visemes to be correct. The top visemes (after silence /v18/) are /v04/, /v12/, /v11/ and /v01/. These are vowels (/v12/, /v11/) and front-of-mouth consonant visemes (/v04/, /v01/).

Figure 6 demonstrates a shallower decline from left to right than the shape graph in Figure 5 but still there is a greater decline for visual features than for audio. The error bars here increase after the 7th position viseme[‡]. The shape of the graph in Figure 6 is similar between audio and video which implies that appearance-based recognition is similar to noisy acoustic recognition for both talkers and hence is less fragile. The top visemes in Figure 6 (not including silence /v18/) are: /v04/ /v12/ /v11/ /v01/ /v7/ i.e. identical for shape-only in the first six positions.

Where the error bars increase, we consider this may be due to the small data available, which makes recognition more unreliable due to less well trained HMM classifiers. We have reduced the impact of this with the

[‡]Note that the order of the audio viseme ordering is identical in both Figures 5 and 6 as this is the same experiment.




Figure 5. Top ten viseme recognition probability in descending order with a shape model.

/garb/ viseme but note with Figure 3 there are similarities between our top performing visemes and those with the most training samples.

Table 4 is the visemes ordered by correctness showing, for example that viseme 18 /v18/ is the best performing viseme overall. It is natural to ask if the differences in ranking are significant. To compare the viseme ordering we compute the Spearman rank correlation coefficient, r . The results are shown in Tables 5 and 6. Also shown is the p -value for the null that $r = 0$ (randomly ordered). Those that are significant at the 5% threshold are underlined. Talker 2 has poor audio performance which tends to degrade the audio correlation. Lip-reading does not depend on audio though so these results confirm the strong relation between shape-only and viseme-only classification. Also note for T1 (Figure 6) the audio ranking is similar to the video ranking although as we have previously noticed there is a more rapid drop-off for video.

In Table 7 we have provided the overall mean and Standard Error Accuracy score for the whole viseme set recognition performance over all five folds. Talker 2 outperforms Talker 1 with all features but for visual features also has a larger degree of error. Appearance features outperform shape for both talkers and audio outperforms appearance for both talkers. As we have seen in Figures 5 and 6 this recognition is based upon a larger spread of visemes than the shape models with the audio having the largest spread of visemes and hence being the most

Table 5. Spearman rank correlation, r and p -value for visemes ranked by performance for Talker 1 and Talker 2

Talker 1	Talker 2	r	p
Audio	Audio	<u>0.43</u>	1.63×10^{-2}
Shape	Shape	<u>0.92</u>	0.00
Appearance	Appearance	<u>0.93</u>	0.00




Figure 6. Top ten viseme recognition probability in descending order with an appearance model.

Table 6. Spearman rank correlation, r and p -values for visemes ordered by feature for Talker 1 (left) and Talker 2 (right)

Talker 1	Talker 1	r	p	Talker 2	Talker 2	r	p
Shape	Appearance	<u>0.90</u>	0.00	Shape	Appearance	<u>0.92</u>	0.00
Audio	Shape	<u>0.85</u>	0.00	Audio	Shape	0.42	0.12
Audio	Appearance	<u>0.74</u>	0.00	Audio	Appearance	0.48	0.07

robust recognition mode.

Table 7. Mean accuracy scores of each feature type by talker

Feature type	Mean	Standard error
T1 Audio	45.6380	2.0086
T2 Audio	75.8780	1.7839
T1 Shape	21.7360	0.7501
T1 Appearance	38.9860	0.4637
T2 Shape	32.1360	1.0339
T2 Appearance	52.9540	1.9996

5. CONCLUSIONS

Our principal observations are:

- Assuming there is enough data to properly train classifiers, then the performance ordering of the visemes is relatively stable across modes of recognition (audio, shape and appearance).

- That said, the visual classifiers are far more dependent on the good performance of a few visemes than the audio.
- Of the video classifiers, shape is the most dependent on a few visemes.

These are important results because they illuminate the often made observation that lip-reading is fragile. In other words if one cannot build classifiers for a few critical visemes then lip-reading is impossible. In a human lip-reading context, humans are often trained to recognise a small number of critical gestures which are then processed via a very sophisticated language and context model to create a transcript.

In audio is it surprisingly rare to see this effect measured even though a good acoustic unit will have accuracies that are at least 10% higher than an average unit (the mean audio viseme performance on T2 is 76% for the whole viseme set).

It is important to acknowledge that most work in this field focuses on improving mean accuracies over the set of all visemes which can conceal the real source of overall performance. A system that achieves a mean viseme accuracy of, say 53% maybe one that contains one supremely accurate viseme classifier or it maybe a system that has a set of classifiers of much more modest performance.

This paper therefore raises two different tactics for improving lip-reading systems. Either one makes the best viseme classifiers better or, one focuses upon improving the worst. At this stage we do not know which tactic is likely to be more successful but we hope this methodology allows future work to focus attention where it is likely to do the most good.

REFERENCES

- [1] Bowden, R., Cox, S., Harvey, R., Lan, Y., Ong, E.-J., Owen, G., and Theobald, B.-J., “Recent developments in automated lip-reading,” in [*SPIE Security+ Defence*], 89010J–89010J, International Society for Optics and Photonics (2013).
- [2] Cappelletta, L. and Harte, N., “Phoneme-to-viseme mapping for visual speech recognition.,” in [*ICPRAM (2)*], 322–329 (2012).
- [3] Davis, S. and Mermelstein, P., “Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences,” *Acoustics, Speech and Signal Processing, IEEE Transactions on* **28**(4), 357–366 (1980).
- [4] Bowden, R., Cox, S. J., Harvey, R. W., Lan, Y., Ong, E.-J., Owen, G., and Theobald, B., “Is automated conversion of video to text a reality?,” in [*Optics and Photonics for Counterterrorism, Crime Fighting, and Defence VIII*], Lewis, C. and Burgess, D., eds., **SPIE 8546**, 85460U–85460U–9, SPIE (2012).
- [5] Petajan, E. D., *Automatic Lipreading to Enhance Speech Recognition*, PhD thesis, University of Illinois, Urbana-Champaign (1984).
- [6] Brooke, N. M. and Summerfield, Q., “Analysis, synthesis and perception of visible articulatory movements,” *Journal of Phonetics* **11**, 63–76 (1983).
- [7] Kaucic, R. and Blake, A., “Accurate, real-time, unadorned lip tracking,” in [*Computer Vision, 1998. Sixth International Conference on*], 370–375, IEEE (1998).
- [8] Z., Z., X., H., and M., Z. G. . P., “A compact representation of visual speech data using latent variables,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **36**(1), 181–187 (2014).
- [9] Bear, H., Harvey, R. W., Theobald, B.-J., and Lan, Y., “Resolution limits on computer lip-reading,” in [*IEEE International Conference on Image Processing*], (2014).
- [10] Lan, Y., Theobald, B.-J., and Harvey, R., “View independent computer lip-reading,” in [*Multimedia and Expo (ICME), 2012 IEEE International Conference on*], 432–437, IEEE (2012).
- [11] McGurk, H. and MacDonald, J., “Hearing lips and seeing voices,” *Nature* **264**, 746–748 (1976).
- [12] Jeffers, J. and Barley, M., [*Speechreading (lipreading)*], Thomas Springfield, IL: (1971).
- [13] Bozkurt, E., Erdem, C., Erzin, E., Erdem, T., and Ozkan, M., “Comparison of phoneme and viseme based acoustic units for speech driven realistic lip animation,” in [*3DTV Conference*], 1–4, IEEE (May 2007).

- [14] Association, I. P., [*Handbook of the International Phonetic Association: A guide to the use of the International Phonetic Alphabet*], Cambridge University Press (1999).
- [15] Chen, T. and Rao, R. R., “Audio-visual integration in multimodal communication,” *Proceedings of the IEEE* **86**(5), 837–852 (1998).
- [16] Fisher, C. G., “Confusions among visually perceived consonants,” *Journal of Speech, Language and Hearing Research* **11**(4), 796 (1968).
- [17] Hazen, T. J., Saenko, K., La, C.-H., and Glass, J. R., “A segment-based audio-visual speech recognizer: Data collection, development, and initial experiments,” in [*Proceedings of the 6th International Conference on Multimodal Interfaces*], *ICMI '04*, 235–242, ACM, New York, NY, USA (2004).
- [18] Binnie, C. A., Jackson, P. L., and Montgomery, A. A., “Visual intelligibility of consonants: A lipreading screening test with implications for aural rehabilitation,” *Journal of Speech and Hearing Disorders* **41**(4), 530 (1976).
- [19] Kricos, P. B. and Lesner, S. A., “Differences in visual intelligibility across talkers.,” *The Volta Review* **84**, 219–226 (1982).
- [20] Nitchie, E. B., [*Lip-Reading, principles and practise: A handbook for teaching and self-practise*], Frederick A Stokes Co, New York (1912).
- [21] Cootes, T., Edwards, G., and Taylor, C., “Active appearance models,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **23**, 681–685 (Jun 2001).
- [22] Matthews, I. and Baker, S., “Active appearance models revisited,” *International Journal of Computer Vision* **60**(2), 135–164 (2004).
- [23] Quinn, P. F., “The critical mind of Edgar Poe: Claude Richard. Edgar Allan Poe: Journaliste et critique.,” *Poe Studies-Old Series* **13**(2), 37–40 (1980).
- [24] Carnegie Mellon University, “CMU pronunciation dictionary,” (2008).
- [25] Massaro, D., [*Perceiving Talking Faces*], The MIT Press (1998).
- [26] Walden, B. E., Prosek, R. A., Montgomery, A. A., Scherr, C. K., and Jones, C. J., “Effects of training on the visual recognition of consonants,” *Journal of Speech, Language and Hearing Research* **20**(1), 130 (1977).
- [27] Young, S., Evermann, G., Gales, M., Hain, T., Kershaw, D., Liu, X. A., Moore, G., Odell, J., Ollason, D., Povey, D., Valtchec, V., and Woodland, P., [*The HTK Book (for HTK Version 3.4)*], Cambridge University Engineering Department (2006).
- [28] Matthews, I., Cootes, T., Bangham, J., Cox, S., and Harvey, R., “Extraction of visual features for lipreading,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **24**, 198–213 (feb 2002).
- [29] Erber, N. P., “Auditory-visual perception of speech,” *Journal of Speech and Hearing Disorders* **40**(4), 481 (1975).
- [30] Stork, D. G. and Hennecke, M. E., [*Speechreading by humans and machines: models, systems, and applications*], vol. 150, Springer (1996).