

A speech-controlled environmental control system for people with severe dysarthria

Mark S. Hawley¹, Pam Enderby², Phil Green³, Stuart Cunningham^{1&4}, Simon Brownsell¹,
James Carmichael³, Mark Parker², Athanassios Hatzis³, Peter O'Neill¹, Rebecca Palmer^{1&2}

*1. Department of Medical Physics and Clinical Engineering, Barnsley District General
Hospital, UK,*

2. Institute of General Practice and Primary Care, University of Sheffield, UK,

3. Department of Computer Science, University of Sheffield, UK

4. Department of Human Communication Science, University of Sheffield, UK

***Abstract*—Automatic speech recognition (ASR) can provide a rapid means of controlling electronic assistive technology. Off-the-shelf ASR systems function poorly for users with severe dysarthria because of the increased variability of their articulations. We have developed a limited vocabulary speaker dependent speech recognition application which has greater tolerance to variability of speech, coupled with a computerised training**

package which assists dysarthric speakers to improve the consistency of their vocalizations and provides more data for recogniser training

We present results of field trials to evaluate the training program and the speech-controlled environmental control system (ECS). The training phase increased the recognition rate from 88.5% to 95.4% ($p < 0.001$). Recognition rates were good for people with even the most severe dysarthria in everyday usage in the home (mean word recognition rate 86.9%). Speech-controlled ECS were less accurate (mean task completion accuracy 78.6% vs 94.8%) but were faster to use than switch-scanning systems, even taking into account the need to repeat unsuccessful operations (mean task completion time 7.7 vs 16.9 seconds, $p < 0.001$). It is concluded that a speech-controlled ECS are a viable alternative to switch-scanning systems for some people with severe dysarthria and would lead, in many cases, to more efficient control of the home.

Index Terms—electronic assistive technology, environmental control system, speech recognition, training

I. INTRODUCTION

A significant proportion of people requiring electronic assistive technology (EAT) have dysarthria, a motor speech disorder, associated with their physical disability. Speech control of EAT is seen as desirable for these people but machine recognition of dysarthric speech is a difficult problem due to the variability of their articulatory output [1]. Large vocabulary adaptive

speech recognition systems have been successfully used for people with mild and moderate dysarthria as a means of inputting text, but these systems are less successful for people with severe dysarthria [2-6]. Speaker independent speech recognition algorithms with the aim of improving the recognition of dysarthric speech patterns have been described [7] but they have not appeared in a widely available form.

Small vocabulary speaker dependent speech recognition for control of assistive technology has been in the literature for more than twenty five years [8,9]. Speaker dependent recognition is arguably more appropriate to severe dysarthria since it allows users to train the system with their own utterances rather than requiring speech that is close to 'normal' [10]. Speaker dependent recognisers can, however, also perform poorly for severely dysarthric speech [3].

There is some evidence to suggest that speech training of dysarthric speakers can improve their ability to accurately use speech recognition [6]. We have therefore taken a two-pronged approach to addressing the problem of reliable use of speech as a control method for people with severe dysarthria. Our aims were:

- To develop a speech recognition system with greater tolerance to variability of speech utterances;
- To develop a computerised training package to assist dysarthric speakers to improve the recognition likelihood and consistency of their vocalisations for a small vocabulary.

This paper briefly describes these two applications. The speech recogniser has been deployed as the control interface for an environmental control system (ECS), which has been tested by

disabled people, who were long-term switch-scanning ECS users. The training, recognition and ECS applications are used as a single package known as STARDUST (for Speech Training and Recognition for Disabled Users of Assistive Technology).

II. SPEECH RECOGNISER

Initially, a small vocabulary recogniser has been specified to allow a limited number of control operations. The recogniser uses isolated words as its recognition units, as there is some evidence that people with severe dysarthria perform better with this type of recogniser than with continuous speech recognisers [10]. Since there is so much variation between individuals, speaker-dependent recognisers are trained for each individual.

The HTK toolkit [11], using Continuous Density Hidden Markov Models [12], has been used for this project. We choose whole words as our modeling units because the phonetic abnormality of severely dysarthric speech prohibits the definition of reliable sub-word units. The models we use are quite standard and take Mel-Frequency Cepstral Coefficients as their acoustic vectors. What is different is the methodology for building recognisers, which is adapted to deal with the scarcity of training data. It is difficult and time-consuming to collect speech samples as our subjects have physical problems that make the production of large amounts of speech on any given occasion tiring. Scarcity of data is problematic [13], since recognition accuracy tends to increase with the size of the training set. In the case of dysarthric speech and its greater variability, the scarce data problem is exacerbated.

We have addressed this problem by closing the loop between recogniser building and user training. We start by training recognisers with a relatively small amount of data for an individual. Speech samples are collected using customised audio recording software and a high-quality microphone, so that utterances are recorded as digitised data onto the computer. These samples are used to prime the user training application. As it is used, the user training program records all examples of utterances used in training and these new speech samples can be used to increase the amount of data in subsequent versions of the recogniser, thus facilitating the collection of larger data sets for each individual.

To further increase the recognition rate, we have introduced statistical confusability measures derived from the recogniser and its training set [13] to identify problematic words that are easily confused by the recogniser. These words can be removed from the vocabulary and replaced by other words that are less easily confused.

III. USER TRAINING

The user training software is based upon the speaker dependent speech recogniser described in section II and runs on a personal computer. The user initially chooses a vocabulary that s/he wishes to train with. In this project the vocabularies were chosen as a command set for the ECS functions required by the user. To set up the program, the recogniser is trained using example utterances for each word in the vocabulary. It was found that for all users of the system, 30 examples of each word was more than sufficient to initially train each recogniser – indeed as few as 10 examples may be sufficient to prime the iterative practice and recogniser-building cycle.

The training data is used to produce hidden Markov models (HMMs) for each word in the vocabulary. A ‘best fit’ utterance for each word is then determined by the software. If one thinks of the HMM as a generative model, this is the utterance in the training set which the model is most likely to produce. The best-fit utterance is not necessarily the most intelligible production of the word, but is the example that best approximates the user’s most likely production. The best fit utterance then becomes the target which the user tries to reproduce: feedback is thus based on something we know the user can achieve, rather than an ideal articulation.

The word to be trained is displayed on the screen (e.g. ‘Lamp’, see figure 1). The user has three options, accessed by using a switch adapted to his/her needs. The user can:

- play the ‘best fit’ example of the word through the computer speakers;
- speak the word;
- or move on to the next word in the vocabulary list.

If the user chooses to speak the displayed word the utterance is recorded and compared to the best-fit example by the recogniser. If the utterance is recognised correctly, the display (see figure 1) then shows two bars. The height of the bar on the right represents the closeness of fit score of the new utterance to the word model. The bar on the left represents the closeness of fit of the ‘best fit’ utterance to the model. The closeness of fit score is derived from the log probability of the model generating the word by the most likely (Viterbi) path [12,13]. The user is thus given an indication of how similar the utterance is to the ‘best fit’ utterance taken from the training set. The user can then carry on practising the word, trying to raise the height of the right hand bar and use the play back facility to hear the best fit example so that s/he has an auditory target to aim for. When the user has completed the attempt, s/he can move on to the

next word.

The aim of the training is three-fold. Firstly, in trying to make each utterance as close as possible to the target (ie to maximise the closeness of fit), the user is increasing the likelihood of his utterances being recognised correctly by the existing recogniser. Secondly, in striving to imitate a stable target, which is a typical example of her/his own speech, the user is expected, through repetition, to reduce the overall variability in the production of these words. This is also expected to have a positive effect on the recognition accuracy. Thirdly, as each utterance is recorded during training, a larger corpus of data is assembled, which is used as training data for new recognisers to improve the robustness and accuracy of recognition.

IV. ENVIRONMENTAL CONTROL SYSTEM

The speech-controlled environmental control system runs on a standard lap-top personal computer connected, via the serial port, to an infra-red remote control unit. The computer is provided with two inputs: a high quality microphone, which can be head-mounted or a remote array; and a switch, adapted to the individual needs of the user. The system requires the user to press the switch to activate the ASR application.

As described in section II, the speech recognizer uses individual words as its recognition units. However, words can be combined into command strings, increasing the range of ECS operations that can be carried out. The speaker says the command phrase, for example, '**TV volume up**' with sufficient pause between the words so that the recogniser does not treat them as one word.

The recogniser identifies the individual words and parses the command, which is then converted to a code and sent via the serial port to the external infra-red sender unit.

The system is also provided with an additional facility, allowing the user to choose to use the switch alone to control the ECS interface. If the switch is held down for a period longer than a pre-set time, the computer displays a scanning interface that can be used to select the desired operation. This facility has been added because it is acknowledged that speech recognition can never be 100% accurate in a home environment. The switch-only operation is important for two reasons: firstly, because users get frustrated with speech recognition if its accuracy is perceived as insufficient. This may be temporary if, for example, the background noise level is too high. Secondly, some ECS equipment is safety critical, such as the need to urgently summon assistance, and a back-up selection method must be provided as an alternative to speech recognition.

V. FIELD TRIALS: METHOD

Field trial participants were recruited to represent a range of severity of dysarthria and ECS use. The main inclusion criteria were for participants to have stable severe dysarthria and physical disability requiring use of ECS. The level of dysarthria was assessed using the Frenchay Dysarthria Assessment [14] and candidates with 25% intelligibility or more for single words were excluded. Identification of participants was through speech and language therapist and clinical engineer caseloads in South Yorkshire, UK after ethical approval had been received. Due to the communication difficulties of some participants a familiar communication partner

was present whilst informed consent to take part in the research was obtained.

To trial the training method, computers running the training software were supplied to each participant for a 6 week period. At the beginning of this period, 30 examples of each word in the participant's control vocabulary were recorded. These were used to train speech recognisers, which were used as the basis for the training feedback. The computers were then left with the participants, who were instructed to use the training program as often as they wished. During the training phase, each utterance was recorded. At the end of the training period, all recorded utterances were used to train a second recogniser. This second recogniser was used in the speech-controlled ECS field trial. The recognition accuracy of the recognisers constructed before and after training were tested with the same unseen data (ie. speech data not used to train the recogniser) consisting of 10 examples of each word.

Following the training trial, the participants were provided with a speech-controlled ECS tailored to their individual needs. All participants chose to use remote array microphones in preference to head-mounted microphones. The computers with infra-red senders were positioned in the home according to individual participants' preferences. Some participants chose to have the computer screen in view and others to rely on audible feedback from the computer, with the screen out of view. Microphones were sited, according to participants' preferences and the layout of the rooms, at distances from the usual sitting positions of the participants ranging from 0.5 – 3.0 metres and were not re-positioned during the trial. No safety-critical control functions were included and most of the functions were for control of audio-visual equipment, eg TV, hi-fi etc. Participants were allowed a 2-week period to familiarise themselves with the systems.

Following this period, participants were encouraged to use the speech-controlled ECS in preference to their usual ECS, for those functions that were available, during a trial period of 6 weeks. During this time, participants used their usual ECS to control all other functions.

During the period of the trial, the system recorded every command issued by the user. To gain an insight into the recognition performance of the system, a random selection of 30 of these recorded commands was selected for each subject. A listener who was familiar with all the subjects then transcribed these utterances blind to recognition results. The transcriptions were then compared with the output from the recogniser for these examples, to give the word recognition and command phrase recognition accuracies of the system during normal use.

At the end of the six week period, the participants were visited and a structured trial of task completion time and task completion accuracy was carried out. The task completion time was measured for speech-controlled and usual switch-scanning ECS for each participant. Each participant was asked to use the speech-controlled ECS to complete each control task available on its interface and also asked to complete the same tasks using their usual switch-scanning ECS systems. Each task was repeated three times and tasks were presented to the user in random order, to avoid any order effects. The time taken to complete each task was recorded. The time measured was between issuing the request to perform the task and the successful completion of the task, no matter how many attempts it took to complete the task. Task completion accuracy was also recorded as the proportion of tasks completed successfully on the first, second, third or subsequent attempt.

A questionnaire was devised to elicit the views of the participants on the speech-controlled ECS and their preferences in relation to the usual ECS system. Opportunities were given for both closed and open responses.

VI. RESULTS

Seventeen people were approached to take part in the study. At initial assessment 5 were found not to fit inclusion criteria, 4 declined further involvement with the study, despite meeting inclusion criteria, and 8 volunteered to go through with the complete trial. The characteristics of the trial participants are shown in Table 1.

One of the original eight (participant 4) was unable to complete the training phase of the project due to deterioration in health. Two further participants (6 and 7) completed the training phase but did not complete the field trial of the speech-controlled ECS, one for health reasons and one because he felt he would not find the speech-controlled ECS helpful. Complete trial data has therefore been gathered on five people, four with cerebral palsy and one with multiple sclerosis.

Table 2 shows the pre-training and post-training recognition accuracy for each of the participants who completed the training phase. These results demonstrate that, for all participants, recognition accuracy increased as a result of training the participants to become more consistent and re-training the recogniser with a larger corpus of data. The overall recognition accuracy increased from 88.5% to 95.4% showing a significant effect of training

(chi-square test $p < 0.001$).

Table 3 shows both the word recognition accuracy and the command phrase recognition accuracy for a random sample of 30 command phrases, for each participant, spoken during the ECS field trial period.

Table 4 shows task completion scores for participants' usual switch-scanning ECS and speech-controlled ECS on first attempt, first or second attempt and first, second or third attempt. For example, for participant 3, the switch-scanning control was not 100% accurate as they selected the incorrect option on four occasions; however on each occasion the task was completed successfully on the second attempt. For the speech ECS, participant 3 successfully completed 83% of tasks on the first attempt, an additional 10% of tasks on the second attempt and the remaining tasks (around 7%) on the third attempt. For all participants it can be seen that the speech-controlled ECS was not as accurate as the switch-scanning system. In all cases the speech-controlled ECS executed the command successfully by the third attempt.

Table 5 shows the average time taken to complete tasks for both speech-controlled ECS systems and the participants' usual ECS. The results show that the speech-controlled system is around twice as fast as the switch-operated systems (t-test, $p < 0.001$).

The responses to the questionnaire were mixed but the majority of participants thought the speech-control system easier to use and faster than their conventional system. The majority of

participants commented that the system was faster as they did not have to use a scanning interface. One said:

“I preferred it as I don't have to wait for it like I do with the [switch-scanning] ECS.”

Two participants particularly noted that they liked the system as it required less physical effort to use, one commenting:

“[it takes] much less effort - I don't get so tired.”

Interestingly, although speech control is not as accurate as conventional control methods, some perceived the system as more accurate than, or as accurate as, their usual system. The majority of participants found the system to be reliable, and three participants said they would be happy to use a speech-controlled system for safety critical operations, such as for an alarm call. Most participants found speech control more frustrating than their usual ECS and this is probably a reflection of lower accuracy. Most users felt more independent using the system and felt they could do more for themselves, requiring less help.

The majority of participants expressed satisfaction with the system. One participant found the experience unsatisfactory as the system was unable to control all of the devices they used regularly. Preferences were balanced, with two people preferring the usual switch-scanning system and two preferring speech control, with one finding both about equal.

VII. DISCUSSION AND CONCLUSIONS

The recogniser's accuracy under test conditions (see Table 2), compares favourably with reported recognition accuracy for dysarthric speech apparently comparable in severity to that of

the users in our trial. Recognition accuracy reported in the literature varies from 22% to 78% under test conditions [2-6]. It is not a straightforward task to compare recognition rates achieved with the STARDUST recogniser to those achievable by other systems, due to the uncertainty in comparisons of severity of dysarthria and the variety of different recognition methods and recogniser training methods used. Nevertheless, it appears that the techniques we have applied in this project have led to a significant advance in recogniser performance. By adopting speaker-dependent whole-word modelling for a small vocabulary we have identified a task for which ASR is viable for these speakers, and by changing recogniser-building methodology we have achieved results good enough for the ECS application to be viable for our clients. Results show an increase in recognition accuracy as a result of participants using the user training software described. These increases in recognition accuracy are considered to be due to closing the loop of data collection, recogniser training and user practice.

Recognition results in real usage show lower accuracy than seen for these same recognition models in test conditions (comparing Table 3 with Table 2). This difference is due to the fact that the results in Table 3 were for uncontrolled domestic acoustic conditions during the field trial, whereas those in Table 2 were obtained in quieter, though not silent, domestic conditions. The field trial was conducted in the participant's home under normal noise conditions, and all systems were used to control devices such TV or radio, which contribute to a higher background level of noise. All results were obtained with remote microphones, the placement of which was determined by individual circumstances, with relatively large subject-to-microphone distances. This is thus a difficult environment for speech recognition, but represents a set of real scenarios and thus tests the system in conditions typical of those found in users' homes.

In all cases, phrase recognition accuracy is lower than word recognition accuracy. This is to be expected as the phrases consist of two or three words and any word being recognised incorrectly would contribute to incorrect phrase recognition.

In the trials of the speech-controlled ECS in comparison to subjects' own switch-scanning ECS, the speech-controlled ECS was consistently faster to operate, taking around half the time to complete an operation, even when taking into account the fact that users sometimes had to repeat the command. The accuracy of the system in normal noise conditions within the home was less than accuracy achieved with the switch based systems with, on average, 79% of commands being recognised on the first attempt and 92% on the second attempt. By the third attempt 100% of commands were recognised for all users.

Experience in providing speech recognition ECS in clinical practice suggests that the level of accuracy we achieved is comparable with the accuracy encountered for commercially available systems in real usage, for non-dysarthric speakers [2]. At these levels of accuracy, some non-dysarthric speakers reject commercially available speech recognition as they find it too frustrating, preferring the more reliable, though slower, switch interface [2]. On the other hand, many non-dysarthric speakers successfully use speech recognition at these accuracy levels.

An additional, and arguably more serious, source of frustration for users of commercially available systems is their susceptibility to environmental noise, which can result in frequent false activations by the system. This has been circumvented with the STARDUST ECS by requiring that the user indicate an imminent command sequence by the single press of a switch.

Although this additional requirement for user action could be seen as a disadvantage of our system, we found that our field trial participants did not find it such.

It appears that in deciding whether they prefer the speech-controlled ECS, individual users balance the increased speed and ease of use of the speech-controlled system against the increased frustration arising from lower accuracy. For those who are very efficient users of switch interfaces (i.e. those who can cope with fast scanning speeds or fast two-switch users), or for those who have direct access (i.e. those who can use a large number of switches with each switch accessing a different function), this balance is more likely to favour the switch-scanning ECS. For those whose switch use is not as efficient, the balance seems to fall on preferring the speech-controlled ECS. There is scope for increasing the accuracy of the speech-controlled ECS as we have the ability to continually collect new speech data. The more the accuracy of speech recognition rises, the more people are likely to accept speech control as a more efficient alternative to switch-based systems.

ACKNOWLEDGEMENTS

This research was sponsored by the UK Department of Health New and Emerging Application of Technology (NEAT) programme and received a proportion of its funding from the UK National Health Service Executive. The views expressed in this publication are those of the authors and not necessarily those of the Department of Health or the NHS Executive.

REFERENCES

1. Blaney B, Wilson J 'Acoustic variability in dysarthria and computer speech recognition', *Clinical Linguistics and Phonetics*, 14(4), 307-327, 2000
2. Hawley MS Speech Recognition as an Input to Electronic Assistive Technology, *British Journal of Occupational Therapy*, 65(1), 15-20, 2002
3. Rosengren E, Raghavendra P, Hunnicut S How does automatic speech recognition handle severely dysarthric speech? Proc 2nd TIDE Congress, Paris, April 1995
4. Thomas-Stonell N, Kotler A-L, Leeper HA, Doyle C Computerized speech recognition: influence of intelligibility and perceptual consistency on recognition accuracy, *Journal of Augmentative and Alternative Communication*, 14, 51-55, 1998
5. Ferrier LJ. Shane HC. Ballard HF. Carpenter T. Benoit A. Dysarthric Speakers' Intelligibility and Speech Characteristics in Relation to Computer Speech Recognition. *Journal of Augmentative and Alternative Communication*. 11:165-174. 1995.
6. Kotler, A., Thomas-Stonell, N Effects of speech training on the accuracy of speech recognition for an individual with a speech impairment, *Journal of Augmentative and Alternative Communication*, 12: 71-80, 1997.

7. Deller Jr JR, Hsu D, Ferrier LJ On the use of hidden Markov modelling for recognition of Dysarthric speech. *Computer Methods & Programs in Biomedicine*, 35(2), 125-139, 1991.
8. Clark JA, Roemer RB Voice controlled wheelchair. *Archives of Physical Medicine & Rehabilitation*. 58(4), 169-75, 1977
9. Cohen A, Graupe D (1980) Speech recognition and control system for the severely disabled. *Journal of Biomedical Engineering*. 2(2), 97-107
10. Rosen, K., Yampolsky, S Automatic Speech Recognition and a Review of Its Functioning with Dysarthric Speech, *Journal of Augmentative and Alternative Communication*, 16: 48-60, 2000
11. Young, S., Evermann, G., Kershaw, D., Moore, G., Odell, J., Ollason, D., Valtcho, V., Woodland, P., *The HTK Book (for Version 3.1)*, Cambridge University Engineering Department, 2002
12. Rabiner, L., A tutorial on hidden Markov models and selected applications in speech recognition, *Proceedings of the IEEE*, 37: 257-286, 1989

13. Green P, Carmichael J, Hatzis A, Enderby P, Hawley M, Parker M (2003) Automatic Speech Recognition with Sparse Data for Dysarthric Speakers, Eurospeech 2003
14. Enderby P.M., Frenchay Dysarthria Assessment, Pro-Ed, Texas, 1983.

FIGURE 1. THE TRAINING AID DISPLAY ON A PERSONAL COMPUTER.



TABLE 1: CHARACTERISTICS OF THE TRIAL PARTICIPANTS

Participant	Sex	Diagnosis	Usual ECS system	Intelligibility rating			
				Individual word		Sentence	
				Range	Mean	Range	Mean
1	F	Cerebral Palsy (CP)	Possum PSU6	0-20	10	0-10	6
2	M	Multiple Sclerosis	Steeper Fox	0-50	22	20-60	34
3	M	CP	Steeper Persona	0	0	0	0
4	F	CP	Steeper Fox	0-10	4	0-10	2
5	M	CP	Possum PSU6	10-40	22	0-20	10
6	M	CP	Gewa Prog II	0-10	5	0-20	8
7	M	CP	Possum Companion	0-10	5	0-10	3
8	F	CP	Steeper Fox	0	0	0-10	3

TABLE 2

RECOGNITION ACCURACY FOR RECOGNISERS TRAINED USING DATA COLLECTED BEFORE USER TRAINING TRIAL (PRE-TRAINING), AND RECOGNISERS TRAINED WITH DATA COLLECTED DURING THE USER TRAINING TRIAL (POST-TRAINING) FOR 7 PARTICIPANTS. THE RECOGNITION ACCURACY WAS DETERMINED USING A TEST SET OF 10 UNSEEN EXAMPLES OF EACH WORD.

Participant	Vocabulary size	Pre-training		Post-training	
		No. of examples of each word	recognition accuracy (%)	No. of examples of each word	recognition accuracy (%)
1	11	30	95.8	103-108	100.0
2	7	13	96.2	32-34	100.0
3	10	30	82.0	51-58	86.0
5	13	30	96.9	79-110	99.7
6	11	30	92.7	46-55	96.4
7	11	30	77.3	35-50	95.5
8	13	30	80.0	56-66	90.8
overall			88.5		95.4

TABLE 3

COMMAND PHRASE AND WORD RECOGNITION ACCURACY DURING THE FIELD TRIAL, FOR UNCONTROLLED USAGE AND CONDITIONS IN THE PARTICIPANTS' HOMES

Participant	Command phrase accuracy (%)	Word accuracy (%)
1	83.3	90.0
2	76.7	81.6
3	86.7	93.0
5	76.7	86.7
8	70.0	83.3
overall	78.7	86.9

TABLE 4

TASK COMPLETION ACCURACY FOR PARTICIPANTS' SWITCH SCANNING ECS VS. SPEECH-CONTROLLED ECS.

Participant	Total no. of tasks	Switch-scanning ECS accuracy (%)		Speech ECS accuracy (%)		
		Attempt 1	Attempt 1 or 2	Attempt 1	Attempt 1 or 2	Attempt 1, 2 or 3
1	30	100.0	100.0	73.3	83.3	100.0
2	18	100.0	100.0	77.8	100.0	100.0
3	30	86.7	100.0	83.3	93.3	100.0
5	48	100.0	100.0	75.0	89.6	100.0
8	24	87.5	100.0	83.3	91.8	100.0

TABLE 5

MEAN TIME TAKEN TO COMPLETE TASKS WITH SWITCH-SCANNING AND SPEECH-CONTROLLED ECS.

Participant	Switch-scanning ECS (s)	Speech-controlled ECS (s)
1	16.3	8.4
2	12.2	6.3
3	19.3	7.2
5	17.7	8.1
8	19.2	8.7
Average Overall	16.9	7.7