Vocal tract and voice source features for monitoring cognitive workload

Manuela Meier[1,2], Michal Borsky[1], Eydis H. Magnusdottir[1], Kamilla R. Johannsdottir[1], and Jon Gudnason[1]

[1] Department of Electrical and Computer Engineering, Technical University of Munich
[2] Center for Analysis and Design of Intelligent Agents, Reykjavik University

manuela.S.meier@mytum.de, michalb@ru.is, eydis07@ru.is, kamilla@ru.is, jg@ru.is

Abstract—**Monitoring cognitive workload from speech signals has received a lot of attention from researchers in the past few years as it has the potential to improve performance and fidelity in human decision making. The bulk of the research has focused on classifying speech from talkers participating in cognitive workload experiments using simple reading tasks, memory span tests and the Stroop test, typically into three levels of low, medium and high. This study focuses on using parameters extracted from the vocal tract and the voice source components of the speech signal for cognitive workload monitoring. The experiment used in this study contains 92 participants, the levels were obtained by using a reading task and three Stroop tasks which were randomly ordered for each participant and an adequate rest time was used in-between tasks to mitigate the effect of cognitive workload from one task affecting the subsequent one. Vocal tract features were obtained from the first three formants and voice source features were extracted using signal analysis on the inverse filtered speech signal. The results show that on their own, the vocal tract features outperform the voice source features. The lowest MCR of 33.92 ± 1.05 was achieved with a SVM classifier.**

Index Terms: speech science, voice source signal, vocal tract features, computational paralinguistic.

I. INTRODUCTION

Monitoring cognitive workload in individuals that are performing safety-critical jobs has huge potential, for example in aviation [1]. The relationship between cognitive workload and performance has been well studied [2] and the connection between cognitive workload and physical health has also been highlighted [3]. Clearly, it is crucial to manage cognitive workload in the modern day work environment, both in terms of performance and health. Speech processing offers the ability to monitor cognitive workload in a non-intrusive way, compared to measures such as blood pressure, heart rate, electroencephalogram or electrocardiogram. Measuring voice is easy and recent work has shown very promising results in classifying cognitive workload levels using the speech signal [4], [5], [6]. Successful design and implementation of such a method would provide a powerful tool to developers of cognitive infocommunications systems [7].

The main objective of this work is to provide an independent verification of whether there is a link between increased cognitive workload and changes in the speech signal and, if this link exists, to characterize what part of voice is mostly affected. To do this we conducted a cognitive workload experiment where the set tasks were primarily solved using speech from 92 participants. Each participant read a standard passage of text on a computer screen and solved Stroop tasks with three difficulty levels [8]. A separate classifier of the three difficulty levels was trained for each speaker based on two sets of voice parameters comprising of the vocal tract and voice source features respectively. The conclusion is that task load does affect the speech signal and that vocal tract parameters are a better indicator of task level than voice source parameters. The context of the work is presented in Sec. II and the speech processing approach is described in Sec. III. The experimental methodology and detailed results are given in Sec. IV and V. The results are summarized in Sec. VI and the discussion is continued in Sec. VII.

II. SPEECH PROCESSING FOR COGNITIVE WORKLOAD

There is a growing body of work that supports the statement that cognitive workload affects the speech signal. For example, a data set of 15 participants was presented in [4] where each participant performed reading and Stroop tasks. The difficulty levels of these tasks were classified using mel-frequency cepstrum coefficients and prosodic features and a classifier based on a speaker-adapted Gaussian mixture model. The feature extraction was extended to include targeted extraction of vocal tract features through sub-band centroids [5]. These classification schemes indicated a strong relationship between cognitive workload and the speech signal within the experimental framework of the studies. The INTERSPEECH 2014 Cognitive Load Challenge (ComParE) was based on this methodology. A data set contained 26 participants providing speech recordings and an electroglottograph during a reading task and a low-, medium-, and high cognitive load level Stroop tasks [9]. The winning entry used an i-vector classification scheme on a combined feature set of fused speech streams, prosody and phone-rate [6]. Other entries used, for example, sub-band centroid features [10] and voice quality features [11] supporting earlier findings.

Other work that does not rely on induced cognitive load levels via tasks includes experiments done in military flight simulation [12]. Mean change in fundamental frequency and speech intensity was shown to increase and their range decreased as cognitive load of the flights increased. Urban search and rescue training operations were recorded and annotated for cognitive load analysis [13] using prosodic features, Teager

energy, voicing strengths and spectral envelope.

This work focuses on two sets of voice features extracted from the vocal tract and the voice source respectively. Voice can be decomposed into the vocal tract and the voice source using the traditional linear source-tract model [14], [15]. The vocal tract can be modeled using an auto-regressive moving average model corresponding to the formants and anti-formants. These were extracted using extended Kalman smoothing [16] and used as vocal tract features for analyzing depression [17].

Voice source features can be extracted indirectly from the speech signal via covariance analysis and cepstrum processing without relying on inverse filtering [18], [19]. Various methods for inverse filtering have however been implemented and evaluated [20] and feature extraction methods based on the glottal flow estimate have been proposed [21], [22], [23]. A set of voice source features was studied in relation to short affect bursts produced by 10 professional actors [24].

In this work vocal tract features are extracted using the Kalman smoothing approach [16] and the features described in [24] were used for the voice source. The two feature sets were combined using the temporal cross-correlation structure described in [17], which has been proven to be effective for detecting emotions and depression in speech and other signals. The cross-correlation structure simultaneously captures the short- and long term relationship between individual features (vocal tract and voice source) in the feature vector. The detail is described in the following section.

III. VOICE FEATURE EXTRACTION

The approach adopted in this work was based on the idea that decomposing the speech signal into the vocal tract characteristics and the voice source signal gives a better insight into the voice changes brought by increased cognitive workload. Two sets of distinct features were extracted from the speech signal, the first set consisted of features describing the vocal tract and the second set consisted of features describing the voice source signal.

A. Vocal Tract Features

The set of vocal tract features used in this work consisted of ordinary formants F1, F2 and F3. These were extracted from the speech signal using the KARMA [16] algorithm. It works on the principle of Kalman-based auto-regressive moving average smoothing. The main advantage of using KARMA is that the algorithm produces smoother formant contours than other methods and they have less erratic spikes around voiced-unvoiced transitions. The authors have demonstrated that this approach exhibits lower overall root-mean-square error and thus can be considered as more reliable. The three formants were obtained for each 20 ms of speech and concatenated in a vocal tract feature vector $\mathbf{x}_{vt}(j)$, where j denotes the frame index. Figure 1b illustrates an output of KARMA algorithm for a given speech signal on Figure 1a.

B. Voice Source Features

The voice source features consisted of 10 different parameters that were extracted either directly from the speech signal or from an estimate of the glottal flow or its derivative (voice
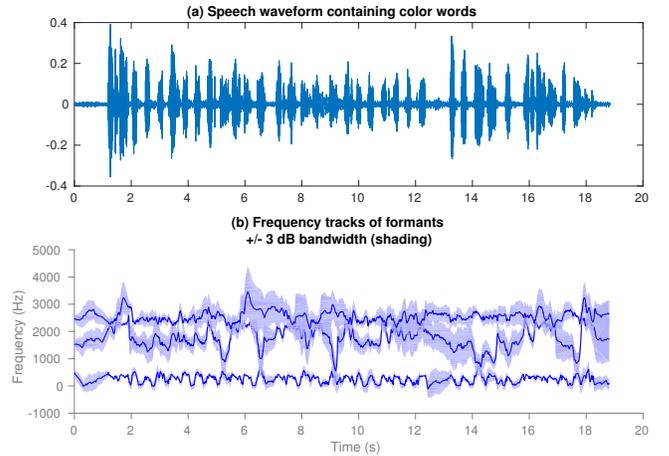

Fig. 1. *The result of formant tracking using KARMA algorithm.*

source signal). The following is a list of the parameters and their time- or frequency domain association:

time - maximum flow declination rate (MFDR), pulse amplitude (PA), normalized amplitude quotient (NAQ) and closed quotient (CQ),

freq. - fundamental frequency (f0), harmonics to noise ratio (HNR), harmonics richness factor (HRF), 1^{st} to 2^{nd} harmonics ratio (H1-H2), jitter (Jitt) and cepstral peak prominence (CPP).

Cepstral peak prominence (CPP) was the only measure extracted directly from the speech signal and has been historically widely used to classify and rate levels of dysphonia [25]. While there is no clear understanding of what the parameter measures, the general findings among the speech pathologists show that the parameter is tied to vocal attributes of breathiness, roughness and hoarseness. A theoretical study on the parameter's nature was done in [26] where the authors concluded that CPP integrates measure of several features describing the aperiodicity and waveform of the acoustic voice signal.

The iterative adaptive inverse filtering algorithm [27] was employed in order to obtain the voice source signal from which the rest of the voice source measures were extracted. Authors in [24] studied the relation of these features to the emotions and found a statistical significance for all of them aside from NAQ. However, other works [28] have reported on increase of this parameter for hypo-functioning (e.g. relived) voice and a decrease for hyper-functioning voice (e.g. angered) [29].

The voice source features were obtained for each 10 ms of speech and concatenated in a voice source feature vector $\mathbf{x}_{vs}(j)$, where j denotes the frame index.

C. Feature level fusion

Feature level fusion was achieved by concatenating the frame-level feature vectors of the vocal tract features ($n_{c-vt} = 3$ formant values) and voice source features ($n_{c-vs} = 10$ parameters) for each frame, so that, instead of having only $\mathbf{x}_{vt}(j)$ or $\mathbf{x}_{vs}(j)$, we now have $\mathbf{x}(j) = [\mathbf{x}_{vt}(j)^T, \mathbf{x}_{vs}(j)^T]^T$

for each frame j. From this a correlation matrix (with $n_c = 13$ parameters) and an utterance-level feature vector $\mathbf{u}$ was made.

D. Temporal correlation structure

The three feature streams $\mathbf{x}_{vt}(j)$, $\mathbf{x}_{vs}(j)$ and $\mathbf{x}(j)$ are processed further and summarized to produce a single feature vector $\mathbf{u}_{vt}$, $\mathbf{u}_{vs}$ or $\mathbf{u}$ for the utterance to be classified. This is done by using the correlation structure of the feature stream [17]. A fixed time scale of 2 frames was used to create a concatenated feature vector of the current feature vector at time j and 13 successive time delays. For the joint vocal tract and voice source vector $\mathbf{x}(j)$ a the new data vector of $N = n_c n_d = 13 \times 14 = 182$ dimensions is created with,

$$\mathbf{y}(j) = [\mathbf{x}^T(j), \mathbf{x}^T(j-2), \mathbf{x}^T(j-4), \ldots \mathbf{x}^T(j-26)]^T. \quad (1)$$

The vocal tract $\mathbf{y}_{vt}(j)$ and voice source feature vectors $\mathbf{y}_{vs}(j)$ therefore have dimensions $N_{vt} = n_{c-vt}n_d = 3 \times 14 = 42$ and $N_{vs} = n_{c-vs}n_d = 10 \times 14 = 140$ dimensions respectively.

The cross-correlation matrix for the utterance is then formed using,

$$\mathbf{R} = \frac{1}{N_s} \sum_j \tilde{\mathbf{y}}(j)\tilde{\mathbf{y}}^T(j) \quad (2)$$

where N_s is the number of frames in a utterance and $\mathbf{y}(j)$ is normalized to have a zero mean and unit variance, denoted as $\tilde{\mathbf{y}}(j)$. The eigenvalues λ_i are obtained from $\mathbf{R}$ in a descending order and each component in the utterance level feature vector $\mathbf{u} = [u_1, u_2, \ldots, u_N]^T$ is computed as the normalized cumulative sum

$$u_n = \frac{\sum_{i=1}^{n} \lambda_i}{\sum_{i=1}^{N} \lambda_i} \quad (3)$$

where $n = 1, 2, \ldots, N$. The component u_n therefore represents the proportion of energy contained in the first n eigenvectors.

Figure 3 illustrates the extracted vocal tract $\mathbf{u}_{vt}$ and voice source $\mathbf{u}_{vs}$ feature vectors from for three cognitive workload classes (explained in next section). The figure demonstrates that for these utterances for the vocal tract features, more than 90% of the energy is captured by the first 8 out of $N_{vt} = 42$ components whereas for the voice source features more than 43 out of $N_{vs} = 140$ components are needed to capture 90% of the energy. The number of eigenvalue components in the figure are truncated for illustrative purposes after the cumulated eigenvalues reached 99.9%.

E. Utterance level fusion

Utterance level fusion was achieved by concatenating the high-level feature vectors $\mathbf{u}_{vt}$ and $\mathbf{u}_{vs}$ into a single vector $\mathbf{u}_f$. This approach is distinctly different from the frame-level fusion. Feature-level fusion assumes that modelling cross-correlations between vocal tract and voice parameters might provide new information about the cognitive workload. Utterance-level fusion, on the other hand, assumes that these two streams are independent of each other and thus can be used in tandem.
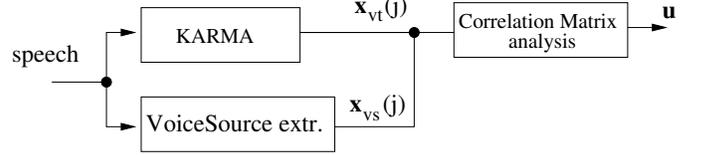

Fig. 2. *Feature extraction block scheme*

IV. METHODOLOGY

A. Experiment

Our experiments were designed around the Stroop test [8], which shows words for colours, such as 'blue' and 'red' on a computer screen in either congruent colour (the word 'red' is shown in red), or incongruent colour (the word 'red' is shown in, for example, blue). Participants are instructed to name the colours in which the words are shown as quickly as possible. If the word is shown in incongruent colour, the task is likely more difficult than when it is shown in congruent colour as the brain is wired to automatically read words when shown, pre-loading the brain with the wrong information (colour name). An additional increase in task difficulty was created by adding a rather strict time limit for naming the colours and only showing a single coloured word on the screen at once, preventing participants from 'looking ahead'. We used the same five colours and corresponding words across the three conditions and the words were shown in the native language of the participants (Icelandic). All speech was recorded using a head-mounted microphone and the speech was sampled at 48 kHz.

A group of 92 participants started with a 5 minute resting period during which baseline values for their cardiovascular activity was determined. Ideally, the subjects would be completely at rest. Then all subjects were asked to read aloud a short passage of text which took around 2 minutes. It aims to engage the subjects in the same physical activity - speaking - as during the Stroop tasks, but performing a task (reading) that they were likely experienced in and comfortable with. After the reading task followed 3 minutes of recovery time (rest), before starting on the Stroop tests. Three Stroop difficulty levels, congruent (L1), mostly incongruent (L2), and time-limited (L3), were presented in random order (Latin square with 6 possible orders) and with 5 minute intervals between each Stroop level. Each set consisted of 6 screen s with 35+1 words each (216 words per Stroop level), where the last word was used to indicate the end of the screen. The time-limited task had an randomly alternating limit of 0.75 or 0.65 seconds per word, which was chosen for being roughly the amount of time participants needed to do the non-time-limited tasks (0.75s) and a bit less so as to ensure the task would be considered extra hard (0.65s). The average time taken to complete a single screen of the congruent task was 23 seconds, increasing to 30 seconds for the incongruent task, and 29 and 25 seconds for the time-limited tasks. Each screen was recorded as a single utterance from which an utterance level feature vector $\mathbf{u}$ is extracted as explained in Sec. III.

B. Classifiers & Evaluation

The primary goal of this paper was to answer the questions: "*Can increased workload be detected in the recordings from a*
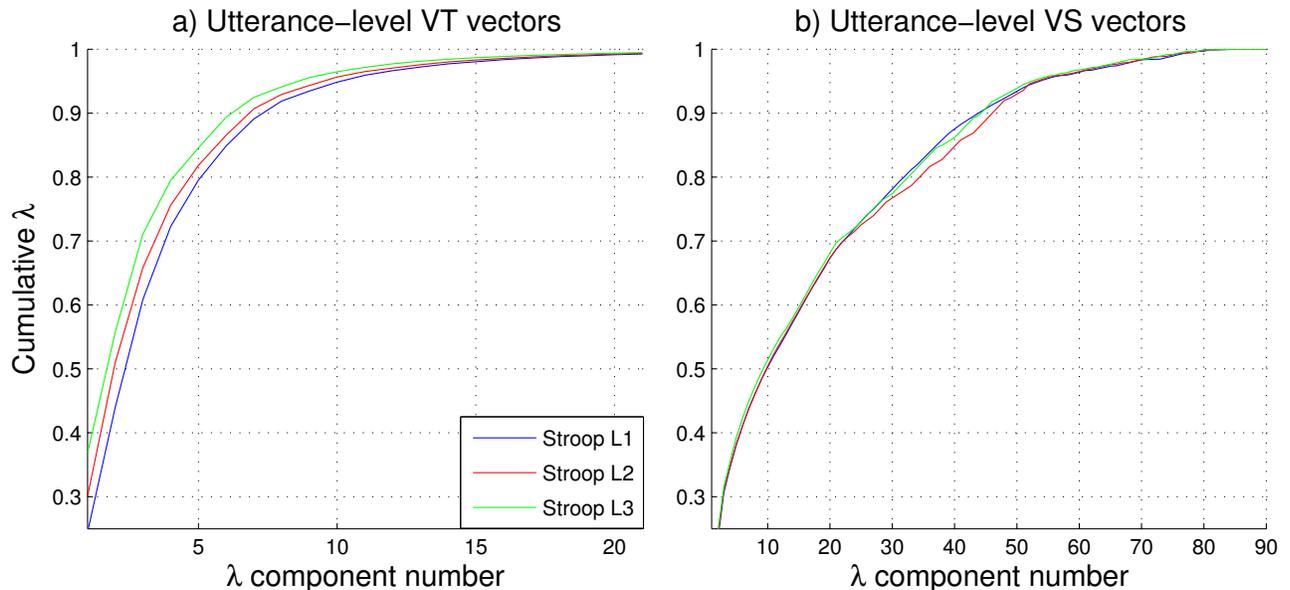
Fig. 3. *An example of a vocal tract and voice source features for all three Stroop tasks.*

series of Stroop tests? And if yes, what features are best suited for this task?" Thus, the experimental task was to construct a trinary classifier that would classify unknown recordings into three different Stroop levels. In order to measure the discriminatory capabilities of the discussed features without introducing additional speaker variability, the classifiers were built as speaker-dependent.

Three distinct types of supervised classifiers were used in this study: Minimum distance (MD) to the class centroid, Support Vector Machine (SVM) and Random Forests (RF). The MD based classifier represents the most intuitive approach to the classification which makes it an ideal choice to gain an insight into data's separability using chosen features. A class centroid was calculated from the training data and the classification was based on the minimum Euclidean distance between the centroids and the test vector. While the SVM is fundamentally a binary classifier, it is also possible to combine multiple SVMs and solve multiclass classification problems. The approach taken in this article was based on constructing all possible two-class SVMs and leaving the MD to resolve the ambiguous cases. The SVMs used linear kernel function, the soft margin optimization method, the hyperparameter C was set to 1 and no observations were allowed to violate the Karush-Kuhn-Tucker conditions. The third classifier was based on random forests, when the final ensemble had 20 trees and the minimum number of observations in a leaf was set to 1.

The amount of available data did not allow us to draw statistically valid conclusions from just a single split of dataset into training and testing sets. As a consequence, we used a leave-one-out cross validation strategy. The testing set consisted of a single signal, while the rest of the signals (20) were put into the training set to train the classifier. And finally, the evaluation results was recorded into a cumulative confusion matrix. This whole process was repeated for each signal separately, which allowed us to train 21 classifiers and to obtain 21 classification scores for every possible set division. The final confusion matrix after the whole run was

used to compute the total misclassification rate (MCR) that is presented in the tables below.

V. RESULTS

The results obtained from the classification tests are shown in Table I. There are several interesting things that can be taken from it. The first thing to look at is the performance of the proposed speech features in the task of classifying the Stroop level and its correlation to the cognitive workload. A trinary classification task with a set of completely random features would theoretically achieve the MCR of 66%, but the results for our best features were below 34%. These results indicate that the cognitive workload causes changes to the voice that can be observed and objectively measured.

The vocal tract features achieved better MCRs than any other studied set of features. The overall best results of 33.92% were obtained with a SVM classifier, while the RF and MD classifiers followed with the MCRs of 43.15% and 43.93% respectively. The second best overall results were achieved with combined features (both utterance and feature-level fused) and the voice source features scored as the last. Another interesting thing to note is the fact that utterance-level fusion outperformed the feature-level fusion by 5.49% for SVMs. In all studied setups, the SVMs proved to consistently outperform other classifiers.

TABLE I.   *Average misclassification rate [%] over all speakers with different sets of features and classifiers.*

Features	Min.Dist	SVM	RF
VTfeat	43.93±1.09	**33.92±1.05**	43.15±1.09
VSfeat	65.65±1.05	47.47±1.1	55.59 ± 1.1
Utt. Fused	53.43±1.08	35.86±1.06	41.77±1.09
Feat. Fused	42.91±1.09	41.35± 1.09	49.08±1.1

A closer look at the confusion matrices gives an insight into the separability of distinct Stroop levels. Table II presents the results for SVM classifiers with the vocal tract features

averaged over all speakers. The actual Stroop L1 utterances were more often confused for Stroop L2 tasks then they were for the L3 tasks. The same trend, although naturally reversed, was observed for actual Stroop L3 utterances. This trend was much more pronounced for actual Stroop L3 tasks. This observation leads us to the conclusion that an increase in the cognitive workload increases the changes in the voice. Another interesting fact is that a classifier trained for Stroop L2 tasks misclassified the Stroop L3 features more often that the Stroop L1. This observation leads to the conclusion that the increased cognitive difficulty of Stroop L2 and Stroop L3 tasks introduces changes to the voice, which can be accurately detected using the vocal tract features.

TABLE II.   *Misclassification rate [%] matrix for the vocal tract features with the SVM classifier.*

	Stroop L1	Stroop L2	Stroop L3
Stroop L1	**72.5**	15.7	11.8
Stroop L2	17.5	**59.7**	22.8
Stroop L3	10.1	24.6	**65.3**

The histogram of MCR for all 92 participants is presented in Fig. 4 and there are several interesting things to look at. A single participant scored at 60%, which is very close to random classification result, while only a single participant scored at 0%. Most of the participants scored around the average MCR of 33%.
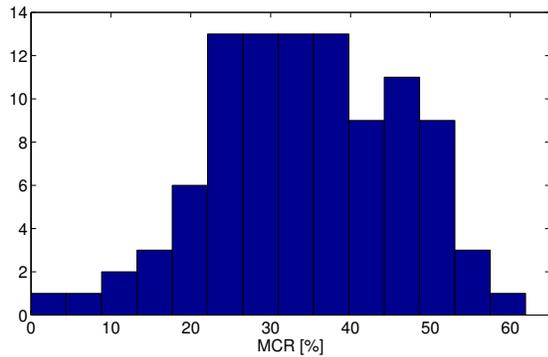

Fig. 4.   *A histogram of MCR [%] for the whole set of speakers.*

VI.   Conclusions

This study gives an independent verification of the relationship between cognitive workload and speech. Speaker independent classifiers trained on 92 participants were able to distinguish between utterances of low, medium and high cognitive workload with $33.92\pm1.05\%$ misclassification rate using vocal tract features and $35.86\pm1.06\%$ when fused together with voice source features. The voice source features only achieved $47.47\pm1.1\%$ misclassification rate on their own.

Specific conclusions of this work are that the particular vocal tract parameters used in this work [16] outperform the voice source parameters that are studied in [24] for emotion detection of which some are also used for cognitive workload classification in [22]. This is in concurrence with studies compare vocal tract and voice source features both for other tasks [18], [19] and cognitive workload [22]. More generally,

the study reinforces previous findings that show that there is a link between cognitive workload and speech [12], [22], [11], [6], [13].

VII.   Discussion

Speaking is a cognitive activity so it is not surprising that a simultaneous task affects the voice. Misclassification rate of over 30% is high however, given that each test utterance is over 20 s in duration and the number of classes is only three. There are two possible reasons behind this discrepancy. The first is related to precisely that of speaker identification but individuals react differently to cognitive workload [2]. This study has tried to overcome this by training individual classifiers for each participant and other studies do similar things, for example using speaker adaptation [22]. Better methods for addressing individual differences for cognitive workload are therefore called for. The second reason that explains the bad classification performance has to do with the dynamic nature of the body's response to cognitive workload [30]. Evidence from studies using cardiovascular reactivity shows that the body's response to cognitive stimulus is strong during the first few seconds of the task but fades as the participant gets used to the task. This dynamic behaviour is not captured in the current experimental framework used for speech classification but there is a strong reason to suggest that a time-varying analysis would improve classification.

In another sense, the results in this paper give reasons to optimism. The results show the voice responds differently to different task load levels. The data gathering setup gives us an opportunity to adapt the methodology to the concerns regarding the time-varying nature of the body's response to cognitive workload. This is indeed the focus of our future research.

REFERENCES

[1] Experimental Centre Note No 18/06, "Evaluation of the human voice for indications of workload induced stress in the aviation environment," Tech. Rep., 2006.

[2] E. Galy, M. Cariou, and C. Mélan, "What is the relationship between mental workload factors and cognitive load types?" *International Journal of Psychophysiology*, vol. 83, pp. 269–275, 2012.

[3] M. A. Kompier, B. Aust, A. M. van den Berg, and J. Siegrist, "Stress prevention in bus drivers: evaluation of 13 natural experiments." *Journal of occupational health psychology*, vol. 5, pp. 11–31, 2000.

[4] B. Yin, F. Chen, N. Ruiz, and E. Ambikairajah, "Speech-based cognitive load monitoring system." in *Proc. IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, 2008, pp. 2041–2044.

[5] P. N. Le, E. Ambikairajah, J. Epps, V. Sethu, and E. H. Choi, "Investigation of spectral centroid features for cognitive load classification," *Speech Communication*, vol. 53, no. 4, pp. 540–551, Apr. 2011. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0167639311000069

[6] M. Van Segbroeck, R. Travadi, C. Vaz, J. Kim, M. P. Black, A. Potamianos, and S. S. Narayanan, "Classification of cognitive load from speech using an i-vector framework," in *Proc. Interspeech*, 2014.

[7] P. Baranyi and A. Csapo, "Cognitive infocommunications: coginfocom," in *Computational Intelligence and Informatics (CINTI), 2010 11th International Symposium on*. IEEE, 2010, pp. 141–146.

[8] J. R. Stroop, "Studies of interference in serial verbal reactions," *Journal of Experimental Psychology*, vol. 18, pp. 643–662, 1935.

[9] B. Schuller, S. Steidl, A. Batliner, J. Epps, F. Eyben, F. Ringeval, E. Marchi, and Y. Zhang, "The interspeech 2014 computational paralinguistics challenge: cognitive & physical load." in *INTERSPEECH*, 2014, pp. 427–431.

[10] J. M. K. Kua, V. Sethu, P. N. Le, and E. Ambikairajah, "The unsw submission to interspeech 2014 compare cognitive load challenge." in *INTERSPEECH*, 2014, pp. 746–750.

[11] M. Huckvale, "Prediction of cognitive load from speech with the voqal voice quality toolbox for the interspeech 2014 computational paralinguistics challenge." 2014.

[12] K. Huttunen, H. Keränen, E. Väyrynen, R. Pääkkönen, and T. Leino, "Effect of cognitive load on speech prosody in aviation: Evidence from military simulator flights," *Applied ergonomics*, vol. 42, no. 2, pp. 348–357, 2011.

[13] M. Charfuelan and G.-J. Kruijff, "Analysis of speech under stress and cognitive load in usar operations," in *Natural Interaction with Robots, Knowbots and Smartphones*. Springer, 2014, pp. 275–281.

[14] G. Fant, *Acoustic Theory of Speech Production*. The Hague, The Netherlands: Mouton, 1960.

[15] D. Y. Wong, J. D. Markel, and J. A. H. Gray, "Least squares glottal inverse filtering from the acoustic speech waveform," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 27, no. 4, pp. 350–355, Aug. 1979.

[16] D. D. Mehta, M. Zañartu, S. W. Feng, H. A. C. II, and R. E. Hillman, "Mobile voice health monitoring using a wearable accelerometer sensor and a smartphone platform," *IEEE Transactions on Biomedical Engineering*, vol. 59, no. 11, pp. 3090–3096, 2012.

[17] J. R. Williamson, T. F. Quatieri, B. S. Helfer, G. Ciccarelli, and D. D. Mehta, "Vocal and facial biomarkers of depression based on motor incoordination and timing," in *Proceedings of the 4th International Workshop on Audio/Visual Emotion Challenge*. ACM, 2014, pp. 65–72.

[18] M. D. Plumpe, T. F. Quatieri, and D. A. Reynolds, "Modeling of the glottal flow derivative waveform with application to speaker identification," *IEEE Trans. Speech Audio Process.*, vol. 7, no. 5, pp. 569–576, Sep. 1999.

[19] J. Gudnason and M. Brookes, "Voice Source cepstrum coefficients for speaker identification," in *Proc. IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, 2008, pp. 4821–4824.

[20] J. Gudnason, D. D. Mehta, and T. F. Quatieri, "Evaluation of speech inverse filtering techniques using a physiologically based synthesizer," in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 4245–4249.

[21] P. N. Le, J. Epps, E. H. Choi, and E. Ambikairajah, "A study of voice source and vocal tract filter based features in cognitive load classification," in *Pattern Recognition (ICPR), 2010 20th International Conference on*. IEEE, 2010, pp. 4516–4519.

[22] T. F. Yap, J. Epps, E. Ambikairajah, and E. H. Choi, "Voice source features for cognitive load classification," in *Proc. IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, 2011, pp. 5700–5703.

[23] J. Gudnason, M. R. P. Thomas, D. P. W. Ellis, and P. A. Naylor, "Data-driven voice source analysis and synthesis," *Speech Communication*, vol. 52, no. 2, pp. 199–211, February 2012.

[24] S. Patel, K. R. Scherer, E. Björkner, and J. Sundberg, "Mapping emotions into acoustic space: the role of voice production," *Biological psychology*, vol. 87, no. 1, pp. 93–98, 2011.

[25] Y. D. Heman-Ackah, R. T. Sataloff, G. Laureyns, D. Lurie, D. D. Michael, R. Heuer, A. Rubin, R. Eller *et al.*, "Quantifying the cepstral peak prominence, a measure of dysphonia," *Journal of Voice*, vol. 28, pp. 783–788, 2014.

[26] R. Fraile and J. I. Godino-Llorente, "Cepstral peak prominence: A comprehensive analysis," *Biomedical Signal Processing and Control*, vol. 14, pp. 42–54, 2014.

[27] P. Alku, "Glottal wave analysis with pitch synchronous iterative adaptive filtering," *Speech Communication*, vol. 11, pp. 109–118, 1992.

[28] P. Alku and T. Backstrom, "Normalized amplitude quotient for parametrization of the glottal flow," *J. Acoust. Soc. Am.*, vol. 112, no. 2, pp. 701–710, Aug. 2002.

[29] T. Waaramaa, P. Alku, and A.-M. Laukkanen, "The role of F3 in the vocal expression of emotions," *Logopedics, Phoniatrics, Vocology*, vol. 31, no. 4, pp. 153–156, 2006.

[30] A. Stuiver, D. De Waard, K. Brookhuis, C. Dijksterhuis, B. Lewis-Evans, and L. Mulder, "Short-term cardiovascular responses to changing task demands," *International Journal of Psychophysiology*, vol. 85, no. 2, pp. 153–160, 2012.