


Real-time tracking of speakers' emotions, states, and traits on mobile platforms

Erik Marchi, Florian Eyben, Gerhard Hagerer, Björn Schuller

Angaben zur Veröffentlichung / Publication details:

Marchi, Erik, Florian Eyben, Gerhard Hagerer, and Björn Schuller. 2016. "Real-time tracking of speakers' emotions, states, and traits on mobile platforms." In *Understanding speech processing in humans and machines: 17th Annual Conference of the International Speech Communication Association (INTERSPEECH 2016), San Francisco, California, USA, 8-12 September 2016; Volume 2*, 1182–83. Baixas: International Society for Conversation Analysis (ISCA).





Real-time Tracking of Speakers' Emotions, States, and Traits on Mobile Platforms

Erik Marchi¹, Florian Eyben¹, Gerhard Hagerer¹, Björn Schuller^{1,2}

¹audEERING GmbH, Gilching, Germany

²Department of Computing, Imperial College London, London, UK

em@audeerling.com

Abstract

We demonstrate audEERING's sensAI technology running natively on low-resource mobile devices applied to emotion analytics and speaker characterisation tasks. A show-case application for the Android platform is provided, where audEERING's highly noise robust voice activity detection based on LSTM-RNN is combined with our core emotion recognition and speaker characterisation engine natively on the mobile device. This eliminates the need for network connectivity and allows to perform robust speaker state and trait recognition efficiently in real-time without network transmission lags. Real-time factors are benchmarked for a popular mobile device to demonstrate the efficiency, and average response times are compared to a server based approach. The output of the emotion analysis is visualized graphically in the arousal and valence space alongside the emotion category and further speaker characteristics.

Index Terms: emotion recognition, mobile application, computational paralinguistics

1. Introduction

It is broadly agreed upon that state and trait recognition from speech bears interest for mobile services, albeit being challenging, e.g., due to noise and low-quality audio [1]. Few advanced products already exist and can be used through nowadays smartphones such as Beyond Verbal's app "Moodies" or EI Technologies's app "Xpression". Both applications analyse speech and indicate changes to allow training and improvement for its users. With such mobile services, manifold applications of computational paralinguistics, i.e., automatic recognition of speakers' states and traits, can be enhanced due to the mobile and ever-present character of off-line-enabled smart-phone analysis. To illustrate this, let us give but three examples:

First, a job applicant could use such an app not only to train for a job interview and improve the aspects of self-assurance and determination in order to have a higher chance of securing the job [2], but could even secretly use such monitoring during the interview. Second, giving online and in situ rhetorical feedback for different purposes becomes a reality. Third, health is another sector for utilising the possibilities of mobile off-line enabled speech monitoring [3]. EI Technologies Xpression app was designed to record mood changes throughout the day for people with anxiety, depression or stress to support the psychologists for their patient's treatment fine-tuning. It records five emotional stages throughout the day and emails the list to the psychologist at the end of the day. With off-line processing ability, a new possibility could be triggering help or warning patients in extreme stress situations [4] by detecting their emo-

tional stage and reacting based on it.

In the remainder of this paper, we introduce our implementation (Section II), and show experimental benchmarking results (Section III) before drawing conclusions (Section IV) from these.

2. Implementation

Most current mobile speech processing applications, such as ASR (e.g., Google's voice search), rely on a frontend-backend architecture and require internet access for the services to operate [5]. This has the great advantage of flexibility, as the frontend can be kept light-weight and generic, and all the specific, heavy weight processing can be done on the backend. The frontend typically includes functionality to record, optionally compress, and finally transmit (or stream while recording) audio to the backend over a network/internet connection. The backend receives the recorded audio snippets or streams, decodes them, and runs the requested analysis, e.g., feature extraction and classification of speech emotions and speaker characteristics in our case. Such a distributed approach has been shown feasible including feature vector compression for computational paralinguistics in [6].

In this demo-showcase we show the world's first working *industrial* prototype of a novel embedded audio analysis engine, where all steps are run entirely on a mobile phone and no connectivity to a backend is required. In fact, only related academic prototypes have previously been introduced in the literature, e.g., [7], where speakers' emotions are detected locally on off-the-shelf mobile phones.

For on device feature extraction we use a modified version of our popular feature extractor openSMILE [8], which we have compiled natively for the Android platform (ARM processor).

The demonstrator Android App is capable of recording speech, either in a hands-free mode or in a push-to-talk mode. In case of the hands-free mode, a robust voice activity detector based on memory enhanced neural networks [9] is used to detect segments which contain speech (separated by speech pauses). The detector described in [9] (also included in the latest version of openSMILE, see [10]) has been optimised computationally to run well faster than real-time on the mobile device. Acoustic features are then extracted only for speech segments and speaker's characteristics and emotional state are classified via Support-Vector Machines (SVMs) and Support Vector Regression (SVR).

The observed vocal emotional state of the speaker as well as its history over the previous speech segments are shown graphically as dots in the upper half of the phone screen in a 2D activation valence coordinate system. Speaker characteristics such

Feature set	# feats.	RTF (i7)	RTF (M9)
Interspeech '09	384	0.01	0.06
Interspeech ComParE	6373	0.04	0.23

Table 1: Benchmarks (real-time-factor – RTF) of feature extraction (from uncompressed waveform audio) on an Intel Core-i7 (2.5 GHz) laptop (i7) and an HTC One M9 Android phone (M9).

as gender and personality are shown in the bottom half. The output is further encoded in EmotionML [11] ready for usage, e. g., in other applications run on the phone.

3. Experiments and Results

In order to verify the performance of the on-device feature extraction and classification, run-time benchmarks were conducted for two different acoustic feature sets which are part of the openSMILE package: the Interspeech 2009 Emotion Challenge baseline feature set [12], and the Interspeech 2013 Computational Paralinguistics Challenge (ComParE) baseline feature set [13]. The motivation of their choice is a) their free-availability and well-defined standard, b) their increasing size from below 400 to several 1000 features (cf. Table 1), and c) the fact that with these sets, one can reproduce the same benchmark results in terms of recognition rates as were reached in the baselines and above of the challenge events the Interspeech sets have been designed for: These tasks include speaker’s cognitive load, physical load, social signals such as laughter and breathing, conflict, up to 12-class emotion, eating condition, degree of nativeness, and for health applications recognition of autism manifestations, and Parkinsons disease [13, 14, 15]. In addition, the former Interspeech challenge task results can be reproduced with improved results due to the larger feature spaces, namely speaker’s age, gender, five dimensional personality profile, degree of likability, intelligibility, intoxication, sleepiness, and level of interest.

Table 1 lists average real-time factors (over 50 runs) for extracting features from a 120 second audio file (uncompressed), both on the mobile phone (ARM processor) and a standard pc platform (Intel processor). The results show that the Interspeech 2009 Emotion Challenge feature set can be extracted easily in real-time, which is visible in the demo app by very low response times. Extracting the other features set on the device is also feasible.

4. Conclusions

We demonstrate the feasibility of mobile recognition of speakers’ emotions and characteristics on a mobile platform such as a smart phone in real-time through a fully functional, standalone Android application.

In future efforts, more emphasis will be given on weakly supervised learning by efficient algorithms from data collected via users of an according service while laying utmost importance on privacy of the data such as by low bit-rate feature vector compression in batches if users want to contribute to improvement of services.

5. Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No. 680883),

6. References

- [1] W.-J. Yoon, Y.-H. Cho, and K.-S. Park, “A study of speech emotion recognition and its application to mobile services,” in *Ubiquitous Intelligence and Computing*. Springer, 2007, pp. 758–766.
- [2] K. Anderson, E. André, T. Baur, S. Bernardini, M. Chollet, E. Chryssafidou, I. Damian, C. Ennis, A. Egges, P. Gebhard *et al.*, “The tardis framework: intelligent virtual agents for social coaching in job interviews,” in *Advances in Computer Entertainment*. Springer, 2013, pp. 476–491.
- [3] B. Schuller, E. Marchi, S. Baron-Cohen, A. Lassalle, H. O’Reilly *et al.*, “Recent developments and results of asc-inclusion: An integrated internet-based environment for social inclusion of children with autism spectrum conditions,” in *Proc. of IDGEI*, Atlanta, GA, 2015, no pagination.
- [4] K.-h. Chang, D. Fisher, J. Canny, and B. Hartmann, “How’s my mood and stress?: an efficient speech analysis library for unobtrusive monitoring on mobile phones,” in *Proceedings of the 6th International Conference on Body Area Networks*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2011, pp. 71–77.
- [5] Z.-H. Tan and B. Lindberg, *Automatic speech recognition on mobile devices and over communication networks*. Springer Science & Business Media, 2008.
- [6] Z. Zhang, E. Coutinho, J. Deng, and B. Schuller, “Distributing Recognition in Computational Paralinguistics,” *IEEE Transactions on Affective Computing*, 2014, 14 pages, to appear.
- [7] K. K. Rachuri, M. Musolesi, C. Mascolo, P. J. Rentfrow, C. Longworth, and A. Aucinas, “Emotionsense: a mobile phones based adaptive platform for experimental social psychology research,” in *Proceedings of the 12th ACM international conference on Ubiquitous computing*. ACM, 2010, pp. 281–290.
- [8] F. Eyben, F. Weninger, F. Gross, and B. Schuller, “Recent developments in opensmile, the munich open-source multimedia feature extractor,” in *Proc. of MM*. ACM, pp. 835–838.
- [9] F. Eyben, F. Weninger, S. Squartini, and B. Schuller, “Real-life Voice Activity Detection with LSTM Recurrent Neural Networks and an Application to Hollywood Movies,” in *Proceedings 38th IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2013*. Vancouver, Canada: IEEE, 2013, pp. 483–487.
- [10] F. Eyben and B. Schuller, “openSMILE: The Munich Open-Source Large-Scale Multimedia Feature Extractor,” *ACM SIGMM Records*, vol. 6, no. 4, December 2014.
- [11] M. Schröder, L. Devillers, K. Karpouzis, J.-C. Martin, C. Pelachaud, C. Peter, H. Pirker, B. Schuller, J. Tao, and I. Wilson, “What should a generic emotion markup language be able to represent?” in *Affective Computing and Intelligent Interaction*. Springer, 2007, pp. 440–451.
- [12] B. Schuller, S. Steidl, A. Batliner, and F. Jurcicek, “The INTERSPEECH 2009 Emotion Challenge,” in *Proc. of INTERSPEECH 2009*, Brighton, UK, Sep. 2009, pp. 312–315.
- [13] B. Schuller, S. Steidl, A. Batliner, A. Vinciarelli, K. Scherer, F. Ringeval, M. Chetouani, F. Weninger, F. Eyben, E. Marchi, M. Mortillaro, H. Salamin, A. Polychroniou, F. Valente, and S. Kim, “The INTERSPEECH 2013 Computational Paralinguistics Challenge: Social Signals, Conflict, Emotion, Autism,” in *Proceedings INTERSPEECH 2013*. Lyon, France: ISCA, 2013, pp. 148–152.
- [14] 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 *Proceedings INTERSPEECH 2014*. Singapore, Singapore: ISCA, 2014.
- [15] B. Schuller, S. Steidl, A. Batliner, S. Hantke, F. Hönl, J. R. Orozco-Arroyave, E. Nöth, Y. Zhang, and F. Weninger, “The INTERSPEECH 2015 Computational Paralinguistics Challenge: Degree of Nativeness, Parkinson’s & Eating Condition,” in *Proceedings INTERSPEECH 2015*. Dresden, Germany: ISCA, September 2015, 5 pages, to appear.