

Assessing the Prosody of Non-Native Speakers of English: Measures and Feature Sets (Abstract)

Eduardo Coutinho¹, Florian Hönig², Yue Zhang^{1, 3}, Simone Hantke³,
Anton Batliner^{2,3}, Elmar Nöth², Björn Schuller^{1,4}

¹Department of Computing, Imperial College London, UK

²Pattern Recognition Lab, Friedrich-Alexander University Erlangen-Nuremberg, Germany

³Machine Intelligence & Signal Processing Group, TUM, Munich, Germany

⁴Chair of Complex & Intelligent Systems, University of Passau, Germany

e.coutinho@imperial.ac.uk

1

Abstract

In this paper, we describe a new database with audio recordings of non-native (L2) speakers of English, and the perceptual evaluation experiment conducted with native English speakers for assessing the prosody of each recording. These annotations are then used to compute the gold standard using different methods, and a series of regression experiments is conducted to evaluate their impact on the performance of a regression model predicting the degree of naturalness of L2 speech. Further, we compare the relevance of different feature groups modelling prosody in general (without speech tempo), speech rate and pauses modelling speech tempo (fluency), voice quality, and a variety of spectral features. We also discuss the impact of various fusion strategies on performance. Overall, our results demonstrate that the prosody of non-native speakers of English as L2 can be reliably assessed using supra-segmental audio features; prosodic features seem to be the most important ones.

Index Terms: English, L2,

non-native speech, prosody, gold standard, feature evaluation

1. Introduction

Non-native (L2) speakers of English diverge from native English speakers in terms of linguistic (e.g. morphology, syntax, lexicon) and phonetic aspects. Phonetic aspects comprise segmental and supra-segmental (prosodic) phenomena. There are two main tasks within automatic speech processing where an assessment of non-native traits can help. First, the performance of Automatic Speech Recognition (ASR) degrades drastically while processing non-native speech. A central aspect here is the mismatch between training data from native speech and test data from non-native speech [1], preventing the acoustic models to capture the variability and ‘innovations’ introduced by L2 speakers. By assessing the degree of nativeness in general, and by specifying the speaker’s L1, it is possible to build speaker-

and/or speaker-group specific acoustic models. Second, such an assessment can be employed within Computer-Aided Pronunciation Training (CAPT) for the automatic screening of learners.

In this paper, we describe a new database with audio recordings of L2 English speakers, and the perceptual evaluation experiment conducted with native English speakers to obtain judgements of the naturalness of the pronunciation of each recording with respect to supra-segmental characteristics. These annotations are then used to compute the gold standard using different methods, and a series of regression experiments is conducted to evaluate their impact on the performance of a regression model predicting the degree of naturalness of L2 speech. We use the publicly available general-purpose feature set openSMILE [2] to compute a large set of acoustic features. Additionally, to add some more specialised (but still text-independent features) we also compute duration, rhythm, and prosodic features derived from a segmentation of the recordings into vocalic and consonantal intervals, and from pseudo-syllables derived from that segmentation.

2. Recording and annotation methods

2.1. Collection of voice recordings

We obtained a set of voice recordings from 54 non-native English speakers with varied degree of proficiency (see Table 1 for details). Each speaker was asked to read aloud a set of 11 sentences from two short stories written in the English language widely used in phonetics and speech pathology research. The data recorded consisted of 5 sentences taken of the fable “The North Wind and the Sun”, and 6 sentences extracted from “The Rainbow”. All recording sessions took place in a room at the Technische Universität München (Germany) using the same hardware and software. The full database comprises 594 recordings (11 sentences * 54 speakers), totalling to 1.4 hours of speech.

2.2. Annotation procedure

The full set of recordings was delivered for annotation to a group of 27 native English speakers. The group of annotators were instructed to rate the naturalness of supra-segmental characteristics by answering the question “*The English language has a characteristic prosody (sentence melody and rhythm, i. e. timing of syllables). This sentence’s prosody sounds ...*,” using a 5-point Likert scale (1 - normal; 2 - acceptable; 3 -

¹This work was funded by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 645378 (ARIA-VALUSPA), the German Federal Ministry of Education, Science, Research and Technology (BMBF) under grant 01IS07014B (C-AuDiT), and the German Ministry of Economics (BMWi) under grant KF2027104ED0 (AUWL), and the European Union’s 7th Framework Program under grant agreement No. 338164 (Starting Grant iHEARu).

Table 1: *Non-native (L2) English speakers characteristics. M - mean; SD - standard deviation*

Mother tongue (L1)
Arabic (4), Bulgarian (1), Chinese (13), French (1), German (22), Italian (2), Indian (2), Japanese (1), Lithuanian (1), Persian (2), Polish (1), Portuguese (1), Romanian (1), Russian (1), Turkish (1)
Age
M = 31 y.o.; SD = 9 y.o.; Range = 20-57 y.o.
Gender
28 female / 26 male

slightly unusual; 4 - unusual; 5 - very unusual). Inter-rater agreement (per session and rater gender) was computed using two common measures – Cronbach’s alpha (α ; computed using “alpha” function from R’s [3] package “coefficientalpha” [4]) and Spearman’s rank correlation coefficient (ρ). Results are shown in the last column of Table 2.

Table 2: *Rating sessions details: number of raters (#raters), number of recordings (#recordings) rated in each session, and inter-rater agreement per session (α - Cronbach’s alpha; ρ - Spearman’s rank correlation coefficient).*

Session	#recordings	#raters	inter-rater agreement	
			α	ρ
1	132	23 (All)	0.97	0.54
		15 (Female)	0.95	0.55
		8 (Male)	0.90	0.51
2	132	16 (All)	0.93	0.47
		10 (Female)	0.89	0.44
		6 (Male)	0.84	0.50
3	165	17 (All)	0.94	0.47
		11 (Female)	0.89	0.43
		6 (Male)	0.88	0.55
4	165	18 (All)	0.95	0.53
		11 (Female)	0.90	0.45
		7 (Male)	0.93	0.64
mean	-	19 (All)	0.95	0.50
		12 (Female)	0.91	0.47
		7 (Male)	0.89	0.55

2.3. Gold standard

We tested three different methods to infer the gold standard from individual annotations. First, we combined the individual ratings for each instance by calculating the median across raters for each instance. Second, we computed the arithmetic mean across all raters for each instance. Third, in order to consider individual (unknown) differences amongst annotators (e.g., level of attention or individual biases when providing the ratings), we attempted to improve the quality of the gold standard calculation by using the evaluator weighted estimator (EWE) method [5]. The histogram of frequencies for each gold standard is shown in Fig. 1.

3. Feature sets

We used five different features sets for acoustic modelling of this database as described in the following paragraphs.




Figure 1: *Histogram of frequencies for each gold standard.*

OpenSMILE Features (oS): openSMILE (version 2.0; [2]) is a framework for extracting general-purpose acoustic and prosodic features, and has applied successfully to a garden-variety of paralinguistic tasks. In this paper we employ the configuration file of the INTERSPEECH 2013 Computational Paralinguistics Challenge (*ComParE*) [6], which includes 6373 static features of functionals of low-level descriptor (LLD) contours. In order to evaluate the nature of acoustic features that permit to better predict the L2 speakers nonverbal proficiency, we subdivided the *ComParE* set into three partitions: prosodic (*oS_PROS*), voice quality (*oS_VQ*), and spectral and cepstral features (*oS_SPEC*). The LLDs in each subset are described in [7].

Prosodic Features (PROS): The set of features extracted (a total of 64) are based on duration, energy, pitch, and pauses, and were applied to locally describe arbitrary units of speech such as words or syllables. The sequence of these local features was then converted into a fixed-length vector using 12 functionals.

Rhythm Features (RHYT): Given the importance of temporal features to the evaluation of L2 speakers proficiency (cf. Section 1), we considered set of specific speech-rhythm related features. Based on the segmentation of the PR (see above) into vocalic and consonantal intervals, we computed Grabe’s raw pairwise variability index (rPVI; [8]) on consonants and vowels plus its (rate-of-speech-) normalised version (nPVI). Additionally, we computed variants of Grabe’s measures using squared instead of absolute differences (root average squared local change). This led to 8 features reflecting local variability in durations. Additionally, another set of five features was added to reflect global variability and proportions. These were calculated using Ramus’ V (percentage of vocalic intervals) and vocalic and consonantal Deltas (global standard deviations of durations) [9], plus Dellwo’s variation coefficient Varco (rate-of-speech-normalised standard deviation of durations) [10].

Duration Features (DUR): In order to include temporal features specifically modelling speech tempo, we created another set of features comprising *articulation rate* (the number of syllables divided by the total length of the speech segments) and the *percentage of pauses* (excluding initial and final silences in each recording). These features were computed using the segmentation of pauses, vowels, consonants, and speaker noise derived from the PR.

4. Experiments and Results

In this section we describe the results of the acoustic modelling of our database using the feature described in Section 3. Additionally, we also present the results from various combinations of the individual feature sets using early and late fusion approaches. In all our experiments we employ Support Vector Regression (SVR) [11] with linear kernel and Sequential Minimal Optimization learning (SMO) [12]. In all experiments (implemented with *sklearn* [13] and *LibSVM* [14]), the performance of the models was estimated using a 10-fold speaker-independent cross-validation scheme (SICV), i. e., the train and test set of each fold was disjoint w. r. t. speakers. The SVR complexity parameter C was optimised in each fold using an inner 10-fold SICV on the training set of that fold. The SVR performance was estimated using Spearman’s rank correlation coefficient between target and predicted values on all test instances (ρ). The details and results of our tests are shown in Tables 3 (individual features sets) and 4 (fusion). The best results using individual features sets were obtained with the *PROS* ($\rho = .600$), *oS_PROS* ($\rho = .542$) and *DUR* ($\rho = .526$) feature set. This confirms our initial expectations (see Section 1) regarding the relevance of prosodic features (including speech tempo) to the perception of naturalness of L2 speech. In relation to the fusion tests (see Table 4), we evaluated various combinations of the individual sets. Chiefly, we found that (1) the late fusion approach performance was globally better than early fusion (av. $\rho(\text{early}) = .562$; av. $\rho(\text{late}) = .586$); (2) the best performance was obtained with the late fusion of all feature sets except *oS_SPEC*; (3) the top performance is only slightly higher (.017 increase in ρ) than the performance of the best individual set *PROS* ($\rho = .600$).

Table 3: Regression experiments results for the EWE gold standard using individual features sets.

Feature set	# features	ρ
<i>oS_PROS</i>	483	<u>0.542</u>
<i>oS_SPEC</i>	5500	0.514
<i>oS_VQ</i>	390	0.513
<i>DUR</i>	2	<u>0.526</u>
<i>PROS</i>	3072	0.600
<i>RHYT</i>	13	0.411

5. Discussion and Conclusions

In this paper, we presented a new database with audio recordings of non-native (L2) speakers of English, and the perceptual evaluation experiment conducted with native English speakers for assessing the prosody of each recording. These annotations were used to compute the gold standard using different methods, and a series of regression experiments to evaluate their impact on the performance of a regression model predicting the degree of naturalness of L2 speech. To this end, we compared the relevance of different feature groups modelling prosody in general (without speech tempo), speech rate and pauses modelling speech tempo (fluency), speech rhythm, voice quality, and a variety of spectral features. Overall, using a speaker-independent schema, we have demonstrated that it is feasible to predict the level of naturalness of L2 speech using acoustic features alone. More specifically, we corroborated the findings from [15, 16, 17] that specialised prosodic features (including speech tempo) seem to be the relevant ones to assess the natural-

Table 4: Regression experiments results for the EWE gold standard using different combinations of the individual feature sets (early and late fusion approaches).

Feature set	ρ	
	early fusion	late fusion
<i>oS_PROS+oS_SPEC</i>	0.508	0.544
<i>oS_PROS+oS_VQ</i>	0.551	0.552
<i>oS_PROS+DUR</i>	0.547	0.571
<i>oS_PROS+PROS</i>	0.600	<u>0.602</u>
<i>oS_PROS+RHYT</i>	0.544	0.573
<i>oS_SPEC+oS_VQ</i>	0.519	0.543
<i>oS_SPEC+DUR</i>	0.515	0.566
<i>oS_SPEC+PROS</i>	0.576	<u>0.603</u>
<i>oS_SPEC+RHYT</i>	0.513	0.546
<i>oS_PROS+DUR</i>	0.526	0.569
<i>oS_VQ+PROS</i>	<u>0.607</u>	<u>0.606</u>
<i>oS_VQ+RHYT</i>	<u>0.607</u>	<u>0.567</u>
<i>DUR+PROS</i>	0.600	<u>0.604</u>
<i>DUR+RHYT</i>	0.554	0.565
<i>PROS+RHYT</i>	0.599	<u>0.604</u>
<i>all</i>	0.572	0.616
<i>all-oS_PROS</i>	0.575	0.614
<i>all-oS_SPEC</i>	0.608	0.617
<i>all-oS_VQ</i>	0.573	0.612
<i>all-DUR</i>	0.572	0.608
<i>all-PROS</i>	0.522	0.599
<i>all-RHYT</i>	0.572	0.604

ness of English produced by non-native speakers. The question remains which other features and methods could be employed in addition, and which might be the ‘empirical’ upper baseline for the performance. For the upper baseline, we computed the averaged gold standard per speaker, under the simplifying assumption that speaker performance is constant within the same recording session. The correlation with the EWE gold standard was $\rho=0.823$. When we employ the default cross-validation procedures of toolboxes such as WEKA or sklearn, we obtain (with 10 folds and late fusion of all six feature groups) $\rho=0.715$. However, these procedures are not speaker-independent. These two figures indicate the range of improvement that we can end up with when employing (1) item specific information obtained via speech recognition and (2) speaker modelling.

6. References

- [1] D. Van Compernelle, “Recognizing speech of goats, wolves, sheep and non-natives,” *Speech Communication*, vol. 35, no. 1, pp. 71–79, 2001.
- [2] F. Eyben, F. Weninger, F. Groß, and B. Schuller, “Recent Developments in openSMILE, the Munich Open-Source Multimedia Feature Extractor,” in *Proceedings of the 21st ACM International Conference on Multimedia, MM 2013*, ACM, Barcelona, Spain: ACM, October 2013, pp. 835–838.
- [3] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2013. [Online]. Available: <http://www.R-project.org/>
- [4] Z. Zhang and K.-H. Yuan, *coefficientalpha: Robust Cronbach’s Alpha and McDonald’s Omega with Missing and Non-normal Data*, 2014, r package version 0.3.3. [Online]. Available: <http://CRAN.R-project.org/package=coefficientalpha>
- [5] M. Grimm and K. Kroschel, “Evaluation of natural emotions using self assessment manikins,” in *Proc. ASRU*. IEEE, 2005, pp. 381–385.

- [6] 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 *Proc. INTERSPEECH 2013*, ISCA. Lyon, France: ISCA, August 2013, pp. 148–152.
- [7] F. Weninger, F. Eyben, B. W. Schuller, M. Mortillaro, and K. R. Scherer, "On the Acoustics of Emotion in Audio: What Speech, Music and Sound have in Common," *Frontiers in Psychology*, vol. 4, pp. 1–12, 2013.
- [8] E. Grabe and E. L. Low, "Durational variability in speech and the rhythm class hypothesis," *Papers in laboratory phonology*, vol. 7, no. 515-546, 2002.
- [9] F. Ramus, "Acoustic correlates of linguistic rhythm: perspectives," in *Proc. of Speech Prosody*. Aix-en-Provence, France: Laboratoire Parole et Langage, 2002, pp. 115–120.
- [10] V. Dellwo, "Influences of speech rate on the acoustic correlates of speech rhythm: An experimental phonetic study based on acoustic and perceptual evidence," *Universität Bonn, Bonn University, Bonn, Germany*, 2010.
- [11] A. Smola and B. Schölkopf, "A tutorial on support vector regression," *Statistics and Computing*, vol. 14, no. 3, pp. 199–222, 2004.
- [12] J. C. Platt, "Fast training of support vector machines using sequential minimal optimization," in *Advances in kernel methods – support vector learning*. MIT Press, 1999, pp. 185–208.
- [13] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [14] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011, software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [15] F. Hönig, A. Batliner, K. Weilhammer, and E. Nöth, "Automatic assessment of non-native prosody for English as L2," in *Proc. Speech Prosody*, Chicago, 2010, no pagination.
- [16] F. Hönig, A. Batliner, and E. Nöth, "Automatic assessment of non-native prosody – annotation, modelling and evaluation," in *Proceedings of IS-ADEPT, International Symposium on Automatic Detection of Errors in Pronunciation Training, June 6-8, Stockholm, Sweden*, 2012.
- [17] F. Hönig, T. Bocklet, K. Riedhammer, A. Batliner, and E. Nöth, "The automatic assessment of non-native prosody: Combining classical prosodic analysis with acoustic modelling," in *Proc. INTERSPEECH 2012*, Portland (Oregon), USA, September 2012, pp. 823–826.