

Automatic Prosodic Variations Modeling for Language and Dialect Discrimination

Jean-Luc Rouas

Abstract—This paper addresses the problem of modeling prosody for language identification. The aim is to create a system that can be used prior to any linguistic work to show if prosodic differences among languages or dialects can be automatically determined. In previous papers, we defined a prosodic unit, the pseudosyllable. Rhythmic modeling has proven the relevance of the pseudosyllable unit for automatic language identification. In this paper, we propose to model the prosodic variations, that is to say model sequences of prosodic units. This is achieved by the separation of phrase and accentual components of intonation. We propose an independent coding of those components on differentiated scales of duration. Short-term and long-term language-dependent sequences of labels are modeled by n-gram models. The performance of the system is demonstrated by experiments on read speech and evaluated by experiments on spontaneous speech. Finally, an experiment is described on the discrimination of Arabic dialects, for which there is a lack of linguistic studies, notably on prosodic comparisons. We show that our system is able to clearly identify the dialectal areas, leading to the hypothesis that those dialects have prosodic differences.

Index Terms—Automatic language identification (ALI), prosody, read and spontaneous speech.

I. INTRODUCTION

THE STANDARD approach to automatic language identification (ALI) considers a phonetic modeling as a front-end. The resulting sequences of phonetic units are then decoded according to language-specific phonotactic grammars [1].

Other information sources can be useful to identify a language, however. Recent studies (see [2] for a review) reveal that humans use different levels of perception to identify a language. Three major kinds of features are employed: segmental features (acoustic properties of phonemes), suprasegmental features (phonotactics and prosody), and high-level features (lexicon). Besides acoustics, phonetics, and phonotactics, prosody is one of the most promising features to be considered for language identification, even if its extraction and modeling are not a straightforward issue.

In the NIST 2003 Language Verification campaign, most systems used acoustic modeling, using Gaussian mixture models adapted from universal background models (a technique derived from speaker verification [3]), and/or phonotactic modeling (parallel phone recognition followed by language modeling (PPRLM), see [1] and [4]). While these techniques gave

the best results, systems using prosodic cues have also been investigated, following research in speaker recognition [5], notably Adami's system [6]. More recently, systems using syllable-scale features have been under research, although their aim is to model acoustic/phonotactic properties of languages [7] or also prosodic cues [8].

Beside the use of prosody to improve the performances of ALI systems, we believe that there is a real linguistic interest in developing an automatic language identification system using prosody and not requiring any *a priori* knowledge (e.g., manual annotations). Hence, we will have the possibility of testing if prosodically unstudied languages can be automatically differentiated. The final aim of our studies is to automatically describe prosodic language typologies.

In this paper, we will describe our prosodic-based language identification system. In Section II, we will recall the main linguistic theories about differences among languages. After reviewing the linguistic and perceptual elements that demonstrate the interest of prosody modeling for language identification, we will address the problem of modeling.

Indeed, modeling prosody is still an open problem, mostly because of the suprasegmental nature of the prosodic features. To address this problem, automatic extraction techniques of sub-phonemic segments are used (Section III). After an activity detection and a vowel localization, a prosodic syllable-like unit adapted to language identification is characterized. Results previously obtained by modeling prosodic features extracted on this unit are briefly described at the end of Section III. We believe, however, that prosodic models should take into account temporal fluctuations, as prosody perception is mostly linked to variations (in duration, energy, and pitch). That is why we propose a prosody coding which enables to consider the sequences of prosodic events in the same way as language models are used to model phone sequences in the PPRLM approach. The originality of our method lies in differentiating phrase accent from local accent, and modeling them separately. The method is described in Section IV. The system is first tested on databases with languages that have known prosodic differences to assess the accuracy of the modeling. These experiments are carried out on both read (Section V) and spontaneous (Section VI) speech, using, respectively, the MULTTEXT [9] and OGI-MLTS [10] corpora. Then, a final experiment is performed on Arabic dialects (Section VII), for which we investigate if prosodic differences between those dialects can be automatically detected.

II. PROSODIC DIFFERENCES AMONG LANGUAGES

The system described in this paper aims at determining to what extent languages are prosodically different. In this section, we will describe what the rhythmic and intonational properties of languages are and how humans perceive them.

Manuscript received November 15, 2006; revised March 12, 2007. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Simon King.

The author is with the Laboratório de Sistemas de Língua Falada (L2F), Instituto de Engenharia de Sistemas e Computadores (INESC), 1000-029 Lisbon, Portugal (e-mail: rouas@l2f.inesc-id.pt).

Digital Object Identifier 10.1109/TASL.2007.900094

A. Rhythm

The rhythm of languages has been defined as an effect involving the isochronous (that is to say at regular intervals) recurrence of some type of speech unit [11]. Isochrony is defined as the property of speech to organize itself in pieces equal or equivalent in duration. Depending on the unit considered, the isochrony theory classifies languages in the following three main sets:

- stress-timed languages (as English and German);
- syllable-timed languages (as French and Spanish);
- mora-timed languages (as Japanese).

Syllable-timed languages share the characteristic of having regular intervals between syllables, while stress-timed languages have regular intervals between stressed syllables, and for mora-timed languages, successive mora are quasi-equivalent in terms of duration. This point of view was made popular by Pike [12] and later by Abercrombie [13]. According to them, distinction between stress-timed and syllable-timed languages is strictly categorical, since languages cannot be more or less stress or syllable-timed. Despite its popularity among linguists, the rhythm class hypothesis is contradicted by several experiments (notably by Roach [14] and Dauer [15]). This forced some researchers (e.g., Beckman [16]) to shift from “objective” to “subjective” isochrony. True isochrony is described as a constraint, and the production of isochronous units is perturbed by phonetic, phonologic, and grammatical rules of the languages. Some other researchers have concluded that isochrony is mainly a perceptual phenomenon (e.g., Lehiste [17]). Isochrony can then be seen as a concept relative to speech perception.

The lack of an empirical proof of isochrony led Dauer [15] to propose a new rhythmic classification system. From her point of view, speakers do not try to keep equal interstress or intersyllabic intervals, but languages are more or less stress or syllable-timed. Nespore [18] introduced the notion of rhythmically intermediate languages, which share properties associated with stress-timed languages and other associated with syllable-timed languages. As an example, she cites Polish—classified as stress-timed, although it does not have vocalic reduction, and Catalan—syllable-timed but having vocalic reduction.

B. Intonation

The following three main groups of languages can be characterized regarding their use of intonation:

- tone languages (as Mandarin Chinese);
- pitch-accent languages (as Japanese);
- stress-accent languages (as English and German).

According to Cummins [19], distinguishing languages using fundamental frequency alone had a moderate success. The explanation is twofold.

- On the one hand, we can imagine a discrimination based on the use of lexical tone (Mandarin) or not (English), but intermediate cases exist (Korean dialects) which are usually considered as representing transitory states between languages of one class and those of another [20].
- On the other hand, phenomena linked to accents and intonation are less easy to handle. There are multiple theories on utterance intonation that do not agree. The situation is

made more complex by studies on the nonlinguistic uses of intonation, as for example to express emotions. Several studies agree on a classification by degrees rather than separate classes [21].

C. Perception

Over the last few decades, numerous experiments have shown the human capability for language identification [2]. The following three major kinds of cues help humans to identify languages:

- 1) segmental cues (acoustic properties of phonemes and their frequency of occurrence);
- 2) suprasegmental cues (phonotactics, prosody);
- 3) high-level features (lexicon, morpho-syntax).

About prosodic features, several perceptual experiments try to shed light on human abilities to distinguish languages keeping only rhythmic or intonation properties. The method is to degrade speech recordings by filtering or resynthesis to remove all segmental cues to the subjects whose task is to identify the language. The subjects are either naive or trained adults, infants or newborns, or even nonhuman primates. For example, all the syllables are replaced by a unique syllable “/sa/” in Ramus’ experiments [22]. In other cases, processing of speech through a low-pass filter (cutoff frequency 400 Hz) is used to degrade the speech signal [23]. Other authors [24] propose different methods to degrade the speech signal in order to keep only the desired information (intensity, intonation, or rhythm). From a general point of view, all those experiments show the notable human capacity to identify to some extent foreign languages after a short period of exposure.

III. SEGMENTATION, VOWEL DETECTION, AND PSEUDOSYLLABLES

The starting point of recent work on rhythm is the rhythm description method proposed by Ramus *et al.* in 1999 [22]. Following him, others have proposed different and more complex rhythm modeling methods (e.g., [11] and [25]). The weak point in many of those approaches is that they have only been applied after a manual segmentation of the speech signal. Consequently, their performance has only been assessed on relatively small corpora.

To overcome this limitation and to model the prosody of languages automatically, we use automatic processing to extract prosodic information. The following three baseline procedures lead to relevant consonant, vocalic, and silence segment boundaries.

- Automatic speech segmentation leading to quasistationary segments: This segmentation results from the “Forward–Backward Divergence” (DFB) algorithm [26], which is based on a statistical study of the signal in the temporal framework. The segmentation achieves a sub-phonemic segmentation, where segments correspond to steady or transient parts of phonemes.
- Vocal activity detection: The vocal activity detection is based on a first-order statistical analysis of the temporal signal [27]. The activity detection algorithm detects the less intense segments of the excerpt (in terms of energy), and the other segments are classified as Silence or Activity according to an adaptive threshold.




Fig. 1. Result of the automatic segmentation, vowel location, and vocal activity detection on a French recording of the MULTEXT database. The sentence is: “les choux avaient été entièrement dévorés par les limaces.” “#” labels are for nonactivity segments, “V” are for vocalic segments, and other segments are labeled “C.” Word boundaries are displayed for illustration purposes.




Fig. 2. Pseudosyllable segmentation performed for the sentence “les choux avaient été entièrement dévorés par les limaces.” Consecutive “C” segments are gathered until a “V” segment is found.

- **Vowel localization:** The vowel location algorithm is based on a spectral analysis (see [27] for more details). The fact that neither labeled data nor supervised learning are necessary constitutes the main advantage of this algorithm. The fact that no learning phase is necessary allows the algorithm to be used on different languages, even if no hand-labeled data is available. However, the consequence is that the algorithm is not optimized for any language even if it behaves correctly when compared to other systems [28].

This front-end processing results in a segmentation into vocalic, consonantal and silence segments. Labels “V,” “C,” or “#” are used to qualify each segments (Fig. 1).

The syllable is a privileged unit for rhythm modeling. However, automatic extraction of syllables (in particular, boundary detection) is a controversial operation that is still in debate among phonologists: the pronunciation quality and the speech rate are factors influencing directly the syllable segmentation [29]. Furthermore, segmenting the speech signal in syllables is a language-specific task [30]. No language-independent algorithm can be easily applied.

We therefore used the notion of pseudosyllable [31]. The basic idea is to articulate the prosodic unit around primordial elements of the syllables—vowels—and to gather the neighboring consonants around those nuclei. We have decided to gather only the preceding consonants. This choice is explained by the fact that syllable boundary detection is not an easy task in a multilingual framework, and that the most frequent syllables correspond to the consonant/vowel structure [15]. An example of this segmentation is shown in Fig. 2.

We showed in previous papers [31] that the pseudosyllable segmentation can be successfully used for language identification. Features characterizing durations and fundamental frequency variations are extracted from each pseudosyllable and are used to learn the parameters of Gaussian mixtures for each language of the database.

With the duration features, the correct identification rate is 67% for the seven languages of the MULTEXT corpus. This result is obtained using a mixture of eight Gaussians. With the intonation features, the correct identification rate is 50%, using a GMM with eight components. A fusion of the results obtained with both duration and intonation features allows to reach

70% correct. The confusions occur mainly across languages belonging to the same groups evoked in linguistic theories.

Identifying stress-timed languages (English, German, and hypothetically Mandarin) versus syllable-timed languages (French, Italian, and Spanish) versus mora-timed languages (Japanese), accuracy is 91% correct using both duration and intonation models.

Nevertheless, the statistical models (Gaussian mixture models) we use to model pseudosyllabic features are intrinsically static models. This does not fit with the perceptive reality of prosody which is continue, dynamic, and suprasegmental. We must use different models to take the temporal variations into account.

IV. MODELING PROSODIC VARIATIONS

Following Adami’s work [6], we used the features computed on each pseudosyllable to label the fundamental frequency and energy trajectories. Two models are used to separate the long-term and short-term components of prosody. The long-term component characterizes prosodic movements over several pseudosyllables, while the short-term component represents prosodic movements inside a pseudosyllable. An overview of the system is displayed in Fig. 3. Fundamental frequency and energy are extracted from the signal using the SNACK Sound toolkit [32].

A. Fundamental Frequency Coding

The fundamental frequency processing is divided into two phases, representing the phrase accentuation and the local accentuation, as in Fujisaki’s work [33].

1) *Baseline Computing and Coding:* The baseline extraction consists in finding all the local minima of the F_0 contour, and linking them. Then, the baseline is labeled in terms of U(p), D(own), and #(silence or unvoiced).

To find local minima, the sound file is automatically divided into “sentences,” defined here as intervals between silent segments of duration over 250 ms. The linear regression of the F_0 curve is then computed on each sentence. Then each part of the curve under the linear regression is used to find a unique minima (Fig. 4).

Successive minima are linked by linear approximation. An example of a resulting baseline curve is displayed in Fig. 5. The slope of the regression is used to label the baseline. We use one

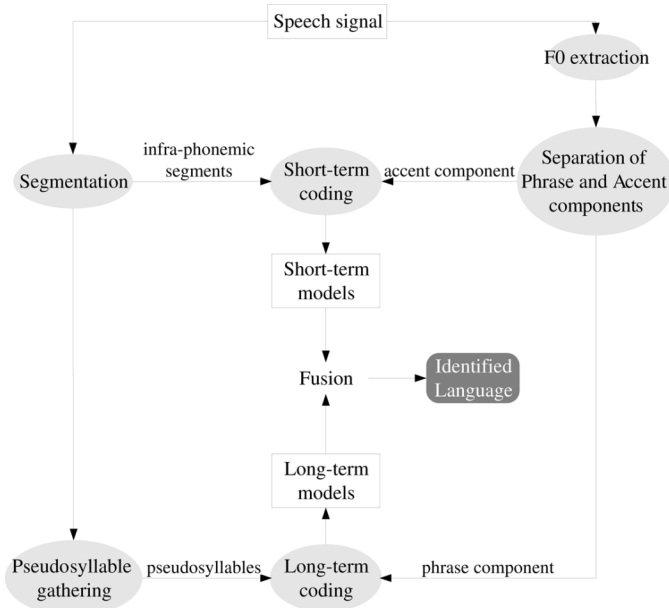


Fig. 3. Description of the system.

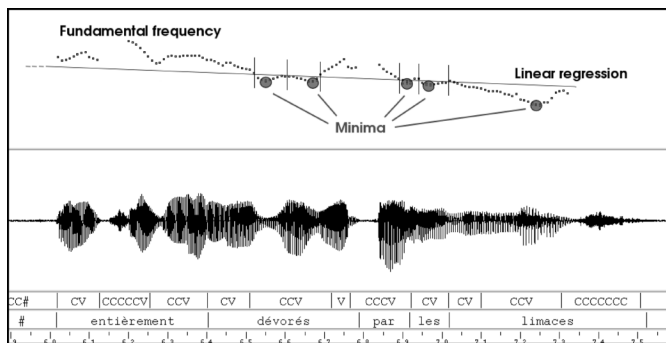


Fig. 4. Finding local minima for the sentence “les choux avaient été entièrement dévorés par les limaces.” Local minima under the linear regression are displayed as a big dot.

label per pseudosyllable. Labels are “U” for a positive slope and “D” for a negative slope. Unvoiced pseudosyllables (less than 70% voiced in duration) are labeled “#.” In this example, the label sequence corresponding to the sentence is

U. U. U. U. D. D. D. D. D. D. #.

2) *Residue Approximation and Coding:* The baseline is subtracted from the original contour. The resulting curve is called residue (Fig. 6). This residue is then approximated for each segment by a linear regression. The slope of the linear regression is used to label the F_0 movement on the unit, according to three available labels (Up, Down, and Silence). The example sentence is labeled in the following way:

D. D. #. #. #. #. #. D. D. U. D. D. D. D. U. U. D. U.
#. #. D. D. U. D. U. D. D. U. D. #. #. #. #. #. #.

B. Energy Coding

The energy curve is approximated by linear regressions for each considered units (subphonemic segments or pseudosylla-

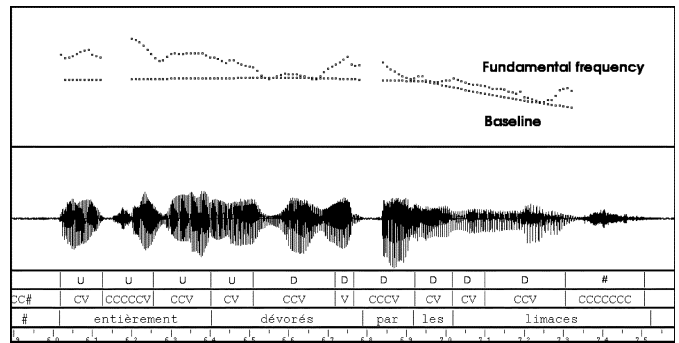


Fig. 5. Extraction of the baseline for the sentence “les choux avaient été entièrement dévorés par les limaces.” Previously found local minima are linked with a straight line.




Fig. 6. Approximation of the residue for the sentence “les choux avaient été entièrement dévorés par les limaces.” The residue is approximated by a linear regression for each segment.




Fig. 7. Approximation of the energy for the sentence “les choux avaient été entièrement dévorés par les limaces.” The energy is approximated by a linear regression for each segment.

bles) (Fig. 7). The process is the same as the one used for the residue coding. As there is no segment with no energy, only two labels are used: Up and Down. In this example, the sequence of labels is

U. D. D. U. U. U. D. U. D. U. D. U. U. D. U. D. D.—
D. U. U. D. D. U. D. U. D. D. D. U. U. D. D. D. D.

C. Duration Coding

Duration labels are computed only for the subphonemic segment units. The labels are assigned considering the mean duration of each kind of segment (vocalic, consonantic, or silence). If the segment to label is a vocalic segment with a duration above

the mean vocalic segments duration computed on the learning part of the corpus, it is labeled “l” (long). If the current vocalic segment duration is below the mean, the “s” (short) label is then used. The duration labels generated on the example sentence are

s.l.s.s.s.s.s.s.s.l.s.l.s.s.l.s.—
s.s.s.l.s.l.s.s.s.s.l.s.s.s.s.s.s.

D. Modeling

To model the prosodic variations, we use classical n-gram language modeling provided by the SRI language modeling toolkit [34]. For each system—long- and short-term—each language is modeled by a n-gram model during the learning procedure. During the test phase, the most likely language is picked according to the model which provides the maximum likelihood. For the long-term models, this modeling is applied at the pseudosyllable level, and n-grams are learnt using baseline labels, eventually combined with energy labels coded at the pseudosyllable scale. The short-term models are learnt using the subphonemic segmentation, and using the residue labels, optionally combined with energy and duration labels. For each segment, the label is then composed of three symbols. For the example sentence, we have

DU_s.DD_l.#D_s.#U_s.#U_s.#U_s.#D_s.DU_s.—
DD_s.UU_s.DD_l.DU_s.DU_l.DD_s.UU_s.UD_l.—
DD_s.UD_s.#U_s.#U_s.DU_l.DD_s.UD_l.DU_s.—
UD_s.DU_s.DD_s.UD_l.DD_s.#U_s.#U_s.#D_s.—
#D_s.#D_s.#D_s.

Several lengths for the n-gram models have been tested (from 3- to 5-grams), but as the best results are obtained with 3-grams, only the results obtained using 3-gram models are displayed.

V. EXPERIMENTS ON READ SPEECH

The first experiments are made on the MULTTEXT corpus [9] (a subset of EUROM1), which contains five languages (English, French, German, Italian, and Spanish), and ten speakers per language, balanced between male and female. We have ten 20-s files per speaker. This baseline corpus has been extended with recordings of Japanese speakers [35]. Mandarin Chinese recordings are also added to the original corpus, thanks to Komatsu [24].

The three theoretical rhythmic classes are represented in this corpus: English, German, and Mandarin Chinese are stress-timed languages; French, Italian, and Spanish are syllable-timed languages, and Japanese is a mora-timed language. Moreover, Mandarin Chinese is a tone language, and Japanese is a pitch-accent language.

For the learning phase, we used eight speakers (four for Japanese), and two (one male and one female) were used for the tests. We have 80 learning files per language (≈ 26 min) and 20 test files per language (≈ 6 min)—except for French (19 files)—resulting in a total of 139 test files.

A. Long-Term Modeling

The sequences of labels computed on each pseudosyllable are modeled by n-gram models. We investigated different combinations of labels to see which features are more useful (Table I).

TABLE I
EXPERIMENTS WITH LONG-TERM LABELS (139 FILES)

Features	# Labels	Model	% correct
Baseline	3	3-grams	50.3% \pm 8.3
Energy	3	3-grams	26.6% \pm 7.4
Baseline + Energy	9	3-grams	29.4% \pm 7.6

TABLE II
LONG-TERM PROSODIC MODEL [CORRECT= 50.3 \pm 8.2% (70/139 FILES)].
RESULTS ARE DISPLAYED AS PERCENTAGES

ref \ id	Eng	Ger	Man	Fre	Ita	Spa	Jap
Eng	25	-	-	35	5	20	15
Ger	5	85	-	5	-	-	5
Man	5	5	10	5	25	25	25
Fre	-	-	-	69	5	26	-
Ita	5	-	15	5	40	20	15
Spa	-	-	-	-	15	80	5
Jap	5	10	-	5	15	20	45

TABLE III
EXPERIMENTS WITH SHORT-TERM LABELS (139 FILES)

Features	# labels	Model	% correct
Duration	2	3-grams	39.6% \pm 7.9
Residue	3	3-grams	52.5% \pm 8.2
Energy	3	3-grams	56.8% \pm 8.3
Residue+Energy+Duration	18	3-grams	79.8% \pm 6.6

TABLE IV
SHORT-TERM PROSODIC MODEL [CORRECT= 79.8 \pm 6.6% (111/139 FILES)].
RESULTS ARE DISPLAYED AS PERCENTAGES

ref \ id	Eng	Ger	Man	Fre	Ita	Spa	Jap
Eng	65	-	15	-	10	10	-
Ger	-	95	5	-	-	-	-
Man	-	20	80	-	-	-	-
Fre	-	-	-	90	5	5	-
Ita	10	-	-	-	85	5	-
Spa	-	-	-	-	45	55	-
Jap	-	-	-	-	10	-	90

The best performance is obtained using only the baseline labels. With these labels, the correct identification rate is 50%. The confusion matrix is drawn in Table II with identified languages in columns and references in rows (confidence intervals are provided in the legend). Table II show that French, Spanish, and German are clearly identified.

B. Short-Term Modeling

The sequences of labels computed on each subphonemic segment are also modeled by n-gram models. Contributions of different features are shown in Table III. The best results are obtained using the combination of residue, energy, and duration labels. The identification rate is 80%.

Detailed results using this configuration are displayed in Table IV. These results allow us to hypothesize that the most characteristic prosodic elements of languages are not pseudosyllable sequences but sequences of elements constituting them.

C. Merging Long and Short-Term Components

The merging of the two systems described previously is addressed. The merging technique is a simple addition of the log-likelihoods. The identification rate obtained with this method is

TABLE V
MERGING OF SHORT- AND LONG-TERM MODELS [CORRECT= 83.5 ± 6.1% (116/139 FILES)]. RESULTS ARE DISPLAYED AS PERCENTAGES

ref \ id	Eng	Ger	Man	Fre	Ita	Spa	Jap
Eng	70	-	15	-	5	10	-
Ger	-	95	5	-	-	-	-
Man	-	-	100	-	-	-	-
Fre	-	-	-	95	5	-	-
Ita	10	-	-	-	80	5	5
Spa	-	-	-	-	45	55	-
Jap	-	-	-	-	10	-	90

TABLE VI
RHYTHM CLASSES IDENTIFICATION TASK [CORRECT= 93.5 ± 4.1% (130/139 FILES)]. RESULTS ARE DISPLAYED AS PERCENTAGES

reference \ identified	Stress-timed	Syllable-timed	Mora-timed
Stress-timed	90	10	-
Syllable-timed	2	98	-
Mora-timed	10	-	90

83%, which is considerably better than the 70% obtained with static modeling (Section III).

Despite the poor performance of the long-term model, merging the long- and short-term results allows to improve the identification rate. Results (in Table V) show that most languages are well identified.

The language classes evoked in Section II have an influence on the correct identification rates: Mandarin (i.e., the only tone language of the corpus) is the most clearly identified language. Japanese, the only mora-timed language, and the only pitch-accent language in our corpus, is also well identified. The main confusions are between Italian and Spanish, both belonging to the syllable-timed and stress-accent groups. Considering rhythmic classes (represented in different strength of grey in the matrix), we can see that most confusions are within languages of the same rhythmic family.

Consequently, a rhythmic classes identification experiment was performed using the same data with the fusion of the short and long-term models. English, German and Mandarin are part of the stress-timed group; French, Italian, and Spanish are in the syllable-timed group, and Japanese constitutes in itself the mora-timed group. The rhythmic classes identification rate is 94% (Table VI).

On this data, our system manages well to classify languages according to their prosodic properties.

VI. EXPERIMENTS ON SPONTANEOUS SPEECH

The same experiments have been made on a spontaneous speech corpus, the OGI Multilingual Telephone Speech Corpus (OGI-MLTS) corpus [10]. This corpus is composed of telephone speech recorded in ten languages: English, Farsi, French, German, Japanese, Korean, Mandarin, Spanish, Tamil, and Vietnamese. Experiments are made on six languages (English, French, German, Japanese, Mandarin, and Spanish) using spontaneous speech utterances of 45 s each, one file per speaker. The data organization is displayed in Table VII.

The tests here are made using the same tuning as for the read speech experiments—same features used (baseline labels for the long-term model, and the combination of residue, energy, and

TABLE VII
DATA ORGANIZATION SUMMARY—OGI

Language	Learning part			Test part		
	# files	Duration (s.)		# files	Duration (s.)	
		Total	Mean		Total	Mean
English	89	2456.69	27.60	18	878.09	48.78
French	88	2438.94	27.71	15	692.29	46.15
German	88	2469.20	28.05	19	909.64	47.87
Japanese	90	2492.78	27.69	19	884.72	46.56
Mandarin	75	1952.62	26.03	18	805.66	44.75
Spanish	88	2553.12	29.01	13	618.14	47.55
Overall	518	14363.30	27.73	102	4788.57	46.94

TABLE VIII
LONG-TERM PROSODIC MODELS [CORRECT= 20.6 ± 8.0% (21/102 FILES)]. RESULTS ARE DISPLAYED AS PERCENTAGES

ref \ id	Eng	Ger	Man	Fre	Spa	Jap
Eng	33	27	6	28	-	6
Ger	26	11	16	21	16	10
Man	28	-	11	17	5	39
Fre	20	20	7	20	20	13
Spa	15	7	8	8	31	31
Jap	26	16	5	16	16	21

TABLE IX
SHORT-TERM PROSODIC MODELS [CORRECT= 40.2 ± 9.5% (41/102 FILES)]. RESULTS ARE DISPLAYED AS PERCENTAGES

ref \ id	Eng	Ger	Man	Fre	Spa	Jap
Eng	55	17	6	5	17	-
Ger	42	37	-	11	10	-
Man	17	33	33	11	-	6
Fre	33	33	-	20	14	-
Spa	38	-	-	16	38	8
Jap	10	10	-	11	16	53

TABLE X
RHYTHM CLASSES IDENTIFICATION TASK [CORRECT= 68.6 ± 9.1% (70/102 FILES)]. RESULTS ARE DISPLAYED AS PERCENTAGES

reference \ identified	Stress-timed	Syllable-timed	Mora-timed
Stress-timed	85	11	4
Syllable-timed	50	46	4
Mora-timed	16	31	53

duration labels for the short-term model), same n-gram configuration (3-grams). Results are displayed in Tables VIII and IX for the long- and short-term models, respectively.

The long-term model only manages to reach performances slightly better than chance (on a six-way classification, the chance level is ≈ 16%).

The short-term model achieves a performance of 40% of correct identifications. The best recognized languages are English (55%) and Japanese (53%). We can observe that German is mainly confused with English.

The performances of the system for the rhythm classes identification task is given in Table X.

The global identification rate is 69%. The identification rates for each class are, respectively, 85% for the stress-timed languages, 46% for the syllable-timed languages, and 53% for the mora-timed language. As expected from language identification results, syllable-timed languages are the least recognized. The identification mistakes are made principally on the French language, responsible for ten errors—that is to say 66% of the

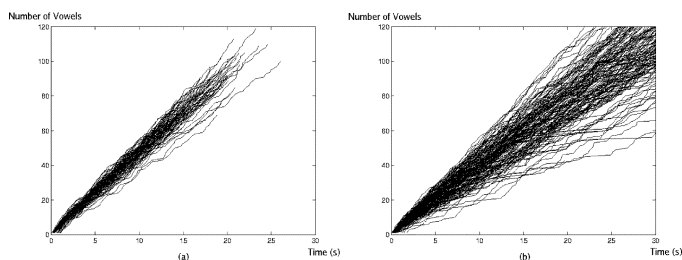


Fig. 8. (a) Number of detected vowels as a function of time on English read speech. (b) Number of detected vowels as a function of time on English spontaneous speech.

TABLE XI

STANDARD DEVIATIONS OF SPEECH RATE (ESTIMATION) FOR DIFFERENT LANGUAGES ON READ AND SPONTANEOUS SPEECH

Lang.	Eng	Fre	Ger	Ita	Jap	Man	Spa	All
Read	0.36	0.35	0.37	0.56	0.42	0.49	0.40	0.50
Spont.	0.77	0.93	0.80	n.d.	0.98	0.77	0.81	0.85

syllable-timed languages identification error. This may be explained by the different varieties of French (Canadian and European) encountered in the corpus.

It is, however, quite difficult to directly transpose our approach, designed on read speech, to spontaneous speech. The main reason may be the intrinsic variability of the spontaneous data. This variability can be linked to the great variations of speech rate observed within each language (see [36] for a study of automatic speech rate estimation on read and spontaneous speech). Fig. 8 displays the accumulated number of vowels per second for the English parts of the MULTEXT and OGI-MLTS corpus. Spontaneous speech leads to much more variation in terms of number of vowels per second than read speech. Table XI shows the standard deviations of the speech rate approximation on spontaneous and read speech for different languages (45-s spontaneous speech utterances of the OGI-MLTS corpus and 20-s read utterances from the MULTEXT corpus). These standard deviation values show that there is a much greater dispersion of spontaneous speech rate within each language than for read speech. These intralanguage variations can explain why our models achieve a poor performance on the test part of the corpus.

VII. EXPERIMENTS ON ARABIC DIALECTS

We believe nonetheless that our system can be applied, at least on groups of languages, to investigate if the groups considered can be automatically distinguished prosodically. We carried out experiments with different Arabic dialects to show how the features described in this paper help to classify them.

The corpus ARABER has been recorded at the *Dynamique du Langage* Laboratory, Lyon, France. It consists of semispontaneous recordings (comments on an image book) from 40 speakers from Maghreb, Middle-East, and an intermediate area (Tunisia, Egypt). The mean duration for each speaker's recording is 5 min, in 40 files of 7.6 s. All the data have been used both for learning and testing, via a cross-validation procedure.

The quality of the recordings made this corpus interesting because it is intermediate between the studio-recorded read speech

TABLE XII
ARABIC DIALECTAL AREAS IDENTIFICATION [CORRECT= $98.0 \pm 0.6\%$ (1563/1592 FILES)]. RESULTS ARE DISPLAYED AS PERCENTAGES

identified \ reference	Occidental	Intermediate	Oriental
Occidental	99.5	-	0.5
Intermediate	1.8	96.0	2.2
Oriental	1.0	0.3	98.7

corpus MULTEXT and the telephone spontaneous speech of the OGI corpus.

As research on the prosody of Arabic dialects is emerging, very few observations are available to hypothesize how much they differ. At this point, Arabic dialects are classified according to intercomprehension between speakers, resulting in three main areas: Occidental (Moroccan, Algerian), Intermediate (Tunisian, Egyptian), and Oriental (Lebanese, Jordanian, Syrian). The aim of this study is to see, with an ALI system that does not require any kind of annotation, if it is possible to identify the hypothesized dialectal areas using only prosody.

The prosodic ALI system—using the combination of the long and short-term models—has been applied to this corpus without any tuning procedure: same features (baseline labels for the long-term model and the combination of residue, energy, and duration features for the short-term model), same n-gram configuration. Experiments are made according to a cross-validation procedure applied on each speaker: learning is done using all speakers except one who is used for the test. This is repeated until all speakers have been used for the test. Results of the cross-validation tests are displayed in Table XII.

The system performs very well on this data, which shows that prosodic differences may be important between the hypothesized dialectal areas of Arabic. Further studies are needed in order to identify the exact prosodic differences between dialectal areas, and if there are differences among dialects of a same area.

VIII. CONCLUSION AND PERSPECTIVES

Experiments on read speech show that our system is able to automatically identify languages using prosody alone. Differences between languages seem more characterized by microprosodic events (short term) than macroprosodic ones. The experiments show that variations of fundamental frequency, energy, and segment duration that occur within a syllable are more characteristic. The dynamic modeling allows to reach 83% of correct identification on a seven language discrimination task. Results tend to confirm the existence of automatically identifiable rhythmic classes (accuracies above 90% for rhythm class identification).

Considering spontaneous speech, accuracies are lower. This can be partly explained by the greater interspeaker variability due to the number of different prosodic realizations allowed by spontaneous speech, especially in terms of changes in speech rate. The results remain, however, interesting for the language class identification experiment.

An applicative example has been shown with the experiments on Arabic dialects, where the speech quality is intermediate between read and spontaneous. This experiment has shown that there exist some automatically detectable prosodic differences between hypothesized dialectal areas of Arabic. A careful study

with the help of linguists is needed in order to define precisely what are those differences and where they appear. The next experiment will be to test if the dialects can be automatically clustered using our system.

The main advantage of our prosodic ALI system lies in the fact it does not require any manual annotations (especially phonetic annotations which are very time-consuming). Hence, the system can be directly applied to unknown data, and be used to evaluate if the prosodic differences between languages or dialects can be automatically detected. Using this system can help linguists to verify that further investigation is needed, leading to new collaborations between the automatic processing and the linguistic communities.

ACKNOWLEDGMENT

The author would like to thank F. Pellegrino for providing helpful comments on this paper and M. Barkat for her help on the Arabic dialects experiments.

REFERENCES

- [1] M. A. Zissman and K. M. Berkling, "Automatic language identification," *Speech Commun.*, vol. 35, no. 1–2, pp. 115–124, 2001.
- [2] M. Barkat and I. Vasilescu, "From perceptual designs to linguistic typology and automatic language identification: Overview and perspectives," in *Eurospeech*, Aalborg, Denmark, 2001, pp. 1065–1068.
- [3] D. A. Reynolds and R. C. Rose, "Robust text-independent speaker identification using Gaussian mixtures speaker models," *IEEE Trans. Speech Audio Process.*, vol. 3, no. 1, pp. 72–83, Jan. 1995.
- [4] T. J. Hazen and V. W. Zue, "Segment-based automatic language identification," *J. Acoust. Soc. Amer.*, pp. 2323–2331, 1997.
- [5] D. Reynolds, W. Andrews, J. Campbell, J. Navratil, B. Peskin, A. Adami, Q. Jin, D. Klusacek, J. Abramson, R. Mihaescu, J. Godfrey, D. Jones, and B. Xiang, "The supersid project: Exploiting high-level information for high-accuracy speaker recognition," in *Proc. ICASSP*, Hong Kong, China, Apr. 2003, vol. 4, pp. 784–787.
- [6] A. Adami, R. Mihaescu, D. A. Reynolds, and J. Godfrey, "Modeling prosodic dynamics for speaker recognition," in *Proc. ICASSP*, Hong Kong, China, 2003, vol. 4, pp. 788–791.
- [7] D. Zhu, M. Adda-Decker, and F. Antoine, "Different size multilingual phone inventories and context-dependent acoustic models for language identification," in *InterSpeech*, Lisbon, Portugal, Sep. 2005, pp. 2833–2836.
- [8] T. Matrin, B. Bake, E. Wong, and S. Sridharan, "A syllable-scale framework for language identification," *Comput. Speech, Lang.*, vol. 20, pp. 276–302, 2006.
- [9] E. Campione and J. Voéronis, "A multilingual prosodic database," in *Proc. ICSLP*, Sidney, Australia, 1998, pp. 3163–3166.
- [10] Y. K. Muthusamy, R. A. Cole, and B. T. Oshika, "The OGI multi-language telephone speech corpus," in *Proc. ICSLP*, Oct. 1992, pp. 895–898.
- [11] E. Grabe and E. L. Low, "Durational variability in speech and the rhythm class hypothesis," *Papers Lab. Phonol.*, vol. 7, pp. 515–546, 2002.
- [12] *The Intonation of American English*. P. Ladefoged, Ed. Ann Arbor, MI: Univ. Michigan Press, 1945.
- [13] *Elements of General Phonetics*. D. Abercrombie, Ed. Edinburgh, U.K.: Edinburgh Univ. Press, 1967.
- [14] P. Roach, "On the distinction between "stress-timed" and "syllable-timed" languages," *Linguist. Controversies*, pp. 73–79, 1982.
- [15] R. M. Dauer, "Stress-timing and syllable-timing reanalysed," *J. Phonetics*, vol. 11, pp. 51–62, 1983.
- [16] M. E. Beckman, "Evidence for speech rhythms across languages," in *Speech Perception, Production and Linguistic Structure*, Y. Tohkura, E. Vatikiotis-Bateson, and Y. Sagisaka, Eds. Amsterdam, The Netherlands: IOS Press, 1992, pp. 457–463.
- [17] I. Lehiste, "Isochrony reconsidered," *J. Phonetics*, vol. 5, pp. 253–263, 1977.
- [18] M. Nespó, "On the rhythm parameter in phonology," in *Logical Issues in Language Acquisition*, I. Roca, Ed. Berlin, Germany: de Gruyter, 1990, pp. 157–175.
- [19] F. Cummins, "Speech rhythm and rhythmic taxonomy," in *Proc. Speech Prosody*, Aix-en-Provence, France, 2002, pp. 121–126.
- [20] L. Hyman, "Word-prosodic typology," Tech. Rep. Univ. California Berkeley, Berkeley, CA, 2005.
- [21] D. R. Ladd and R. Morton, "The perception of intonation emphasis: Continuous or categorical?," *J. Phonetics*, vol. 25, pp. 313–342, 1997.
- [22] F. Ramus, M. Nespó, and J. Mehler, "Correlates of linguistic rhythm in the speech signal," *Cognition*, vol. 73, no. 3, pp. 265–292, 1999.
- [23] S. Frota, M. Vigarío, and F. Martins, "Language discrimination and rhythm classes: Evidence from portuguese," in *Proc. Speech Prosody*, Aix-en-Provence, France, 2002, pp. 319–322.
- [24] M. Komatsu, T. Arai, and T. Sugawara, "Perceptual discrimination of prosodic types," in *Proc. Speech Prosody*, Nara, Japan, 2004, pp. 725–728.
- [25] A. Galves, J. Garcia, D. Duarte, and C. Galves, "Sonority as a basis for rhythmic class discrimination," in *Proc. Speech Prosody*, Aix en Provence, France, Apr. 2002, pp. 323–326.
- [26] R. André-Obrecht, "A new statistical approach for the automatic segmentation of continuous speech signals," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, no. 1, pp. 29–40, Jan. 1988.
- [27] F. Pellegrino and R. André-Obrecht, "Vocalic system modeling: A VQ approach," in *Proc. IEEE Digital Signal Process.*, Santorini, Greece, Jul. 1997, pp. 427–430.
- [28] F. Pellegrino, "Une approche phonétique en identification automatique des langues: La modélisation acoustique des systèmes vocaliques," Ph.D. dissertation, Univ. Paul Sabatier, Toulouse, France, 1998.
- [29] I. Kopeček, "Syllable based approach to automatic prosody detection; Applications for dialogue systems," in *Proc. Workshop Dialogue Prosody*, Eindhoven, The Netherlands, Sep. 1999, pp. 89–93.
- [30] H. Pfitzinger, S. Burger, and S. Heid, "Syllable detection in read and spontaneous speech," in *Proc. 4th ICSLP*, Philadelphia, PA, Oct. 1996, vol. 2, pp. 1261–1264.
- [31] J.-L. Rouas, J. Farinas, F. Pellegrino, and R. André-Obrecht, "Rhythmic unit extraction and modeling for automatic language identification," *Speech Commun.*, vol. 47, no. 4, pp. 436–456, 2005.
- [32] K. Sjölander, "The Snack Sound Toolkit," [Online]. Available: <http://www.speech.kth.se/snack/>.
- [33] H. Fujisaki, "Prosody, information and modeling – With emphasis on tonal features of speech," in *Proc. ISCA Workshop Spoken Lang. Process.*, Mumbai, India, Jan. 2003, pp. 5–14.
- [34] A. Stolcke, "SRILM—An extensible language modeling toolkit," in *Proc. ICSLP*, 2002, pp. 901–904.
- [35] S. Kitazawa, "Periodicity of Japanese accent in continuous speech," in *Proc. Speech Prosody*, Aix en Provence, France, Apr. 2002, pp. 435–438.
- [36] F. Pellegrino, J. Farinas, and J.-L. Rouas, "Automatic estimation of speaking rate in multilingual spontaneous speech," in *Proc. Speech Prosody*, Nara, Japan, Mar. 23–26, 2004, pp. 517–520.



Jean-Luc Rouas received the Ph.D. degree in computer science from the Computed Science Research Institute (IRIT), Toulouse, France, and Paul Sabatier University, Toulouse, France, in 2005. His thesis was "Languages' characterisation and identification."

He had a one-year Postdoctoral position at the French National Institute for Safety and Transports (INRETS). He is now with the Laboratório de Sistemas de Língua Falada (L2F), Instituto de Engenharia de Sistemas e Computadores (INESC), Lisbon, Portugal. His research interests include

signal analysis, speech processing, and prosody.