

Similarity Search of Acted Voices for Automatic Voice Casting

Nicolas Obin, *Member, IEEE*, Axel Roebel, *Member, IEEE*

Abstract—This paper presents a large-scale similarity search of professionally acted voices for computer-aided voice casting. The proposed voice casting system explores GMM-based acoustic models and multi-label recognition of perceived para-linguistic content (speaker states and speaker traits, e.g., age/gender, voice quality, emotion) for the voice casting of professionally acted voices. First, acoustic models (universal background model, super-vector, i-vector) are constructed to model the acoustic space of voices, from which the similarity between voices can be measured directly in the acoustic space. Second, multiple binary classification of speaker traits and states is added to the acoustic models in order to represent the vocal signature of a voice, which is then used to measure the similarity between voices in the para-linguistic space. Finally, a similarity search is processed in order to determine the set of target actors that are the most similar to the voice of a source actor. In a subjective experiment conducted in the real-context of cross-language voice casting, the multi-label scoring system significantly outperforms the acoustic scoring system. This constitutes a proof of concept for the role of perceived para-linguistic categories in the perception of voice similarity.

Index Terms: voice casting, voice similarity, speaker recognition, speaker traits and states, para-linguistics, multi-label classification.

I. INTRODUCTION

THE production of multi-media content (films, series, video-games) available to various countries requires the translation of the speech content from a source language (typically, English) to a set of target languages (typically, French, German, Spanish, Japanese, Mandarin). Translation of the speech content can be simply obtained by subtitles, but very often the original speech content is totally replaced by the corresponding speech content in the target language. This process, referred to as dubbing, is obtained by first translating the text from the source to the target language, then selecting actors in the target language, and finally recording the actors synced to the original speech content. These actors must be selected so as to preserve as much as possible the voice and the acting of the original actors. Voice casting denotes the selection of a voice in a target language that is the most similar to a voice in a source language, and is usually performed by human experts who manually select actors according to a database of available voices in the target language. Beyond voice casting stands the open scientific issue on the perception and the measurement of voice similarity: the closer/farther a source voice is perceived from a target voice, the smaller/larger the distance should be measured. What defines the perception voice similarity remains vague: common expressions (i.e., gender: male/female, age: young/old) are generally used to

describe the main traits of a voice/speaker [1], and the role of voice quality in the perception of voice similarity has been recently addressed [2]. Also, some recent research in speaker clustering (e.g., speech retrieval [3], [4], and speech synthesis [5], [6]) have addressed to some extent the measurement of speaker/voice similarities.

To the best of our knowledge, this paper is the first scientific investigation into the measurement of voice similarity for the voice casting of professionally acted voices. Two alternative solutions are investigated and compared:

- Intuitively, the use of speaker recognition techniques [7]–[9] for voice casting appears seducing: the scoring used in speaker recognition system can be interpreted straightforwardly as a similarity measure between voices, and this similarity can be directly measured in the acoustic space. In particular, the similarity measure as determined for speaker recognition has been proven to be extremely accurate in the local acoustic neighbourhood of a speaker: a speaker can be authenticated in the presence of close impostor speakers. However, there is no evidence that this similarity measure remains valid in the entire acoustic space, and actually reflects the perception of the similarity between voices.
- Alternatively, the description of a voice by perceived para-linguistic categories (speaker states and traits [10], [11] e.g., age/gender, voice quality, emotions) may efficiently capture the perception of a voice, and then serve to measure the similarity between voices. Furthermore, common expressions are widely used by human experts in voice casting to stereotype a role. For instance, Albus Dumbledore from the movie “Harry Potter” can be described as a male, old, wise, and breathy voice. Accordingly, the stereotype may be more important for voice casting than the actual acoustic similarities.

This paper explores and compares the use of GMM-based acoustic models and multi-label classification of perceived para-linguistic categories (e.g., age/gender, voice quality, emotion) for the voice casting of professionally acted voices. This extends the preliminary work presented in [12] by 1) presenting the complete details of the two contributions proposed for voice casting, one is based on acoustic models derived from speaker recognition and one is based on multi-label recognition of perceived para-linguistic categories (speaker traits and speaker states); and 2) providing a detailed evaluation of para-linguistic recognition for all considered categories, including a comparison of MFCC, super-vector, and i-vector acoustic representations for some classic and novel para-linguistic categories. First, GMM-based acoustic models and multi-label classification of speaker traits and states are presented in order to score the similarity between voices in the acoustic and para-

Nicolas Obin and Axel Roebel are with IRCAM, UPMC-Sorbonne Universités, CNRS, Paris, France.

linguistic spaces (Section II). Second, the set of categories used to describe the traits and states of a speaker for the multi-label scoring is presented (Section III). The performance of the acoustic and the multi-label scoring is first evaluated in an objective experiment, and then compared in a subjective experiment in the real context of professional voice casting (Section IV).

II. SIMILARITY SCORING FOR VOICE CASTING

This section presents the details of the acoustic models and scoring derived from speaker recognition, and the multi-label classification and scoring from perceived para-linguistic categories for voice casting. The implementation is based on *ircamClassifier* [13], a system developed in the context of Music Information Retrieval (MIR) [14], [15]. This system includes the *Alizée 3.0* speaker recognition [16] and the *LibSvm* [17] SVM libraries.

A. Acoustic Space Modeling: Universal Background Model and GMM supervector

The Universal Background Model (UBM) is used to model the distribution of the entire acoustic space [7]. This modelling is usually achieved with a standard Gaussian Mixture Model (GMM-UBM). The likelihood of the $(D \times 1)$ feature vector \mathbf{o} describing the acoustic characteristics of speech is defined as:

$$p(\mathbf{o}|\boldsymbol{\lambda}) = \sum_{i=1}^M \alpha_i p_i(\mathbf{o}) \quad (1)$$

where M is the number of mixture components, $\boldsymbol{\lambda} = \{\alpha_i, \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i\}_{i \in [1, M]}$ represents the weights, means, and variances of the i -th Gaussian, and $p_i(\mathbf{o}) = \mathcal{N}(\mathbf{o}|\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$ denotes the likelihood of the i -th mixture component, where \mathcal{N} denotes a Gaussian distribution.

Then, the mean parameters $\boldsymbol{\mu} = \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_M\}$ of the UBM are adapted to each speech recording by using relevance maximum a posteriori (MAP) adaptation [7]. This is achieved by updating the means of the mixture components to a sequence of acoustic observations $\mathbf{o} = [\mathbf{o}_1, \dots, \mathbf{o}_T]$, of length T . Finally, each speech recording is represented by the mean vectors of the adapted mixture components:

$$\boldsymbol{\mu}^{adapt} = [\boldsymbol{\mu}_1^{adapt \top}, \dots, \boldsymbol{\mu}_M^{adapt \top}]^\top \quad (2)$$

where \top denotes the transposition operator, and $\boldsymbol{\mu}^{adapt}$, referred to as a GMM-supervector, is the concatenation of all the mean vectors of the adapted mean parameters of the UBM.

B. Factor Analysis: Total Variability Space and i-vector

An i-vector is the compact representation of a high-dimensional speech recording into a low-dimensional space called Total Variability space [9], assuming an affine linear model (i.e., factor analysis):

$$\boldsymbol{\mu}' = \boldsymbol{\mu} + \mathbf{T}\mathbf{x} \quad (3)$$

where $\boldsymbol{\mu}'$ is the GMM-supervector of a speech recording, $\boldsymbol{\mu}$ is the GMM-supervector corresponding to the UBM mean

parameters, \mathbf{T} is the $(DM \times q)$ total variability matrix, and \mathbf{x} is a q -dimensional vector assuming a prior normal distribution, referred to as an i-vector. The total variability matrix \mathbf{T} is estimated by Expectation-Maximization (EM) [9]. The i-vector of a speech recording is determined as a MAP point estimate of the latent variable \mathbf{x} [9].

C. Inter-Session Compensation: i-vector Transformation

The i-vector transformation is used to project the total variability of the high-dimensional acoustic space (i.e., speaker/class information and session/channel information) in a low-dimensional space in which the i-vectors distribution is assumed to be normal for each speaker/class. In order to compensate for the session/channel information, and to constrain the i-vector distribution to be normally distributed for each speaker/class, a large number of methods have been proposed from Linear Discriminant Analysis [9] (LDA) for inter-session compensation, to Within-Class Covariance Normalization (WCCN, [18]), Length Normalization (LN, [19]), Eigen Factor Radial Normalization (EFR, [20]), and Sphere Nuisance Normalization (SN, [16], [20]) for speaker/class normalization.

The Length Normalization (LN) is a simple normalization:

$$\hat{\mathbf{x}} = \frac{\mathbf{x}}{\|\mathbf{x}\|} \quad (4)$$

where $\|\cdot\|$ denotes the L-2 norm.

The Within-Class Covariance Normalization (WCCN) whitens the covariance matrix of each class:

$$\hat{\mathbf{x}} = \mathbf{W}^{-\frac{1}{2}} \mathbf{x} \quad (5)$$

$$\mathbf{W} = \sum_{k=1}^K p(k) \boldsymbol{\Sigma}_x^{(k)} \quad (6)$$

where \mathbf{W} is the covariance matrix defined as the weighted sum of the within-class covariance matrices $\boldsymbol{\Sigma}_x^{(k)}$, where K is the number of classes, and $p(k)$ is the prior probability of class k .

The Eigen Factor Radial Normalization (EFR) processes recursively standardization and normalization:

$$\mathbf{x}^{(i+1)} = \frac{\boldsymbol{\Sigma}_x^{(i)-\frac{1}{2}} (\mathbf{x}^{(i)} - \boldsymbol{\mu}_x^{(i)})}{\|\boldsymbol{\Sigma}_x^{(i)-\frac{1}{2}} (\mathbf{x}^{(i)} - \boldsymbol{\mu}_x^{(i)})\|} \quad (7)$$

where $\boldsymbol{\mu}_x^{(i)}$ and $\boldsymbol{\Sigma}_x^{(i)}$ denote the mean vector and covariance matrix of all i-vectors at iteration i .

The Sphere Nuisance Normalization (SN) is similar to the Eigen Factor Radial Normalization, except that the covariance matrix is replaced by the within-class covariance matrix \mathbf{W} .

D. Acoustic Scoring

The first contribution of this paper is to investigate the acoustic scoring derived from speaker recognition to measure

the similarity between voices for voice casting (Figure 1). While the Support Vector Machine (SVM) is historically a milestone in speaker recognition [8], some recent advances have been proposed to score the similarity between speakers.

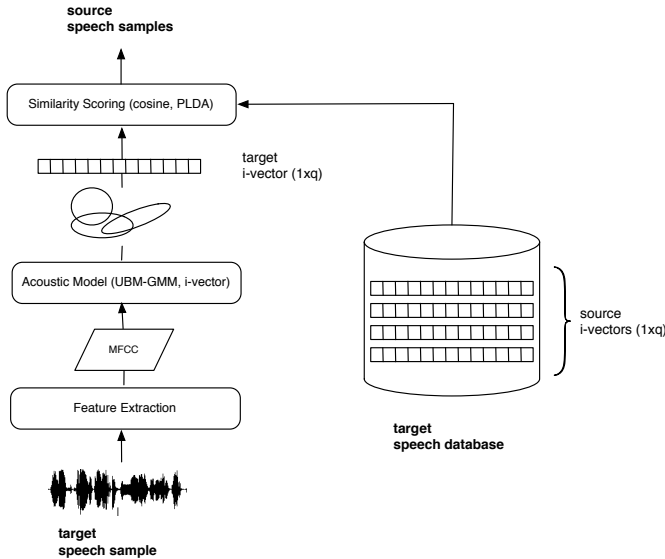


Fig. 1. Architecture of the acoustic voice casting system. On bottom, unsupervised acoustic extraction and modelling; on top, supervised/unsupervised acoustic scoring.

1) *Direct Scoring: Cosine Similarity:* First, direct cosine similarity [21] has been proven to be extremely efficient for speaker recognition. The cosine similarity measures the similarity between two speech recordings \mathbf{x}_{src} and \mathbf{x}_{tgt} in the i-vector acoustic space:

$$s(\mathbf{x}_{src}, \mathbf{x}_{tgt}) = \frac{\langle \mathbf{x}_{src}, \mathbf{x}_{tgt} \rangle}{\|\mathbf{x}_{src}\| \|\mathbf{x}_{tgt}\|} \quad (8)$$

where $\langle \cdot, \cdot \rangle$ is the scalar product operator.

Importantly, the cosine similarity assumes that only the angle between two i-vectors provides information about the similarity between speech recordings. Furthermore, the cosine similarity can be computed directly in the acoustic space, without any prior training.

2) *Generative Model: PLDA:* One of the last advances is the introduction of generative models for speaker recognition [22]. Among them, the Probabilistic Linear Discriminant Analysis (PLDA) [23] is the most popular generative model currently used for speaker recognition. In the original form, PLDA linearly decomposes an i-vector in eigen-speaker and eigen-channel subspaces (respectively of rank $N_{speaker}$ and $N_{channel}$). In the case where the eigen-channel is assumed to be full-rank ($N_{channel} = q$) (Gaussian PLDA [24] or simplified PLDA [19]), each i-vector \mathbf{x}_s of a speaker s can be expressed as:

$$\mathbf{x}_s = \boldsymbol{\mu}_x + \mathbf{S}\mathbf{h}_s + \epsilon \quad (9)$$

where $\boldsymbol{\mu}_x$ is the total i-vectors mean vector, \mathbf{S} is the ($N_{speaker} \times q$) eigen-speaker matrix, \mathbf{h}_s is the position of the

i-vector within the eigen-speaker space \mathbf{S} (the latent speaker vector, assumed to be normally distributed), and ϵ is the q residual vector with a full covariance matrix. Maximum-Likelihood (ML) estimation of the PLDA parameters is described in [23].

Then, the similarity between two speech recordings \mathbf{x}_{src} and \mathbf{x}_{tgt} can be computed as the likelihood ratio [24]:

$$s(\mathbf{x}_{src}, \mathbf{x}_{tgt}) = \frac{p(\mathbf{x}_{src}, \mathbf{x}_{tgt} | \mathcal{H}_1)}{p(\mathbf{x}_{src} | \mathcal{H}_0) p(\mathbf{x}_{tgt} | \mathcal{H}_0)} \quad (10)$$

where the hypothesis \mathcal{H}_1 indicates that both vectors come from the same latent speaker (respectively, class), and \mathcal{H}_0 indicates they come from different latent speakers (respectively, classes). A closed form solution can be computed as detailed in [24], [25].

The acoustic scores derived from speaker recognition can be straightforwardly turned into a similarity measure between two speech recordings for voice casting, by ignoring the identification, recognition, verification hard decision of a speaker identity. The main advantages of the acoustic scoring for voice casting is that the scoring can be performed directly in the acoustic space, and the similarity measure has been proven to be extremely efficient for speaker recognition. Also, the similarity metric may be strictly unsupervised (e.g., cosine distance), or supervised by available information (e.g., PLDA). Since the speaker's identity is generally the only available information in the context of voice casting and assuming the accuracy of speaker recognition systems, this first contribution will address the use of a similarity metric supervised by speaker's identity for voice casting.

E. Multi-Label Scoring

The second contribution of this paper is the use of a multi-label scoring based on the "semantic" description of a voice with perceived para-linguistic categories (speaker states and speaker traits). This multi-label scoring is presented as an alternative to the acoustic scoring as used for speaker recognition. First, a multi-label classifier is added on top of GMM-based acoustic models to assign the labels corresponding to a speech recording. Then, the posterior probabilities of each label are concatenated to form a vector that represents the signature of a voice, which is used to measure the similarity between voices (Figure 2).

Multi-label classification [26] is commonly used for the indexing, retrieval, and similarity search of multi-media content (e.g., [27], [28] for text, [29], [30] for music, [31] for image, and [32] for video). Multi-label classification assumes that a media content (text, image, video, sound) can be described with a set of labels that are independent to each other. Multi-label classification is opposed to the multi-class classification commonly used in para-linguistic classification: the classification does not result to a single label, but to a vector of multiple co-occurring and non-exclusive labels

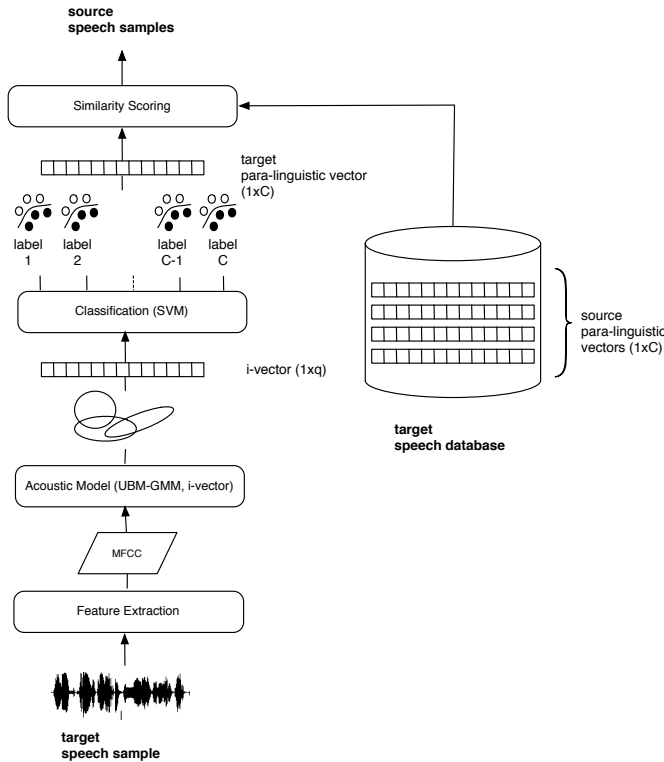


Fig. 2. Architecture of the multi-label voice casting system. On bottom, unsupervised acoustic extraction and modelling; on top, supervised multi-label recognition and scoring.

(Figure 3). In consequence, multiple labels can be assigned to a media content. Typically, the emotion content of a speech recording can be a mix of “sadness” and “fear” and even of “sadness” and “joy”, whereas the speech recording would be classified only as “sadness” with multi-class classification. This provides an extended description of a speech content that can be further used for similarity search.

For voice casting, a multi-label scoring is constructed by converting the classification of multiple labels into multiple binary classifications [14], [26]. First, each label of the speech description (e.g., the speech recording is creaky) is turned into a binary representation (i.e., yes/no). Then, a classifier is trained for each label separately, which results into C independent one-versus-all classifiers [33]. A complete description of the labels used in this paper is provided in Section III.

Here, the Support Vector Machine ([17]) is used for multi-label classification. For each label c , the classification of a vector \mathbf{x} (e.g., supervector, i-vector) corresponding to a speech recording is obtained with regard to the decision function:

$$f_c(\mathbf{x}) = \sum_{i=1}^N \omega_c^i K(\mathbf{x}, \mathbf{x}_c^i) + b_c \quad (11)$$

where $\Theta_c = \{\omega_c^i, \mathbf{x}_c^i, b_c\}_{i=1}^N$ are the parameters of the maximum-margin hyperplane determined during training (respectively weights, support vectors, and offset), and $K(\cdot, \cdot)$ the SVM kernel [34].

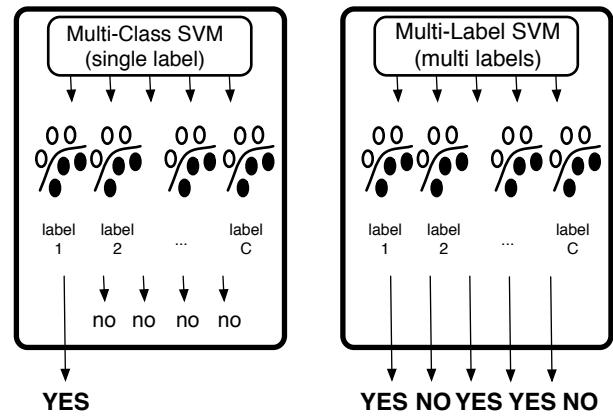


Fig. 3. Multi-class vs. multi-label classification.

In a standard SVM, the binary label corresponding to the observation vector \mathbf{x} is assigned with regard to the sign of the decision function:

$$\hat{y}_c = \text{sign}(f_c(\mathbf{x})) \quad (12)$$

where $\text{sign}(x) = +1$ for $x \geq 0$ and -1 otherwise, so that y_c is 1 when the label is positive, and 0 when the label is negative.

Here, the decision function is converted into a posterior probability estimate for each label c , as detailed in [35]:

$$\psi_c = p_c(y = 1|\mathbf{x}) = p(y = 1|\mathbf{x}, \Theta_c), \quad c \in [1, \dots, C] \quad (13)$$

Then, the posterior probabilities of each label are concatenated to form a vector that represents the vocal signature of a speech recording:

$$\Psi = [\psi_1, \dots, \psi_C]^T \quad (14)$$

where ψ_c is the posterior probability of the c -th label conditionally to the observation vector \mathbf{x} . Similarly to the GMM-supervector and the i-vector, the vector Ψ representing the vocal signature of a speech recording is a single vector summarizing each speech recording.

Finally, the similarity of a source to a target speech recording is defined as the distance d of their vocal signatures:

$$s(\mathbf{x}_{src}, \mathbf{x}_{tgt}) = d(\Psi_{src}, \Psi_{tgt}) \quad (15)$$

where Ψ_{src} and Ψ_{tgt} denotes the vocal signature of the source and target speech recordings, respectively. Here, $d(\cdot, \cdot)$ is defined as the Kullback-Leibler divergence $KL(\cdot||\cdot)$ which is a natural distance measure between posterior probabilities [36], [37].

The main advantage of the multi-label scoring for voice casting lies on the assumption that the para-linguistic content of a voice (speaker traits and states) may reflect more explicitly the *perceived* similarity between voices. Furthermore, the multi-label classification system can be used to automatically tag and search voices based on their perceived para-linguistic content within large speech databases.

III. DESCRIPTION OF SPEAKER TRAITS AND STATES

This section presents the details of the para-linguistic categories used for the multi-label scoring system presented in Section II-E. A short review on standard description and recognition of a voice content is first provided, followed by a specification of the selected description for voice casting.

Common linguistic expressions can be associated with a voice to describe the perceived “quality” of the voice [1]. Among them, the age (young/old), sex (male/female), and emotions are the most widely used expressions to describe a voice. These expressions can be directly associated with speaker traits and states: speaker traits denote persistent/external traits of a speaker (e.g., personality), and speaker states denote the temporary/internal states of a speaker (e.g., emotions). The definition and the description of speaker traits and states has been widely studied in the literature, from biological speaker traits primitives (e.g., age and gender), individual, social and cultural speaker traits (e.g., voice quality, [38]), to speaker states (e.g., emotions [39]). Also, research on the automatic recognition of perceived para-linguistics categories in speech has considerably increased over the past few years, pushed by the emergence and needs for human-robot interaction and multi-media retrieval applications (computational paralinguistics challenges: speaker age and gender [40], speaker state [10], [41], speaker traits [11]). The recognition scores obtained for perceived para-linguistics recognition significantly varies depending on the task: the classification of the gender of a speaker is accurate (around 90% for adult speakers [42]), the age can be reasonably determined (within 10 years on telephone speech, [43]), while emotion remains an open issue (from around 80% [44] for acted speech [45], to only 60% for spontaneous speech [46], [47]). More recently, the recognition and modeling of voice quality has raised as a novel topic in para-linguistics recognition [48]–[51].

For voice casting, one must first define a comprehensive set of para-linguistics categories that can be used by expert voice casting operators and exploited for voice similarity search. This set must cover the main traits and states of a speaker, and fulfill specific ad-hoc needs of expert voice casting operators (mostly related to acting and stereotypes). The first set of speaker traits and states comprises standard para-linguistic categories:

- biological speaker traits: sex (male, female), and age (child, teenager, young adult, adult, old, very old);
- speaker state: emotion (tender, excited, happy, neutral, sad, angry, fear, stressed, surprise, other);

A second set comprises categories associated with perceived acoustic characteristics of the voice:

- phonation: voice quality (breathy, creaky, hoarse), tension (relaxed, normal, tensed, pressed), vocal effort (whispered/soft, normal, loud/shouted);
- articulation: articulation (hypo, normal, hyper);
- timbre: timbre (clear, dark);




Fig. 4. Multi-label tagging of a speech recording.

- prosody: F0 register (extreme-low, low, medium, high, extreme-high), F0 range (flat, normal, extended), and speech rate (slow, normal, fast);

A last set comprises categories associated with the role and the situation of acting:

- attitude/modality: affirmation, confirmation, exclamation, interrogation, order, other;
- situation: action, conversation, information, monologue, other;
- archetype: announcer, artificial intelligence, basic soldier, brute, commander, hero, neutral, old wise, rookie soldier, sensual, suffer, veteran soldier, other.

The selected description of a speaker includes 14 classes (e.g., gender, age, emotion voice quality), and 68 labels (e.g., for voice quality: breathy, creaky, hoarse). For clarity, the terms “class” and “label” are here used by analogy to multi-class and multi-label classification (see Table II): a class denotes a group containing multiple instances (e.g., the emotion class contains multiple instances: angry, happy, neutral, sad), and a label denote each particular instance (e.g., angry, happy, neutral, sad are labels). The multiple labeling of a speech recording results in a binary vector which represents the voice signature of the speech recording (Figure 4).

A preliminary phase of manual labeling was conducted in order to train multi-label classifiers for the recognition of speaker traits and speaker states, and to process multi-label scoring for voice casting. Beforehand, a guideline was created to define each class and each label, accompanied by a set of representative speech samples, and a PHP web interface was designed to allow easy and fast on-line annotation of a speech database. The manual labeling was produced by a non-expert individual, preliminary trained by two speech experts (the author, and an expert voice casting operator). First, pilot campaigns were conducted on small sets of speech recordings (around 50-100) by the non-expert annotator and the two expert annotators, until the non-expert annotator presents a sufficiently satisfactory agreement with the expert annotators. The final inter-annotator agreement for coding speaker traits and states has an average Krippendorff’s alpha of $\alpha = 0.52$ [52], which represents a fairly reliable agreement regarding the ambiguity and the diversity of the classes considered for

labeling. Then, a large-scale annotation was conducted on a selection of 4,000 speech recordings extracted from the 20,000 speech recordings of the French version of the Mass Effect 3 video game, covering 54 speakers interpreting 500 roles, with a maximum of 10 speech recordings for each role (see section IV for a detailed description of the speech database).

IV. EXPERIMENTS

Two experiments were conducted to compare acoustic and multi-label similarity scoring in the context of professional voice casting. First, an objective experiment was conducted to determine the parameters of the optimal configurations of the acoustic and multi-label scoring systems. Then, a subjective experiment was conducted to compare the optimal acoustic and multi-label similarity scoring systems in the real context of professional voice casting. For all comparisons, the acoustic and multi-label similarity scoring systems share the same unsupervised acoustic space representation (MFCC, super-vector, i-vector). The systems differ only by the way the similarity measure is constructed: for the acoustic scoring, the similarity metric is defined in the acoustic space, and supervised with respect to the speaker's identity; for the multi-label scoring, the similarity metric is defined in the para-linguistic space, and supervised with respect to each para-linguistic label.

A. Objective Experiment

The purpose of the objective experiment is to determine the optimal configurations of the acoustic and multi-label scoring systems, in order to select the configurations that will be used for the subjective comparison. Accordingly, the objective experiment is only concerned with separate optimization of the acoustic scoring and multi-label scoring systems. The acoustic scoring system is optimized with respect to a speaker recognition experiment, and the multi-label system is optimized with respect to a perceived para-linguistic classification experiment. Besides optimization, the objective experiment explores the use of advanced acoustic modeling (super-vector, i-vector) for the recognition of a large set of para-linguistic speech categories, which extends preliminary research (for age recognition, [43]), and includes novel para-linguistic categories (e.g., attitude/modality, situation, archetypes).

The objective experiment was conducted on the French version of the Mass Effect 3 video game containing 20,000 speech recordings, around 500 roles, around 50 speakers, and around 20 hours of speech of professional actors. A subset of 4,000 speech recordings was used for the manual annotation of perceived para-linguistic categories. All speech recordings were recorded in professional conditions (professional studio recordings, same recording material, same supervision), and encoded into a 48 kHz-16 bits high-quality format. The duration of speech recording varies from 0.1 s to 15 s. Speech recordings shorter than 1 s were removed from the speech database. The front-end processing consisted in the extraction of short-term (20 ms. Hanning window with 50%

overlapping) Mel-frequency cepstral coefficients (MFCC, 13 cepstral coefficient determined with 40 Mel-frequency bands), without delta and delta-delta. The system setups were defined as follows: $N_{\text{GMM}} = 8$ to 2048 (number of GMM-UBM mixture components), $q = 10$ to 800 (dimension of i-vector), and shared among the acoustic scoring and multi-label scoring systems. For the acoustic scoring system, $N_{\text{LDA}} = 10$ to 200 (dimension of LDA reduction), $N_{\text{it}} = 1$ for EFR (length normalization), $N_{\text{it}}=3$ for sphNorm (number of iterations), $N_{\text{speaker}} = 10$ to 400 and $N_{\text{channel}} = q$ (dimension of the speaker and channel spaces for PLDA). For the cosine and PLDA scoring, the scoring was performed by using the mean i-vector of the speaker [16]. For the multi-label scoring, a SVM classifier with a Gaussian kernel [53] was used for binary classification of each label (Figure 2), each trained on the subset of manually annotated speech recordings.

The experiment was conducted in the form of a 2-fold cross-validation for speaker recognition and 5-fold cross-validation for para-linguistic classification. In k -fold cross-validation, the dataset is first partitioned into k subsets of equal size, then $k-1$ subsets are used for training the model parameters, and the remaining subset is used for testing the model. This process is repeated for the k folds. The main advantage of cross-validation is the explicit consideration of the performance variability, which can be then be used to assess the statistical difference between different model configurations (acoustic space modeling, inter-session compensation, and scoring). Here, the subsets are constructed by randomly partitioning the available speech recordings regardless of the speakers. For speaker recognition, the standard Equal Error Rate (EER) was used to measure the performance, as determined from the detection error trade-off (DET) curve [54] by following the NIST SRE 2012 guidelines [55]. For para-linguistic classification, the balanced accuracy (BA%) is used to measure the recognition performance [56]. The balanced accuracy is the equivalent for binary classification of the unweighted average recall (WA%, [57]) for multi-class classification, a well-established measure for emotion and other para-linguistics recognition ([10], [58], [59]). Indeed, the balanced accuracy is simply defined as the unweighted average of true and false recalls of a binary classification. For one label, it is computed as:

$$BA = \frac{R_P + R_N}{2} = \frac{1}{2} \left(\frac{T_P}{T_P + F_N} + \frac{T_N}{F_P + T_N} \right) \quad (16)$$

where R_P and R_N are the positive and negative recalls, and T_P , T_N , F_P , and F_N are the true positive, true negative, false positive, and true negative counts.

The main idea of these measures is to compensate for imbalanced datasets when computing the accuracy score. This is particularly true for para-linguistic binary classification, where the class of interest (the positive one) is highly under-represented as compared to the other (e.g., creaky = yes vs. creaky = no, emotion = sad vs. emotion = not sad, etc...). For each model configuration, the average balanced accuracy is obtained by averaging the balanced accuracy over all folds to compute the per label score, and over

all labels to compute the overall score. A grid search is processed to determine the optimal model configuration. First, cross-validation is processed for each model configuration: model parameters are estimated on the training set and the corresponding performance is evaluated on the test set. Then, the hyper-parameters are tuned for each configuration as the one maximizing the overall score.

The performance obtained for speaker recognition is presented in Table I. The optimal performance was obtained with the i-vector + sphNorm + PLDA scoring method with the following configuration, 512 GMM (UBM), $q = 400$ (i-vector), $N_{\text{speaker}} = 50$ and $N_{\text{channel}} = 400$ (full-rank) (PLDA). The speaker recognition performance (EER = 2.50%) indicates the robustness of GMM-based acoustic models to the expressive variability of the speaker, and to the variability in duration of the speech recordings.

METHOD	EER (%)
i-vector + cosine	4.04
i-vector + LDA/WCCN + cosine	3.02
i-vector + PLDA	2.80
i-vector + EFR + PLDA	2.73
i-vector + sphNorm + PLDA	2.50

TABLE I
PERFORMANCE OF SPEAKER RECOGNITION SYSTEMS (EER (%)).

The performance obtained for para-linguistic classification is presented in details in Table II (with the exception of 4 "other" labels and 4 other minor labels, only for the sake of space purpose), and in summary with 95% confidence intervals in Figure 5. The 95% confidence interval is computed by assuming a normal distribution of the cross-validation scores, and is equal to 1.96 times the standard deviation of the cross-validation scores, divided by the square root of the number of folds. The optimal performance was obtained with the i-vector + EFR + SVM method with the following configuration: 512 GMM (UBM), and $q = 50$ (i-vector). In all cases, the i-vector recognition (from 73.45% to 74.62%) has a greater recognition rate than the super-vector recognition (from 71.08% to 72.32%), and the MFCC recognition (70.09%). Also, the inter-session compensation improves the recognition performance, from 71.08% to 72.32% for super-vectors, and from 73.45% to 74.62% for i-vectors. A statistical comparison (one-way ANOVA [60]) shows that the super-vector recognition rate is significantly higher than the MFCC recognition rate ($F(1, 18) = 20.52$, $p - \text{value} \leq 10^{-4}$), and that the i-vector recognition rate is significantly higher than the super-vector recognition rate ($F(1, 18) = 35.20$, $p - \text{value} \leq 10^{-5}$). Also, the i-vector EFR recognition rate is significantly higher than the average i-vector recognition rate ($F(1, 58) = 6.78$, $p - \text{value} \leq 10^{-2}$), and is higher but not significantly with the LN and WCCN recognition rates (respectively, $F(1, 18) = 2.52$, $p - \text{value} = 0.12$ and

$F(1, 18) = 4.18$, $p - \text{value} = 0.05$). For details, the optimal configuration corresponds to 95.4% for gender, 80.4% for age, 76.8% for voice quality, 73.6% for tension, 80.3% for vocal effort, 73.3% for timbre, 62.9% for articulation, 76.7% for F0 (range and register), 65.6% for speech rate, 66.7% for attitude/modality, 66.4% for emotion, 78.5% for situation, and 78.2% for archetypes. These scores correspond to one single and globally optimal configuration for super-vectors, i-vectors, inter-session compensation, and SVM hyper-parameters in order to figure out a computationally realistic scenario, though all individual performances might be improved through dedicated optimizations. From these observations, some para-linguistic categories can be consistently recognized (age, gender, voice quality, tension, vocal effort, timbre, situation, archetype) while some others remain an open issue (articulation, F0, speech rate, attitude/modality, and emotion). In particular, some novel para-linguistic categories specific to multi-media applications (situation and archetype) are more recognized than some standard para-linguistic categories (attitude/modality, emotions). Also, extreme para-linguistic labels are generally more recognized than standard ones (e.g., normal, medium, neutral) which are more ambiguous. As a conclusion, this generalizes the role of advanced acoustic modeling (i-vector and inter-session compensation) for para-linguistic recognition, as preliminarily reported for age estimation in [43]. Moreover, these constitute encouraging performances for further similarity search for voice casting.




Fig. 5. Overall average recognition score (BA%) and 95% confidence interval obtained for the multi-label classification.

The optimal configurations were further retained for the subjective comparison of acoustic scoring and multi-label scoring systems in the real context of professional voice casting.

B. Subjective Experiment

The real context of voice casting consists in selecting the actors of a target language (e.g., French, German, Spanish, Japanese, Mandarin) whose voice is the most similar to actors

of a source language (typically, English). Accordingly, a subjective experiment was conducted in order to address the ability of acoustic and multi-label similarity scores to estimate the perceived similarity between voices for cross-language voice casting. The objective is to compare the role of acoustic and para-linguistic information in the human perception of voice similarity.

The subjective experiment consisted in the comparison of the two optimal configurations previously determined for a voice casting from American-English to French. The American-English (source language) and the French (target language) versions of the Mass Effect 3 video game were used for the experiment. First, 50 speech samples were selected from the American-English version, one speech recording for each of 50 speakers (50% male, 50% female, around 5 sec. in duration). For each source speech sample, the 3 most similar samples were determined in the target speech database for each scoring system. Then, the source speech sample and the 3 target speech samples determined by the 2 scoring systems were presented to the listener. For each source speech sample, the listener was asked to rate the overall similarity of the target speech samples to the source speech sample on a 5 degree scale: very dissimilar (-2), fairly dissimilar (-1), slightly similar (0), fairly similar (+1), very similar (+2). 30 French native individuals participated in the experiment (20 males/ 10 females, 20-35 years old, same headphones, same professional listening room, paid experiment).




Fig. 6. Mean similarity score and 95% confidence interval for the 2 systems. The similarity scale is very dissimilar (-2), fairly dissimilar (-1), slightly similar (0), fairly similar (+1), very similar (+2).

The comparison of the 2 scoring systems is presented in Figure 6. The multi-label scoring system significantly outperforms the acoustic scoring system in the similarity judgement for voice casting. For comparison, the target speech samples determined by the acoustic scoring (i-vector + sphNorm + PLDA) and the multi-label scoring (i-vector + EFR (noNorm) + SVM) systems are considered as *slightly similar* and *fairly similar* to the source sample in average, respectively. This constitutes almost a one degree difference on the 5 degree scale. A statistical comparison

(one-way ANOVA [60]) shows that the multi-label scoring system is judged as significantly more similar than the acoustic scoring system ($F(1, 298) = 10.86$, $p - value \leq 10^{-3}$).

This experiment provides instructive information about the role of acoustic and para-linguistic information in the perception of voice similarity. First, acoustic information appears necessary but not sufficient to fully capture the perceived similarity between voices. Second, para-linguistic information, as abstractions extracted from the speech content, provides some valuable information about the perception of voice similarity. These observations suggest that the abstraction of a voice into categories (speaker traits and states) play an important role in the human perception of voice similarity, which may prevail over pure acoustic similarity. Beyond, this highlights the role of stereotypes in the human perception of voice similarity, which might be particularly true for professionally acted voices that are generally more stereotyped than everyday speech.

V. CONCLUSION

In this paper, a large-scale similarity search of voices was presented to measure the perceived similarity between voices for computer-aided voice casting. The proposed voice casting system explored and compared GMM-based acoustic models and multi-label recognition of perceived para-linguistic content (e.g., age/gender, voice quality, emotion) to measure the perceived similarity between voices. In a subjective experiment, the multi-label scoring significantly outperformed acoustic scoring in the real-context of voice casting, which constitutes evidence for the role of perceived para-linguistic content in the perception of voice similarity. This constitutes a preliminary research on voice similarity search for the voice casting of professional acted voices. Further research will investigate the use of short- and long-term speech characteristics (glottal source [48], [49], prosody [61]) during acoustic modeling, the elaboration of acoustic scoring more specific to voice casting and less constrained by speaker's identity, and the construction of a similarity scoring that covers the entire expressive range of actors instead on being based on the particular expression of a single speech recording. Finally, human experts in voice casting will be added into the subjective evaluation procedure in order to compare the judgements of naive and expert listeners, and to define guidelines for the validation of a voice casting system.

VI. ACKNOWLEDGEMENTS

This research was supported by the European FEDER project VOICE4GAMES (2011-2014).

Copyright (c) 2015 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

REFERENCES

- [1] H. Kido and H. Kasuya, "Everyday Expressions associated with Voice Quality of Normal Utterance — Extraction by Perceptual Evaluation," *Journal of the Acoustic Society of Japan*, vol. 57, no. 5, pp. 337–344, 2001.
- [2] F. Nolan, P. French, K. McDougall, L. Stevens, and T. Hudson, "The Role of Voice Quality 'Settings' in Perceived Voice Similarity," in *International Association for Forensic Phonetics and Acoustics*, Vienna, Austria, 2011.
- [3] Z. Karam, W. M. Campbell, and N. Dehak, "Graph Relational Features for Speaker Recognition and Mining," in *IEEE Statistical Signal Processing Workshop*, 2011, p. 525–528.
- [4] W. M. Campbell and E. Singer, "Query-by-Example using Speaker Content Graphs," in *Interspeech*, Portland, USA, 2012.
- [5] R. Dall, C. Veaux, J. Yamagishi, and S. King, "Analysis of Speaker Clustering Strategies for HMM-based Speech Synthesis," in *Interspeech*, Portland, USA, 2012.
- [6] H. Zen, N. Braunschweiler, S. Buchholz, M. J. F. Gales, K. Knill, S. Krstulovic, and J. Latorre, "Statistical Parametric Speech Synthesis Based on Speaker and Language Factorization," *IEEE transactions on Audio, Speech, and Language Processing*, vol. 20, no. 6, pp. 1713–1724, 2012.
- [7] D. A. Reynolds, T. F. Quatieri, and R. B. Dunn, "Speaker Verification Using Adapted Gaussian Mixture Models," *Digital Signal Processing*, vol. 10, no. 1-3, pp. 19–41, 2000.
- [8] W. M. Campbell, D. E. Sturim, and D. A. Reynolds, "Support Vector Machines using GMM Supervectors for Speaker Verification," *IEEE Signal Processing Letters*, vol. 13, no. 5, pp. 308–311, 2006.
- [9] N. Dehak, P. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-End Factor Analysis for Speaker Verification," *IEEE transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2011.
- [10] B. Schuller, S. Steidl, A. Batliner, F. Schiel, and J. Krajewski, "The INTERSPEECH 2011 Speaker State Challenge," in *Interspeech*, Florence, Italy, 2012.
- [11] B. Schuller, S. Steidl, A. Batliner, E. Nöth, A. Vinciarelli, F. Burkhardt, R. Van Son, F. Weninger, F. Eyben, T. Bocklet, G. Mohammadi, and B. Weiss, "The INTERSPEECH 2012 Speaker Trait Challenge," in *Interspeech*, Portland, USA, 2012.
- [12] N. Obin, A. Roebel, and G. Bachman, "On Automatic Voice Casting for Expressive Speech: Speaker Recognition vs. Speech Classification," in *International Conference on Acoustics, Speech, and Signal Processing*, Florence, Italy, 2014.
- [13] G. Peeters, "A Generic System for Audio Indexing: Application to Speech/Music Segmentation and Music Genre," in *International Conference on Digital Audio Effects*, Bordeaux, France, 2007.
- [14] J.-J. Burred and G. Peeters, "An Adaptive System for Music Classification and Tagging," in *International Workshop on Learning the Semantics of Audio Signals*, Graz, Austria, 2009.
- [15] C. Charbuillet, D. Tardieu, and G. Peeters, "GMM-Supervector for Content based Music Similarity," in *International Conference on Digital Audio Effects*, Paris, France, 2011, pp. 425–428.
- [16] A. Larcher, J.-F. Bonastre, B. Fauve, K. A. Lee, C. Levy, H. Li, J. S. Mason, and J.-Y. Parfait, "ALIZE 3.0 - Open Source Toolkit for State-of-the-Art Speaker Recognition," in *Interspeech*, Lyon, France, 2013.
- [17] C.-C. Chang and C.-J. Lin, "LIBSVM: a Library for Support Vector Machines," 2001.
- [18] A. Hatch, S. Kajarekar, and A. Stolcke, "Within-Class Covariance Normalization for SVM-based Speaker Recognition," in *International Conference on Spoken Language Processing*, Pittsburgh, USA, 2006, pp. 1471–1474.
- [19] D. Garcia-Romero and C. Espy-Wilson, "Analysis of i-vector Length Normalization in Speaker Recognition Systems," in *Interspeech*, 2011, p. 249–252.
- [20] P.-M. Bousquet, A. Larcher, D. Matrouf, J.-F. Bonastre, and O. Plchot, "Variance-Spectra based Normalization for I-vector Standard and Probabilistic Linear Discriminant Analysis," in *Odyssey: The Speaker and Language Recognition Workshop*, Singapore, Singapore, 2012, pp. 157–164.
- [21] N. Dehak, R. Dehak, P. Kenny, N. Brummer, P. Ouellet, and P. Dumouchel, "Support Vector Machines versus Fast Scoring in the Low-Dimensional Total Variability Space for Speaker Verification," in *Interspeech*, 2009, pp. 4237–4240.
- [22] N. Dehak, "Discriminative and Generative Approaches for Long- and Short-Term Speaker Characteristics Modeling: Application to Speaker Verification," PhD. Thesis, Ecole de Technologie Supérieure, 2009.
- [23] S. J. D. Prince and J. H. Elder, "Probabilistic Linear Discriminant Analysis for Inferences about Identity," in *IEEE International Conference on Computer Vision*, Rio de Janeiro, Brazil, 2007, pp. 1751–1758.
- [24] N. Dehak, R. Dehak, J. Glass, D. Reynolds, and P. Kenny, "Bayesian Speaker Verification with Heavy-Tailed Priors," in *Odyssey: The Speaker and Language Recognition Workshop*, Brno, Czech Republic, 2010.
- [25] S. Prince, *Computer Vision: Models Learning and Inference*. Cambridge University Press, 2012.
- [26] M.-L. Zhang and Z.-H. Zhou, "A Review on Multi-Label Learning Algorithms," *IEEE Transactions on Knowledge and Data Engineering*, vol. 26, no. 8, pp. 1819 – 1837, 2014.
- [27] R. Schapire and Y. Singer, "Boostexter: a Boosting-based System for Text Categorization," *Machine Learning*, vol. 39, p. 135–168, 2000.
- [28] I. Katakis, G. Tsoumakas, and I. Vlahavas, "Multilabel Text Classification for Automated Tag Suggestion," in *ECML/PKDD Discovery Challenge*, Antwerp, Belgium, 2008, p. 75–83.
- [29] K. Trohidis, G. Tsoumakas, G. Kalliris, and I. Vlahavas, "Multilabel Classification of Music into Emotions," in *International Conference on Music Information Retrieval (ISMIR)*, Philadelphia, USA, 2008.
- [30] D. Tardieu, C. Charbuillet, F. Cornu, and G. Peeters, "Mirex-2011 Single-Label and Multi-Label Classification Tasks: Ircamclassification2011 Submission," in *Music Information Retrieval Evaluation eXchange (MIREX)*, Miami, USA, 2011.
- [31] M. Boutell, J. Luo, X. Shen, and C. Brown, "Learning Multi-Label Scene Classification," *Pattern Recognition*, vol. 37, no. 9, p. 1757–1771, 2004.
- [32] G.-J. Qi, X.-S. Hua, Y. Rui, J. Tang, T. Mei, and H.-J. Zhang, "Correlative Multi-Label Video Annotation," in *International Conference on Multimedia (MULTIMEDIA)*, New York, USA, 2007, p. 17–26.
- [33] R. Rifkin and A. Klautau, "In Defense of One-Vs-All Classification," *Journal of Machine Learning Research*, vol. 5, pp. 101–141, 2004.
- [34] C. Cortes and V. Vapnik, "Support-Vector Networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [35] T.-F. Wu, C.-J. Lin, and R. C. Weng, "Probability Estimates for Multi-Class Classification by Pairwise Coupling," *Journal of Machine Learning Research*, vol. 5, p. 975–1005, 2004.
- [36] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley, 1991.
- [37] J. V. G. Aradilla and H. Bourlard, "Using Posterior-based Features in Template Matching for Speech Recognition," in *Interspeech*, Pittsburgh, USA, 2006.
- [38] J. Laver, *The Phonetic Description of Voice Quality*. Cambridge: Cambridge University Press, 1980.
- [39] K. R. Scherer, R. Banse, H. G. Wallbott, and T. Goldbeck, "Vocal Cues in Emotion Encoding and Decoding," *Motivation and Emotion*, vol. 15, p. 123–148, 1991.
- [40] B. Schuller, S. Steidl, A. Batliner, F. Burkhardt, L. Devillers, C. Müller, and S. Narayanan, "The INTERSPEECH 2010 Paralinguistic Challenge," in *Interspeech*, Makuhari, Japan, 2010.
- [41] B. Schuller, S. Steidl, and A. Batliner, "The INTERSPEECH 2009 Emotion Challenge," in *Interspeech*, Brighton, UK, 2009.
- [42] M. Kockmann, L. Burget, and J. H. Cernocky, "Brno University of Technology System for Interspeech 2010 Paralinguistic Challenge," in *Interspeech*, Makuhari, Japan, 2010, pp. 2822–2825.
- [43] M. H. Bahari, M. McLaren, H. Van Hamme, and D. Van Leeuwen, "Age Estimation from Telephone Speech using i-vectors," in *Interspeech*, Portland, USA, 2012.
- [44] A. Hassan and R. I. Dampier, "Multi-Class and Hierarchical SVMs for Emotion Recognition," in *Interspeech*, Makuhari, Japan, 2010, pp. 2354–2357.
- [45] F. Burkhardt, A. Paeschke, M. Rolfes, W. Sendmeier, and B. Weiss, "A Database of German Emotional Speech," in *Interspeech*, Lisbon, Portugal, 2005.
- [46] C. Busso, M. Bulut, C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. Chang, S. Lee, and S. Narayanan, "IEMOCAP: Interactive Emotional Dyadic Motion Capture Database," *Journal of Language Resources and Evaluation*, vol. 42, no. 4, pp. 335–359, 2008.
- [47] B. Schuller, M. Valstar, F. Eyben, G. McKeown, R. Cowie, and M. Pantic, "AVEC 2011 – The First International Audio/Visual Emotion Challenge," in *International HUMAINE Association Conference on Affective Computing and Intelligent Interaction*, Memphis, USA, 2011, pp. 415–424.
- [48] N. Obin, "Cries and Whispers - Classification of Vocal Effort in Expressive Speech," in *Interspeech*, Portland, USA, 2012.
- [49] N. Obin and M. Liuni, "On the Generalization of Shannon Entropy for Speech Recognition," in *IEEE workshop on Spoken Language Technology*, Miami, USA, 2012.

- [50] S. Scherer, J. Kane, C. Gobl, and F. Schwenker, "Investigating Fuzzy-Input Fuzzy-Output Support Vector Machines for Robust Voice Quality Classification," *Speech Communication*, vol. 27, no. 1, p. 263–287, 2013.
- [51] G. Degottex and N. Obin, "Phase Distortion Statistics as a Representation of the Glottal Source: Application to the Classification of Voice Qualities," in *Interspeech*, Singapore, Singapore, 2014.
- [52] A. F. Hayes and K. Krippendorff, "Answering the Call for a Standard Reliability Measure for Coding Data," *Communication Methods and Measures*, vol. 1, pp. 77–89, 2007.
- [53] N. Dehak and G. Chollet, "Support Vector GMMs for Speaker Verification," in *Odyssey: The Speaker and Language Recognition Workshop*, San Juan, Puerto Rico, pp. 1–4.
- [54] A. Martin, G. Doddington, T. Kamm, M. Ordowski, and M. Przybocki, "The DET curve in Assessment of Detection Task Performance," in *European Conference on Speech Communication and Technology*, 1997, p. 1895–1898.
- [55] "The NIST Year 2012 Speaker Recognition Evaluation Plan," Tech. Rep., 2012. [Online]. Available: http://www.nist.gov/itl/iad/mig/upload/NIST_SRE12_evalplan-v17-r1.pdf
- [56] D. R. Velez, B. C. White, A. A. Motsinger, W. S. Bush, M. D. Ritchie, S. M. Williams, and J. H. Moore, "A Balanced Accuracy Function for Epistasis Modeling in Imbalanced Datasets using Multifactor Dimensionality Reduction," *Genetic Epidemiology*, vol. 31, no. 4, p. 306–315, 2007.
- [57] H. Wanger, "Measuring Performance in Category Judgment Studies on Nonverbal Behavior," *Journal of Nonverbal Behavior*, vol. 17, no. 1, p. 3–28, 1993.
- [58] C.-C. Lee, E. Mower, C. Busso, S. Lee, and S. Narayanan, "Emotion Recognition using a Hierarchical Binary Decision Tree Approach," *Speech Communication*, vol. 53, p. 1162–1171, 2011.
- [59] B. Schuller and A. Batliner, *Computational Paralinguistics: Emotion, Affect and Personality in Speech and Language Processing*. Wiley, 2013.
- [60] J. Hair, R. Anderson, M. Tatham, and W. Black, *Multivariate data analysis*. Prentice-Hall, 1995.
- [61] N. Obin, "MeLos: Analysis and Modelling of Speech Prosody and Speaking Style," PhD. Thesis, Ircam - UPMC, 2011.



Axel Roebel (M'08) received the Diploma in electrical engineering from Hanover University in 1990 and the Ph.D. degree (summa cum laude) in computer science from the Technical University of Berlin in 1993. In 1994 he joined the German National Research Center for Information Technology (GMD-First) in Berlin where he continued his research on adaptive modeling of time series of nonlinear dynamical systems. In 1996 he became assistant professor for digital signal processing in the communication science department of the Technical University of Berlin. In 2000 he was visiting researcher at CCRMA Stanford University, where he worked on adaptive sinusoidal modeling. In the same year he joined the IRCAM to work on sound analysis, synthesis and transformation algorithms. In summer 2006 he was Edgar-Varese guest professor for computer music at the Electronic studio of the Technical University of Berlin and currently he is head of the Sound Analysis and Synthesis team at IRCAM. His current research interests are related to music and speech signal analysis and transformation.



Nicolas Obin (M'13) is Associate Professor at the Institute for Research and Coordination in Acoustics & Music (IRCAM), and the University of Pierre and Marie Curie (UPMC) - Sorbonne Universités. He received a MSc degree in Acoustics, Signal Processing, and Computer science applied to Music and a Ph.D. in Computer Sciences from the University of Paris VI in 2006 and 2011, respectively. Also, he received a MSc degree in Musicology, Arts, and Aesthetic from the University of Paris VIII in 2007. In 2006, he was a visiting researcher at CNMAT (Center for New Music and Audio Technologies) at the University of California, Berkeley. He conducted his Ph.D. at IRCAM on the modeling of speech prosody and speaking style for text-to-speech synthesis, for which he received the award for the best French Ph.D. thesis in computational sciences from "La Fondation Des Treilles" in 2011. His primary research interests cover audio signal processing and machine learning with applications to voice conversion, speech synthesis, computational linguistics/para-linguistics, and computational auditory scene analysis. His research is devoted to the elaboration of audio technologies for creation, art, and culture.