

Robust Recognition of Simultaneous Speech by a Mobile Robot

Jean-Marc Valin, *Member, IEEE*, Shun'ichi Yamamoto, *Student Member, IEEE*, Jean Rouat, *Senior Member, IEEE*, François Michaud, *Member, IEEE*, Kazuhiro Nakadai, *Member, IEEE*, and Hiroshi G. Okuno, *Senior Member, IEEE*

Abstract—This paper describes a system that gives a mobile robot the ability to perform automatic speech recognition with simultaneous speakers. A microphone array is used along with a real-time implementation of geometric source separation (GSS) and a postfilter that gives a further reduction of interference from other sources. The postfilter is also used to estimate the reliability of spectral features and compute a missing feature mask. The mask is used in a missing feature theory-based speech recognition system to recognize the speech from simultaneous Japanese speakers in the context of a humanoid robot. Recognition rates are presented for three simultaneous speakers located at 2 m from the robot. The system was evaluated on a 200-word vocabulary at different azimuths between sources, ranging from 10° to 90° . Compared to the use of the microphone array source separation alone, we demonstrate an average reduction in relative recognition error rate of 24% with the postfilter and of 42% when the missing features approach is combined with the postfilter. We demonstrate the effectiveness of our multisource microphone array postfilter and the improvement it provides when used in conjunction with the missing features theory.

Index Terms—Cocktail party, geometric source separation (GSS), microphone array, missing feature theory, robot audition, speech recognition.

I. INTRODUCTION

THE human hearing sense is very good at focusing on a single source of interest and following a conversation even when several people are speaking at the same time. This ability is known as the *cocktail party effect* [1]. To operate in human and natural settings, autonomous mobile robots should be able

Manuscript received November 30, 2006. This paper was recommended for publication by Associate Editor Hirai and Editor H. Arai upon evaluation of the reviewers' comments. This work was supported in part by the Canada Research Chair (CRC) Program, by the Natural Sciences and Engineering Research Council of Canada (NSERC), by the Canadian Foundation for Innovation (CFI), by the Ministry of Education, Culture, Sports, Science and Technology (MEXT) and Japan Society for the Promotion of Science (JSPS) under Grant 15200015, and by the Informatics Research Center for Development of Knowledge Society Infrastructure under Grant 16016251. The work of J.-M. Valin was supported by the NSERC, the Quebec Fonds de recherche sur la nature et les technologies, and the JSPS short-term exchange student scholarship. The work of J. Rouat was supported by NSERC, Canada.

J.-M. Valin is with the Commonwealth Scientific and Industrial Research Organization Information and Communication Technologies (CSIRO ICT) Centre, Sydney, Australia (e-mail: jean-marc.valin@csiro.au).

J. Rouat and F. C. Michaud are with the Department of Electrical and Computer Engineering, Université de Sherbrooke, Sherbrooke, QC J1K 2R1, Canada (e-mail: jean.rouat@usherbrooke.ca; francois.michaud@usherbrooke.ca).

S. Yamamoto and H. G. Okuno are with the Department of Intelligence Science and Technology, Graduate School of Informatics, Kyoto University, Kyoto 606-8501, Japan (e-mail: shunichi@kuis.kyoto-u.ac.jp; okuno@kuis.kyoto-u.ac.jp).

K. Nakadai is with the Honda Research Institute Japan Company, Ltd., Saitama 351-0114, Japan (e-mail: nakadai@jp.honda-ri.com).

Digital Object Identifier 10.1109/TRO.2007.900612

to do the same. This means that a mobile robot should be able to separate and recognize all sound sources present in the environment at any moment. This requires the robots not only to detect sounds, but also to locate their origin, separate the different sound sources (since sounds may occur simultaneously), and process all of this data to be able to extract useful information about the world from these sound sources.

Recently, studies on robot audition have become increasingly active [2]–[8]. Most studies focus on sound source localization and separation. Recognition of separated sounds has not been addressed as much, because it requires integration of sound source separation capability with automatic speech recognition, which is not trivial. Robust speech recognition usually assumes source separation and/or noise removal from the feature vectors. When several people speak at the same time, each separated speech signal is severely distorted in spectrum from its original signal. This kind of interference is more difficult to counter than background noise because it is nonstationary and similar to the signal of interest. Therefore, conventional noise reduction techniques such as spectral subtraction [9], used as a front-end of an automatic speech recognizer, usually do not work well in practice.

We propose the use of a microphone array and a sound source localization system integrated with an automatic speech recognizer using the missing feature theory [10], [11] to improve robustness against nonstationary noise. In previous work [5], the missing feature theory was demonstrated using a mask computed from clean (nonmixed) speech. The system we now propose can be used in a real environment by computing the missing feature mask only from the data available to the robot. To do so, a microphone array is used and a missing feature mask is generated based only on the signals available from the array postfiltering module.

This paper focuses on the integration of speech/signal processing and speech recognition techniques into a complete system operating in a real (nonsimulated) environment, demonstrating that such an approach is functional and can operate in real-time. The novelty of this approach lies in the way we estimate the missing feature mask in the speech recognizer and in the tight integration of the different modules.

More specifically, we propose an original way of computing the missing feature mask for the speech recognizer that relies on a measure of frequency bin's quality, estimated by our proposed postfilter. In opposition to most missing feature techniques, our approach does not need estimation of prior characteristics of the corrupting sources or noise. This leads to new capabilities in robot speech recognition with simultaneous speakers. As an

example, for three simultaneous speakers, our system can allow at least three speech recognizers running simultaneously on the three separated speaker signals.

It is one of the first systems that runs in real-time on real robots while performing simultaneous speech recognition. The real-time constraints guided us in the integration of signal and speech processing techniques that are sufficiently fast and efficient. We, therefore, had to reject signal processing techniques that are too complex, even if potentially yielding better performance.

The paper is organized as follows. Section II discusses the state of the art and limitations of speech enhancement and missing feature-based speech recognition. Section III gives an overview of the system. Section IV presents the linear separation algorithm and Section V describes the proposed postfilter. Speech recognition integration and computation of the missing feature mask are shown in Section VI. Results are presented in Section VII, followed by the conclusion.

II. AUDITION IN MOBILE ROBOTICS

Artificial hearing for robots is a research topic still in its infancy, at least when compared to the work already done on artificial vision in robotics. However, the field of artificial audition has been the subject of much research in recent years. In 2004, the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) included, for the first time, a special session on robot audition. Initial work on sound localization by Irie [12] for the Cog [13] and Kismet robots can be found as early as 1995. The capabilities implemented were, however, very limited, partly because of the necessity to overcome hardware limitations.

The SIG robot¹ and its successor SIG2,² both developed at Kyoto University, have integrated increasing auditory capabilities [14]–[20] over the years (from 2000 to now). Both robots are based on binaural audition, which is still the most common form of artificial audition on mobile robots. Original work by Nakadai *et al.* [14], [15] on active audition has made it possible to locate sound sources in the horizontal plane using binaural audition and active behavior to disambiguate front from rear. Later work has focused more on sound source separation [18], [19] and speech recognition [5], [6].

The ROBITA robot, designed at Waseda University, uses two microphones to follow a conversation between two people, originally requiring each participant to wear a headset [21], although a more recent version uses binaural audition [22].

A completely different approach is used by Zhang and Weng [23] in the SAIL robot with the goal of making a robot develop auditory capabilities autonomously. In this case, the *Q-learning* unsupervised learning algorithm is used instead of supervised learning, which is most commonly used in the field of speech recognition. The approach is validated by making the robot learn simple voice commands. Although current speech recognition accuracy using conventional methods is usually higher than the results obtained, the advantage is that the robot learns words autonomously.

More recently, robots have started taking advantage of using more than two microphones. This is the case of the Sony QRIO SDR-4XII robot [24] that features seven microphones. Unfortunately, little information is available regarding the processing done with those microphones. A service robot by Choi *et al.* [25] uses eight microphones organized in a circular array to perform speech enhancement and recognition. The enhancement is provided by an adaptive beamforming algorithm. Work by Asano, Asoh, and others [2], [26], [27] also uses a circular array composed of eight microphones on a mobile robot to perform both localization and separation of sound sources. In more recent work [28], particle filtering is used to integrate vision and audition in order to track sound sources.

In general, human–robot interface is a popular area of audition-related research in robotics. Work on robot audition for human–robot interface has also been done by Prodanov *et al.* [29] and Theobalt *et al.* [30], based on a single microphone near the speaker. Even though human–robot interface is the most common goal of robot audition research, there is research being conducted for other goals. Huang *et al.* [31] use binaural audition to help robots navigate in their environment, allowing a mobile robot to move toward sound-emitting objects without colliding with those objects. The approach even works when those objects are not visible (i.e., not in line of sight), which is an advantage over vision.

III. SYSTEM OVERVIEW

One goal of the proposed system is to integrate the different steps of source separation, speech enhancement, and speech recognition as closely as possible to maximize recognition accuracy by using as much of the available information as possible and with a strong real-time constraint. We use a microphone array composed of omnidirectional elements mounted on the robot. The missing feature mask is generated in the time–frequency plane since the separation module and the postfilter already use this signal representation. We assume that all sources are detected and localized by an algorithm such as [32], [33], although our approach is not specific to any localization algorithm. The estimated location of the sources is used by a linear separation algorithm. The separation algorithm we use is a modified version of the geometric source separation (GSS) approach proposed by Parra and Alvino [34], designed to suit our needs for real-time- and real-life applications. We show that it is possible to implement the separation with relatively low complexity that grows linearly with the number of microphones. The method is interesting for use in the mobile robotics context because it makes it easy to dynamically add or remove sound sources as they appear or disappear. The output of the GSS still contains residual background noise and interference, which we further attenuate through a multichannel postfilter. The novel aspect of this postfilter is that, for each source of interest, the noise estimate is decomposed into stationary and transient components assumed to be due to leakage between the output channels of the initial separation stage. In the results, the performance of that postfilter is shown to be superior to those obtained when considering each separated source independently.

¹<http://winnie.kuis.kyoto-u.ac.jp/SIG/oldsig/>

²<http://winnie.kuis.kyoto-u.ac.jp/SIG/>

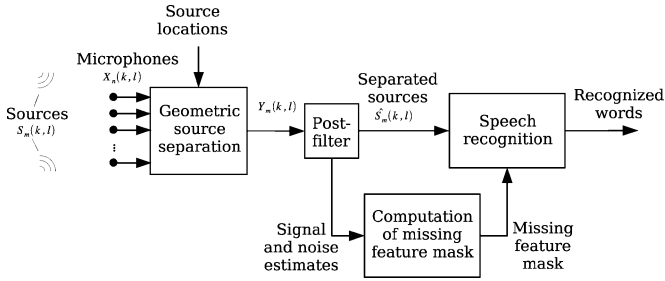


Fig. 1. Overview of the separation system with the postfilter being used both to improve the audio quality and to estimate the missing feature mask.

The postfilter that we use can not only reduce the amount of noise and interference, but its behavior provides useful information, which is used to evaluate the reliability of different regions of the time–frequency plane for the separated signals. Based also on the ability of the postfilter to model independently background noise and interference, we propose a novel way to estimate the missing feature mask to further improve speech recognition accuracy. This also has the advantage that acoustic models trained on clean data can be used and that no multicondition training is required.

The structure of the proposed system is shown in Fig. 1 and its four main parts are:

- 1) linear separation of the sources, implemented as a variant of the GSS algorithm;
- 2) multichannel postfiltering of the separated output;
- 3) computation of the missing feature mask from the postfilter output;
- 4) speech recognition using the separated audio and the missing feature mask.

IV. GEOMETRIC SOURCE SEPARATION

Although the work we present can be adapted to systems with any linear source separation algorithm, we propose to use the GSS algorithm because it is simple and well suited to a mobile robotics application. More specifically, the approach has the advantage that it can make use of the location of the sources. In this work, we only make use of the direction information, which can be obtained with a high degree of accuracy using the method described in [3]. It was shown in [32] that distance can be estimated as well. The use of location information is important when new sources are observed. In that situation, the system can still provide acceptable separation performance (at least equivalent to the delay-and-sum beamformer), even if the adaptation has not yet taken place.

The method operates in the frequency domain using a frame length of 21 ms (1024 samples at 48 kHz). Let $S_m(k, \ell)$ be the real (unknown) sound source m at time-frame ℓ and for discrete frequency k . We denote as $\mathbf{s}(k, \ell)$ the vector of the sources $S_m(k, \ell)$ and matrix $\mathbf{A}(k)$ as the transfer function from the sources to the microphones. The signal received at the microphones is, thus, given by

$$\mathbf{x}(k, \ell) = \mathbf{A}(k)\mathbf{s}(k, \ell) + \mathbf{n}(k, \ell) \quad (1)$$

where $\mathbf{n}(k, \ell)$ is the noncoherent background noise received at the microphones. The matrix $\mathbf{A}(k)$ can be estimated using the result of a sound localization algorithm by assuming that all transfer functions have unity gain and that no diffraction occurs. The elements of $\mathbf{A}(k)$ are, thus, expressed as

$$a_{ij}(k) = e^{-j2\pi k\delta_{ij}} \quad (2)$$

where δ_{ij} is the time delay (in samples) to reach microphone i from source j .

The separation result is then defined as $\mathbf{y}(k, \ell) = \mathbf{W}(k, \ell)\mathbf{x}(k, \ell)$, where $\mathbf{W}(k, \ell)$ is the separation matrix that must be estimated. This is done by providing two constraints (the index ℓ is omitted for the sake of clarity):

- 1) decorrelation of the separation algorithm outputs (second-order statistics are sufficient for nonstationary sources), expressed as $\mathbf{R}_{\mathbf{y}\mathbf{y}}(k) - \text{diag}[\mathbf{R}_{\mathbf{y}\mathbf{y}}(k)] = \mathbf{0}$.
- 2) geometric constraint $\mathbf{W}(k)\mathbf{A}(k) = \mathbf{I}$, which ensures unity gain in the direction of the source of interest and places zeros in the direction of interferences.

In theory, constraint 2) could be used alone for separation (the method is referred to as LS-C2 [34]), but this is insufficient in practice, as the method does not take into account reverberation or errors in localization. It is also subject to instability if $\mathbf{A}(k)$ is not invertible at a specific frequency. When used together, constraints 1) and 2) are too strong. For this reason, we use a “soft” constraint (referred to as GSS-C2 in [34]) combining 1) and 2) in the context of a gradient descent algorithm.

Two cost functions are created by computing the square of the error associated with constraints 1) and 2). These cost functions are defined as, respectively

$$J_1(\mathbf{W}(k)) = \|\mathbf{R}_{\mathbf{y}\mathbf{y}}(k) - \text{diag}[\mathbf{R}_{\mathbf{y}\mathbf{y}}(k)]\|^2 \quad (3)$$

$$J_2(\mathbf{W}(k)) = \|\mathbf{W}(k)\mathbf{A}(k) - \mathbf{I}\|^2 \quad (4)$$

where the matrix norm is defined as $\|\mathbf{M}\|^2 = \text{trace}[\mathbf{M}\mathbf{M}^H]$ and is equal to the sum of the square of all elements in the matrix. The gradient of the cost functions with respect to $\mathbf{W}(k)$ is equal to [34]

$$\frac{\partial J_1(\mathbf{W}(k))}{\partial \mathbf{W}^*(k)} = 4\mathbf{E}(k)\mathbf{W}(k)\mathbf{R}_{\mathbf{x}\mathbf{x}}(k) \quad (5)$$

$$\frac{\partial J_2(\mathbf{W}(k))}{\partial \mathbf{W}^*(k)} = 2[\mathbf{W}(k)\mathbf{A}(k) - \mathbf{I}]\mathbf{A}(k) \quad (6)$$

where $\mathbf{E}(k) = \mathbf{R}_{\mathbf{y}\mathbf{y}}(k) - \text{diag}[\mathbf{R}_{\mathbf{y}\mathbf{y}}(k)]$.

The separation matrix $\mathbf{W}(k)$ is then updated as follows:

$$\mathbf{W}^{n+1}(k) = \mathbf{W}^n(k) - \mu \left[\alpha(k) \frac{\partial J_1(\mathbf{W}(k))}{\partial \mathbf{W}^*(k)} + \frac{\partial J_2(\mathbf{W}(k))}{\partial \mathbf{W}^*(k)} \right] \quad (7)$$

where $\alpha(f)$ is an energy normalization factor equal to $\|\mathbf{R}_{\mathbf{x}\mathbf{x}}(k)\|^{-2}$ and μ is the adaptation rate.

The difference between our implementation and the original GSS algorithm described in [34] lies in the way the correlation matrices $\mathbf{R}_{\mathbf{x}\mathbf{x}}(k)$ and $\mathbf{R}_{\mathbf{y}\mathbf{y}}(k)$ are computed. Instead of using several seconds of data, our approach uses instantaneous estimates, as used in the stochastic gradient adaptation of the least-mean square (LMS) adaptive filter [35]. We, thus, assume

that

$$\mathbf{R}_{\mathbf{x}\mathbf{x}}(k) = \mathbf{x}(k)\mathbf{x}(k)^H \quad (8)$$

$$\mathbf{R}_{\mathbf{y}\mathbf{y}}(k) = \mathbf{y}(k)\mathbf{y}(k)^H. \quad (9)$$

It is, then, possible to rewrite (5) as

$$\frac{\partial J_1(\mathbf{W}(k))}{\partial \mathbf{W}^*(k)} = 4 [\mathbf{E}(k)\mathbf{W}(k)\mathbf{x}(k)] \mathbf{x}(k)^H \quad (10)$$

which only requires matrix-by-vector products, greatly reducing the complexity of the algorithm. Similarly, the normalization factor $\alpha(k)$ can also be simplified as $[\|\mathbf{x}(k)\|^2]^{-2}$. With a small update rate, it means that the time averaging is performed implicitly. In early experiments, the instantaneous estimate of the correlation was found to have no significant impact on the performance of the separation, but is necessary for real-time implementation.

The weight initialization we use corresponds to a delay-and-sum beamformer, referred to as the I1 (or C1) initialization method in [34]. Such initialization ensures that prior to adaptation, the performances are at worst equivalent to a delay-and-sum beamformer. In fact, if only a single source is present, our algorithm is strictly equivalent to a delay-and-sum beamformer implemented in the frequency domain.

V. MULTICHANNEL POSTFILTER

To enhance the output of the GSS algorithm presented in Section IV, we derive a frequency-domain postfilter that is based on the optimal estimator originally proposed by Ephraim and Malah [36], [37]. Several approaches to microphone array postfiltering have been proposed in the past. Most of these postfilters address reduction of stationary background noise [38], [39]. Recently, a multichannel postfilter taking into account nonstationary interferences was proposed by Cohen [40]. The novelty of our postfilter resides in the fact that, for a given channel output of the GSS, the transient components of the corrupting sources are assumed to be due to leakage from the other channels during the GSS process. Furthermore, for a given channel, the stationary and the transient components are combined into a single noise estimator used for noise suppression, as shown in Fig. 2. In addition, we explore different suppression criteria (α values) for optimizing speech recognition instead of perceptual quality. Again, when only one source is present, this postfilter is strictly equivalent to standard single-channel noise suppression techniques.

A. Noise Estimation

This section describes the estimation of noise variances that are used to compute the weighting function $G_m(k, \ell)$ by which the outputs $Y_m(k, \ell)$ of the GSS are multiplied to generate a clean signal whose spectrum is denoted $\hat{S}_m(k, \ell)$. The noise variance estimation $\lambda_m(k, \ell)$ is expressed as

$$\lambda_m(k, \ell) = \lambda_m^{\text{stat.}}(k, \ell) + \lambda_m^{\text{leak}}(k, \ell) \quad (11)$$

where $\lambda_m^{\text{stat.}}(k, \ell)$ is the estimate of the stationary component of the noise for source m at frame ℓ for frequency k , and $\lambda_m^{\text{leak}}(k, \ell)$ is the estimate of source leakage.

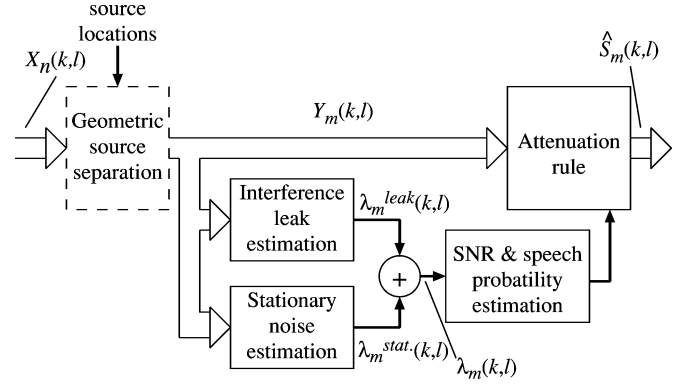


Fig. 2. Overview of the postfilter. $X_n(k, \ell)$, $n = 0 \dots N - 1$: Microphone inputs, $Y_m(k, \ell)$, $m = 0 \dots M - 1$: Inputs to the postfilter, $\hat{S}_m(k, \ell) = G_m(k, \ell)Y_m(k, \ell)$, $m = 0 \dots M - 1$: Postfilter outputs.

We compute the stationary noise estimate $\lambda_m^{\text{stat.}}(k, \ell)$ using the minima-controlled recursive average (MCRA) technique proposed by Cohen [41].

To estimate λ_m^{leak} , we assume that the interference from other sources has been reduced by a factor η (typically $-10 \text{ dB} \leq \eta \leq -3 \text{ dB}$) by the separation algorithm (GSS). The leakage estimate is, thus, expressed as

$$\lambda_m^{\text{leak}}(k, \ell) = \eta \sum_{i=0, i \neq m}^{M-1} Z_i(k, \ell) \quad (12)$$

where $Z_m(k, \ell)$ is the smoothed spectrum of the m^{th} source, $Y_m(k, \ell)$, and is recursively defined (with $\alpha_s = 0.7$) as

$$Z_m(k, \ell) = \alpha_s Z_m(k, \ell - 1) + (1 - \alpha_s) |Y_m(k, \ell)|^2. \quad (13)$$

It is worth noting that if $\eta = 0$ or $M = 1$, then the noise estimate becomes $\lambda_m(k, \ell) = \lambda_m^{\text{stat.}}(k, \ell)$ and our multisource postfilter is reduced to a single-source postfilter.

B. Suppression Rule

From here on, unless otherwise stated, the m index and the ℓ arguments are omitted for clarity and the equations are given for each m and for each ℓ . The proposed noise suppression rule is based on minimum mean-square error (MMSE) estimation of the spectral amplitude in the $(|X(k)|^\alpha)$ domain. The power coefficient α is chosen to maximize the recognition results.

Assuming that speech is present, the spectral amplitude estimator is defined by

$$\hat{A}(k) = (E[|S(k)|^\alpha | Y(k)])^{\frac{1}{\alpha}} = G_{H_1}(k) |Y(k)| \quad (14)$$

where $G_{H_1}(k)$ is the spectral gain assuming that speech is present.

The spectral gain for arbitrary α is derived from [37, eq. (13)]

$$G_{H_1}(k) = \frac{\sqrt{v(k)}}{\gamma(k)} \left[\Gamma\left(1 + \frac{\alpha}{2}\right) M\left(-\frac{\alpha}{2}; 1; -v(k)\right) \right]^{\frac{1}{\alpha}} \quad (15)$$

where $M(a; c; x)$ is the confluent hypergeometric function, $\gamma(k) \triangleq |Y(k)|^2/\lambda(k)$ and $\xi(k) \triangleq E[|S(k)|^2]/\lambda(k)$ are,

respectively, the a posteriori signal-to-noise ratio (SNR) and the a priori SNR. We also have $v(k) \triangleq \gamma(k)\xi(k)/(\xi(k) + 1)$ [36].

The a priori SNR $\xi(k)$ is estimated recursively as [36]

$$\hat{\xi}(k, \ell) = \alpha_p G_{H_1}^2(k, \ell - 1) \gamma(k, \ell - 1) + (1 - \alpha_p) \max\{\gamma(k, \ell) - 1, 0\}. \quad (16)$$

When taking into account the probability of speech presence, we obtain the modified spectral gain

$$G(k) = p^{1/\alpha(k)} G_{H_1}(k) \quad (17)$$

where $p(k)$ is the probability that speech is present in the frequency band k and given by

$$p(k) = \left\{ 1 + \frac{\hat{q}(k)}{1 - \hat{q}(k)} (1 + \xi(k)) \exp(-v(k)) \right\}^{-1}. \quad (18)$$

The a priori probability of speech presence $\hat{q}(k)$ is computed as in [41] using speech measurements on the current frame for a local frequency window, a larger frequency, and for the whole frame.

VI. INTEGRATION WITH SPEECH RECOGNITION

Robustness against noise in conventional³ automatic speech recognition (ASR) is being extensively studied, in particular, in the AURORA project [42], [43]. To realize noise-robust speech recognition, *multicondition training* (training on a mixture of clean speech and noises) has been studied [44], [45]. This is currently the most common method for vehicle and telephone applications. Because an acoustic model obtained by multicondition training reflects all expected noises in specific conditions, recognizer's use of the acoustic model is effective as long as the noise is stationary. This assumption holds, for example, with background noise in a vehicle and on a telephone. However, multicondition training is not effective for mobile robots, since those usually work in dynamically changing noisy environments and, furthermore, multicondition training requires an important amount of data to learn from.

Source separation and speech enhancement algorithms for robust recognition are another potential alternative for automatic speech recognition on mobile robots. However, their common use is to maximize the perceptual quality of the resulting signal. This is not always effective since most preprocessing source separation and speech enhancement techniques distort the spectrum and, consequently, degrade features, reducing the recognition rate (even if the signal is perceived to be cleaner by naïve listeners [46]). For example, the work of Seltzer *et al.* [47] on microphone arrays addresses the problem of optimizing the array processing specifically for speech recognition (and not for a better perception). Recently, Araki *et al.* [48] have applied ICA to the separation of three sources using only two microphones. Aarabi and Shi [49] have shown speech-enhancement feasibility, for speech recognition, using only the phase of the signals from an array of microphones.

³We use conventional in the sense of speech recognition for applications where a single microphone is used in a static environment such as a vehicle or an office.

A. Missing Features Theory and Speech Recognition

Research of confidence islands in the time–frequency plane representation has been shown to be effective in various applications and can be implemented with different strategies. One of the most effective is the missing feature strategy. Cooke *et al.* [50], [51] propose a probabilistic estimation of a mask in regions of the time–frequency plane where the information is not reliable. Then, after masking, the parameters for speech recognition are generated and can be used in conventional speech recognition systems. They obtain a significant increase in recognition rates without any explicit modeling of the noise [52]. In this scheme, the mask is essentially based on the dominance speech/interference criteria and a probabilistic estimation of the mask is used.

Conventional missing feature theory-based ASR is a hidden Markov model (HMM)-based recognizer, whose output probability (emission probability) is modified to keep only the reliable feature distributions. According to the work by Cooke *et al.* [51], HMMs are trained on clean data. Density in each state S_i is modeled using mixtures of M Gaussians with diagonal-only covariance.

Let $f(\mathbf{x}|S)$ be the output probability density of feature vector \mathbf{x} in state S_i , and $P(j|S_i)$ represent the mixture coefficients expressed as a probability. The output probability density is defined by

$$f(\mathbf{x}|S_i) = \sum_{j=1}^M P(j|S_i) f(\mathbf{x}|j, S_i). \quad (19)$$

Cooke *et al.* [51] propose to transform (19) to take into consideration the only reliable features x_r from \mathbf{x} and to remove the unreliable features. This is equivalent to using the marginalization probability density functions $f(x_r|j, S_i)$ instead of $f(\mathbf{x}|j, S_i)$ by simply implementing a binary mask. Consequently, only reliable features are used in the probability calculation, and the recognizer can avoid undesirable effects due to unreliable features.

Hugo van Hamme [53] formulates the missing feature approach for speech recognizers using conventional parameters such as mel frequency cepstral coefficients (MFCC). He uses data imputation according to Cooke [51] and proposes a suitable transformation to be used with MFCC for missing features. The acoustic model evaluation of the unreliable features is modified to express that their clean values are unknown or confined within bounds. In a more recent paper, Hugo van Hamme [54] presents speech recognition results by integrating harmonicity in the SNR for noise estimation. He uses only static MFCC as, according to his observations, dynamic MFCC do not increase sufficiently the speech recognition rate when used in the context of missing features framework. The need to estimate pitch and voiced regions in the time–space representation is a limit to this approach. In a similar approach, Raj *et al.* [55] propose to modify the spectral representation to derive cepstral vectors. They present two missing feature algorithms that reconstruct spectrograms from incomplete noisy spectral representations (*masked* representations). Cepstral vectors can be derived from the reconstructed spectrograms for missing feature recognition. Seltzer

et al. [56] propose the use of a Bayesian classifier to determine the reliability of spectrographic elements. Ming, Jancovic, and others [57], [58] propose the *probabilistic union model* as an alternative to the missing feature framework. According to the authors, methods based on the missing feature framework usually require the identification of the noisy bands. This identification can be difficult for noise with unknown, time-varying band characteristics. They designed an approach for speech recognition involving partial, unknown corrupted frequency-bands. In their approach, they combine the local frequency-band information based on the union of random events, to reduce the dependence of the model on information about the noise. Cho and Oh [59] apply the *union model* to improve robust speech recognition based on frequency bands selection. From this selection, they generate “channel-attentive” mel frequency cepstral coefficients. Even if the use of missing features for robust recognition is relatively recent, many applications have already been designed.

To avoid the use of multicondition training, we propose to merge a multimicrophone source separation and speech enhancement system with the missing feature approach. Very little work has been done with arrays of microphones in the context of missing feature theory. To our knowledge, only McCowan *et al.* [60] apply the missing feature framework to microphone arrays. Their approach defines a missing feature mask based on the input-to-output ratio of a postfilter, but is, however, only validated on stationary noise.

Some missing feature mask techniques can also require the estimation of prior characteristics of the corrupting sources or noise. They usually assume that the noise or interference characteristics vary slowly with time. This is not possible in the context of a mobile robot. We propose to estimate quasi-instantaneously the mask (without preliminary training) by exploiting the postfilter outputs along with the local gains (in the time–frequency plane representation) of the postfilter. These local gains are used to generate the missing feature mask. Thus, the speech recognizer with clean acoustic models can adapt to the distorted sounds by consulting the postfilter feature missing masks. This approach is also a solution to the automatic generation of simultaneous missing feature masks (one for each speaker). It allows the use of simultaneous speech recognizers (one for each separated sound source) with their own mask.

B. Reliability Estimation

The postfilter uses adaptive spectral estimation of background noise and interfering sources to enhance the signal produced during the initial separation. The main idea lies in the fact that, for each source of interest, the noise estimate is decomposed into stationary and transient components assumed to be due to leakage between the output channels of the initial separation stage. It also provides useful information concerning the amount of noise present at a certain time, for each particular frequency. Hence, we use the postfilter to estimate a missing feature mask that indicates how reliable each spectral feature is when performing recognition.

C. Computation of Missing Feature Masks

The missing feature mask is a matrix representing the reliability of each feature in the time–frequency plane. More specifically, this reliability is computed for each frame and for each mel-frequency band. This reliability can be either a continuous value from 0 to 1, or a discrete value of 0 or 1. In this paper, discrete masks are used. It is worth mentioning that computing the mask in the mel-frequency band domain means that it is not possible to use MFCC features, since the effect of the DCT cannot be applied to the missing feature mask.

For each mel-frequency band, the feature is considered reliable if the ratio of the postfilter output energy over the input energy is greater than a threshold T . The reason for this choice is that it is assumed that the more noise is present in a certain frequency band, the lower the postfilter gain will be for that band.

One of the dangers of computing missing feature masks based on an SNR measure is that there is a tendency to consider all silent periods as nonreliable, because they are dominated by noise. This leads to large time–frequency areas where no information is available to the ASR, preventing it from correctly identifying silence (we made this observation from practice). For this reason, it is desirable to consider as reliable at least some of the silence, especially when there is no nonstationary interference.

The missing feature mask is computed in two steps: for each frame ℓ and for each mel frequency band i

- 1) We compute a continuous mask $m_\ell(i)$ that reflects the reliability of the band

$$m_\ell(i) = \frac{S_\ell^{\text{out}}(i) + N_\ell(i)}{S_\ell^{\text{in}}(i)} \quad (20)$$

where $S_\ell^{\text{in}}(i)$ and $S_\ell^{\text{out}}(i)$ are, respectively, the postfilter input and output energy for frame ℓ at mel-frequency band i , and $N_\ell(i)$ is the background noise estimate. The values $S_\ell^{\text{in}}(i)$, $S_\ell^{\text{out}}(i)$, and $N_\ell(i)$ are computed using a mel-scale filterbank with triangular bandpass filters, based on linear-frequency postfilter data.

- 2) We deduce a binary mask $M_\ell(i)$. This mask will be used to remove the unreliable mel frequency bands at frame ℓ

$$M_\ell(i) = \begin{cases} 1, & m_\ell(i) > T \\ 0, & \text{otherwise} \end{cases} \quad (21)$$

where T is the mask threshold. We use the value $T = 0.25$, which produces the best results over a range of experiments. In practice the algorithm is not very sensitive to T and all values in the $[0.15, 0.30]$ interval generally produce equivalent results.

In comparison to McCowan *et al.* [60], the use of the multisource postfilter allows a better reliability estimation by distinguishing between interference and background noise. We include the background noise estimate $N_\ell(i)$ in the numerator of (20) to ensure that the missing feature mask equals 1 when no speech source is present (as long as there is no interference). Using a more conventional postfilter as proposed by McCowan *et al.* [60] and Cohen *et al.* [40] would not allow the mask to preserve




Fig. 3. Spectrograms for separation of three speakers, 90° apart with postfilter. (a) Signal as captured at microphone #1. (b) Separated right speaker. (c) Separated center speaker. (d) Separated left speaker. (e)–(g) Corresponding mel-frequency missing feature mask for static features with *reliable features* ($M_\ell(i) = 1$) shown in black. Time is represented on the x -axis and frequency (0–8 kHz) on the y -axis.

silence features, which is known to degrade ASR accuracy. The distinction between background noise and interference also reflects the fact that background noise cancellation is generally more efficient than interference cancellation.

An example of a computed missing feature mask is shown in Fig. 3. It is observed that the mask indeed preserves the silent periods and considers unreliable the regions of the spectrum dominated by other sources. The missing feature mask for delta-features is computed using the mask for the static features. The dynamic mask $\Delta M_\ell(i)$ is computed as

$$\Delta M_\ell(i) = \prod_{k=-2}^2 M_{\ell-k}(i) \quad (22)$$

and is nonzero only when all the mel features used to compute the delta-cepstrum are deemed reliable.

D. Speech Analysis for Missing Feature Masks

Since MFCC cannot be easily used directly with a missing feature mask and as the postfilter gains are expressed in the time–frequency plane, we use spectral features that are derived from MFCC features with the inverse discrete cosine transform (IDCT). The detailed steps for feature generation are as follows.

- 1) [FFT] The speech signal sampled at 16 kHz is analyzed using a 400-sample FFT with a 160-sample frame shift.
- 2) [Mel] The spectrum is analyzed by a 24th-order mel-scale filter bank.

- 3) [Log] The mel-scale spectrum of the 24th order is converted to log-energies.
- 4) [DCT] The log mel-scale spectrum is converted by discrete cosine transform to the Cepstrum.
- 5) [Lifter] Cepstral features 0 and 13–23 are set to zero so as to make the spectrum smoother.
- 6) [CMS] Convolutional effects are removed using Cepstral mean subtraction.
- 7) [IDCT] The normalized Cepstrum is transformed back to the log mel-scale spectral through an inverse DCT.
- 8) [Differentiation] The features are differentiated in the time, producing 24 delta features in addition to the static features.

The [CMS] step is necessary to remove the effect of convolutional noise, such as reverberation and microphone frequency response.

The same features are used for training and evaluation. Training is performed on clean speech, without any effect from the postfilter. In practice, this means that the acoustic model does not need to be adapted in any way to our method. During evaluation, the only difference with a conventional ASR is the use of the missing feature mask as represented in (19).

E. The Missing Feature-Based Automatic Speech Recognizer

Let $f(x|s)$ be the output probability density of feature vector x in state S . The output probability density is defined by (19),




Fig. 4. SIG 2 robot with eight microphones (two are occluded).

and becomes

$$f(x | S) = \sum_{k=1}^M P(k | S) f(x_r | k, S) \quad (23)$$

where M is the dimensionality of the Gaussian mixture, and x_r are the reliable features in x . This means that only reliable features are used in probability calculation, and, thus, the recognizer can avoid undesirable effects due to unreliable features. We used two speech recognizers. The first one is based on the CASA Tool Kit (CTK) [52] hosted at Sheffield University, U.K.⁴ and the second one is the Julius open-source Japanese ASR [61] that we extended to support the previously mentioned decoding process.⁵ According to our preliminary experiments with these two recognizers, CTK provides slightly better recognition accuracy, while Julius runs much faster.

VII. RESULTS

Our system is evaluated on the SIG2 humanoid robot, on which eight omnidirectional (for the system to work in all directions) microphones are installed as shown in Fig. 4. The microphone positions are constrained by the geometry of the robot because the system is designed to be fitted on any robot. All microphones are enclosed within a 22 cm \times 17 cm \times 47 cm bounding box. To test the system, three Japanese speakers (two males, one female) are recorded simultaneously: one in front, one on the left, and one on the right. In nine different experiments, the angle between the center speaker and the side speakers is varied from 10 to 90° degrees. The speakers are placed 2-m away from the robot, as shown in Fig. 5. The distance between the speakers and the robot was not found to have a significant impact on the performance of the system. The only exception is for short distances (<50 cm) where performance decreases due to the far-field assumption we make in this particular work. The position of the speakers used for the GSS algorithm is computed automatically using the algorithm described in [3].

The room in which the experiment took place is 5 m \times 4 m and has a reverberation time (−60 dB) of approximately 0.3 s.




Fig. 5. Position of the speakers relative to the robot in the experimental setup.

The postfilter parameter $\alpha = 1$ (corresponding to a short-term spectral amplitude (STSA) MMSE estimator) is used since it was found to maximize speech recognition accuracy.⁶ When combined together, the GSS, postfilter, and missing feature mask computation require 25% of a 1.6-GHz Pentium-M to run in real-time when three sources are present.⁷ Speech recognition complexity is not reported as it usually varies greatly between different engines and settings.

A. Separated Signals

Spectrograms showing separation of the three speakers⁸ are shown in Fig. 3, along with the corresponding mask for static features. Even though the task involves nonstationary interference with the same frequency content as the signal of interest, we observe that our postfilter is able to remove most of the interference. Informal subjective evaluation has confirmed that the postfilter has a positive impact on both quality and intelligibility of the speech. This is confirmed by improved recognition results.

B. Speech Recognition Accuracy

We report speech recognition experiments obtained using the CTK toolkit. Isolated word recognition on Japanese words is performed using a triphone acoustic model. We use a speaker-independent three-state model trained on 22 speakers (10 males, 12 females), not present in the test set. The test set includes 200 different ATR phonetically-balanced isolated Japanese words (300 s) for each of the three speakers and is used on a 200-word vocabulary (each word spoken once). Speech recognition accuracy on the clean data (no interference, no noise) varies between 94 and 99%.

Speech recognition accuracy results are presented for five different conditions:

- 1) single-microphone recording;
- 2) GSS only;
- 3) GSS with postfilter (GSS+PF);
- 4) GSS with postfilter using MFCC features (GSS+PF w/ MFCC);
- 5) GSS with postfilter and missing feature mask (GSS+PF+MFT).

⁶The difference between $\alpha = 1$ and $\alpha = 2$ on a subset of the test set was less than one percent in recognition rate

⁷Source code for part of the proposed system is available at <http://manyyears.sourceforge.net/>

⁸Audio signals and spectrograms for all three sources are available at: <http://www.gel.usherbrooke.ca/laborius/projects/Audible/sap/>

⁴<http://www.dcs.shef.ac.uk/research/groups/spandh/projects/respite/ctk/>

⁵<http://julius.sourceforge.jp/>




Fig. 6. Speech recognition accuracy results for intervals ranging from 10 to 90° averaged over the three speakers.

Results are shown in Fig. 6 as a function of the angle between sources and averaged over the three simultaneous speakers. As expected, the separation problem becomes more difficult as sources are located closer to each other because the difference in the transfer functions becomes smaller. We find that the proposed system (GSS+PF+MFT) provides a reduction in relative error rate compared to GSS alone that ranges from 10 to 55%, with an average of 42%. The postfilter provides an average of 24% relative error-rate reduction over use of GSS alone. The relative error-rate reduction is computed as the difference in errors divided by the number of errors in the reference setup. The results of the postfilter with MFCC features (4) are included to show that the use of mel spectral features only has a small effect on the ASR accuracy.

While they seem poor, the results with GSS can only be explained by the highly nonstationary interference coming from the two other speakers (especially when the speakers are close to each other) and the fact that the microphones' placement is constrained by the robot dimensions. The single-microphone results are provided only as a baseline. The results are very low because a single omnidirectional microphone does not have any acoustic directivity.

In Fig. 7 we compare the accuracy of the multisource postfilter to that of a "classic" (single-source) postfilter that removes background noise but does not take interference from other sources into account ($\eta = 0$). Because the level of background noise is very low, the single-source postfilter has almost no effect and most of the accuracy improvement is due to the multisource version of the postfilter, which can effectively remove part of the interference from the other sources. The proposed multisource postfilter was also shown in [62] to be more effective for multiple sources than the multichannel approach in [40].

VIII. CONCLUSION

In this paper we demonstrate a complete multimicrophone speech recognition system capable of performing speech recognition on three simultaneous speakers. The system closely in-

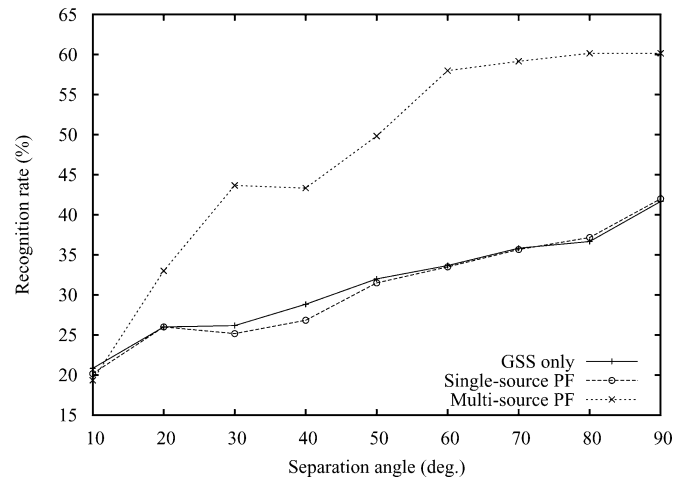


Fig. 7. Effect of the multisource postfilter on speech recognition accuracy.

tegrates all stages of source separation and missing features recognition so as to maximize accuracy in the context of simultaneous speakers. We use a linear source separator based on a simplification of the GSS algorithm. The nonlinear postfilter that follows the initial separation step is a short-term spectral amplitude MMSE estimator. It uses a background noise estimate as well as information from all other sources obtained from the GSS algorithm.

In addition to removing part of the background noise and interference from other sources, the postfilter is used to compute a missing feature mask representing the reliability of mel spectral features. The mask is designed so that only spectral regions dominated by interference are marked as unreliable. When compared to the GSS alone, the postfilter contributes to a 24% (relative) reduction in the word error rate while the use of the missing feature theory-based modules yields a reduction of 42% (also when compared to GSS alone). The approach is specifically designed for recognition on multiple sources and we did not attempt to improve speech recognition of a single source with background noise. In fact, for a single sound source, the proposed work is strictly equivalent to commonly used single-source techniques.

We have shown that robust simultaneous speakers speech recognition is possible when combining the missing feature framework with speech enhancement and source separation with an array of eight microphones. To our knowledge, there is no work reporting multispeaker speech recognition using missing feature theory. This is why this paper is meant more as a proof of concept for a complete auditory system than a comparison between algorithms for performing specific signal processing tasks. Indeed, the main challenge here is the adaptation and integration of the algorithms on a mobile robot so that the system can work in a real environment (moderate reverberation) and that real-time speech recognition with simultaneous speakers be possible.

In future work, we plan to perform the speech recognition with moving speakers and adapt the postfilter to work even in highly reverberant environments, in the hope of developing new capabilities for natural communication between robots and

humans. Also, we have shown that the cepstral-domain speech recognition usually performs slightly better. Hence, it would be desirable for the technique to be generalized for the use of cepstral features instead of spectral features.

REFERENCES

- [1] E. Cherry, "Some experiments on the recognition of speech, with one and with two ears," *J. Acoust. Soc. Am.*, vol. 25, no. 5, pp. 975–979, 1953.
- [2] F. Asano, M. Goto, K. Itou, and H. Asoh, "Real-time source localization and separation system and its application to automatic speech recognition," in *Proc. Eurospeech*, 2001, pp. 1013–1016.
- [3] J.-M. Valin, F. Michaud, B. Hadjou, and J. Rouat, "Localization of simultaneous moving sound sources for mobile robot using a frequency-domain steered beamformer approach," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2004, vol. 1, pp. 1033–1038.
- [4] J.-M. Valin, J. Rouat, and F. Michaud, "Enhanced robot audition based on microphone array source separation with post-filter," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2004, pp. 2123–2128.
- [5] S. Yamamoto, K. Nakadai, H. Tsujino, and H. Okuno, "Assessment of general applicability of robot audition system by recognizing three simultaneous speeches," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2004, pp. 2111–2116.
- [6] S. Yamamoto, K. Nakadai, H. Tsujino, T. Yokoyama, and H. Okuno, "Improvement of robot audition by interfacing sound source separation and automatic speech recognition with missing feature theory," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2004, pp. 1517–1523.
- [7] Q. Wang, T. Ivanov, and P. Aarabi, "Acoustic robot navigation using distributed microphone arrays," *Inf. Fusion (Spec. Issue Robust Speech Process.)*, vol. 5, no. 2, pp. 131–140, 2004.
- [8] B. Mungamuru and P. Aarabi, "Enhanced sound localization," *IEEE Trans. Syst., Man, Cybern., B, Cybern.*, vol. 34, no. 3, pp. 1526–1540, 2004.
- [9] S. F. Boll, "A spectral subtraction algorithm for suppression of acoustic noise in speech," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 1979, pp. 200–203.
- [10] P. Renevey, R. Vetter, and J. Kraus, "Robust speech recognition using missing feature theory and vector quantization," in *Proc. Eurospeech*, 2001, pp. 1107–1110.
- [11] J. Barker, L. Josifovski, M. Cooke, and P. Green, "Soft decisions in missing data techniques for robust automatic speech recognition," in *Proc. IEEE Int. Conf. Spoken Lang. Process.*, 2000, vol. 1, pp. 373–376.
- [12] R. Irie, "Robust sound localization: An application of an auditory perception system for a humanoid robot" Master's thesis, Dept. Elect. Eng. Comput. Sci., Massachusetts Inst. Technol., 1995.
- [13] R. Brooks, C. Breazeal, M. Marjanovic, B. Scassellati, and M. Williamson, "The Cog project: Building a humanoid robot," in *Computation for Metaphors, Analogy, and Agents*, C. Nehaniv, Ed. Berlin, Germany, 1999, pp. 52–87, Springer-Verlag.
- [14] K. Nakadai, T. Lourens, H. G. Okuno, and H. Kitano, "Active audition for humanoid," in *Proc. Natl. Conf. Artif. Intell.*, 2000, pp. 832–839.
- [15] K. Nakadai, T. Matsui, H. G. Okuno, and H. Kitano, "Active audition system and humanoid exterior design," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2000, pp. 1453–1461.
- [16] K. Nakadai, K. Hidai, H. G. Okuno, and H. Kitano, "Real-time multiple speaker tracking by multi-modal integration for mobile robots," in *Proc. Eurospeech*, 2001, pp. 1193–1196.
- [17] H. G. Okuno, K. Nakadai, K.-I. Hidai, H. Mizoguchi, and H. Kitano, "Human–robot interaction through real-time auditory and visual multiple-talker tracking," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2001, pp. 1402–1409.
- [18] K. Nakadai, H. G. Okuno, and H. Kitano, "Real-time sound source localization and separation for robot audition," in *Proc. IEEE Int. Conf. Spoken Lang. Process.*, 2002, pp. 193–196.
- [19] K. Nakadai, H. G. Okuno, and H. Kitano, "Exploiting auditory fovea in humanoid–human interaction," in *Proc. Natl. Conf. Artif. Intell.*, 2002, pp. 431–438.
- [20] K. Nakadai, D. Matsuura, H. G. Okuno, and H. Kitano, "Applying scattering theory to robot audition system: Robust sound source localization and extraction," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2003, pp. 1147–1152.
- [21] Y. Matsusaka, T. Tojo, S. Kubota, K. Furukawa, D. Tamiya, K. Hayata, Y. Nakano, and T. Kobayashi, "Multi-person conversation via multi-modal interface—A robot who communicates with multi-user," in *Proc. Eurospeech*, 1999, pp. 1723–1726.
- [22] Y. Matsusaka, S. Fujie, and T. Kobayashi, "Modeling of conversational strategy for the robot participating in the group conversation," in *Proc. Eurospeech*, 2001, pp. 2171–2176.
- [23] Y. Zhang and J. Weng, "Grounded auditory development by a developmental robot," in *Proc. INNS/IEEE Int. Joint Conf. Neural Netw.*, 2001, pp. 1059–1064.
- [24] M. Fujita, Y. Kuroki, T. Ishida, and T. Doi, "Autonomous behavior control architecture of entertainment humanoid robot SDR-4X," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2003, pp. 960–967.
- [25] C. Choi, D. Kong, J. Kim, and S. Bang, "Speech enhancement and recognition using circular microphone array for service robots," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2003, pp. 3516–3521.
- [26] H. Asoh, S. Hayamizu, I. Hara, Y. Motomura, S. Akaho, and T. Matsui, "Socially embedded learning of the office-conversant mobile robot jijo-2," in *Proc. Int. Joint Conf. Artif. Intell.*, 1997, vol. 1, pp. 880–885.
- [27] F. Asano, H. Asoh, and T. Matsui, "Sound source localization and signal separation for office robot jijo-2," in *Proc. Int. Conf. Multisens. Fusion Integr. Intell. Syst.*, 1999, pp. 243–248.
- [28] H. Asoh, F. Asano, K. Yamamoto, T. Yoshimura, Y. Motomura, N. Ichimura, I. Hara, and J. Ogata, "An application of a particle filter to Bayesian multiple sound source tracking with audio and video information fusion," in *Proc. Int. Conf. Inf. Fusion*, 2004, pp. 805–812.
- [29] P. J. Prodanov, A. Drygajlo, G. Ramel, M. Meisser, and R. Siegwart, "Voice enabled interface for interactive tour-guided robots," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2002, pp. 1332–1337.
- [30] C. Theobalt, J. Bos, T. Chapman, A. Espinosa-Romero, M. Fraser, G. Hayes, E. Klein, T. Oka, and R. Reeve, "Talking to Godot: Dialogue with a mobile robot," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2002, pp. 1338–1343.
- [31] J. Huang, T. Supaongprapa, I. Terakura, F. Wang, N. Ohnishi, and N. Sugie, "A model-based sound localization system and its application to robot navigation," *Robots Auton. Syst.*, vol. 27, no. 4, pp. 199–209, 1999.
- [32] J.-M. Valin, F. Michaud, and J. Rouat, "Robust 3D localization and tracking of sound sources using beamforming and particle filtering," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2006, pp. 841–844.
- [33] J.-M. Valin, F. Michaud, and J. Rouat, "Robust localization and tracking of simultaneous moving sound sources using beamforming and particle filtering," *Robot. Auton. Syst.*, vol. 55, no. 3, pp. 216–228, 2007.
- [34] L. C. Parra and C. V. Alvino, "Geometric source separation: Merging convolutive source separation with geometric beamforming," *IEEE Trans. Speech Audio Process.*, vol. 10, no. 6, pp. 352–362, Sep. 2002.
- [35] S. Haykin, *Adaptive Filter Theory*, 4th ed.: Prentice-Hall, 2002.
- [36] Y. Ephraim and D. Malah, "Speech enhancement using minimum mean-square error short-time spectral amplitude estimator," *IEEE Trans. Acoust. Speech Signal Process.*, vol. ASSP-32, no. 6, pp. 1109–1121, 1984.
- [37] Y. Ephraim and D. Malah, "Speech enhancement using minimum mean-square error log-spectral amplitude estimator," *IEEE Trans. Acoust. Speech Signal Process.*, vol. ASSP-33, no. 2, pp. 443–445, 1985.
- [38] R. Zelinski, "A microphone array with adaptive post-filtering for noise reduction in reverberant rooms," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 1988, vol. 5, pp. 2578–2581.
- [39] I. McCowan and H. Bourlard, "Microphone array post-filter for diffuse noise field," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2002, vol. 1, pp. 905–908.
- [40] I. Cohen and B. Berdugo, "Microphone array post-filtering for non-stationary noise suppression," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2002, pp. 901–904.
- [41] I. Cohen and B. Berdugo, "Speech enhancement for non-stationary noise environments," *Signal Process.*, vol. 81, no. 2, pp. 2403–2418, 2001.
- [42] AURORA. [Online] Available: "<http://www.elda.fr/proj/aurora1.html>" "<http://www.elda.fr/proj/aurora2.html>".
- [43] D. Pearce, "Developing the ETSI Aurora advanced distributed speech recognition front-end & what next," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2001, pp. 131–134.
- [44] R. P. Lippmann, E. A. Martin, and D. B. Paul, "Multi-style training for robust isolated-word speech recognition," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 1987, pp. 705–708.
- [45] M. Blanchet, J. Boudy, and P. Lockwood, "Environment adaptation for speech recognition in noise," in *Proc. Eur. Signal Process. Conf.*, 1992, vol. vol. VI, pp. 391–394.
- [46] D. O'Shaughnessy, "Interacting with computers by voice: Automatic speech recognition and synthesis," in *Proc. IEEE*, Sep. 2003, vol. 91, no. 9, pp. 1272–1305.
- [47] M. L. Seltzer and R. M. Stern, "Subband parameter optimization of microphone arrays for speech recognition in reverberant environments," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2003, pp. 408–411.

- [48] S. Araki, S. Makino, A. Blin, R. Mukai, and H. Sawada, "Underdetermined blind separation for speech in real environments with sparseness and ICA," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process*, 2004, pp. 881–884.
- [49] P. Aarabi and G. Shi, "Phase-based dual-microphone robust speech enhancement," *IEEE Trans. Syst., Man, Cybern., B, Cybern.*, vol. 34, no. 4, pp. 1763–1773, Aug. 2004.
- [50] M. Cooke, P. Green, and M. Crawford, "Handling missing data in speech recognition," in *Proc. IEEE Int. Conf. Spoken Lang. Process*, 1994, pp. 1555–1558.
- [51] M. Cooke, P. Green, L. Josifovski, and A. Vizinho, "Robust automatic speech recognition with missing and unreliable acoustic data," *Speech Commun.*, vol. 34, pp. 267–285, 2001.
- [52] J. Barker, M. Cooke, and P. Green, "Robust ASR based on clean speech models: An evaluation of missing data techniques for connected digit recognition in noise," in *Proc. Eurospeech*, 2001, pp. 213–216.
- [53] H. van Hamme, "Robust speech recognition using missing feature theory in the cepstral or LDA domain," in *Proc. Eurospeech*, 2003, pp. 1973–1976.
- [54] H. van Hamme, "Robust speech recognition using cepstral domain missing data techniques and noisy masks," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process*, 2004, pp. 213–216.
- [55] B. Raj, M. L. Seltzer, and R. M. Stern, "Reconstruction of missing features for robust speech recognition," *Speech Commun.*, vol. 43, no. 4, pp. 275–296, 2004.
- [56] M. L. Seltzer, B. Raj, and R. M. Stern, "A Bayesian framework for spectrographic mask estimation for missing feature speech recognition," *Speech Commun.*, vol. 43, no. 4, pp. 379–393, 2004.
- [57] J. Ming, P. Jancovic, and F. J. Smith, "Robust speech recognition using probabilistic union models," *IEEE Trans. Speech Audio Process.*, vol. 10, no. 6, pp. 403–414, 2004.
- [58] J. Ming and F. J. Smith, "Speech recognition with unknown partial feature corruption—a review of the union model," *Comput. Speech Lang.*, vol. 17, pp. 287–305, 2003.
- [59] H.-Y. Cho and Y.-H. Oh, "On the use of channel-attentive MFCC for robust recognition of partially corrupted speech," *IEEE Signal Process. Lett.*, vol. 11, no. 6, pp. 581–584, 2004.
- [60] I. McCowan, A. Morris, and H. Bourlard, "Improved speech recognition performance of small microphone arrays using missing data techniques," in *Proc. IEEE Int. Conf. Spoken Lang. Process*, 2002, pp. 2181–2184.
- [61] A. Lee, T. Kawahara, and K. Shikano, "Julius—an open-source real-time large vocabulary recognition engine," in *Proc. Eurospeech*, 2001, pp. 1691–1694.
- [62] J.-M. Valin, J. Rouat, and F. Michaud, "Microphone array post-filter for separation of simultaneous non-stationary sources," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process*, 2004, pp. 1–224.




Jean-Marc Valin (S'03–M'05) received the B.Eng., M.A.Sc., and Ph.D. degrees in electrical engineering from the University of Sherbrooke, QC, Canada, in 1999, 2001, and 2005, respectively.

Since 2005, he is a Postdoctoral Fellow at the Commonwealth Scientific and Industrial Research Organization Information and Communication Technologies (CSIRO ICT) Centre, Sydney, Australia. His current research interests include acoustic echo cancellation and microphone array processing.

Dr. Valin is a member of the IEEE Signal Process-

ing Society.



Shun'ichi Yamamoto (S'04) received the B.S. and M.S. degrees in engineering and informatics from Kyoto University, Kyoto, Japan, in 2003 and 2005, respectively. He is currently working toward the Ph.D. degree in the Department of Intelligence Science and Technology, Graduate School of Informatics, Kyoto University.

His current research interests include automatic speech recognition, sound source separation, and sound source localization for robot audition.

Dr. Yamamoto is the recipient of several awards including the IEEE Robotics and Automation Society Japan Chapter Young Award.



Jean Rouat (S'83–M'88–SM'xx) received the Master's degree in physics from Univ. de Bretagne, Bretagne, France, in 1981, the E. & E. Master's degree in speech coding and speech recognition from Université de Sherbrooke, Sherbrooke, QC, Canada, in 1984, and the E. & E. Ph.D. in cognitive and statistical speech recognition jointly with the Université de Sherbrooke and McGill University, Montreal, QC, in 1988.


He is now with the Université de Sherbrooke, Sherbrooke, QC, Canada, where he founded the Computational Neuroscience and Intelligent Signal Processing Research Group. His current research interests include audition, speech, and signal processing in relation with networks of spiking neurons. He has been a Reviewer for various speech, neural networks, and signal processing journals. Dr. Rouat is a member of several international scientific associations (ASA, ISCA, IEEE senior, ARO, etc.) and was on the IEEE Technical Committee on Machine Learning for Signal Processing from 2001 to 2005.



François Michaud (S'89–M'92) received the Bachelor's, Master's, and Ph.D. degrees in electrical engineering from the Université de Sherbrooke, Sherbrooke, QC, Canada, in 1992, 1993, and 1996, respectively.

He is currently a Professor in the Department of Electrical and Computer Engineering, Université de Sherbrooke. He was a Postdoctoral Fellow at Brandeis University, Waltham, MA. He founded LABORIOUS, a research laboratory working on designing intelligent autonomous systems that can assist humans in living environments. His current research interests include architectural methodologies for intelligent decision-making, design of autonomous mobile robotics, social robotics, robots for children with autism, and robot learning and intelligent systems.

Prof. Michaud is a member of the Association for the Advancement of Artificial Intelligence (AAAI) and the Ordre des ingénieurs du Québec (OIQ). He is the Canada Research Chairholder in Autonomous Mobile Robots and Intelligent Systems. He was the recipient of the 2003 Young Engineer Achievement Award from the Canadian Council of Professional Engineers.



Kazuhiro Nakadai (M'04) received the B.E. degree in electrical engineering in 1993, the M.E. degree in information engineering in 1995, and the Ph.D. degree in electrical engineering in 2003, all from the University of Tokyo, Tokyo, Japan.

He is currently a Senior Researcher at the Honda Research Institute Japan, Co., Ltd., Saitama, Japan. From 1995 to 1999, he was with Nippon Telegraph and Telephone and NTT Comware Corporation, Tokyo, Japan. From 1999 to 2003, he was a Researcher for the Kitano Symbiotic Systems Project.

Since 2006, he is also a Visiting Associate Professor at the Tokyo Institute of Technology, Tokyo. His current research interests include signal and speech processing, artificial intelligence and robotics, computational auditory scene analysis, multimodal integration, and robot audition.

Prof. Nakadai is a member of the Robotics Society of Japan. He was the recipient of the 2001 Best Paper Award from the International Society for Applied Intelligence.



Hiroshi G. Okuno (M'03–SM'06) received the B.A. and Ph.D. degrees from the University of Tokyo, Tokyo, Japan, in 1972 and 1996, respectively.

He is currently a Professor in the Department of Intelligence Science and Technology, Graduate School of Informatics, Kyoto University, Kyoto, Japan. He was a Visiting Scholar at Stanford University, Stanford, CA, and Visiting Associate Professor at the University of Tokyo. He was with the Nippon Telegraph and Telephone, the Kitano Symbiotic Systems Project, and the Tokyo University of Science, Tokyo.

His current research interests include computational auditory scene analysis, music scene analysis, and robot audition. He is the Co-Editor (with D. Rosenthal) of *Computational Auditory Scene Analysis* (Princeton, NJ: Erlbaum, 1998) and (with T. Yuasa) of *Advanced LISP Technology* (London, U.K.: Taylor and Francis, 2002).

Prof. Okuno is a member of the Robotics Society of Japan, the Association for the Advancement of Artificial Intelligence, the Association for Computing Machinery (ACM), and the American Standards Association (ASA).