



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/133240/>

Version: Accepted Version

Article:

Ma, N., Gonzalez, J. and Brown, G.J. (2018) Robust binaural localization of a target sound source by combining spectral source models and deep neural networks. *IEEE/ACM Transactions on Audio, Speech and Language Processing*, 26 (11). pp. 2122-2131. ISSN: 2329-9290

<https://doi.org/10.1109/TASLP.2018.2855960>

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other users, including reprinting/ republishing this material for advertising or promotional purposes, creating new collective works for resale or redistribution to servers or lists, or reuse of any copyrighted components of this work in other works. Reproduced in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Robust binaural localisation of a target sound source by combining spectral source models and deep neural networks

Ning Ma, Jose A. Gonzalez and Guy J. Brown

Abstract—Despite there being clear evidence for top-down (e.g., attentional) effects in biological spatial hearing, relatively few machine hearing systems exploit top-down model-based knowledge in sound localisation. This paper addresses this issue by proposing a novel framework for binaural sound localisation that combines model-based information about the spectral characteristics of sound sources and deep neural networks (DNNs). A target source model and a background source model are first estimated during a training phase using spectral features extracted from sound signals in isolation. When the identity of the background source is not available, a universal background model can be used. During testing, the source models are used jointly to explain the mixed observations and improve the localisation process by selectively weighting source azimuth posteriors output by a DNN-based localisation system. To address the possible mismatch between training and testing, a model adaptation process is further employed on-the-fly during testing, which adapts the background model parameters directly from the noisy observations in an iterative manner. The proposed system therefore combines model-based and data-driven information flow within a single computational framework. The evaluation task involved localisation of a target speech source in the presence of an interfering source and room reverberation. Our experiments show that by exploiting model-based information in this way, sound localisation performance can be improved substantially under various noisy and reverberant conditions.

Index Terms—binaural source localisation, machine hearing, reverberation, sound source combination, masking

I. INTRODUCTION

IT has been long established that human listeners use both bottom-up (data-driven) and top-down (model-based) information in order to understand an acoustic scene (e.g., [1]). More specifically, both kinds of information are needed in order to answer ‘what’ and ‘where’ questions about the acoustic scene, i.e. *what* the identities of the sound sources are, and *where* they are in space. Many machine hearing studies have proposed computational approaches for answering these questions, by combining techniques for source separation, classification and sound localisation [2]. However, in such machine systems, data-driven and model-based mechanisms are typically much less well-integrated than they appear to be in biological hearing. The current paper addresses this issue, by proposing a novel machine hearing system that tightly

integrates knowledge about source spectral characteristics with a mechanism for binaural localisation. The resulting system improves binaural sound localisation under challenging acoustic conditions in which multiple sound sources and room reverberation are present.

Psychophysical studies of human hearing have found evidence for top-down effects in sound localisation. For example, listeners take less time to localise a target sound when it is preceded by a cue on the same side of the head, suggesting that sound localisation is enhanced by covert orienting [3]. More recently, based on a review of psychophysical data, Bronkhorst [4] proposed a conceptual model of hearing in which top-down models inform the selection of binaural cues needed for localisation of specific sounds. Physiological studies in animals have also shown that sound localisation is modulated by top-down mechanisms. Studies of the barn owl have shown that selective attention influences sound localisation, including orienting behaviour such as body and head movements. More specifically, neural responses in the midbrain of the owl are enhanced if they are associated with the location of behaviourally relevant stimuli, such as a source of food [5]. Likewise, neural circuitry for gaze control can exert a top-down effect on the responses of auditory neurons that are turned to particular spatial locations [6]. In summary, these psychological and physiological findings suggest that mechanisms of spatial hearing are tightly integrated with top-down, cross-modal processing in the brain.

In contrast, relatively few systems for machine hearing exploit top-down model-based information in source localisation. In [7], Ma et al. proposed a robust binaural localisation system based on deep neural network (DNN) for localisation of multiple sources. However, no source models were used and the system reported azimuth estimates for all directional sources. In [8], Christensen et al. exploit a pitch cue to identify local spectro-temporal fragments that are considered to be dominated by a single source, before integrating binaural cues over such fragments. The integrated binaural cues were more robust for localising multiple speakers, but no explicit model-based information was exploited. Mandel et al. [9] have proposed a probabilistic model for joint sound source localisation and separation based on interaural spatial cues and binary masking. Given the number of sources, the system iteratively updates Gaussian mixture models (GMMs) of spatial cues using the expectation-maximization (EM) algorithm. However, the benefit of the system was demonstrated in terms of source separation rather than localisation. Two systems for

Ning Ma and Guy J. Brown are with the Department of Computer Science, University of Sheffield, Sheffield, UK, S1 4DP (email: {n.ma, g.j.brown}@sheffield.ac.uk). Jose A. Gonzalez is with the Department of Languages and Computer Sciences, University of Malaga, Campus de Teatinos, Malaga 29071, Spain (email: j.gonzalez@uma.es). This work was carried out while Gonzalez was working at the University of Sheffield.

binaural localisation of multiple sources recently proposed by [10], [11] use statistical frameworks to jointly perform localisation and pitch-based segregation. However, they do not use information about source characteristics other than through statistical models of pitch dynamics and binaural cues. In [12], [13], model-based information is explicitly employed together with binaural cues, but the task in those studies was automatic speech recognition rather than sound localisation. Closest to the approach proposed here is the attention-driven model of sound localisation proposed by [14], in which top-down connections from a cortical model are used to potentiate responses to attended locations. However, attentional control in their model is driven by a simple neural circuit that fixates on sounds arriving from the same spatial location, and their approach is currently unable to localise multiple sound sources.

The current paper proposes a framework for sound localisation in which model-based information about the spectral characteristics of sound sources in the acoustic scene is used to selectively weight binaural cues. Models for the target and the background sources are first estimated in a training stage, using spectral features computed from source signals in isolation. When the identity of the background source is not available, a universal background model can be used. The source models are then used during the testing stage to jointly explain the mixed observations in terms of which spectral features belong to each source (target or masker), and thereby improve the localisation process by selectively weighting binaural cues in each time-frequency bin. In our previous work [15] model-based information was incorporated in a GMM-based localisation system, whereas in this study it was used in a state-of-the-art DNN-based system [7].

The proposed system also addressed the possible mismatch due to power level differences between training and testing. An iterative adaptation process is employed on-the-fly to estimate a frequency-dependent scaling factor for the background model parameters directly from the mixed signals. The proposed system therefore combines data-driven and model-based information flow within a single computational framework. We show that by exploiting source models in this way, sound localisation performance can be improved under conditions in which multiple sources and room reverberation are present.

The rest of the paper is organised as follows. In the next section we give a system overview and describe features used for source localisation and for modelling source spectral characteristics. A source localisation framework is then presented in Section III which allows model-based information of source spectral characteristics to be incorporated. Evaluation and experiments are described in Section IV. Section V discusses the results in detail. Finally, Section VI concludes the paper and makes suggestions for future work.

II. SYSTEM OVERVIEW

Figure 1 shows a schematic diagram of the proposed binaural sound localisation system. A target source and an interfering source (masker) are spatialised in a virtual environment, which allows sounds to be placed in the full 360°




Fig. 1. Schematic diagram of the proposed system.

azimuth range around a virtual listener. In this study, the prior knowledge of the target source is assumed to be known so that the system can “attend” to the target for localisation. No prior knowledge of the masker is assumed when a universal background model (UBM) is used, but the knowledge of the masker can be exploited by the framework when available (see Section III-B for more details).

The binaural ear signals are analysed by a binaural auditory model, which consists of a bank of 32 overlapping Gammatone filters with centre frequencies uniformly spaced on the equivalent rectangular bandwidth (ERB) scale between 80 Hz and 8 kHz [2]. Inner hair cell function was approximated by half-wave rectification. Subsequently, a cross-correlation between the right and left ears was computed independently for each of the 32 frequency bands using overlapping frames of 20 ms with a shift of 10 ms. The cross-correlation function (CCF) was further normalised by the autocorrelation value at lag zero as described in [16] and evaluated for time lags in the range of ± 1.1 ms.

Two binaural features, interaural time differences (ITDs) and interaural level differences (ILDs), are typically used in binaural localisation systems [17]. ITD is estimated as the lag corresponding to the maximum in the CCF. However, it has been shown that the CCF outperforms the ITDs in state-of-the-art DNN-based localisation systems [7], [18]. Thus, in this work we use the entire CCF as localisation features. For signals sampled at a rate of 16 kHz, the CCF with a lag range of ± 1 ms produced a 33-dimensional binaural feature space for each frequency band. This was supplemented by the ILD, which corresponds to the energy ratio between the left and right ears within each analysis window, expressed in dB. The final 34-dimensional (34D) feature vectors were passed to a DNN-based localisation system which estimates posterior probabilities for each azimuth in the full 360° azimuth range. DNNs were trained using sound mixtures consisting of the target and masker sources rendered in a virtual acoustic environment.

Ratemap features [19] were used to model source spectral characteristics. A ratemap is a spectro-temporal representation of auditory nerve firing rate, extracted from the inner hair cell output of each frequency channel by leaky integration and downsampling (see Fig. 2 for examples). For the binaural signals used here, the ratemap features were computed for each ear channel separately and then averaged across the two ears. They were finally log-compressed to form 32D feature vectors. The source ratemap features were exploited together with a target source model and a background source model

to estimate a set of localisation weights, in order to bias the system towards localising the target source. In the next section we describe the localisation framework in detail.

III. LOCALISATION WITH TOP-DOWN SOURCE MODELS

In the following we use P for probability, p for density functions and C for cumulative distribution functions.

A. Localisation model

DNNs were used to model statistical relationships between the binaural features and corresponding azimuth angles [7], [18]. Unlike many other binaural sound localisation systems (e.g. [20], [21]), this study does not assume that sound sources are located only in the frontal-hemifield. Instead it considers 72 azimuth angles ϕ in the full 360° azimuth range (5° steps).

A separate DNN was trained for each of the 32 frequency bands, largely because this study is concerned with localisation of multiple sources. Although simultaneous sources overlap in time, within a local time frame each frequency band is mostly dominated by a single source¹. By adopting a separate DNN for each frequency band, it is possible to train the DNNs using single-source data without having to construct multi-source training data.

The DNN model was largely based on a state-of-art binaural localisation system described in [7]. In brief, each DNN consists of an input layer, two hidden layers, and an output layer. The input layer contained 34 nodes and each node was assumed to be a Gaussian random variable with zero mean and unit variance. The hidden layers had sigmoid activation functions with 128 hidden nodes. The output layer contained 72 nodes corresponding to the 72 azimuth angles in the full 360° azimuth range, with a 5° step. A ‘softmax’ activation function was applied at the output layer.

Given the observed localisation feature vector \mathbf{o}_{tf} (the 34D localisation features described in Section II) at time frame t and frequency band f , the posterior probability of azimuth angle $P(\phi|\mathbf{o}_{tf})$ is estimated using the DNN trained for each frequency band f . The posteriors are then integrated across frequency to produce the probability of azimuth ϕ given features $\mathbf{o}_t = [\mathbf{o}_{t1}^\top, \dots, \mathbf{o}_{t32}^\top]^\top$ of the entire frequency range at time t ,

$$P(\phi|\mathbf{o}_t) = \frac{\prod_f P(\phi|\mathbf{o}_{tf})^{\omega_{tf}}}{P(\mathbf{o}_t)} \quad (1)$$

where

$$P(\mathbf{o}_t) = \sum_\phi \prod_f P(\phi|\mathbf{o}_{tf})^{\omega_{tf}} \quad (2)$$

Here $\omega_{tf} \in [0, 1]$ is introduced for selectively weighting the contribution of binaural cues from each time-frequency bin in order to localise the attended target source in the presence of competing sources. When ω_{tf} is 0 the time-frequency bin will be excluded from localisation of the target source. The next section will discuss in detail how model-based knowledge about sound sources can be used to jointly estimate the weighting factors.

¹See Bregman’s notion of ‘exclusive allocation’ [1]. Note in reverberant environments this is only a crude approximation

Assuming that the target sound source is stationary, the frame posteriors are further averaged across time to produce a posterior distribution $P(\phi)$ of sound source activity given a segment of signal consisting of T time frames

$$P(\phi|\mathbf{o}_{1..T}) = \frac{1}{T} \sum_t^{t+T-1} P(\phi|\mathbf{o}_t) \quad (3)$$

The target location is considered to be the azimuth $\hat{\phi}$ that maximises Eq. (3)

$$\hat{\phi} = \operatorname{argmax}_\phi P(\phi|\mathbf{o}_{1..T}) \quad (4)$$

B. Exploiting model-based information

In the presence of multiple sources, the binaural cues computed from spectro-temporal regions dominated by the masking sources will bias the localisation decision towards the location of the maskers, thus leading to localisation errors. To address this issue, we propose an ‘attentional’ mechanism that exploits prior information about the spectral characteristics of the acoustic sources in order to weight the binaural cues towards localising the target source. First, the prior source models are used to estimate the probability of each time-frequency (T-F) bin of the observed signal being dominated by the energy of the target source. These probabilities are then employed during sound localisation to selectively weight the binaural cues: cues extracted from T-F regions dominated by the masking sources are penalised during localisation.

Let $\mathbf{y}_t = [y_{t1}, \dots, y_{tD}]$ denote the ratemap feature vector (see Section II) extracted from the observed signals at frame t , where D is the feature dimension. For simplification we will omit the dependence on the time index t for the remainder of this section. Similarly, let \mathbf{x} and \mathbf{n} denote the ratemap feature vectors of the underlying target and masker, respectively. In the log-ratemap domain, the relationship between these variables can be approximated as

$$y_f \approx \max(x_f, n_f) \quad (5)$$

where f is the frequency band within $[1 \dots D]$. This is known as the *log-max* model [22], [23], which is an approximation of the exact interaction model between two sound sources when they are expressed in a log-compressed spectral domain. According to this model, the effect of the masker sound on the target can be modelled as a kind of spectral masking.

Let ω denote the localisation weights to be estimated, and each of its elements $\omega_f \in [0, 1]$ indicates whether y_f is dominated by the energy of the target source x_f ($\omega_f \approx 1$) or the masker source n_f ($\omega_f \approx 0$). From a probabilistic point of view, and under the restrictions imposed by the *log-max* model, ω_f corresponds to the following *a posteriori* probability

$$\omega_f \triangleq P(x_f \geq n_f | \mathbf{y}) \equiv P(x_f = y_f, n_f \leq y_f | \mathbf{y}) \quad (6)$$

To estimate this probability, prior models for the spectral characteristics of the sound sources are used. Each sound source $s = 1, \dots, S$ is modelled using a GMM λ_s with K_s mixtures. Because the identity of the target source is known,

the probability of the observation given the target source GMM λ_x is computed simply as

$$p(\mathbf{y}|\lambda_x) = \sum_k^{K_x} P(k|\lambda_x) \mathcal{N}(\mathbf{y}; \boldsymbol{\mu}_x^{(k)}, \boldsymbol{\Sigma}_x^{(k)}) \quad (7)$$

where $\boldsymbol{\mu}_x^{(k)}$ and $\boldsymbol{\Sigma}_x^{(k)}$ are the mean and covariance of the k^{th} component for the target source GMM.

For the case of the masker source we will distinguish two cases. First, if there is only one masker source in the acoustic scene and its identity is known *a priori*, $p(\mathbf{y}|\lambda_n)$ can be computed as in (7) but using the GMM for the known masker, λ_n . If the identity of the masker is unknown, we estimate a UBM λ_{ubm} using signals from various sources (see Section IV-D for more details about the estimation of the UBM).

Using $p(\mathbf{y}|\lambda_x)$ and $p(\mathbf{y}|\lambda_n)$, Eq. (6) for computing the localisation weights ω_f can be rewritten as follows (see [24]–[26]):

$$\omega_f = \sum_{k_x} \sum_{k_n} \gamma^{(k_x, k_n)} P(x_f = y_f, n_f \leq y_f | k_x, k_n, \mathbf{y}) \quad (8)$$

Here, k_x and k_n denote the index for the mixture components in λ_x and λ_n , $P(x_f = y_f, n_f \leq y_f | k_x, k_n, \mathbf{y})$ is the target presence probability (TPP) and $\gamma^{(k_x, k_n)}$ is defined as the posterior probability

$$\gamma^{(k_x, k_n)} \triangleq P(k_x, k_n | \mathbf{y}) = \frac{p(\mathbf{y}|k_x, k_n)P(k_x)P(k_n)}{\sum_{k'_x, k'_n} p(\mathbf{y}|k'_x, k'_n)P(k'_x)P(k'_n)} \quad (9)$$

Here we have omitted the dependence on the models λ_x and λ_n in order to simplify the notation.

Assuming the frequency bands are conditionally independent given the mixture components, $p(\mathbf{y}|k_x, k_n)$ can be expressed as

$$p(\mathbf{y}|k_x, k_n) = \prod_{f=1}^D p(y_f | k_x, k_n) \quad (10)$$

with $p(y_f | k_x, k_n)$ being the following marginal distribution

$$p(y_f | k_x, k_n) = \iint p(y_f | x_f, n_f) p(x_f | k_x) p(n_f | k_n) dx_f dn_f \quad (11)$$

The terms $p(x_f | k_x)$ and $p(n_f | k_n)$ of the equation can be directly calculated using the GMMs λ_x and λ_n . Using the *log-max* model, $p(y_f | x_f, n_f)$ can be expressed as

$$\begin{aligned} p(y_f | x_f, n_f) &= \delta(y_f - \max(x_f, n_f)) \\ &= \delta(y_f - x_f) \mathbb{1}_{n_f \leq x_f} + \delta(y_f - n_f) \mathbb{1}_{x_f < n_f} \end{aligned} \quad (12)$$

where $\delta(\cdot)$ is the Dirac delta function and $\mathbb{1}_{\mathcal{C}}$ is an indicator function which equals 1 when the condition \mathcal{C} is true and 0 otherwise. As shown in [26], (11) then becomes

$$p(y_f | k_x, k_n) = p_x(y_f | k_x) C_n(y_f | k_n) + p_n(y_f | k_n) C_x(y_f | k_x) \quad (13)$$

where p_x and p_n are the Gaussian probability functions and C_x and C_n are the corresponding cumulative distribution functions.

Using Bayes' rule, the TPP in (8) can be written as

$$\begin{aligned} P(x_f = y_f, n_f \leq y_f | k_x, k_n, \mathbf{y}) &= \frac{p(x_f = y_f, n_f \leq y_f | k_x, k_n)}{p(x_f = y_f, n_f \leq y_f | k_x, k_n) + p(n_f = y_f, x_f < y_f | k_x, k_n)} \\ &= \frac{p_x(y_f | k_x) C_n(y_f | k_n)}{p(y_f | k_x, k_n)} \end{aligned} \quad (14)$$

Finally, using (9) and (14), the localisation weight (8) becomes a weighted average of the TPPs in (14) for all possible combinations of (k_x, k_n)

$$\omega_f = \sum_{k_x, k_n} \frac{\gamma^{(k_x, k_n)} p_x(y_f | k_x) C_n(y_f | k_n)}{p_x(y_f | k_x) C_n(y_f | k_n) + p_n(y_f | k_n) C_x(y_f | k_x)} \quad (15)$$

C. Adaptation to power level differences

The mismatch in the power level of acoustic sources between training and testing conditions may cause the source GMM to inaccurately represent the spectral characteristics of the sources in the testing condition. To alleviate this problem, we introduce an algorithm that adapts the source GMMs directly from the noisy signals on-the-fly, before estimating the localisation weights from the same signal. To simplify the discussion, here we only consider adaptation of the GMM for the interfering sources (i.e. we assume that the power level of the target source does not change). However, the extension of the procedure to also adapt the target source GMM is straightforward.

We assume that all the mixture components in the GMM are adapted by the same level as follows:

$$p(\mathbf{y}|\lambda_n) = \sum_k P(k|\lambda_n) \mathcal{N}(\mathbf{y}; \boldsymbol{\mu}_n^{(k)} + \boldsymbol{\beta}, \boldsymbol{\Sigma}_n^{(k)}) \quad (16)$$

where $\boldsymbol{\beta} = [\beta_1, \dots, \beta_D]$ is the vector that accounts for the power level differences in each frequency channel. Note that only the means are adapted and variances are kept the same. To determine $\boldsymbol{\beta}$ directly from the noisy observations, we resort to an iterative approach based on the EM algorithm [27]. Let $\hat{\boldsymbol{\beta}}$ denote the current estimate of $\boldsymbol{\beta}$. The function to be optimised is

$$Q(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}) = \sum_t \sum_{k_x} \sum_{k_n} \gamma_t^{(k_x, k_n)} \sum_f \log p(y_{tf} | k_x, k_n, \boldsymbol{\beta}_f) \quad (17)$$

where $\gamma_t^{(k_x, k_n)}$ is the posterior probability in (9) and is computed using the current estimate $\hat{\boldsymbol{\beta}}$. $p(y_{tf} | k_x, k_n, \boldsymbol{\beta}_f)$ is given by (13).

By setting the derivative $\partial Q(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}) / \partial \beta_f$ equal to zero, we obtain the following updating equation for β_f :

$$\beta_f = \frac{1}{T} \sum_t \sum_{k_x} \sum_{k_n} \gamma_t^{(k_x, k_n)} \left(\bar{n}_{tf}^{(k_x, k_n)} - \mu_{nf}^{(k_n)} \right) \quad (18)$$

where $\bar{n}_{tf}^{(k_x, k_n)}$ is the estimate of n_{tf} ,

$$\bar{n}_{tf}^{(k_x, k_n)} = \alpha_{tf}^{(k_x, k_n)} \hat{n}_{tf}^{(k_n)} + \left(1 - \alpha_{tf}^{(k_x, k_n)} \right) y_{tf}. \quad (19)$$

In the last equation we use the short-hand notation $\alpha_{tf}^{(k_x, k_n)} \triangleq P(x_f = y_f, n_f \leq y_f | k_x, k_n, \hat{\beta}_f, \mathbf{y})$ for the TPP in (14). $\tilde{n}_{tf}^{(k_n)}$ represents the estimate for n_{tf} under the assumption that n_{tf} is masked by the target source energy x_{tf} . In this case, n_{tf} is upper-bounded by the observation y_{tf} . Thus, $\tilde{n}_{tf}^{(k_n)}$ corresponds to the following expected value,

$$\tilde{n}_{tf}^{(k_n)} = \int_{-\infty}^{y_{tf}} n_{tf} p(n_{tf} | k_n, \hat{\beta}_f) dn_{tf}. \quad (20)$$

From (18), it can be seen that the updated value for β_f is a weighted average of the power level differences between the estimates $\tilde{n}_{tf}^{(k_x, k_n)}$ and the GMM means $\mu_{nf}^{(k_n)}$. This equation is iteratively applied until β_f converges.

IV. EVALUATION

To evaluate the proposed binaural localisation system, a virtual acoustic environment was created which contained a target speech source and a masking source selected from a number of different noise types.

A. Binaural simulations

Binaural audio signals were created by convolving monaural signals with head related impulse responses (HRIRs) or binaural room impulse responses (BRIRs). An HRIR catalog based on the Knowles Electronic Manikin for Acoustic Research (KEMAR) dummy head [28] was used for simulating the anechoic training signals. The evaluation stage used the Surrey BRIR database [29] to simulate reverberant room conditions. The Surrey database was captured using a Cortex head and torso simulator (HATS) and includes four room conditions with various amounts of reverberation (see Table I for the reverberation time (T_{60}) and the direct-to-reverberant ratio (DRR) of each room). Binaural mixtures of two simultaneous sources were created by convolving each source signal with BRIRs separately before adding them together in each of the two binaural channels.

TABLE I
ROOM CHARACTERISTICS OF THE SURREY BRIR DATABASE [29].

	Room A	Room B	Room C	Room D
T_{60} (s)	0.32	0.47	0.68	0.89
DRR (dB)	6.09	5.31	8.82	6.12

B. Target and masker signals

The target source signals were drawn from the GRID speech corpus [30]. Each GRID sentence is approximately 2 sec long and has a six-word form (e.g., “lay red at G 9 now”). Both a male talker and a female talker were selected as the target speech source. All target signals were sampled at 16 kHz.

In our previous study [15], all noise types were used to estimate the UBM. In order to investigate how well the system is able to deal with unseen noise types, two sets of noise signals were used as the masker source. Noise Set A consists of six masker types with various degrees of spectro-temporal complexities, and were used to estimate the UBM. Noise Set B

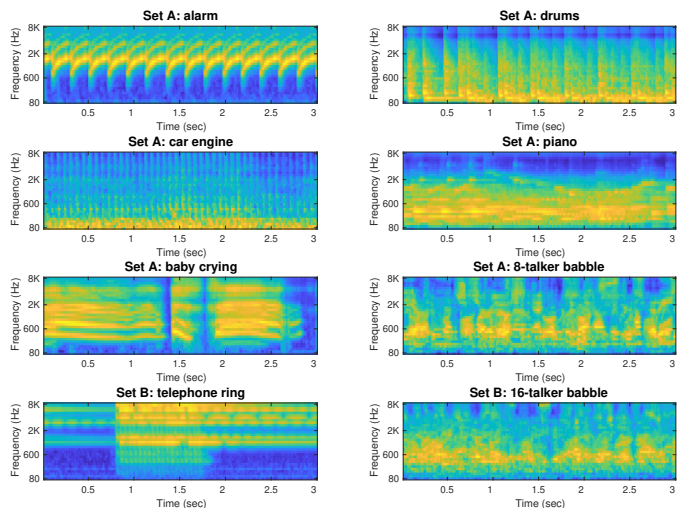


Fig. 2. Ratemap representations of various masker sounds used in this study.

consists of two additional masker types as the “unseen noise” set, i.e. they were excluded for training the UBM. The details of both noise sets are summarised in Tables II and III. In all cases, noise signals were 90 seconds long and were sampled at 16 kHz. Fig. 2 shows example ratemap representations of these noise signals.

TABLE II
DESCRIPTIONS OF MASKER SOUNDS USED IN NOISE SET A.

alarm	car alarm sound, rhythmic moderate-narrow-band signal between 600 Hz and 3 kHz.
drums	drumming sound, strong onsets synchronised across frequency, and with significant energy distributed between 100–300 Hz.
car engine	highly modulated at 15 Hz with most energy distributed below 300 Hz.
piano	fast playing solo piano sound, with most energy distributed below 2 KHz.
baby crying	crying baby sound, high pitch, less rhythmic with harmonics often lasting longer than 1 second.
16-talker babble	created from 16 talkers randomly selected from the TIMIT corpus [31], mostly overlapping, speech frequency range.

TABLE III
DESCRIPTIONS OF MASKER SOUNDS USED IN NOISE SET B.

telephone ring	telephone ring, periodic and narrow-band sound with significant energy at around 1 kHz and above 3 kHz.
32-talker babble	created from 32 talkers randomly selected from the TIMIT corpus, mostly overlapping, speech frequency range.

C. Localisation DNN training

As shown in previous studies [10], [21], [32], multi-conditional training (MCT) can increase the robustness of localisation systems in reverberant multi-source conditions. To train the DNNs, in this study a MCT dataset was created by mixing the training signals at a specified azimuth with diffuse

noise as described in [32]. The diffuse noise consisted of 72 uncorrelated, white Gaussian noise sources that were placed across the full 360° azimuth range in steps of 5° . Both the target signals and the diffuse noise were created by using the same anechoic HRIR recorded using a KEMAR dummy head [28].

Following [7], the localisation training data consisted of speech sentences from the TIMIT database [31]. A set of 30 speech sentences was randomly selected for each of the 72 azimuth locations, equally distributed in the 360° range. For each spatialised training sentence, the anechoic signal was corrupted with diffuse noise at three signal-to-noise ratios (SNRs) (20, 10 and 0 dB). The corresponding binaural features (ITDs, CCF, and ILDs) were then extracted. Only those features for which the *a priori* SNR between the target and the diffuse noise exceeded -5 dB were used for training. This negative SNR criterion ensured that the multi-modal clusters in the binaural feature space at higher frequencies, which are caused by periodic ambiguities in the cross-correlation function, were properly captured. The localisation DNNs were not retrained for the target source drawn from the GRID speech corpus.

D. Source model training

This study is concerned with binaural sound localisation. In order to learn a source spectral model in a binaural setting, binaural source signals were created by convolving monaural signals with the Room A BRIR from the Surrey database [29] for azimuths ranging between $[-90^\circ, 90^\circ]$ in 5° steps. Ratemap features were first extracted from each of the binaural channels and then averaged across the two channels. The ratemap features for all the azimuths were used to train source models using the EM algorithm as described in Section III-B.

The identity of the attended target source is assumed known *a priori* in this study. For the target speech source, 90 seconds of training data were used to train a 16-mixture target GMM.

For each noise source in Noise Set A, 4/5 of the 90-second signal was used as the training data and the remaining 1/5 was used as the test data. When the identity of the masker source is unavailable, the system uses a universal background model and performs adaptation on-the-fly during testing in order to better match the spectral profile of the masker source in a mixed test signal. All the training data from Noise Set A were used to train a 16-mixture UBM. Noises in Set B were excluded in this process.

If the identity of the masker source is also available *a priori*, the system can directly exploit the corresponding masker model. The number of mixtures for each GMM was selected based on its spectro-temporal complexity and is listed in Table IV.

E. Experimental setup

The evaluation set contained 50 GRID sentences from the target speech source which were not included in training. For each target signal, a masker signal that matched the length of the target was randomly selected from the test set. The target

TABLE IV
THE NUMBER OF GAUSSIAN MIXTURE COMPONENTS USED FOR EACH SOURCE MODEL.

Target speech	Noise Set A						UBM
	alarm	drum	car	piano	baby	8-talker	
16	2	2	2	3	3	4	8




Fig. 3. Schematic diagram of the virtual listener configuration, showing a typical arrangement of target source (here at $+30^\circ$) and masker (at -15°). Target source positions were limited to the range $[-90^\circ, +90^\circ]$ as indicated by the gray arrows. Potentially, a target source in front of the head could be incorrectly attributed to a location behind the head – a *front-back error* – as shown by the open circle.

source was then mixed with the masker signal in a binaural setting. As shown in Fig. 3, the target source varied in azimuth within the range of $[-90^\circ, 90^\circ]$ in 10° steps. However, this knowledge was not available to the systems and the full 360° azimuth range could be reported (and hence front-back errors could occur, as shown in Fig. 3).

The azimuth of the masker was randomly selected each time from the same azimuth range, while ensuring an angular distance of at least 10° between the two competing sources. Source locations were limited to this azimuth range because the Surrey database only includes azimuths in the frontal hemifield. However, our localisation system was not provided in any case with information that the azimuth of the source lay within this range; it was free to report the azimuth within the full 360° range. Both target and masker signals were normalised by their root mean square (RMS) amplitudes prior to spatialisation at two target-to-masker ratios (TMRs): 0 dB and -6 dB.

The baseline system was a state-of-the-art binaural localisation system using DNNs [7], with no model-based knowledge about the sources. The proposed localisation system exploiting model-based information employed the same localisation DNNs, but was given prior knowledge of the target source so that it could be “attended to” for more accurate localisation. The system was evaluated in three scenarios:

1) *The identity of the masker source was assumed to be available a priori*: The corresponding masker source GMM was used as the background model in order to estimate the set of localisation weights ω_{tf} in Eq. (1). Noise Set A was used for evaluating this scenario.

2) *The masker source identity was unavailable but the source was used for training the UBM*: In the second scenario, the UBM estimated from all the noise types in Noise Set A was used as the background model.

3) *The masker source identity was unavailable and the source was not used for training the UBM*: Similar to scenario 2, in this scenario the UBM estimated from all the noise types in Noise Set B was used as the background model. However, to test how well the system could generalise unseen noise types, Noise Set B was used for evaluation, which was not used for estimating the UBM.

In all scenarios, the background model (either the correct masker model or the UBM) was employed with or without adaptation.

For all the evaluated systems, the number of competing sources (two in this study) was assumed to be known *a priori*. Therefore each system reported two estimated source azimuths. The localisation performance was measured for the target source only by comparing true target source azimuths with the estimated azimuths. The *target localisation error rate* was measured by counting the number of sentences for which one of the azimuth estimates was outside a predefined grace boundary of $\pm 5^\circ$ with respect to the true target azimuths.

V. RESULTS AND DISCUSSION

The target localisation error rates of various binaural localisation systems are shown in Figs. 4 and 5 for the 0dB and the -6dB TMR cases, respectively. The results are averaged between the male speech source case and the female speech source case.

For each of the known masker types (top 6 maskers), the results using both UBM and the respective masker model were shown while for the unknown masker types in Set B only the results using the UBM are shown. Furthermore, since in this study the system did not assume that the sound sources are located in the frontal hemifield, the front-back error rates are shown as white bars in each figure.

The performance of the baseline system, which uses no model-based information for localisation, varied greatly across different masker conditions. The results show that the baseline system performed worse in masker conditions such as ‘phone’, ‘alarm’, and N-talker babble. While it is straightforward to understand why localisation of the target speech is less reliable in N-talker babble, as the masker spectrum overlaps that of the target speech, the poor performance in the ‘phone’ and ‘alarm’ conditions requires further explanation. This is likely due to two reasons. Firstly, these sounds are narrowband as shown in Fig. 2, and the target speech was completely masked at the frequencies where the maskers’ energy dominated. Hence, in these frequency bands only small glimpses were available to localise the target source and the baseline system tended to report high probabilities for the masker azimuths. Secondly, most of the energy for the narrowband sounds is located at high frequencies. The baseline system employed cross-correlation features and ILDs as localisation features, and these features are known to be more reliable in high frequencies [7]. In particular, the ILDs are more pronounced

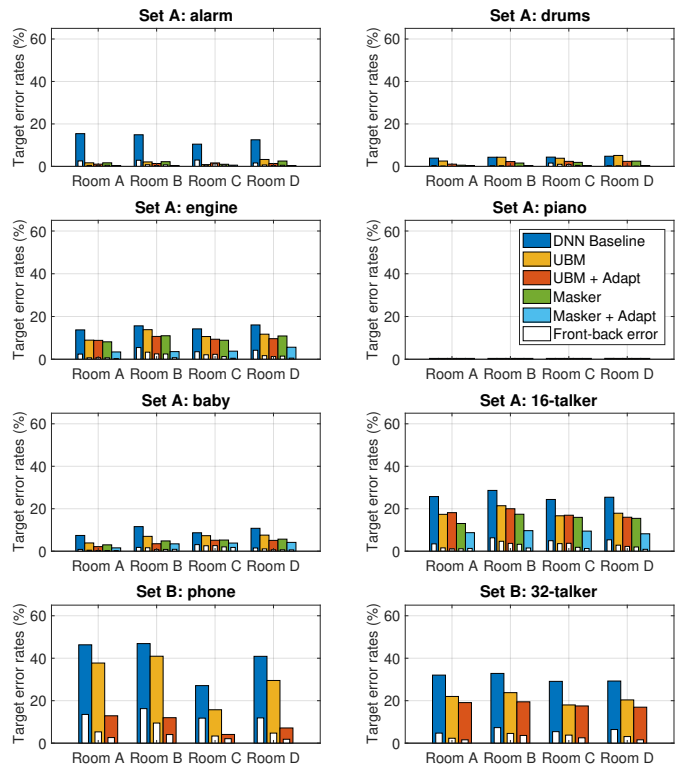


Fig. 4. Localisation error rates for localising the target source in the presence of various maskers, at a target-to-masker ratio (TMR) of 0dB. The proportions of front-back errors are indicated as white bars.

at high frequencies due to the size of the head compared to the wavelength of incoming sounds [17]. When these frequency bands were dominated by the masker, the system was more likely to report the location of the masker, especially when the global TMR was negative. Detailed analysis of the results shows that most localisation errors were due to incorrectly reporting the masker locations.

In contrast, the target localisation error rates were lower when the masker source dominates low frequencies. For example, in the ‘piano’ and ‘drums’ conditions, the target localisation error rates were below 10% at the TMR of 0dB, and even at the TMR of -6dB, the localisation error rates were still below 5% in the ‘piano’ condition. This is largely because the maskers dominated frequency regions that were less reliable for localisation with the DNN system, and thus the performance remained robust in the presence of the masker.

When model-based source knowledge was employed (in Figs. 4 and 5, UBM or Masker), the localisation system was able to identify the spectro-temporal regions dominated by the target speech. Therefore the system was more likely to report the location of the target source by weighing those regions more. Figs. 4 and 5 show that the results significantly improve when model-based information was used for localisation, particularly for narrowband sounds such as ‘alarm’.

A. Effect of using the masker model vs. the UBM

For each masker source in Noise Set A, a corresponding masker model was created. If the masker type is assumed

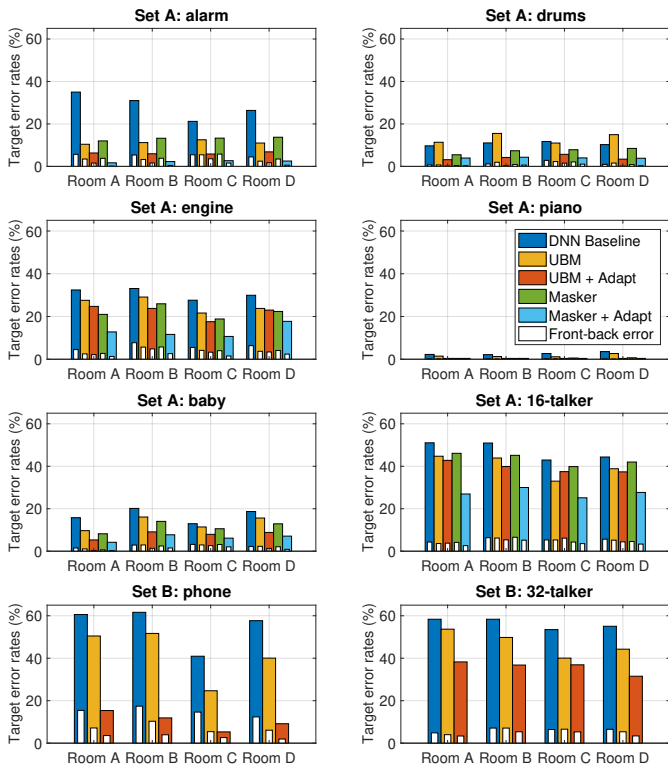


Fig. 5. Localisation error rates for localising the target source in the presence of various maskers, at a target-to-masker ratio (TMR) of -6 dB. The proportions of front-back errors are indicated as white bars.

to be known, the correct masker model was used as the background model, otherwise the UBM was used. Across different masker conditions, the use of a background model greatly improved the target localisation accuracy. Comparing the results in Fig. 4 using the UBM and the correct masker model, without adaptation, one can see that at 0 dB TMR using the UBM the system performed comparably with that using the correct masker model. This is largely because the UBM was trained using signals of the masker types from Noise Set A. The use of the GMM was effective to capture the spectral profiles of all the maskers and this helped localise the target source using the proposed framework.

As can be seen in Fig. 5, at -6 dB TMR, the localisation error reduction using the correct masker model becomes larger. This is expected, as a more detailed masker model will produce more reliable localisation weights than the general UBM. However, the use of the UBM minimises the assumptions made about the active masker sources. Such a system is potentially more suitable for an attention-driven model of sound localisation, in which the attended target source may be switched and the localisation weights can be dynamically recomputed in order to localise the newly attended source.

B. Effect of unseen masker types

In Noise Set A, when the identity of the masker was assumed unknown, the UBM was used as the background model. However, the parameters of the UBM were estimated using all the masker types from Noise Set A. To evaluate how

well the system could generalise to unseen masker types, Noise Set B was excluded from training the UBM. Comparing the UBM results for Noise Sets A and B, it is clear that using a UBM the system did not improve the target localisation accuracy for Noise Set B as much as for Noise Set A. This is especially the case for narrowband maskers. In the unseen ‘phone’ condition, the average target localisation error rate was reduced from 41% to 31% at the TMR of 0 dB, and was reduced from 55% to 42% at the TMR of -6 dB. In contrast, in the ‘alarm’ condition, the localisation error rate was reduced from 14% to 2% at 0 dB TMR and from 30% to 12% at -6 dB TMR. The error reduction was more similar between the 16-talker and 32-talker conditions, probably because the two speech babbles are more similar in spectral shapes. This suggests that the UBM worked less effectively for masker conditions that were not used for training the UBM, but it is still beneficial to use the UBM in unseen masker conditions.

C. Effect of background-model adaptation

As demonstrated by the results, background-model adaptation had a major contribution in reducing the target localisation error rates across various conditions. At the TMR of 0 dB the energy level mismatch between the trained background models and the test signals was minimal, but the adaptation process was still beneficial for Noise Set A as it can accommodate the level difference of individual test signals. The improvement was larger when the background model was the correct masker model (as opposed to the UBM), particularly for masker types that are difficult to model, such as ‘engine’ (many unpredictable events) and ‘16-talker babble’ (less stationary). For Noise Set B, the benefit of adaptation is more apparent, especially for the narrowband ‘phone’ condition. This suggests that the adaptation process was also able to adapt the model parameters to match better the spectral profile of an unseen masker.

Background-model adaptation benefitted the systems more at the TMR of -6 dB where there is a mismatch between the model level and the test signal level. Similarly to the 0 dB TMR case, the improvement appeared larger when the correct masker model was used as the background model. This is likely because with the correct masker model only the energy level needs to be adapted, whereas with the UBM the spectral profile also needs to be adapted. As the model parameters were adapted on-the-fly using a short test signal (less than 2-sec long) that was mixed with a masker sound, it is more difficult to reestimate the model parameters correctly.

To illustrate the effect of various stages in the proposed system, Fig. 6 shows an example of the estimated localisation weights. Here the target speech was mixed with an alarm sound at a TMR of -6 dB. The UBM was used when estimating the localisation weights. The ‘oracle’ mask shows the spectro-temporal regions dominated by the target speech using *a priori* information of the pre-mixed signals. The localisation weights estimated without adaptation bear some resemblance to the oracle mask, but incorrectly give more weight to the high frequency regions above 2 kHz that are dominated by the masker. With adaptation these masker-dominated regions are

given less weight, and the estimated mask is closer to the oracle mask.

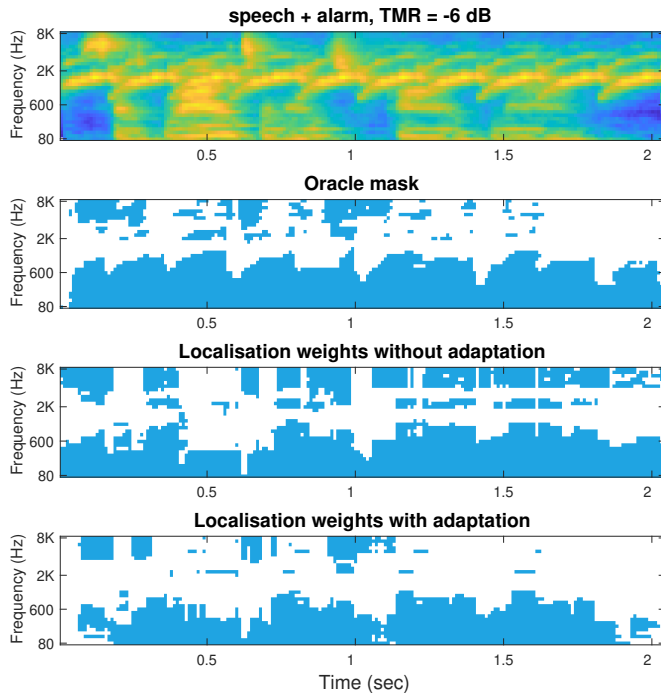


Fig. 6. Localisation weights estimated for target speech mixed with alarm sound at a target-to-masker ratio (TMR) of -6 dB using the UBM with and without adaptation. The ‘oracle’ mask was shown to indicate the spectrotemporal regions dominated by the target speech (blue regions). The localisation weights larger than 0.5 were shown in blue in the bottom 2 panels.

VI. CONCLUSIONS

This paper has presented a novel computational framework for binaural sound localisation that combines data-driven and model-based information flow. By jointly exploiting model-based knowledge about the source spectral characteristics in the acoustic scene, the system is able to selectively weight binaural cues in order to more reliably localise the attended source. To address the mismatch between the training and testing conditions, a model adaptation process is employed on-the-fly which re-estimates the background model parameters directly from the noisy observations in an iterative manner. Evaluation using masker sources with varying spectrotemporal complexity, including masker types that are not seen during the training stage, showed that by exploiting source models in this way, sound localisation performance can be improved substantially under conditions where multiple sources and room reverberation are present.

One of the advantages of the proposed system is that information fusion across frequency bands is done after the DNNs estimate the azimuth posteriors. The late fusion of information allows use of single-source data for training, as otherwise the amount of required training data would increase exponentially with the number of sources. In this way, the framework is also able to generalise to localisation of other sounds and is not tailored to the type of sound used for training.

The current evaluation involved only two sound sources (a target source and an interfering source). In general, the proposed approach is able to generalise to the scenario where there are more than one interfering source. In this case, one would consider all interfering sources as a single masker and use the universal background model to model the combination of all interfering sources. The complexity of the framework stays the same in this case.

The proposed sound localisation framework could be combined with source identification in order to estimate the identity of the target source that the system ‘attends’ to. Such an attention-driven model could be used to localise an attended source whose identity is not available *a priori*, e.g. a talker that speaks a keyword in an acoustic mixture. Source localisation and source identification could then interact in an ongoing iterative process. In addition, the framework described here could be integrated with an approach that uses source models to ‘perceptually restore’ parts of the target sound that have been masked [33]. Another future direction is the extension to cross-modal control. The proposed system is a general framework through which other modalities could also be incorporated, such as a vision system on a mobile robot.


ACKNOWLEDGEMENTS

This work was supported by the EC FP7 project TWO!EARS under grant agreement No. 618075.

REFERENCES

- [1] A. Bregman, *Auditory Scene Analysis*. Cambridge, MA: MIT Press, 1990.
- [2] D. L. Wang and G. J. Brown, Eds., *Computational Auditory Scene Analysis: Principles, Algorithms and Applications*. Wiley/IEEE Press, 2006.
- [3] G. C. Spence and J. Driver, “Covert spatial orienting in audition: exogenous and endogenous mechanisms,” *Journal of Experimental Psychology*, vol. 20, pp. 555–574, 1994.
- [4] A. Bronkhorst, “The cocktail-party problem revisited: early processing and selection of multi-talker speech,” *Attention, Perception, and Psychophysics*, vol. 77, no. 5, pp. 1465–1487, 2015.
- [5] B. H. Gaese and H. Wagner, “Precognitive and cognitive elements in sound localization,” *Zoology*, vol. 105, pp. 329–339, 2002.
- [6] D. E. Winkowski and E. I. Knudsen, “Top-down gain control of the auditory space map by gaze control circuitry in the barn owl,” *Nature*, vol. 439, no. 7074, pp. 336–339, 2006.
- [7] N. Ma, T. May, and G. J. Brown, “Exploiting deep neural networks and head movements for robust binaural localisation of multiple sources in reverberant environments,” *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 25, no. 12, pp. 2444–2453, 2017.
- [8] H. Christensen, N. Ma, S. Wrigley, and J. Barker, “A speech fragment approach to localising multiple speakers in reverberant environments,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Taipei, 2009, pp. 4593–4596.
- [9] M. Mandel, R. Weiss, and D. Ellis, “Model-based expectation maximization source separation and localization,” *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 18, no. 2, pp. 382–394, 2010.
- [10] J. Woodruff and D. L. Wang, “Binaural localization of multiple sources in reverberant and noisy environments,” *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 20, no. 5, pp. 1503–1512, 2012.
- [11] J. Woodruff and D. Wang, “Binaural detection, localization, and segregation in reverberant environments based on joint pitch and azimuth cues,” *IEEE T. Audio. Speech.*, vol. 21, no. 4, pp. 806–815, 2013.
- [12] J. Barker, M. Cooke, and D. Ellis, “Decoding speech in the presence of other sources,” *Speech Commun.*, vol. 45, pp. 5–25, 2005.
- [13] N. Ma, J. Barker, H. Christensen, and P. Green, “A hearing-inspired approach for distant-microphone speech recognition in the presence of multiple sources,” *Comput. Speech. Lang.*, vol. 27, no. 3, pp. 820–836, 2013.

- [14] J. Liu, H. Erwin, and G.-Z. Yang, "Attention driven computational model of the auditory midbrain for sound localization in reverberant environments," in *Neural Networks (IJCNN), The 2011 International Joint Conference on*, July 2011, pp. 1251–1258.
- [15] N. Ma, G. J. Brown, and J. A. Gonzalez, "Exploiting top-down source models to improve binaural localisation of multiple sources in reverberant environments," in *Proc. Interspeech*, 2015, pp. 160–164.
- [16] T. May, S. van de Par, and A. Kohlrausch, "A probabilistic model for robust localization based on a binaural auditory front-end," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 19, no. 1, pp. 1–13, 2011.
- [17] J. Blauert, *Spatial hearing - The psychophysics of human sound localization*. Cambridge, MA, USA: The MIT Press, 1997.
- [18] N. Ma, G. J. Brown, and T. May, "Robust localisation of multiple speakers exploiting deep neural networks and head movements," in *Proc. Interspeech*, 2015, pp. 3302–3306.
- [19] G. Brown and M. Cooke, "Computational auditory scene analysis," *Comput. Speech. Lang.*, vol. 8, pp. 297–336, 1994.
- [20] J. Woodruff and D. L. Wang, "Binaural localization of multiple sources in reverberant and noisy environments," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 20, no. 5, pp. 1503–1512, 2012.
- [21] T. May, S. van de Par, and A. Kohlrausch, "Binaural localization and detection of speakers in complex acoustic scenes," in *The technology of binaural listening*, J. Blauert, Ed. Berlin–Heidelberg–New York NY: Springer, 2013, ch. 15, pp. 397–425.
- [22] A. Varga and R. Moore, "Hidden Markov model decomposition of speech and noise," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 1990, pp. 845–848.
- [23] S. J. Rennie, J. R. Hershey, and P. A. Olsen, "Single-channel multitalker speech recognition," *IEEE Signal Process. Mag.*, vol. 27, pp. 66–80, 2010.
- [24] A. M. Reddy and B. Raj, "Soft mask estimation for single channel speaker separation," in *ISCA Tutorial and Research Workshop (ITRW) on Statistical and Perceptual Audio Processing*, 2004.
- [25] —, "Soft mask methods for single-channel speaker separation," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 15, no. 6, pp. 1766–1776, 2007.
- [26] J. A. Gonzalez, A. M. Gómez, A. M. Peinado, N. Ma, and J. Barker, "Spectral reconstruction and noise model estimation based on a basking model for noise robust speech recognition," *Circ. Sys. Signal Pr.*, pp. 1–30, 2017.
- [27] A. Dempster, N. Laird, and D. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *Journal of the Royal Statistical Society. Series B*, vol. 39, no. 1, pp. 1–38, 1977.
- [28] H. Wierstorf, M. Geier, A. Raake, and S. Spors, "A free database of head-related impulse response measurements in the horizontal plane with multiple distances," in *Proc. 130th Conv. Audio Eng. Soc.*, 2011.
- [29] C. Hummerson, R. Mason, and T. Brookes, "Dynamic precedence effect modeling for source separation in reverberant environments," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 18, no. 7, pp. 1867–1871, 2010.
- [30] M. Cooke, J. Barker, S. Cunningham, and X. Shao, "An audio-visual corpus for speech perception and automatic speech recognition," *J. Acoust. Soc. Am.*, vol. 120, pp. 2421–2424, 2006.
- [31] J. S. Garofolo, L. F. Lamel, W. M. Fisher, J. G. Fiscus, D. S. Pallett, and N. L. Dahlgren, "DARPA TIMIT Acoustic-phonetic continuous speech corpus CD-ROM," *National Inst. Standards and Technol. (NIST)*, 1993.
- [32] T. May, N. Ma, and G. J. Brown, "Robust localisation of multiple speakers exploiting head movements and multi-conditional training of binaural cues," in *Proc. ICASSP*, 2015, pp. 2679–2683.
- [33] U. Remes, A. López, L. Juvela, K. Palomäki, G. J. Brown, P. Alku, and M. Kurimo, "Comparing human and automatic speech recognition in a perceptual restoration experiment," *Computer Speech and Language*, vol. 35, pp. 14–31, June 2015.



Ning Ma obtained a MSc (distinction) in Advanced Computer Science in 2003 and a PhD in hearing-inspired approaches to automatic speech recognition in 2008, both from the University of Sheffield. He has been a visiting research scientist at the University of Washington, Seattle, and a Research Fellow at the MRC Institute of Hearing Research, working on auditory scene analysis with cochlear implants. Since 2015 he is a Research Fellow at the University of Sheffield, working on computational hearing. His research interests include robust automatic speech recognition, computational auditory scene analysis, and hearing impairment. He has authored or coauthored over 40 papers in these areas.



a lecturer at the Department of Languages and Computer Science, University of Malaga, Spain.

Jose A. Gonzalez received the B.Sc. and Ph.D. degrees in Computer Science, both from the University of Granada, Spain, in 2006 and 2013, respectively. During his Ph.D. he did two research stays at the Speech and Hearing Research Group, University of Sheffield, U.K., studying missing data approaches for noise robust speech recognition. From 2013 to 2017 he was a Research Associate at the University of Sheffield working on silent speech technology with special focus on direct speech synthesis from speech-related biosignals. Since October 2017, he is



editor for medical applications. He has authored more than 100 papers and is the co-editor (with Prof. DeLiang Wang) of the IEEE book "Computational auditory scene analysis: principles, algorithms and applications".

Guy J. Brown obtained a BSc (Hons) Applied Science from Sheffield City Polytechnic in 1984 and a PhD in Computer Science from the University of Sheffield in 1992. He was appointed to a Chair in the Department of Computer Science, University of Sheffield, in 2013. He has held visiting appointments at LIMSI-CNRS (France), Ohio State University (USA), Helsinki University of Technology (Finland) and ATR (Japan). His research interests include computational auditory scene analysis, speech perception, hearing impairment and acoustic monitoring