



HAL
open science

A Cascaded Multiple-Speaker Localization and Tracking System

Xiaofei Li, Yutong Ban, Laurent Girin, Xavier Alameda-Pineda, Radu Horaud

► **To cite this version:**

Xiaofei Li, Yutong Ban, Laurent Girin, Xavier Alameda-Pineda, Radu Horaud. A Cascaded Multiple-Speaker Localization and Tracking System. IWAENC - LOCATA Challenge Workshop - a satellite event of IWAENC 2018, Sep 2018, Tokyo, Japan. pp.1-5. hal-01957137

HAL Id: hal-01957137

<https://inria.hal.science/hal-01957137>

Submitted on 17 Dec 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A CASCADED MULTIPLE-SPEAKER LOCALIZATION AND TRACKING SYSTEM

Xiaofei Li,¹ Yutong Ban,¹ Laurent Girin,^{1,2} Xavier Alameda-Pineda,¹ Radu Horaud,¹

¹ Inria Grenoble Rhône-Alpes, France

² Univ. Grenoble Alpes, Grenoble-INP, GIPSA-lab, France

ABSTRACT

This paper presents an online multiple-speaker localization and tracking method, as the INRIA-Perception contribution to the LOCATA Challenge 2018. First, the recursive least-square method is used to adaptively estimate the direct-path relative transfer function as an interchannel localization feature. The feature is assumed to associate with a single speaker at each time-frequency bin. Second, a complex Gaussian mixture model (CGMM) is used as a generative model of the features. The weight of each CGMM component represents the probability that this component corresponds to an active speaker, and is adaptively estimated with an online optimization algorithm. Finally, taking the CGMM component weights as observations, a Bayesian multiple-speaker tracking method based on the variational expectation maximization algorithm is used. The tracker accounts for the variation of active speakers and the localization miss measurements, by introducing speaker birth and sleeping processes. The experiments carried out on the development dataset of the challenge are reported.

Index Terms— sound-source localization, multiple moving speakers, tracking, reverberant environments, LOCATA Challenge

1. INTRODUCTION

For multiple-speaker localization, the W-disjoint orthogonality (WDO) [1] assumption is widely used. It assumes that the audio signal is dominated by only one speaker in each small region of the time-frequency (TF) domain, because of the natural sparsity of speech signals in this domain. After applying the short-time Fourier transform (STFT), interchannel localization features, e.g. interaural phase difference, can be extracted. To assign the interchannel features to multiple speakers, a mixture of Gaussian mixture models (GMMs) is used as a generative model of the interchannel features of multiple speakers in [2] with each GMM representing one speaker, and then feature assignment was done based on the maximum likelihood criteria. In [3], instead of setting one GMM for each speaker, a single complex GMM (CGMM) is used for all speakers with each component representing one candidate speaker location. After maximizing the likelihood of the features, with an expectation-maximization (EM) algorithm, the weight of each component represents the probability that there is an active speaker at the corresponding candidate location. To localize moving speakers, in [4], based on a CGMM model similar to [3], a recursive EM algorithm was proposed to update online the CGMM component weights. Counting and localization of active speakers can be jointly carried out by selecting the components with large weights. Taking

the instantaneous outputs of a localization method as observations and using a speaker dynamic model, Bayesian tracking techniques estimate the posterior distribution of source locations, e.g. [5, 6]. To tackle the tracking problem with an unknown and time-varying number of speakers, additional model features such as observation-to-speaker assignments, speaker track birth/death processes and a model of speech activity can be included [7, 8, 9, 10].

In the present paper, we present our online multiple-speaker localization and tracking method contributed to the LOCATA Challenge, which is composed of three modules:

- A recursive DP-RTF (direct-path relative transfer function) estimation module. In the STFT domain, the room impulse response (RIR) can be approximated by the convolutive transfer function (CTF) [11, 12]. DP-RTF is defined by the ratio between the first taps of the CTF of two microphones, thus encodes the direct-path information and is used as an interchannel feature being robust against reverberation. The CTF estimation used for DP-RTF extraction was formulated in batch mode in [13, 14] with speakers considered as static. Based on recursive least-square (RLS), the online CTF estimation method was proposed in [15, 16], and will be briefly presented in this paper, for moving speaker localization.
- An online multiple-speaker localization module. In [14], we adopt the above-mentioned CGMM model [3] to assign the DP-RTF features to speakers, in addition, an entropy-based regularization term was used to impose the spatial sparsity of the estimated component weights. However, [14] only considered the batch mode and static speakers. The recursive EM algorithm proposed in [4] was adopted in [15] for online likelihood maximization without using the entropy regularization. Furtherly, an online optimization method, i.e. exponentiated gradient (EG) [17], was used in [16] for simultaneously online likelihood maximization and entropy minimization. This method will be briefly presented in this paper.
- A multiple-speaker tracking module. In [16], a multiple-speaker tracking method was proposed. The results of the above localization module are taken as inputs and efficiently exploited in a Bayesian framework. The problem is efficiently solved by a variational expectation maximization (VEM) algorithm. Speaker birth and sleeping processes are included in the tracking process. The sleeping process is efficient to tackle a missed detection by the localization procedure. This paper will briefly present this multiple-speaker tracker, please refer to [8, 10, 18, 16] for more detailed algorithmic derivation and description.

In Fig. 1, the three layers of temporal evolutions depicts these three modules, respectively. In the following, we present them one by one from Section 2 to 4, and then gives the experiments on the development dataset of the LOCATA challenge [19] in Section 5.

This work was supported by the ERC Advanced Grant VHIA #340113.

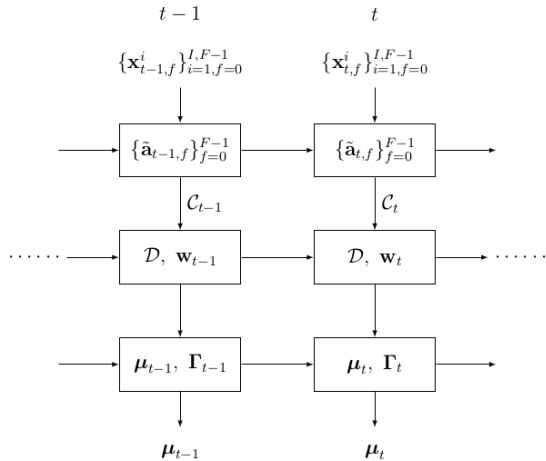


Figure 1: Flowchart of the three localization and tracking modules.

2. RECURSIVE DP-RTF ESTIMATION

To simplify the presentation, let us first consider the noise-free single-speaker case. In the time domain, the i -th, $i = 1, \dots, I$, microphone signal is: $x^i(n) = a^i(n) \star s(n)$, where n is the time index, $s(n)$ is the source signal, $a^i(n)$ is the RIR from the source to the i -th microphone, and \star denotes the convolution. Applying the STFT, and using the CTF approximation, we have for each frequency index $f = 0, \dots, F-1$: $x_{t,f}^i = a_{t,f}^i \star s_{t,f}$, where $x_{t,f}^i$ and $s_{t,f}$ are the STFT coefficients of the corresponding signals, and the CTF $a_{t,f}^i$ is a subband representation of $a^i(n)$. Here, the convolution is executed w.r.t the frame index t . The first CTF coefficient $a_{0,f}^i$ mainly consists of the direct-path information, and DP-RTF is defined as the ratio between the first CTF coefficients of two channels: $a_{0,f}^i/a_{0,f}^r$, where channel r is the reference channel.

Based on the cross-relation method [20], for one microphone pair (i, j) , we have $\mathbf{x}_{t,f}^{i\top} \mathbf{a}_f^j = \mathbf{x}_{t,f}^{j\top} \mathbf{a}_f^i$, where \top denotes matrix/vector transpose, the convolution vectors are $\mathbf{a}_f^i = [a_{0,f}^i, \dots, a_{Q-1,f}^i]^\top$ and $\mathbf{x}_{t,f}^i = [x_{t,f}^i, \dots, x_{t-Q+1,f}^i]^\top$, with Q denoting the CTF length. We concatenate the CTF vector of all channels as $\mathbf{a}_f = [\mathbf{a}_f^\top, \dots, \mathbf{a}_f^\top]^\top$. For each microphone pair (i, j) , we construct a cross-relation equation in terms of \mathbf{a}_f . For this aim, define:

$$\mathbf{x}_{t,f}^{ij} = \underbrace{[0, \dots, 0]_{(i-1)Q}}_{(i-1)Q}, \underbrace{\mathbf{x}_{t,f}^{j\top}}_{(j-i-1)Q}, \underbrace{[0, \dots, 0, -\mathbf{x}_{t,f}^{i\top}, 0, \dots, 0]_{(I-j)Q}}_{(I-j)Q}. \quad (1)$$

Then we have $\mathbf{x}_{t,f}^{ij\top} \mathbf{a}_f = 0$. The CTF vector \mathbf{a}_f can be estimated by solving this equation. To avoid a trivial solution, i.e. $\mathbf{a}_f = \mathbf{0}$, we constrain the first CTF coefficient of the reference channel, say $r = 1$, to be equal to 1. This leads to the following equation

$$\tilde{\mathbf{x}}_{t,f}^{ij\top} \tilde{\mathbf{a}}_f = y_{t,f}^{ij}, \quad (2)$$

where $-y_{t,f}^{ij}$ is the first entry of $\mathbf{x}_{t,f}^{ij}$, $\tilde{\mathbf{x}}_{t,f}^{ij}$ is $\mathbf{x}_{t,f}^{ij}$ with the first entry removed, and $\tilde{\mathbf{a}}_f$ is the relative CTF vector:

$$\tilde{\mathbf{a}}_f = \left[\frac{\tilde{\mathbf{a}}_f^\top}{a_{0,f}^1}, \frac{\mathbf{a}_f^{2\top}}{a_{0,f}^1}, \dots, \frac{\mathbf{a}_f^{I\top}}{a_{0,f}^1} \right]^\top. \quad (3)$$

where $\tilde{\mathbf{a}}_f^1 = [a_{1,f}^1, \dots, a_{Q-1,f}^1]^\top$, is \mathbf{a}_f^1 with the first entry removed. For $i = 2, \dots, I$, the DP-RTFs appear in (3) as the first entries of $\frac{\mathbf{a}_f^{i\top}}{a_{0,f}^1}$.

The DP-RTF estimation amounts to solving the linear problem Eq. (2). Eq. (2) is defined for one microphone pair at one frame. For the online case, we would like to update the estimate of $\tilde{\mathbf{a}}_f$ using the current frame, say t . There is a total of $M = I(I-1)/2$ distinct microphone pairs, For notational convenience, instead of using ij , we use $m = 1, \dots, M$ denote the index of microphone pair. The fitting error of (2) is $e_{t,f}^m = y_{t,f}^m - \tilde{\mathbf{x}}_{t,f}^{m\top} \tilde{\mathbf{a}}_f$. At the current frame t , for the microphone pair m , RLS aims to minimize the error

$$J_{t,f}^m = \sum_{t'=1}^{t-1} \sum_{m'=1}^M \lambda^{t-t'} |e_{t',f}^{m'}|^2 + \sum_{m'=1}^m |e_{t,f}^{m'}|^2, \quad (4)$$

which sums up the fitting error of all the microphone pairs for the past frames and the microphone pairs up to m for the current frame. The forgetting factor $\lambda \in (0, 1]$ gives exponentially lower weight to older frames, whereas at one given frame, all microphone pairs have the same weight. In RLS, this minimization problem can be recursively solved. For the detailed recursion procedure please refer to Algorithm 1 in [16]. For each frame t , let $\tilde{\mathbf{a}}_{t,f}$ denote the CTF estimate, and $\tilde{c}_{t,f}^i$, $i = 2, \dots, I$ denote the DP-RTF estimates extracted from $\tilde{\mathbf{a}}_{t,f}$. Note that implicitly we have $\tilde{c}_{t,f}^1 = 1$.

We now introduce how to extend the above method to the noisy multiple-speaker case. We assume that, over a short time, the speakers are static and only one source is active at each frequency bin. Therefore, the CTF can be estimated using a small number of recent frames. This can be done by adjusting the forgetting factor λ : To approximately have a memory of P frames, we can set $\lambda = \frac{P-1}{P+1}$. P should be empirically set to achieve a good tradeoff between the validity of the above assumptions and a robust CTF estimate. To suppress the noise, we use the inter-frame spectral subtraction algorithm proposed in [21, 15]. At each frequency bin f , frames are first classified into speech frames and noise frames, then inter-frame spectral subtraction is applied between them. In the RLS process, only the speech frames (after spectral subtraction) are used, and the noise frames are skipped. In practice, a DP-RTF estimate can sometimes be unreliable due to the imperfect performance of the above-mentioned methods. We use the consistency test method proposed in [14, 16] to detect the unreliable estimates. Finally, at frame t , we obtain a set of features $\mathcal{C}_t = \{\{\tilde{c}_{t,f}^i\}_{i \in \mathcal{I}_f}\}_{f=0}^{F-1}$, where $\mathcal{I}_f \subseteq \{2, \dots, I\}$ denotes the set of microphone indices that pass the consistency test. Each of the features is assumed to be associated with a single speaker.

3. ONLINE MULTIPLE-SPEAKER LOCALIZATION

In order to assign the DP-RTF features in \mathcal{C}_t to speakers, the CGMM generative model proposed in [3] is adopted. We define a set \mathcal{D} of D candidate source locations. Let $d = 1, \dots, D$ denote the location index. The probability, that an observed feature $\tilde{c}_{t,f}^i$ is emitted by candidate locations, is modelled by a CGMM:

$$P(\tilde{c}_{t,f}^i | \mathcal{D}) = \sum_{d=1}^D w^d \mathcal{N}_c(\tilde{c}_{t,f}^i; c_f^{i,d}, \sigma^2), \quad (5)$$

where the mean $c_f^{i,d}$ is precomputed based on the direct-path propagation model from the d -th candidate location to the microphone i . The variance σ^2 is empirically set as a constant value. The component weight $w^d \geq 0$, with $\sum_{d=1}^D w^d = 1$, is the only free model parameter. Let us denote the vector of weights with $\mathbf{w} = [w^1, \dots, w^D]^\top$.

Let \mathcal{L}_t denote the log-likelihood (normalized by the number of features) of the features in \mathcal{C}_t , as a function of w^d . Once \mathcal{L}_t is maximized, the component weight w^d represents the probability that there exist an active speaker at the d -th candidate location. In addition, taking into account the fact that the number of actual active speakers is much lower than the number of candidate locations, an entropy minimization is taken as a regularization term to impose a sparse distribution of w^d . The entropy is defined as $H = -\sum_{d=1}^D w^d \log(w^d)$. Then the overall cost function is $-\mathcal{L}_t + \gamma H$. In the present online framework, we proceed to a recursive update of the weight vector at current frame t , hence now denoted \mathbf{w}_t , based on the estimate at previous frame \mathbf{w}_{t-1} and on the DP-RTF features of the current frame. This can be formulated as an online optimization problem [17]:

$$\mathbf{w}_t = \underset{\mathbf{w}}{\operatorname{argmin}} \chi(\mathbf{w}, \mathbf{w}_{t-1}) + \eta(-\mathcal{L}_t + \gamma H), \quad (6)$$

with the constraints that w^d , $d = 1, \dots, D$ are positive and sum to 1. $\chi(\mathbf{w}, \mathbf{w}_{t-1})$ is some distance measure between \mathbf{w} and \mathbf{w}_{t-1} . The positive constant η controls the parameter update rate. To exploit the fact that w^d are probability masses, we use the Kullback-Leibler divergence, i.e. $\chi(\mathbf{w}, \mathbf{w}_{t-1}) = \sum_{d=1}^D w^d \log \frac{w^d}{w_{t-1}^d}$, which results in the exponentiated gradient algorithm [17]. Let $\Delta_{t-1} = \frac{\partial(\eta(-\mathcal{L}_t + \gamma H))}{\partial w^d} \Big|_{w_{t-1}^d}$ denote the partial derivatives of $\eta(-\mathcal{L}_t + \gamma H)$ w.r.t w^d at the point w_{t-1}^d . Then, the exponentiated gradient $r_{t-1}^d = e^{-\Delta_{t-1}}$ is used to update the weights

$$w_t^d = \frac{r_{t-1}^d w_{t-1}^d}{\sum_{d'=1}^D r_{t-1}^{d'} w_{t-1}^{d'}}, \quad \forall d \in \mathcal{D}. \quad (7)$$

It is obvious from (7) that the parameter constraints, namely positivity and summation to 1, are automatically satisfied.

4. MULTIPLE-SPEAKER TRACKING

In the following, upper case letters denote random variables while lower case letters denote their realizations. Let N be the maximum number of tracks (speakers), and let n be the speaker index. Moreover, let $n = 0$ denote *no speaker*, or background noise. Let \mathbf{S}_{tn} be a latent (or state) variable associated with speaker n at frame t , and let $\mathbf{S}_t = (\mathbf{S}_{t1}, \dots, \mathbf{S}_{tn}, \dots, \mathbf{S}_{tN})$. \mathbf{S}_{tn} is composed of two parts: the speaker direction and the speaker velocity. In this work, speaker direction is defined by an azimuth θ_{tn} . To avoid phase (circular) ambiguity we describe the direction with the unit vector $\mathbf{U}_{tn} = (\cos(\theta_{tn}), \sin(\theta_{tn}))^\top$. Moreover, let $V_{tn} \in \mathbb{R}$ be the angular velocity. Altogether we define a realization of the state variable as $\mathbf{s}_{tn} = [\mathbf{u}_{tn}; v_{tn}]$ where the notation $[\cdot; \cdot]$ stands for vertical vector concatenation.

As mentioned above, the CGMM component weight w^d represents the probability that there exist an active speaker at the d -th candidate location. The frame-wise localization of active speakers

can be carried out by peak picking over w^d . However, to fully use the weight information, without applying peak picking, all the candidate locations and their associated weights are used. Formally, let $\mathbf{O}_t = (\mathbf{O}_{t1}, \dots, \mathbf{O}_{td}, \dots, \mathbf{O}_{tD})$ be the observed variables at frame t . Each realization \mathbf{o}_{td} of \mathbf{O}_{td} is composed of a candidate location, or azimuth $\theta_{td} \in \mathcal{D}$, and a weight w_{td} . As above, let the azimuth be described by a unit vector $\mathbf{b}_{td} = (\cos(\theta_{td}), \sin(\theta_{td}))^\top$. In summary we have $\mathbf{o}_{td} = [\mathbf{b}_{td}; w_{td}]$. Moreover, let Z_{td} be a (latent) assignment variable associated with each observed variable \mathbf{O}_{td} , such that $Z_{td} = n$ means that the observation indexed by d at frame t is assigned to speaker $n \in \{0, \dots, N\}$.

4.1. Bayesian Tracking Model

The problem at hand can be cast into the estimation of the filtering distribution $p(\mathbf{s}_t, \mathbf{z}_t | \mathbf{o}_{1:t})$, and further inference of \mathbf{s}_t and \mathbf{z}_t . By applying the Bayes rule, the filtering distribution is proportional to:

$$p(\mathbf{s}_t, \mathbf{z}_t | \mathbf{o}_{1:t}) \propto p(\mathbf{o}_t | \mathbf{s}_t, \mathbf{z}_t) p(\mathbf{z}_t) p(\mathbf{s}_t | \mathbf{o}_{1:t-1}), \quad (8)$$

which contains the following three terms.

The audio observation model $p(\mathbf{o}_t | \mathbf{s}_t, \mathbf{z}_t)$ describes the distribution of the observations given speakers state and assignment. We assume the different observations are independent conditionally to speakers state and assignment. For each observation, we adopt the weighted-data GMM model of [22]: $p(\mathbf{b}_{td} | Z_{td} = n, \mathbf{s}_{tn}; w_{td}) = \mathcal{N}(\mathbf{b}_{td}; \mathbf{M} \mathbf{s}_{tn}, \frac{1}{w_{td}} \boldsymbol{\Sigma})$ for $n \in \{1, \dots, N\}$, where the matrix $\mathbf{M} = [\mathbf{I}_{2 \times 2}, \mathbf{0}_{2 \times 1}]$ projects the state variable onto the space of source directions and $\boldsymbol{\Sigma}$ is a covariance matrix (set empirically to a fixed value). Note that the weight plays the role of a precision: The higher the weight w_{td} , the more reliable the source direction \mathbf{b}_{td} . The case $Z_{td} = 0$ follows a uniform distribution over the volume of the observation space, i.e. $p(\mathbf{b}_{td} | Z_{td} = 0) = \mathcal{U}(\operatorname{vol}(\mathcal{G}))$.

The prior distribution of the assignment variable, i.e. $p(\mathbf{z}_t)$, is independent over observations and is assumed to be uniformly distributed over all the speakers, i.e. $p(Z_{td} = n) = \pi_{dn} = \frac{1}{N+1}$.

To calculate the state predictive distribution $p(\mathbf{s}_t | \mathbf{o}_{1:t-1})$, we marginalize it over \mathbf{s}_{t-1} :

$$p(\mathbf{s}_t | \mathbf{o}_{1:t-1}) = \int p(\mathbf{s}_t | \mathbf{s}_{t-1}) p(\mathbf{s}_{t-1} | \mathbf{o}_{1:t-1}) d\mathbf{s}_{t-1}. \quad (9)$$

We model the state dynamics $p(\mathbf{s}_t | \mathbf{s}_{t-1})$ as a linear-Gaussian first-order Markov process, independent over the speakers, i.e. $p(\mathbf{s}_{t,n} | \mathbf{s}_{t-1,n}) = \mathcal{N}(\mathbf{s}_{t,n}; \mathbf{D}_{t-1,n} \mathbf{s}_{t-1,n}, \boldsymbol{\Lambda}_{tn})$, where $\boldsymbol{\Lambda}_{tn}$ is the dynamics' covariance matrix and $\mathbf{D}_{t-1,n}$ is the state transition matrix. Importantly, since the state position subvector \mathbf{u}_{tn} lies on the unit circle, the dynamic model is designed for circular motion. Given the estimated azimuth angle $\theta_{t-1,n}$ at frame $t-1$, the state transition matrix can be written as:

$$\mathbf{D}_{t-1,n} = \begin{pmatrix} 1 & 0 & -\sin(\theta_{t-1,n}) \\ 0 & 1 & \cos(\theta_{t-1,n}) \\ 0 & 0 & 1 \end{pmatrix}. \quad (10)$$

In the following $\mathbf{D}_{t-1,n}$ is written as \mathbf{D} for notational simplicity.

4.2. Variational Expectation Maximization Algorithm

To estimate the model parameters Θ , and infer the posterior distribution of $(\mathbf{s}_t, \mathbf{z}_t)$, we adopt an EM algorithm that alternates between computing and maximizing the expected complete-data log-likelihood:

$$J(\Theta, \Theta^\circ) = \mathbf{E}_{p(\mathbf{z}_t, \mathbf{s}_t | \mathbf{o}_{1:t}, \Theta^\circ)} [\log p(\mathbf{z}_t, \mathbf{s}_t, \mathbf{o}_{1:t} | \Theta)], \quad (11)$$

where \mathbf{E} denotes expectation, Θ° are the ‘‘old’’ model parameter estimates (obtained at previous iteration). Given the hybrid combinatorial-continuous nature of the latent space, it is computationally heavy to operate with the exact a posteriori distribution $p(\mathbf{z}_t, \mathbf{s}_t | \mathbf{o}_{1:t})$. We thus use a variational approximation to solve the problem [8], which factorizes the posterior distribution as:

$$p(\mathbf{z}_t, \mathbf{s}_t | \mathbf{o}_{1:t}) \approx q(\mathbf{z}_t, \mathbf{s}_t) = q(\mathbf{z}_t) \prod_{n=0}^N q(\mathbf{s}_{tn}). \quad (12)$$

It is seen that the posterior distribution factorizes across speakers. This principle is also valid for time $t-1$. As a result, the predictive distribution (9) also factorizes across speakers. If the posterior distribution $q(\mathbf{s}_{t-1,n})$ is assumed to be a Gaussian with mean $\boldsymbol{\mu}_{t-1,n}$ and variance $\boldsymbol{\Gamma}_{t-1,n}$, then the per-speaker predictive distribution $p(\mathbf{s}_{tn} | \mathbf{o}_{1:t-1})$ is a Gaussian:

$$p(\mathbf{s}_{tn} | \mathbf{o}_{1:t-1}) = \mathcal{N}(\mathbf{s}_{tn}; \mathbf{D}\boldsymbol{\mu}_{t-1,n}, \mathbf{D}\boldsymbol{\Gamma}_{t-1,n}\mathbf{D}^\top + \boldsymbol{\Lambda}_{tn}). \quad (13)$$

In E-S step, it can be derived that the variational posterior distribution of $q(\mathbf{s}_{tn})$ is a Gaussian distribution, with variance and mean respectively as

$$\boldsymbol{\Gamma}_{tn} = \left[\left(\sum_{d=1}^D \alpha_{tdn} w_t^d \right) \mathbf{H}^\top \boldsymbol{\Sigma}^{-1} \mathbf{H} + \left(\boldsymbol{\Lambda}_{tn} + \mathbf{D}\boldsymbol{\Gamma}_{t-1,n}\mathbf{D}^\top \right)^{-1} \right]^{-1},$$

$$\boldsymbol{\mu}_{tn} = \boldsymbol{\Gamma}_{tn} \left[\mathbf{H}^\top \boldsymbol{\Sigma}^{-1} \left(\sum_{d=1}^D \alpha_{tdn} w_t^d \mathbf{b}_{td} \right) + \left(\boldsymbol{\Lambda}_{tn} + \mathbf{D}\boldsymbol{\Gamma}_{t-1,n}\mathbf{D}^\top \right)^{-1} \mathbf{D}\boldsymbol{\mu}_{t-1,n} \right],$$

where $\alpha_{tdn} = q(Z_{td} = n)$ is the variational posterior distribution of the assignment variable, which is derived in E-Z step, the model parameters $\boldsymbol{\Lambda}_{tn}$ is updated in M-step, please refer to [16] for detailed descriptions. At time t , VEM converges after a few iterations, e.g. 5 iterations used in this work. After convergence, the posterior mean of \mathbf{s}_{tn} , i.e. $\boldsymbol{\mu}_t = (\boldsymbol{\mu}_{t1}, \dots, \boldsymbol{\mu}_{tn}, \dots, \boldsymbol{\mu}_{tN})$, is output as the result of speaker tracking, and together with the posterior variance $\boldsymbol{\Gamma}_t = (\boldsymbol{\Gamma}_{t1}, \dots, \boldsymbol{\Gamma}_{tn}, \dots, \boldsymbol{\Gamma}_{tN})$ are transmitted to time $t+1$.

4.3. Speaker Track Birth Process and Activity Detection

A track birth process is used to initialize new tracks, i.e. new speakers that enter the scenario. The general principle is the following. In a short period of time, say from frame $t-L$ to frame t , we assume that at most one new (yet untracked) speaker appears. For each frame from $t-L$ to t , we select the observation assigned to the background noise with the highest weight, and obtain an observation sequence $\tilde{\mathbf{o}}_{t-L:t}$. We then calculate the marginal likelihood of this sequence according to our model, $\tau_0 = p(\tilde{\mathbf{o}}_{t-L:t})$. If these observations have been generated by a new speaker, they exhibit smooth trajectories, and τ_0 will be high. Therefore, the birth process is conducted by comparing τ_0 with a threshold. The posterior distribution of the assignment variable, i.e. α_{tdn} , can be

used for multiple-speaker activity detection. This can be formalized as testing, for each frame t and each speaker n , whether the weighted assignments $\sum_{d=1}^D \alpha_{tdn} w_t^d$ (averaged over a small number of frames) is larger than a threshold.

5. EXPERIMENTS ON LOCATA DEVELOPMENT DATA

We report the results of on the LOCATA development corpus for tasks #3 and #5 with a single moving speaker, and tasks #4 and #6 with two moving speakers, each task comprising three recorded sequences. In this work, we use four microphones with indices {5, 8, 11, 12} out of the twelve microphones of a spherical array built in the head of a humanoid robot, i.e. NAO, to perform azimuth localization and tracking. These four microphones are mounted on the top of the robot head, and they approximately lie in a horizontal plane. We perform 360°-wide azimuth estimation and tracking: $D = 72$ azimuth directions at every 5° in $[-175^\circ, 180^\circ]$ are used as candidate directions. The TDOAs are computed based on the coordinate of microphones, which are then used to compute the phase of the CGMM means, while the magnitude of the CGMM means are set to a constant for all the frequencies. All the recorded signals are resampled to 16 kHz. The STFT uses the Hamming window with length of 16 ms and shift of 8 ms. The CTF length is $Q = 8$ frames. The RLS forgetting factor λ is computed using $\rho = 1$. The exponentiated gradient update factor is $\eta = 0.07$. The entropy regularization factor is $\gamma = 0.1$. For the tracker, the covariance matrix is set to be isotropic $\boldsymbol{\Sigma} = 0.03\mathbf{I}_2$.

A localization and tracking example Fig. 2 shows an example for a LOCATA sequence. Two speakers are moving and continuously speaking with short pauses. Fig. 2 (b) shows that the localization method achieves a heatmap with negligible interferences and smooth peak evolution. From Fig. 2 (c), it is seen that the tracking method further smooths the speaker moving trajectories. Even when the observations have a low weight, the tracker is still able to give the correct speaker trajectories. This is ensured by exploiting the source dynamics model. As a result, the tracker is able to preserve the identity of speakers in spite of the (short) speech pauses. In the presented sequence example, the estimated speaker identities are quite consistent with the ground truth.

Quantitative results The following metrics are used for quantitative evaluation. The detected speaker is considered to be successfully localized if the azimuth difference is not larger than 15°. The absolute error is calculated for the successfully localized sources. The mean absolute error (MAE) is computed by averaging the absolute error of all speakers and frames. For the unsuccessful localizations, we count the miss detections (MD) (speaker active but not detected) and false alarms (FA) (speaker detected but not active). Then the MD and FA rates are computed, using all the frames, as the percentage of the total MDs and FAs out of the total number of actual speakers, respectively. In addition to these localization metrics, we also count the identity switches (IDs) to evaluate the tracking continuity. IDs represents the number of the identity changes in the tracks for a whole test sequence. Table 1 gives the quantitative localization and tracking results. The localization results are obtained by applying peak picking on the GMM weights. The tracker slightly reduces the FA rate compared to the localization method alone mainly by eliminating some spurious peaks that are present in the localization outputs. It also reduces the MD rate since some correct speaker trajectories can be recovered even when the observations have (very) weak weights, as explained above. The identity

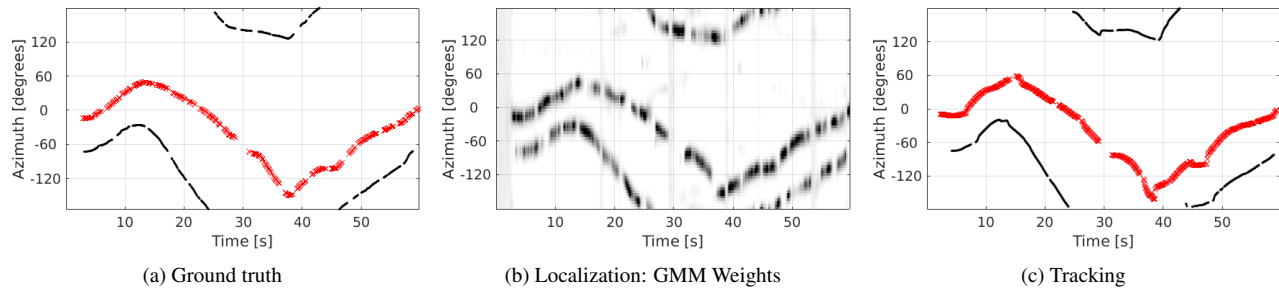


Figure 2: Results of speaker localization and tracking for Recording #1 / Task #6. (a) Ground truth trajectory and voice activity (red for speaker 1, black for speaker 2). Intervals in the trajectories are speaking pauses. (b) Result for localization, i.e. GMM Weights. (f) Result for tracker. Black and red colors demonstrate a successful tracking, i.e. continuity of the tracks despite of speech pauses.

Table 1: Localization and tracking results.

	MD rate (%)	FA rate (%)	MAE (°)	IDs
Localization	24.1	12.7	4.0	-
Tracking	22.7	12.4	4.1	10

switches are mainly due to the crossing of speaker trajectories, a hard case for the source dynamics model.

6. CONCLUSION

In this paper, we presented the INRIA-Perception contribution to the LOCATA Challenge 2018. We combined i) a recursive DP-RTF feature estimation method, ii) an online multiple-speaker localization method, and iii) an multiple-speaker tracking method. The resulting framework provides online speaker counting, localization and consistent tracking (i.e. preserving speaker identity over a track in spite of intermittent speech production).

7. REFERENCES

- [1] O. Yilmaz and S. Rickard, "Blind separation of speech mixtures via time-frequency masking," *IEEE Transactions on Signal Processing*, vol. 52, no. 7, pp. 1830–1847, 2004.
- [2] M. I. Mandel, R. J. Weiss, and D. P. Ellis, "Model-based expectation-maximization source separation and localization," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 2, pp. 382–394, 2010.
- [3] Y. Dorfan and S. Gannot, "Tree-based recursive expectation-maximization algorithm for localization of acoustic sources," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 10, pp. 1692–1703, 2015.
- [4] O. Schwartz and S. Gannot, "Speaker tracking using recursive EM algorithms," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 2, pp. 392–402, 2014.
- [5] J. Vermaak and A. Blake, "Nonlinear filtering for speaker tracking in noisy and reverberant environments," in *Acoustics, Speech, and Signal Processing, 2001. Proceedings. (ICASSP'01). 2001 IEEE International Conference on*, vol. 5, IEEE, 2001, pp. 3021–3024.
- [6] J.-M. Valin, F. Michaud, and J. Rouat, "Robust localization and tracking of simultaneous moving sound sources using beamforming and particle filtering," *Robotics and Autonomous Systems*, vol. 55, no. 3, pp. 216–228, 2007.
- [7] C. Evers, A. H. Moore, P. A. Naylor, J. Sheaffer, and B. Rafaely, "Bearing-only acoustic tracking of moving speakers for robot audition," in *IEEE International Conference on Digital Signal Processing (DSP)*, 2015, pp. 1206–1210.
- [8] S. Ba, X. Alameda-Pineda, A. Xompero, and R. Horaud, "An on-line variational bayesian model for multi-person tracking from cluttered scenes," *Computer Vision and Image Understanding*, vol. 153, pp. 64–76, 2016.
- [9] I. Gebru, S. Ba, X. Li, and R. Horaud, "Audio-visual speaker diarization based on spatiotemporal Bayesian fusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [10] Y. Ban, L. Girin, X. Alameda-Pineda, and R. Horaud, "Exploiting the complementarity of audio and visual data in multi-speaker tracking," in *ICCV Workshop on Computer Vision for Audio-Visual Media*, vol. 3, 2017.
- [11] Y. Avargel and I. Cohen, "System identification in the short-time Fourier transform domain with crossband filtering," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 4, pp. 1305–1319, 2007.
- [12] R. Talmon, I. Cohen, and S. Gannot, "Relative transfer function identification using convolutive transfer function approximation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 17, no. 4, pp. 546–555, 2009.
- [13] X. Li, L. Girin, R. Horaud, and S. Gannot, "Estimation of the direct-path relative transfer function for supervised sound-source localization," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 24, no. 11, pp. 2171–2186, 2016.
- [14] —, "Multiple-speaker localization based on direct-path features and likelihood maximization with spatial sparsity regularization," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 10, pp. 1997–2012, 2017.
- [15] X. Li, B. Mourgue, L. Girin, S. Gannot, and R. Horaud, "Online localization of multiple moving speakers in reverberant environments," in *The Tenth IEEE Workshop on Sensor Array and Multichannel Signal Processing*, 2018.
- [16] X. Li, Y. Ban, L. Girin, X. Alameda-Pineda, and R. Horaud, "Online Localization and Tracking of Multiple Moving Speakers in Reverberant Environment," Jul. 2018, submitted to *Journal on Selected Topics in Signal Processing*. [Online]. Available: <https://hal.inria.fr/hal-01851985>
- [17] J. Kivinen and M. K. Warmuth, "Exponentiated gradient versus gradient descent for linear predictors," *Information and Computation*, vol. 132, no. 1, pp. 1–63, 1997.
- [18] Y. Ban, X. Alameda-Pineda, F. Badeig, S. Ba, and R. Horaud, "Tracking a varying number of people with a visually-controlled robotic head," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 4144–4151.
- [19] H. W. Löllmann, C. Evers, A. Schmidt, H. Mellmann, H. Barfuss, P. A. Naylor, and W. Kellermann, "The LOCATA challenge data corpus for acoustic source localization and tracking," in *IEEE Sensor Array and Multichannel Signal Processing Workshop (SAM)*, Sheffield, UK, July 2018.
- [20] G. Xu, H. Liu, L. Tong, and T. Kailath, "A least-squares approach to blind channel identification," *IEEE Transactions on signal processing*, vol. 43, no. 12, pp. 2982–2993, 1995.
- [21] X. Li, L. Girin, R. Horaud, and S. Gannot, "Estimation of relative transfer function in the presence of stationary noise based on segmental power spectral density matrix subtraction," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2015, pp. 320–324.
- [22] I. D. Gebru, X. Alameda-Pineda, F. Forbes, and R. Horaud, "Em algorithms for weighted-data clustering with application to audio-visual scene analysis," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 12, pp. 2402–2415, 2016.