

Information-based source localization with distinct binaural cues

Patrick DANÈS⁽¹⁾

⁽¹⁾LAAS-CNRS, Université de Toulouse, CNRS, UPS, Toulouse, France, patrick.danes@laas.fr

Abstract

Audio-motor binaural localization algorithms, which combine directional cues extracted from the sensed signals with the motion of the sensor, are known to overcome shortcomings such as front-back ambiguity and source range non-observability. They can be improved by closing the loop from their output to the control inputs of the sensor, i.e., the sensor motor commands. This paper presents an approach, coined “information-based feedback control”, which drives in real time a binaural head so as to gather information on the location of a static source. On the one hand, a “greedy” approach moves the head to its next best position. On the other hand, a multi-step-ahead scheme determines its most effective path over a receding horizon of size N , by reasoning on average over yet uncollected audio data. Both methods internally entail the prediction of binaural cues, e.g., ITDs, ILDs or a combination of both. Some results can be given an elegant commonsense interpretation.

Keywords: Robot audition, Binaural hearing, Source localization, Audio-motor localization, Information-based feedback control

1 INTRODUCTION

Since the early days of robot audition, sound source localization has been widely acknowledged as the early fundamental stage required to endow robots with Cocktail Party ability in real environments, be this problem tackled from the Computational Auditory Scene Analysis viewpoint or by means of standard signal processing techniques [14]. First approaches relied on binaural sensors; in view of their moderate robustness to changing experimental conditions and of the difficulty to get and exploit accurate models of the scatterers holding the microphones, microphone arrays were later envisaged [1]. However, the potential benefits of so-called “active” approaches, which seamlessly combine into audio-motor schemes the sensed audio stream with the motor commands of the sensor, had long been identified [13]. This gave rise to a renewal of interest for binaural setups, which avoid the engineering cost and burden of multichannel synchronous acquisition. The binaural active audition paradigm also complies with psychology and neurology approaches which consider binaural audition in humans as an exploratory and purposive process, entailing the combination of hearing and motion, as well as multiple feedback loops: e.g., from the cognitive level to low-level binaural processing stages (“shaping the ears”), or from the cognitive level to planned motions (in order to reinforce/weaken hypotheses about the meaning of the scene) [3]. Feedback loops do also exist at the very low reflex level. As a matter of fact, the unconscious interweave of binaural sensing and head motion can well be an information-based sensorimotor feedback which moves the head so as to decrease the spatial uncertainty on the acoustic environment.

This paper outlines an engineering approach to information-based binaural localization with robots. The aim is to augment an audio-motor localization scheme, which computes a belief on the relative position of a sound source with respect to a moving binaural head, with a feedback control that computes in real time the head motion that maximizes, on average, the spatial information on the source. Before entering into details, we overview some neighboring problems addressed in the literature.

The early approaches to binaural audio-motor localization on artificial agents noticed the obvious impact of the head trajectory on the uncertainty associated to the source position estimate. For instance, [11] propose a simulated comparison of various head displacements, which are combined with audio cues inside a particle filter. More recently, still relying on a particle filter, a binaural robot endowed with vision actively moves so as to maximize the decrease of entropies between the current and next beliefs on the hidden source position. In [2], the motions of a binaural head obtained by reinforcement learning on the basis of a frequency analysis of the

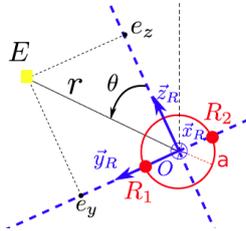


Figure 1. Overview.

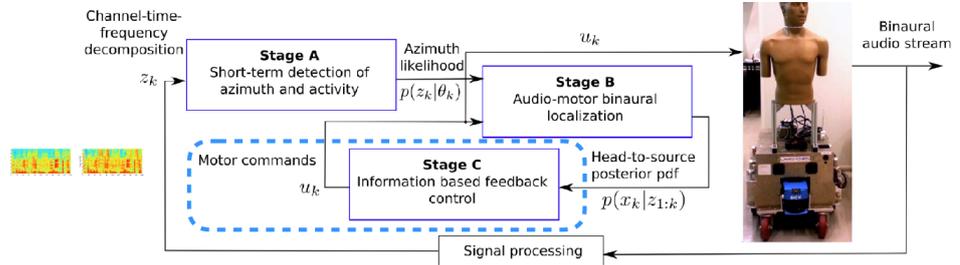


Figure 2. Decomposition of active binaural localization into three stages.

sensed signals show that the expected source should be located in the median plane, or “auditory fovea”. The Monte Carlo approach of [18] drives an agent towards a specific goal while lowering the spatial uncertainty on a sound source by optimizing a criterion which combines the distance to the goal and the entropy of the next source position belief. All these strategies can be termed “greedy” or “myopic”, as their aim is to get the next best head motion. Active motion has also been considered in [8] and [12]. In the context of robotics Simultaneous Localization and Mapping (SLAM) with other sensing modalities, greedy active motion strategies have also been investigated, entailing Shannon entropy, mutual information, or Fisher information matrix, e.g., [4]. Besides, efforts have been made to cope with long-term criteria, that depict the gathered information after the application of a sequence of N motor commands. For instance, [20] designs a dynamic programming based approximate solution to the minimization of the expectation, over the next unknown measurements, of a probabilistic occupation grid. Recently, a Monte Carlo Tree Search solution [15] minimizes the expectation, over the measurements to appear on a receding horizon, of a criterion made up with approximations of the entropies of the source position beliefs over this horizon (themselves expressed as Gaussian sums by means of an extension of the Gaussian sum filter proposed in [16]). Though this complete and convincing method assumes a countable admissible set of control inputs, its major drawback is its computation time, so that it comes within planning rather than feedback control.

The paper is organized as follows. Section 2 provides a mathematical statement of the considered information-based binaural localization problem. It relies on a three-layer framework made up with short-term directional source localization, audio-motor localization, information-based feedback control. The exploration is guided by virtual exploratory audio cues that can be related to the hidden head-to-source direction by means of an analytical equation with constant-variance additive noise. Section 3 sketches the proposed mathematical solution. It is shown that the complexity fundamentally differs depending on whether a greedy (one-step-ahead) or N -step-ahead problem is addressed. Rather than entering into the mathematical details of the solution detailed elsewhere, qualitative insights are given. Section 4 presents some results, in the form of sample trajectories corresponding to the selected exploratory audio cues. Conclusions and prospects end the paper.

2 PROBLEM STATEMENT

2.1 Notations

Consider a binaural head of approximate radius a , endowed with two microphones R_1, R_2 (Figure 1). Its associated frame is $\mathcal{F} = (O, \vec{x}_R, \vec{y}_R, \vec{z}_R)$, with \vec{x}_R pointing downwards, $\vec{y}_R = \frac{R_2 R_1}{\|R_2 R_1\|}$ supporting the interaural axis, and \vec{z}_R pointing frontwards. An motionless omnidirectional sound source E emits continuously. The unknown head-to-source position is characterized by the (hidden) state vector $x = (e_y, e_z)^T \in \mathbb{R}^2$ such that $\vec{OE} = e_y \vec{y}_R + e_z \vec{z}_R$, or, equivalently, by the polar coordinates (r, θ) , with $r = \sqrt{e_y^2 + e_z^2}$ and $\theta = -\text{atan2}(e_y, e_z)$. As time evolves, these

variables are subscripted by the discrete index k .

The measurement vector z_k gathers the binaural information collected by R_1, R_2 over a short temporal snippet ending at time k , typically the the left and right channel-time-frequency decompositions. The measurement model between x and z , though unknown and very involved, is assumed static, i.e., z_k depends solely on x_k .

Between consecutive times k and $k+1$, the head undergoes a rigid motion. Its control input vector, or “motor command”, is defined as $u_k = (T_k^T, \phi_k)^T \in \mathbb{R}^3$, with $T_k = (u_{1k}, u_{2k})^T = (\overrightarrow{O_k O_{k+1} \cdot y_{Rk}}, \overrightarrow{O_k O_{k+1} \cdot z_{Rk}})^T$ the 2D translation and $\phi_k = u_{3k} = \widehat{(\overrightarrow{z_{Rk}}, \overrightarrow{z_{Rk+1}})}_{\text{around } \overrightarrow{x_{Rk}}}$ the finite rotation which, when applied in sequence, turn \mathcal{F}_k into \mathcal{F}_{k+1} . The prior dynamics pdf $p(x_{k+1}|x_k)$ can be equivalently described by the stochastic state equation

$$x_{k+1} = R^T(\phi_k)x_k - R^T(\phi_k)T_k + w_k, \quad w_k \sim \mathcal{N}(0, Q_k) \quad (1)$$

where $R(\phi) = \begin{pmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{pmatrix}$ stands for the nontrivial diagonal block of the 3×3 rotation matrix of angle ϕ around $\overrightarrow{x_R}$, and the additive dynamic noise w_k is assumed white zero-mean Gaussian with known covariance Q_k . Mutual independence of w_k , the random initial state x_0 and of all measurement noises is also assumed.

2.2 A three-layer framework to active binaural localization

Active binaural localization is addressed through the framework outlined on Figure 2. At each time k , the following sequence is completed. Stage A processes the measurement z_k so as to detect the activity of the source and computes a HRTF-based “pseudo-likelihood” $p(z_k|\theta_k)$ of the head-to-source azimuth θ_k , along [17]. Stage B assimilates all these cues up to time k and combines them with the sensor motor commands into a Gaussian sum Unscented Kalman filter, resulting in the state (head-to-source position) filtering pdf (or “belief”) $p(x_k|z_{1:k}) \approx \sum_{j=1}^{J_k} \gamma_k^j \mathcal{N}(x_k; \hat{x}_{k|k}^j, P_{k|k}^j)$ expressed as a Gaussian mixture. Observe that if $Q_k = 0$ in (1), then the current filtering $p(x_k|z_{1:k})$ and next prediction $p(x_{k+1}|z_{1:k})$ pdfs describe the same belief on the source location in the world frame, but expressed from the respective viewpoints of \mathcal{F}_k and \mathcal{F}_{k+1} . This audio-motor localization, described in [16], enables front-back disambiguation as well as faithful source range recovery. This paper addresses the feedback controller featured in Stage C, which processes the state filtering pdf from Stage B so as to move the head in the most informative way.

2.3 Information-based optimum control strategy

The fundamental problem defining Stage C can now be stated.

Multi-step-ahead information-based feedback control. At time k , given the state filtering pdf $p(x_k|z_{1:k}) \approx \sum_{j=1}^{J_k} \gamma_k^j \mathcal{N}(x_k; \hat{x}_{k|k}^j, P_{k|k}^j)$, find the sequence $\bar{u}_N^* = u_k^* : u_{k+N-1}^*$ of control inputs to be applied over a sliding window of size N starting at k (i.e., from time k to $k+N-1$) such that the expected value $\mathbb{E}_{z_{k+1:k+N}|z_{1:k}} h(x_{k+N}|z_{1:k+N})$ of the entropy $h(x_{k+N}|z_{1:k+N})$ of the state filtering pdf $p(x_{k+N}|z_{1:k+N})$ at final time $k+N$ is minimum. \square

This problem makes sense for two reasons. On the one hand, the entropy of a pdf is a convenient measure of the information it conveys. On the other hand, as $z_{k+1:k+N}$ are not yet known at time k , the expectation of the entropy of the state filtering pdf at the end of the receding horizon over these future measurements (conditioned known measurements up to k) is considered. Basically, $N=1$ corresponds to the greedy strategy.

2.4 Simplifying assumptions

Binaural signals to be sensed in the future are hardly predictable, and so is the pdf $p(z_{k+1:k+N}|z_{1:k})$ involved in the criterion to be minimized at time k . This is why the following is assumed to hold.

Fundamental assumption. In Stages A and B of the three-stage framework to active binaural localization (Figure 2), z stands for the measurement variable (channel-time-frequency decomposition) depicted above. However, in the above fundamental problem defining Stage C, z_{k+1}, \dots, z_{k+N} stand for “exploratory audio cues” instead. A variable z is said to be an exploratory audio cue iff it is linked to the hidden state x by a known closed-form measurement

model, the measurement noise of which is additive white zero-mean Gaussian (independent of all other noises) and has a (co)variance unrelated to the hidden state value. In other words, the exploration (computation of \bar{u}_N^*) is driven by kind of virtual measurements

$$z_{k+i} = l(x_{k+i}) + v_{k+i}, \quad i = 1, \dots, N, \quad v_{k+i} \sim \mathcal{N}(0, R_{k+i}) \text{ with } R_{k+i} \text{ independent of hidden } x_{k+i}, \quad (2)$$

and $v_{1:N}$ white independent of $x_0, w_{0:k+N}$, etc. \square

Other less fundamental assumptions are also in effect.

- The state filtering pdf $p(x_k|z_{1:k})$ is reduced to a single Gaussian, so that $p(x_k|z_{1:k}) \approx \mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$.
- Farfield exploratory audio cues are considered, i.e., $l(x_{k+i}) = \bar{l}(\theta_{k+i})$, with $\bar{l}(\cdot)$ given analytically.
- Those cues are assumed scalar, i.e., $z_{k+i} \in \mathbb{R}, R_{k+i} \in \mathbb{R}$.

3 SKETCH OF THE MATHEMATICAL SOLUTION

As aforementioned, the one-step-ahead (or “greedy”) problem is equivalent to $N = 1$. As it will be shown to be significantly easier than the general N -step-ahead problem ($N > 1$), it deserves a specific treatment. Recall that the entropy $h(\mathcal{N}(x; \bar{x}, P))$ of the Gaussian pdf $\mathcal{N}(x; \bar{x}, P)$ is an affine increasing function of $\det(P)$, so that, for the purpose of its optimization it will be replaced (with a slight notation misuse) with $\det(P)$. The higher the value of $\det(P)$, the higher the volume of the 99%-probability confidence ellipsoid of $\mathcal{N}(x; \bar{x}, P)$.

3.1 One-step-ahead problem

In view of the Unscented Kalman filter (UKF) equations [10], when starting from $p(x_k|z_{1:k}) \approx \mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$ the next state filtering pdf comes as $p(x_{k+1}|z_{1:k+1}) \approx \mathcal{N}(x_{k+1}; \hat{x}_{k+1|k+1}, P_{k+1|k+1})$, with $P_{k+1|k+1}$ independent of z_{k+1} . Hence, the expectation $\mathbb{E}_{z_{k+1}|z_{1:k}} h(x_{k+1}|z_{1:k+1})$ of the entropy $h(x_{k+1}|z_{1:k+1})$ of $p(x_{k+1}|z_{1:k+1})$ amounts to $h(x_{k+1}|z_{1:k+1})$. Furthermore, if the dynamic noise covariance Q_k is neglected, then the optimum control is [7, 5]

$$\bar{u}_1^* = (T_{y_k}^*, T_{z_k}^*, \phi_k^*) = \arg \min_{\substack{\bar{u}_1 = u_k = (T_{y_k}, T_{z_k}, \phi_k) \\ ((T_{y_k}, T_{z_k}), \phi_k) \in (\mathcal{T} \times \mathcal{R})^N}} J_1(\bar{u}_1), \quad \text{with } J_1(\bar{u}_1) = h(x_{k+1}|z_{1:k+1}) = K'_1 - \underbrace{h(z_{k+1}|z_{1:k})}_{=: F_1(\bar{u}_1)} \text{ and } K'_1 \text{ independent of } \bar{u}_1, \quad (3)$$

where

- \mathcal{T} and \mathcal{R} term the admissible translation and rotation domains (henceforth, respectively, a disk centered on 0 of given radius for the 2D translation and a 1D line segment centered on 0 for the rotation);
- $h(z_{k+1}|z_{1:k})$ stands for the entropy of the next measurement prediction pdf $p(z_{k+1}|z_{1:k})$; it depends on $\bar{u}_1 = u_k = (T_{y_k}, T_{z_k}, \phi_k)$ and is thus denoted by $F_1(\bar{u}_1)$.

If the Gaussian approximation $p(z_{k+1}|z_{1:k}) \approx \mathcal{N}(z_{k+1}; \hat{z}_{k+1|k}, S_{k+1|k})$ produced by the UKF is used for $p(z_{k+1}|z_{1:k})$, then the information-based one-step-ahead optimum control simplifies as

$$\bar{u}_1^* = (T_{y_k}^*, T_{z_k}^*, \phi_k^*) = \arg \max_{((T_{y_k}, T_{z_k}), \phi_k) \in \mathcal{T} \times \mathcal{R}} F_1(T_{y_k}, T_{z_k}, \phi_k), \quad \text{with } F_1(T_{y_k}, T_{z_k}, \phi_k) = h(z_{k+1}|z_{1:k}) \approx \log \det(S_{k+1|k}). \quad (4)$$

Therein, the criterion $F_1(\bar{u}_1)$ comes as a combination of the image, through the function $l(\cdot)$ entailed in the exploratory audio cues model (2), of sigma-points of the next state filtering pdf $p(x_{k+1}|z_{1:k})$. These sigma-points are themselves functions of the decision variable \bar{u}_1 . Hence, though the genuine $F_1(\bar{u}_1)$ has no closed form, it can easily be approximated. So can its gradient, through Taylor expansion. Consequently, (3)-(4) can be solved by means of the projected gradient algorithm.

3.2 N-step-ahead problem

The corresponding information-based optimum control input sequence is given by [6, 5].

$$\bar{u}_N^* = u_k^* : u_{k+N-1}^* = \arg \min_{\substack{\bar{u}_N = u_k : u_{k+N-1} \\ \bar{u}_N \in (\mathcal{T} \times \mathcal{R})^N}} J_N(\bar{u}_N), \text{ with } J_N(\bar{u}_N) = K'_N - \underbrace{h(z_{k+1}|z_{1:k})}_{=: F_1(\bar{u}_1)} - \sum_{i=2}^N \mathbb{E}_{z_{k+1:k+i-1}|z_{1:k}} \left[\underbrace{h(z_{k+i}|z_{1:k+i-1})}_{=: F_i(\bar{u}_i, z_{k+1:k+i-1})} \right] \quad (5)$$

and K'_N independent of \bar{u}_N , where the entropy $h(z_{k+i}|z_{1:k+i-1})$ of the i -step-ahead measurement prediction pdf $p(z_{k+i}|z_{1:k+i-1})$ depends on the sequence \bar{u}_i of the i next control inputs to be determined and on the $i-1$ future exploratory cues $z_{k+1:k+i-1}$, so that it is denoted by $F_i(\bar{u}_i, z_{k+1:k+i-1})$. Its conditional expectation over these yet unknown virtual measurements must be evaluated, what makes the N -step-ahead problem significantly harder.

A way out is to use the approximation [19]

$$\int p(z_{k+1:k+i-1}|z_{1:k}) F_i(\bar{u}_i, z_{k+1:k+i-1}) dz_{k+1:k+i-1} \approx \sum_{j=0}^{2(i-1)} w_j F_i(\bar{u}_i, Z_{j,i}(\bar{u}_{i-1})), \quad (6)$$

where $\{Z_{j,i}\}$ are the $2(i-1)+1$ sigma-points, deterministically drawn by the Unscented Transform (UT), of the Gaussian approximation $\mathcal{N}(z_{k+1:k+i-1}; \hat{z}_{k+1:k+i-1}|k, S_{k+1:k+i-1}|k)$ of the joint measurement prediction pdf $p(z_{k+1:k+i-1}|z_{1:k})$, and $\{w_j\}$ are the corresponding UT weights. $\mathcal{N}(z_{k+1:k+i-1}; \hat{z}_{k+1:k+i-1}|k, S_{k+1:k+i-1}|k)$ is readily obtained by means of the UT and the exploratory audio cues model (2) on the basis of the Gaussian approximation $\mathcal{N}(x_{k+1:k+i-1}; \hat{x}_{k+1:k+i-1}|k, P_{k+1:k+i-1}|k)$ of the joint state prediction pdf $p(x_{k+1:k+i-1}|z_{1:k})$. $\mathcal{N}(x_{k+1:k+i-1}; \hat{x}_{k+1:k+i-1}|k, P_{k+1:k+i-1}|k)$ itself can be obtained from $\mathcal{N}(x_k; \hat{x}_k|k, P_k|k) \approx p(x_k|z_{1:k})$ by repeated application of the prior dynamics (1) and use of the UT. Hence, $\hat{x}_{k+1:k+i-1}|k, P_{k+1:k+i-1}|k, \hat{z}_{k+1:k+i-1}|k, S_{k+1:k+i-1}|k$ depend on the sequence of decisions variables $\bar{u}_{i-1} = u_k : u_{k+i-2}$. So do the sigma-points $\{Z_{j,i}\}$, which are denoted $\{Z_{j,i}(\bar{u}_{i-1})\}$. In summary, the information-based N -step-ahead optimum control simplifies as

$$\bar{u}_N^* = u_k^* : u_{k+N-1}^* = \arg \max_{\bar{u}_N \in (\mathcal{T} \times \mathcal{R})^N} F_1(\bar{u}_1) + \sum_{i=2}^N \sum_{j=0}^{2(i-1)} w_j F_i(\bar{u}_i, Z_{j,i}(\bar{u}_{i-1})). \quad (7)$$

In spite of the apparent simplicity of (7), it must be kept in mind that for each index i in $\{2, \dots, N\}$, and for each j^{th} sigma-point $Z_{j,i}(\bar{u}_{i-1})$ associated to i , $F_i(\bar{u}_i, Z_{j,i}(\bar{u}_{i-1}))$ is the value $[F_i(\bar{u}_i, z_{k+1:k+i-1})]_{z_{k+1:k+i-1}=Z_{j,i}(\bar{u}_{i-1})}$ of the entropy $F_i(\bar{u}_i, z_{k+1:k+i-1}) = h(z_{k+i}|z_{1:k+i-1})$ of $p(z_{k+i}|z_{1:k+i-1})$, when evaluated at the sequence of virtual measurements $z_{k+1:k+i-1} = Z_{j,i}(\bar{u}_{i-1})$. For each sigma-point $Z_{j,i}$, which implicitly depends on the sequence \bar{u}_{i-1} of decision variables under optimization, a Gaussian approximation $[p(z_{k+i}|z_{1:k+i-1})]_{z_{k+1:k+i-1}=Z_{j,i}(\bar{u}_{i-1})} \approx \mathcal{N}(z_{k+i}; \hat{z}_{(k+i|k+i-1),j}, S_{(k+i|k+i-1),j})$ can be produced by a UKF, leading to $F_i(\bar{u}_i, Z_{j,i}(\bar{u}_{i-1})) \approx \log \det(S_{(k+i|k+i-1),j})$, in the vein of (4). The obtained approximation of the genuine criterion of (5) is quite impossible to write by hand. Nevertheless it can be programmed, and even if Taylor expansions are hardly conceivable, automatic differentiation in the framework of dual numbers algebra [9] is still possible to get its gradient. This enables again a solution by means of the projected gradient algorithm.

3.3 Qualitative insights

Commonsense characterization of one-step-ahead (greedy) control optimality can be established. Figure 3-(a) displays the binaural head at time k with its attached frame, the genuine position of the source E (yellow), as well as the 99%-probability confidence ellipse (grey) corresponding to $p(x_k|z_{1:k}) \approx \mathcal{N}(x_k; \hat{x}_k|k, P_k|k)$. The set of radial lines rigidly linked to the head and intersecting at its center are the loci of the head-to-source positions which correspond to evenly-spaced values of the exploratory measurement at time k . These drawn on the picture typically correspond to the Woodworth-Schlosberg approximation of ITDs: the locus $z = 0$ is supported by \vec{z}_{Rk} ; it splits the 2D plane into two halves onto which z keeps the same sign and evolves monotonically; around \vec{z}_{Rk} is the auditory fovea, for which the sensitivity of z to the head-to-source position is maximum.

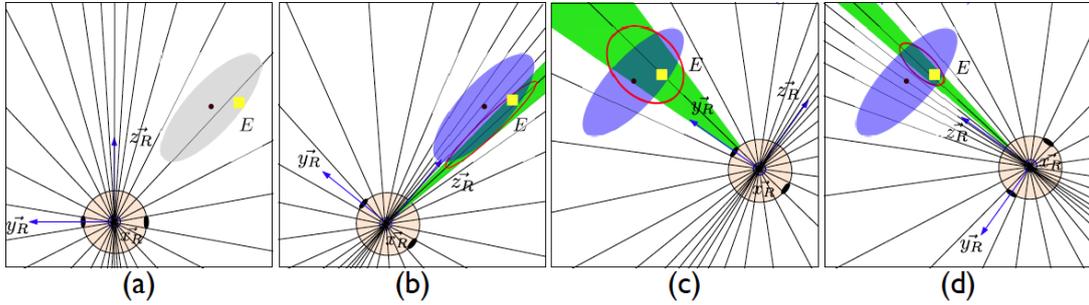


Figure 3. Geometric interpretation of $h(z_{k+1}|z_{1:k})$ for various motions, when using ITD-like exploratory cues.

Let the motor command u_k applied to the head between times k and $k+1$ correspond to a clockwise rotation, and consider the situation at $k+1$ (Figure 3-(b)). The blue 99%-probability confidence ellipse corresponds to $p(x_{k+1}|z_{1:k}) \approx \mathcal{N}(x_{k+1}; \hat{x}_{k+1|k}, P_{k+1|k})$. The green cone represents the 99%-probability confidence set of the pdf $p(x_{k+1}|z_{k+1})$, i.e., the spatial uncertainty sector on the head-to-source position extracted from a single hypothetical measurement z_{k+1} . As the measurement noise (co)variance R_{k+1} is constant, the denser are the iso- z radial lines (which are still rigidly linked to the head) around z_{k+1} , the smaller is the extent of this confidence cone. The red 99%-probability confidence ellipse corresponds to the next state filtering pdf $p(x_{k+1}|z_{1:k+1}) \approx \mathcal{N}(x_{k+1}; \hat{x}_{k+1|k+1}, P_{k+1|k+1})$ obtained by fusing the two last pdfs. As aforementioned, the entropy of $p(x_{k+1}|z_{1:k+1})$ is all the smaller as the volume of this ellipse is low. Heuristically, the entropy of the next measurement prediction pdf $p(z_{k+1}|z_{1:k})$ is all the bigger (resp. smaller) as the number of iso- z radial lines intersected by the blue 99%-probability confidence ellipse of $p(x_{k+1}|z_{1:k})$ is high (resp. low).

By comparing Figures 3-(b)-(c)-(d), the following heuristical rules of thumb can be drawn so that u_k is optimum when using approximate Woodworth-Schlosberg ITDs as exploratory audio cues (virtual measurements): while keeping into its admissible set $\mathcal{F} \times \mathcal{R}$, u_k must

- orient the auditory fovea, rather than the interaural axis, towards the confidence ellipsoid associated to the initial belief $p(x_k|z_{1:k})$;
- drive the head center on the line supporting the minor axis, rather than the major axis, of this ellipsoid;
- make the head get closer to this ellipsoid;

so that after the head motion this ellipsoid intersects as many iso- z radial lines as possible.

4 RESULTS WITH DISTINCT BINAURAL CUES

An experimental setup comprising a KEMAR head-and-torso-simulator mounted on a nonholonomic mobile platform and endowed with a motorized neck has been designed in such a way that its binaural head is omnidirectional. Simulations and live tests were conducted, which implement the Woodworth-Schlosberg approximation of ITDs for a spherical head of same average radius. The simulations can be considered realistic, as they rely on the binaural simulator developed in the framework of the EU FP7 Two!Ears project (<http://docs.twoears.eu/en/latest/binsim/>) which can faithfully emulate the sensed binaural signals. As aforementioned, the Woodworth-Schlosberg ITDs iso- z loci looks like the radial lines sketched on Figure 3-(a). Due to space reasons, only simulated results with such virtual measurements are presented.

Nevertheless, exact farfield ITDs and ILDs for a spherical binaural head have also been implemented in simulation. To this aim, the developments onto Legendre polynomials of the farfield left and right HRTFs as well as their derivatives have been carefully implemented so as to keep within predefined tolerances. The results with farfield spherical ITDs look like these sketched below for their Woodworth-Schlosberg approximation. Farfield ITDs constitute relevant exploratory cues in view of their limited sensitivity to the assumed source frequency and because they constitute a convenient approximation of the genuine range-dependent ITDs for an impor-

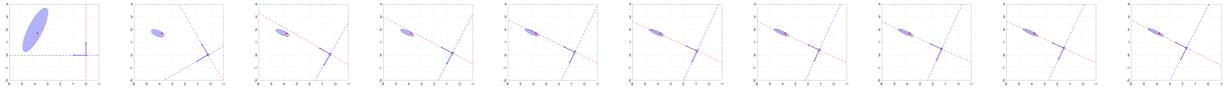


Figure 4. One-step-ahead optimum head motion with Woodworth-Schlosberg ITD exploratory cues (Case 2)

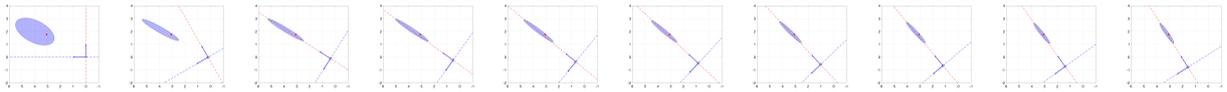


Figure 5. One-step-ahead optimum head motion with Woodworth-Schlosberg ITD exploratory cues

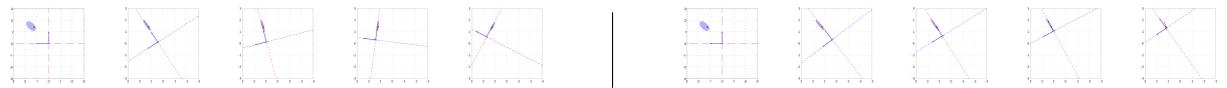


Figure 6. One- (left) vs Five- (right) step-ahead optimum head motion with Woodworth-Schlosberg ITD cues

tant set of head-to-source positions. However, controlling the head motion on the basis of virtual farfield ILD measurements can give rise to bad results if precautions are not taken. Unsurprisingly, iso- z loci of genuine spherical ILDs strongly depend on the assumed source frequency. Moreover, their sensitivity to the source range is such that for a broad interval of source azimuths outside the auditory fovea, they may strongly differ from iso- z radial lines of their farfield limits. Last, farfield ILDs do not vary monotonically on each side of the head, so that they may lead to quite surprising and unefficient motions in some circumstances.

Figures 4 and 5 display the one-step-ahead optimum head motion. It can be seen that the obtained behavior of the head complies with the heuristic rules of thumb sketched in Section 3.3. Notice that Figure 4 corresponds to an advantageous initial belief such that the further fusion with measurements is very efficient.

Figure 6 compares the greedy strategy with the 5-step optimum head motion. It can be seen that the second policy leads to a more “cautious” behavior of the head, in that it tends to first rotate and then follow a straight path. The greedy policy can be transiently better in terms of performance (information criterion), but is less efficient in the long term.

5 CONCLUSION AND PROSPECTS

The information-based feedback for binaural localization outlined in this paper aims at minimizing the expected entropy of the head-to-source belief at the end of a receding N -element horizon. A tractable approximate solution is obtained through the unscented transform, automatic differentiation, and the projected gradient algorithm. With no specific code optimization, it can run in real time on the binaural mobile robot for small N 's. It has been extended to a multi-objective problem similar to [18]. Prospects consist in the incorporation of this stochastic and optimization method with machine learning in order to consider multiple sources in unanechoic/dynamic environments.

ACKNOWLEDGEMENTS

The foundations of this work were developed in collaboration with A. Portello and G. Bustamante during their PhD theses. They are gratefully acknowledged.

REFERENCES

- [1] S. Argentieri, P. Danès, and P. Souères. A survey on sound source localization in robotics: From binaural to array processing methods. *Computer Speech and Language*, 34(1), 2015.

- [2] E. Berglund and J. Sitte. Sound source localisation through active audition. *IEEE/RSJ IROS'2005*, Edmonton, Canada.
- [3] J. Blauert. Reading the world with two ears. *ICSV'2017*, London.
- [4] F. Bourgault, A. Makarenko, S. Williams, B. Grocholsky, and H. Durrant-Whyte. Information based adaptive robotic exploration. *IEEE/RSJ IROS'2002*, Lausanne, Switzerland.
- [5] G. Bustamante. *Active Motion for Binaural Localization of Sound Source in Robotics*. PhD, Toulouse, 2017.
- [6] G. Bustamante and P. Danès. Multi-step-ahead information-based feedback control for active binaural localization. *IEEE/RSJ IROS'2017*, Vancouver, Canada.
- [7] G. Bustamante, P. Danès, F. Fergue, A. Podlubne, and J. Manhès. An information based feedback control for audio-motor binaural localization. *Autonomous Robots*, 42(2):477–490, 2018.
- [8] C. Evers, A.H. Moore and P.A. Naylor. Towards informative path planning for acoustic SLAM. *DAGA'2016*, Aachen, Germany.
- [9] J. E. Horwedel. GRESS, a preprocessor for sensitivity studies of Fortran programs. In *Automatic Differentiation of Algorithms: Theory, Implementation, and Application*. SIAM, Philadelphia, PA, 1991.
- [10] S. J. Julier and J. K. Uhlmann. Unscented filtering and nonlinear estimation. *Proc. IEEE*, 92(3), 2004.
- [11] Y. Lu and M. Cooke. Motion strategies for binaural localisation of speech sources in azimuth and distance by artificial listeners. *Speech Communication*, 2010.
- [12] N. Ma, T. May, and G. Brown. Exploiting deep neural networks and head movements for robust binaural localization of multiple sources in reverberant environments. *IEEE/ACM Trans. on Audio, Speech and Language Processing*, 25(12), 2017.
- [13] K. Nakadai, T. Lourens, H. Okuno, and H. Kitano. Active audition for humanoid. *AAAI'2000*, Austin, TX.
- [14] K. Nakadai and K. Nakamura. Sound source localization and separation. *Wiley Encyclopedia of Electrical and Electronics Engineering*.
- [15] Q. V. Nguyen, F. Colas, E. Vincent, and C. F. Long-term robot motion planning for active sound source localization with monte carlo tree search. *HSCMA'2017*, San Francisco, CA.
- [16] A. Portello, G. Bustamante, P. Danès, J. Piat, and J. Manhès. Active localization of an intermittent sound source from a moving binaural sensor. *Forum Acusticum (FA'2014)*, Krakow, Poland.
- [17] A. Portello, P. Danès, S. Argentieri, and S. Pledel. HRTF-based source azimuth estimation and activity detection from a binaural sensor. *IEEE/RSJ IROS'2013*, Tokyo, Japan.
- [18] C. Schymura, J. Diego, R. Grajales, and D. Kolossa. Monte Carlo exploration for active binaural localization. *IEEE ICASSP'2017*, New Orleans, LA.
- [19] S. Särkkä and J. Hartikainen and L. Svensson and F. Sandblom. Gaussian process quadratures in nonlinear sigma-point filtering and smoothing. *IEEE FUSION'2014*.
- [20] E. Vincent, A. Sini, and F. Charpillat. Audio source localization by optimal control of a mobile robot. *IEEE ICASSP'2015*, Brisbane, Australia.