An Information Based Feedback Control for Audio-Motor Binaural Localization

Gabriel Bustamante, Patrick Danès, Thomas Forgue, Ariel Podlubne, Jérôme Manhes

An Information Based Feedback Control for Audio-Motor Binaural Localization

Gabriel Bustamante · Patrick Danès · Thomas Forgue ·
Ariel Podlubne · Jérôme Manhès

Abstract In static scenarios, binaural sound localization is fundamentally limited by front-back ambiguity and distance non-observability. Over the past few years, "active" schemes have been shown to overcome these shortcomings, by combining spatial binaural cues with the motor commands of the sensor. In this context, given a Gaussian prior on the relative position to a source, this paper determines an admissible motion of a binaural head which leads, on average, to the one-step-ahead most informative audio-motor localization. To this aim, a constrained optimization problem is set up, which consists in maximizing the entropy of the next predicted measurement probability density function over a cylindric admissible set. The method is appraised through geometrical arguments, and validated in simulations and on real-life robotic experiments.

Keywords Robot audition · Binaural audition · Active localization · Information theory · Information based control

1 Introduction

The advent of auditory robots has led to the emergence of binaural audio-motor localization schemes which, by combining binaural perception and motor commands, can disambiguate front from back and recover source range (Cooke et al 2007; Nakadai et al 2000). Some of these can cope with a moving and intermittent sound

The authors are with LAAS-CNRS, Université de Toulouse, CNRS, UPS, Toulouse, France.
Tel: +33 (0)5 61 33 78 25, Fax: +33 (0)5 61 33 64 55,
E-mail: patrick.danes@laas.fr.

source (Portello et al 2012). However, the question remains how to drive a binaural head so as to maximize the spatial information on a source extracted from the sensorimotor flow.

In Robotics, Simultaneous Localization and Mapping (SLAM) techniques have been extended to make robots move in order to improve their knowledge about the environment (Thrun et al 2005). Control policies could be found by maximizing information criteria related to the robot situation, *e.g.*, by determining the direction of maximum local information improvement. Shannon entropy or mutual information have often been used (Bourgault et al 2002), as well as the Fisher information matrix (FIM) (Feder et al 1999). It has been shown that a mapping robot guided by a mutual information based controller can be "attracted" towards unexplored areas (Julian 2013). Similar strategies have been used to coordinate multiple sensor platforms (Grocholsky et al 2003). Information-theoretic controllers can address different objectives such as the control of a robot-mounted camera to optimize depth estimation (Forster et al 2014), or the selection of sensor parameters (*e.g.*, zoom or attitude) for scene analysis (Denzler and Brown 2002; Sommerlade and Reid 2008).

In the bearings-only tracking problem, optimum observer actions can be determined by maximizing a cost functional involving FIM determinants (Le Cadre and Laurent-Michel 1999). When the problem is the reduction of the mean square tracking error, the minimization of the posterior Cramér-Rao lower bound—*i.e.*, the inverse of the Bayesian extension of the FIM—has been addressed (Ristic and Arulampalam 2003).

In robot audition, the problem of auditory scenes exploration has also been investigated (Martinson and Schultz 2009). A mobile robot has been equipped with a microphone array to localize
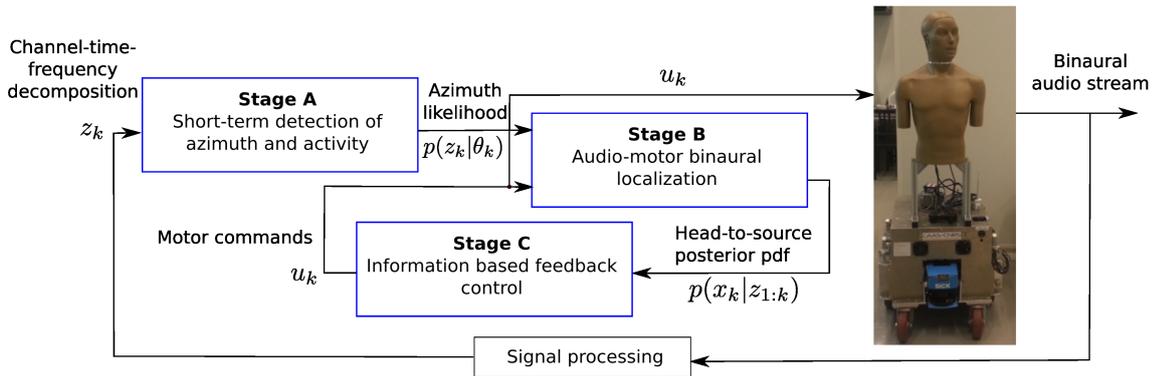
Fig. 1: The three-stage framework to active binaural localization. This paper addresses Stage C.

sound sources and estimate its own position in a known geometric map (Sasaki et al 2010). Motion planning based on audio situation has been proposed to improve speech recognition by a monaural robot (Kumon et al 2010). In (Martinson et al 2011), sound source localization was improved by optimizing the position of microphones deployed in the environment. Recently, a robot equipped with a microphone array was controlled to locate a sound source by minimizing a criterion based on the entropy of an occupancy grid used to represent the source position belief (Vincent et al 2015).

Given a prior knowledge on the relative position of a static sound source with respect to a binaural head, this paper deals with the determination of an admissible finite motion of the sensor which leads, on average, to the minimum uncertainty in the one-step-ahead localization. It is organized as follows. First, the three-stage approach to binaural active localization (Bustamante et al 2015) which has motivated this work is recalled (Section 2). Then, a constrained optimization problem is defined, so as to get the next best position of the sensor (Section 3). A numerical solution scheme is proposed. Further, useful insights into the geometry of the problem are provided (Section 4), when the exploration is guided by directional cues such as the interaural time difference (ITD) between two microphones placed antipodally on a spherical binaural head. Evaluations are then conducted in simulation and on a binaural robotic platform (Section 5). Therein, a comparison is made with some open-loop motion policies. Conclusions and prospects end the paper.

2 A three-stage framework to active binaural localization

This work took place within the EU FET Two!Ears project (www.twoears.eu) whose aim was to develop a

computational model of auditory perception and experience in humans. Listeners are regarded as multi-modal agents that develop their concept of the world by active, exploratory, interaction, and, in the course of this process, interpret percepts, collect knowledge and develop concepts accordingly. To enable this, the Two!Ears model includes not only bottom-up—signal-driven—processing but also top-down—hypothesis-driven—feedbacks. Some of these feedbacks come from the cognitive level, *e.g.*, the context-dependent adjustment of bottom-up processing parameters, or the hypothesis-driven activation of specific low-level processing procedures. Other feedbacks operate at the sensorimotor level—with no cognition in between—at much shorter time scales, *e.g.*, "turn-to-reflex" exploratory movements to dispel localization ambiguities.

Such sensorimotor feedbacks for single-source active binaural localization can be adressed through the three-stage framework depicted on Figure 1. Stage A implements the maximum likelihood estimation of the source azimuth and the information-theoretic detection of its activity from the short-term channel-time-frequency decomposition of the binaural stream (Portello et al 2013). Stage B assimilates these azimuths over time and combines them with the motor commands into a stochastic filter, leading to the posterior probability density function (or "belief") of the head-to-source relative position (Portello et al 2014b). Stage C is the topic of this paper. It consists in a feedback controller which, on the basis of the output from Stage B, can move the head so as to improve the quality of the localization. Stage A has been extended to the multiple-source case (Portello et al 2014a), and Stage B can cope with a moving and/or intermittent source (Portello et al 2012), but this is not considered here.

First an overview of Stages A and B is proposed. Then, the paper focuses on Stage C.
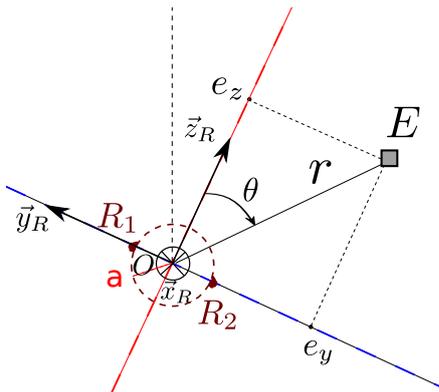
Fig. 2: Sketch of the planar problem. The dashed circle depicts the binaural head, with R_1 and R_2 its left and right microphones, and a its radius. The sound source E is located on the plane $(O, \overrightarrow{y_R}, \overrightarrow{z_R})$ at the distance r and azimuth θ. By convention, θ is zero along along $\overrightarrow{z_R}$ and increases clockwise (so, $\theta > 0$ here). In this plot and all subsequent plots depicting the motion of the binaural head, the front and interaural axes are colored in red and blue, respectively.

2.1 Terminology

A binaural head is fitted with the left and right microphones R_1 and R_2. A frame $\mathcal{F} = (O, \overrightarrow{x_R}, \overrightarrow{y_R}, \overrightarrow{z_R})$ is attached to it, with $\overrightarrow{R_1 O} = \overrightarrow{O R_2}$ (Figure 2). R_1, R_2 and the pointwise emitter E lie on a common horizontal plane defined by $(O, \overrightarrow{y_R}, \overrightarrow{z_R})$, where $\overrightarrow{y_R} = \frac{\overrightarrow{R_2 R_1}}{\|\overrightarrow{R_2 R_1}\|}$ supports the interaural axis and $\overrightarrow{z_R}$ is oriented towards the front direction. So, $\overrightarrow{x_R}$ is vertical and points downwards. a terms the radius of a sphere approximating the head.

Throughout the paper, geometric vectors are denoted with arrows. Scalar, vector or matrix variables are written in normal font. Whether they are deterministic or stochastic can be straightly inferred.

2.2 Stage A: Short-term extraction of directional cues

The interaural transfer function is assumed known over an adequate range of source azimuths and frequencies. The source signal and sensor noises are modeled as jointly Gaussian, zero-mean, individually and jointly "locally stationary" random processes (Portello et al 2013). Then, on the basis of the channel-time-frequency decomposition z_k of the binaural signal on a sliding window ending at time k, the short-term maximum likelihood $\hat{\theta}_k$ of the source azimuth θ_k comes as the argmax of a "pseudo likelihood" $p(z_k | \theta_k)$. This

pseudo likelihood is obtained by replacing in the genuine likelihood of the unknown variables the most likely spectral parameters of the source as a function of its azimuth, by means of a notable separation property.

2.3 Stage B: Combination with motor commands

A discrete-time stochastic state space equation is set up, uniting the motor commands to the head-to-source position $x_k = (e_y, e_z)^T$ to be estimated (Figure 2). A theoretically sound Gaussian mixture square-root unscented Kalman filter (GMsrUKF) is defined so as to incorporate the above pseudo likelihood $p(z_k | \theta_k)$, where θ_k comes as a static function of x_k, and compute a Gaussian mixture approximation of the posterior probability density function (pdf), or "belief",

$$p(x_k | z_{1:k}) = \sum_{i=1}^{I_k} w_k^i \mathcal{N}(x_k; \hat{x}_{k|k}^i, P_{k|k}^i), \tag{1}$$

where $(w_k^i, \hat{x}_{k|k}^i, P_{k|k}^i)$ are the weight, mean and covariance of each hypothesis (Portello et al 2014b). Empirical tests show that self-initialization as well as posterior covariance consistency are generally ensured, so that front and back are disambiguated, and both range and azimuth are faithfully recovered.

2.4 Stage C: Problem statement

Let $\mathcal{F}_k = (O_k, \overrightarrow{x_{Rk}}, \overrightarrow{y_{Rk}}, \overrightarrow{z_{Rk}})$ and $X_k = (e_x, e_y, e_z)^T$ be the frame $\mathcal{F}$ at time k and the Cartesian coordinates of the—static—source in $\mathcal{F}_k$. If between times k and $k+1$ the sensor undergoes the translation $T_y \overrightarrow{y_{Rk}} + T_z \overrightarrow{z_{Rk}}$ followed by the rotation of angle $\phi \triangleq (\widehat{\overrightarrow{z_{Rk}}, \overrightarrow{z_{Rk+1}}})$ around $\overrightarrow{x_{Rk}}$, then the vector X_{k+1} of the source coordinates in $\mathcal{F}_{k+1} = (O_{k+1}, \overrightarrow{x_{Rk+1}}, \overrightarrow{y_{Rk+1}}, \overrightarrow{z_{Rk+1}})$ writes as

$$X_{k+1} = R^T(\phi) X_k - R^T(\phi) T + w_k, \tag{2}$$

with $T = (0, T_y, T_z)^T$, $R(\phi)$ the rotation matrix corresponding to ϕ, and w_k the dynamic noise (if present).

A belief on the sensor-to-source position $x_k = (e_y, e_z)^T$ at time k is given in terms of the 2D Gaussian pdf $\mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$, with $\hat{x}_{k|k}$ the estimate of x_k and $P_{k|k}$ the associated error covariance matrix. The problem consists in determining the motion (T, ϕ) of the sensor which best improves, on average, the next localization of the sound source. First, a metric is described, uniting the belief on the state at time k and the rigid motion applied over $[k; k+1]$, to the expected information obtained after a measurement update at time $k+1$. The exploration

is assumed to be guided by a scalar closed-form observation model such as

$$z_k = l(x_k) + v_k = \bar{l}(\theta_k) + v_k, \ z_k \in \mathbb{R}, \ v_k \sim \mathcal{N}(0, R_k), \quad (3)$$

with v_k the measurement noise and R_k its (co)variance. In the above, z_k is assumed to be a directional cue, in that it solely depends on the source relative azimuth $\theta_k = -\mathrm{atan2}(e_y, e_z)$. Note that by convention, θ_k is 0 along $\overrightarrow{z_R}$ and increases clockwise. Assuming a farfield sound source, $\bar{l}(\theta_k)$ in (3) could express the time delay of arrival (TDOA) between two microphones in free field. Unless otherwise stated, in the sequel $\bar{l}(\theta_k)$ stands for the Woodworth-Schlosberg farfield approximation of the ITD between two antipodal microphones placed on a spherical head. That is, (Aaronson and Hartmann 2014)

$$\bar{l}(\theta_k) = \frac{a}{c}(\theta_k + \sin(\theta_k)) \text{ for } |\theta_k| \in [0, \frac{\pi}{2}],$$
$$\bar{l}(\theta_k) = \frac{a}{c}(\pi - \theta_k + \sin(\theta_k)) \text{ for } \theta_k \in [\frac{\pi}{2}, \pi],$$
$$\bar{l}(\theta_k) = \frac{a}{c}(-\pi - \theta_k + \sin(\theta_k)) \text{ for } \theta_k \in [-\pi, -\frac{\pi}{2}], \quad (4)$$

with c the velocity of sound.

3 Feedback control of the binaural sensor

On the basis of Stages A and B introduced above, the main topic of the paper is now adressed, namely, the development of Stage C. The information based feedback control is first stated, then turned into a constrained optimization problem. A geometric interpretation is discussed. A numerical solution is obtained by means of the projected gradient algorithm.

3.1 Information-theoretic constrained optimization

Let x, y be continuous random variables with joint and marginal pdfs $p(x, y)$ and $p(x), p(y)$. The differential entropy

$$h(x) = -\int p(x) \log p(x) \mathrm{d}x \quad (5)$$

and the mutual information (nonnegative by definition)

$$I(x, y) = \int p(x, y) \log \frac{p(x, y)}{p(x)p(y)} \mathrm{d}x \mathrm{d}y \quad (6)$$

respectively embody the uncertainty in x and measure the amount of information that x contains about y (Cover and Thomas 1991).

When conditioned on the event that a random variable z takes a given value, they will henceforth be denoted by $h(x|z)$, $h(y|z)$ and $I(x, y|z)$. The Bayes rule

underlying the measurement update stage relates the next filtered state pdf $p(x_{k+1}|z_{1:k+1})$, the next predicted state pdf $p(x_{k+1}|z_{1:k})$, the observation model $p(z_{k+1}|x_{k+1})$ and the next predicted measurement pdf $p(z_{k+1}|z_{1:k})$. Consequently, entropies and mutual information of these distributions can be connected with an entropy update rule of the same kind as (Manyika 1993). The expectation

$$\int -\log p(x_{k+1}|z_{1:k+1}) p(x_{k+1}, z_{k+1}|z_{1:k}) \mathrm{d}x_{k+1} \mathrm{d}z_{k+1}, \quad (7)$$

of $-\log p(x_{k+1}|z_{1:k+1})$ conditioned on $z_{1:k}$, which is also equal to $\mathsf{E}_{z_{k+1}|z_{1:k}}\{h(x_{k+1}|z_{1:k+1})\}$, satisfies (Bustamante et al 2016)

$$\mathsf{E}_{z_{k+1}|z_{1:k}}\{h(x_{k+1}|z_{1:k+1})\} = h(x_{k+1}|z_{1:k}) - I,$$
$$\mathsf{E}_{x_{k+1}|z_{1:k}}\{h(z_{k+1}|x_{k+1})\} = h(z_{k+1}|z_{1:k}) - I, \quad (8)$$
$$I = I(x_{k+1}, z_{k+1}|z_{1:k}),$$

with I the conditional mutual information of the next state and measurement. Due to the nonnegativity of I, $\mathsf{E}_{z_{k+1}|z_{1:k}}\{h(x_{k+1}|z_{1:k+1})\} \leq h(x_{k+1}|z_{1:k})$ holds, which highlights the information gain brought by the measurement update.

Between times k and $k+1$, (linear) Kalman time update equations turn the Gaussian belief $p(x_k|z_{1:k}) = \mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$ into the next predicted state pdf $p(x_{k+1}|z_{1:k}) = \mathcal{N}(x_{k+1}; \hat{x}_{k+1|k}, P_{k+1|k})$. Then, (nonlinear) Kalman measurement update equations incorporate the measurement z_{k+1} so as to compute a Gaussian approximation $p(x_{k+1}|z_{1:k+1}) \approx \mathcal{N}(x_{k+1}; \hat{x}_{k+1|k+1}, P_{k+1|k+1})$ of the next filtered state pdf. They involve a Gaussian approximation $p(z_{k+1}|z_{1:k}) \approx \mathcal{N}(z_{k+1}; \hat{z}_{k+1|k}, S_{k+1|k})$ of the next predicted measurement pdf. Let $|.|$ term the determinant of a matrix. If w_k is negligible in (2), then $h(x_{k+1}|z_{1:k}) = \frac{1}{2}\log[(2\pi e)^{n_x}|P_{k+1|k}|]$ is also equal to $\frac{1}{2}\log[(2\pi e)^{n_x}|P_{k|k}|]$ in view of the fact that the sensor undergoes a rigid motion[1], with $n_x = 2$. In addition, $h(z_{k+1}|x_{k+1}) = \frac{1}{2}\log[(2\pi e)^{n_z}|R_{k+1}|]$, with $n_z = 1$, is also independent of the control variables (T, ϕ). Besides, both $h(x_{k+1}|z_{1:k+1}) = \frac{1}{2}\log[(2\pi e)^{n_x}|P_{k+1|k+1}|]$ and $h(z_{k+1}|z_{1:k}) = \frac{1}{2}\log[(2\pi e)^{n_z}|S_{k+1|k}|]$ do not depend on the measurement z_{k+1}. Consequently, the following rule is in effect.

---

[1] Consider again the dynamic equation (2) with no dynamic noise, and assume that the posterior covariance $\overline{P}_{k|k}$ of the full state X_k (defined in $\mathbb{R}^3$) is $\overline{P}_{k|k} = \mathrm{diag}(0, P_{k|k})$. As the vector $R^T(\phi)T$ is constant, the next "full" predicted covariance $\overline{P}_{k+1|k}$ writes as $\overline{P}_{k+1|k} = R^T(\phi)\overline{P}_{k|k}R(\phi)$, with $R(\phi) = \mathrm{diag}(1, r(\phi))$, and $|R(\phi)| = |r(\phi)| = 1$. Consequently, $\overline{P}_{k+1|k} = \mathrm{diag}(0, P_{k+1|k})$ with $|P_{k+1|k}| = |r^T(\phi)P_{k|k}r(\phi)| = |P_{k|k}|$.

Theorem 1 *Finding the next best sensor position which minimizes the entropy $h(x_{k+1}|z_{1:k+1})$ of the next filtered state pdf—which is also its expected value w.r.t. z_{k+1}—is equivalent to maximizing the mutual information $I(x_{k+1}, z_{k+1}|z_{1:k})$ of the next predicted state and measurement, or to maximizing the entropy $h(z_{k+1}|z_{1:k})$ of the next predicted measurement pdf. In other words, the optimum rigid body motion (T^*, ϕ^*) to be applied to the sensor is the solution of*

$$(\mathcal{P}) \begin{cases} (T^*, \phi^*) = \arg \min_{(T,\phi)\in\mathcal{T}\times\mathcal{R}} h(x_{k+1}|z_{1:k+1}) \\ \qquad = \arg \max_{(T,\phi)\in\mathcal{T}\times\mathcal{R}} I(x_{k+1}, z_{k+1}|z_{1:k}) \\ \qquad = \arg \max_{(T,\phi)\in\mathcal{T}\times\mathcal{R}} h(z_{k+1}|z_{1:k}), \end{cases}$$

$$(9)$$

where $\mathcal{T}$ and $\mathcal{R}$ respectively term the sets of admissible translations and rotations.

3.2 Interpretation

A fundamental feature of the observation model (3) is that the spreading of the measurement noise—*i.e.*, the (co)variance R_k—is assumed known and constant independently of the value of the hidden state vector. If the exploration is guided by TDOAs/ITDs, then this assumption is valid since the standard deviation of the noise associated to their extraction is typically a fraction of the audio sampling period.

From (3), the loci of the sensor-to-source positions x corresponding to given values of the measurement z in the absence of noise—or "iso-z loci"—are radial lines rigidly linked to frame $\mathcal{F}$ and passing through O. For TDOA/ITD measurements, because of the nonlinearity of the measurement equation—see for instance (4)—these lines are not uniformly distributed along the azimuths. They are more concentrated along the direction of $\overrightarrow{z_R}$ which defines the auditive fovea, while they are sparser around the interaural axis $\overrightarrow{y_R}$. Given a belief $\mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$ on the head-to-source position at time k, Figure 3 sketches the 2D Gaussian approximation of the next filtered state pdf $\mathcal{N}(x_{k+1}; \hat{x}_{k+1|k+1}, P_{k+1|k+1})$ after applying various rigid motions (T, ϕ) to the sensor. All the involved normal distributions are depicted by related 99%-probability confidence ellipses. Importantly, if the dynamic noise is neglected in (2), then the next predicted state pdf $\mathcal{N}(x_{k+1}; \hat{x}_{k+1|k}, P_{k+1|k})$ is basically described by the same ellipse as for the initial belief, but "viewed" from the sensor once it has completed its motion. Besides, (3) implies that the pdf of the head-to-source position deduced from the sole measurement
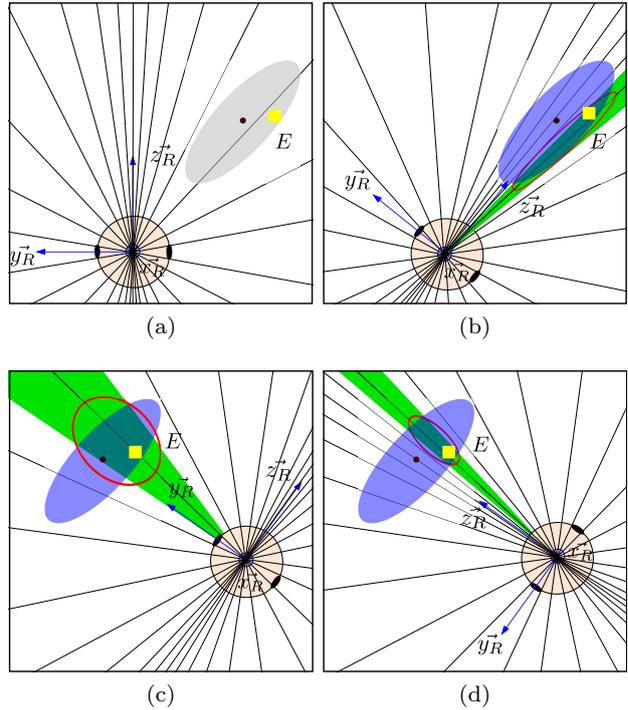

Fig. 3: Iso-z loci and measurement update for various scenarios. (a): Frame $\mathcal{F}_k$ attached to the binaural head (blue); sound source genuine position (yellow square); confidence ellipse associated to the belief at time k (grey); iso-z_k loci depicting the measurement space (grey radial lines). (b)-(c)-(d): Frame $\mathcal{F}_{k+1}$ (blue); confidence ellipse associated to the next predicted state pdf at time $k+1$ (blue); iso-z_{k+1} loci (grey); confidence cone associated to the measurement (green); confidence ellipse associated to the next filtered state pdf (belief at $k+1$) after the incorporation of z_{k+1} (red).

z_{k+1} can be described by a 99%-probability confidence cone tapering to the apex O_{k+1}. For a given variance R_k of the measurement noise, the extent of this cone on each side of the iso-z locus corresponding to the genuine azimuth of the source is all the more important as the iso-z loci are sparse. The measurement update fuses these two last pdfs so as to get the next belief $\mathcal{N}(x_{k+1}; \hat{x}_{k+1|k+1}, P_{k+1|k+1})$. Qualitatively, the fusion is all the more efficient as the overlap of the respective confidence ellipse and cone occurs around the modes of the pdfs and has a limited spatial extent.

From the initial configuration depicted in Figure 3a, the head first undergoes a pure rotational motion so that the auditive fovea (supported by $\overrightarrow{z_{Rk+1}}$) becomes oriented towards the major axis of the confidence ellipse associated to the next predicted state pdf (Figure 3b). On Figure 3c, a translation is applied so as to drive

O_{k+1} on the line supported by the minor axis of that ellipse, and a subsequent rotation makes $\overrightarrow{y_{Rk+1}}$ point towards its center. Last, in Figure 3d, the auditive fovea $\overrightarrow{z_{Rk+1}}$ is driven towards the minor axis of that ellipse.

Equation (9) in Theorem 1 states that the next best sensor position must maximize the (determinant of) the (co)variance $S_{k+1|k}$ of the next predicted measurement pdf $p(z_{k+1}|z_{1:k})$. Here, the scalar value of $S_{k+1|k}$ is heuristically related to the number of iso-z_{k+1} loci intersecting the confidence ellipse associated to the next predicted state pdf. The more iso-z_{k+1} loci intersect that ellipse, the higher is $S_{k+1|k}$.

As aforementioned, the confidence cone describing the spatial uncertainty on the head-to-source position due to the noisy measurement is wide if the source lies along the interaural axis (Figure 3c). In this case, a small number of iso-z loci intersect the confidence ellipse associated to the predicted state pdf, so that the measurement update cannot significantly improve the information in the next filtered state pdf. When the auditive fovea is oriented towards the confidence ellipse, the confidence cone is narrower, so the measurement update is more efficient (Figures 3b and 3d). The variance $S_{k+1|k}$ is also higher than in the above case. Further, if the fovea points to the minor axis of the confidence ellipse, then the measurement update is improved (Figure 3d).

Importantly, the closer the sensor gets to the source, the smaller is the spatial uncertainty on the head-to-source position given a TDOA/ITD measurement. Then, a greater number of iso-z loci cross the confidence ellipse associated to the predicted state pdf, so that the predicted measurement variance $S_{k+1|k}$ increases, what is beneficial.

3.3 Numerical solution

In view of the above, starting from the head-to-source position belief $\mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$ at time k, the desired optimum finite translations T_y^*, T_z^* and rotation ϕ^* maximize the log-determinant—as $z \in \mathbb{R}$, just the log— of the (co)variance the next predicted measurement pdf $p(z_{k+1}|z_{1:k})$, i.e., maximize $F_k(T_y, T_z, \phi) = \log S_{k+1|k}$ with $F_k : \mathbb{R}^3 \to \mathbb{R}$. Then, the optimization problem $(\mathcal{P})$ defined in (9) can be stated as

$$(\mathcal{P}) \begin{cases} (T_y^*, T_z^*, \phi^*) = \arg\max_{(T_y, T_z, \phi) \in (\mathcal{T} \times \mathcal{R})} F_k(T_y, T_z, \phi) \end{cases} \quad (10)$$

with $\mathcal{T} = \{(T_y, T_z) \in \mathbb{R} \times \mathbb{R} \mid T_y^2 + T_z^2 \le r_{max}^2\}$ and $\mathcal{R} = \{\phi \in \mathbb{R} \mid |\phi| \le \phi_{max}\}$ the sets of admissible translations and rotations. $\mathcal{T} \times \mathcal{R}$ thus constitutes a cylinder
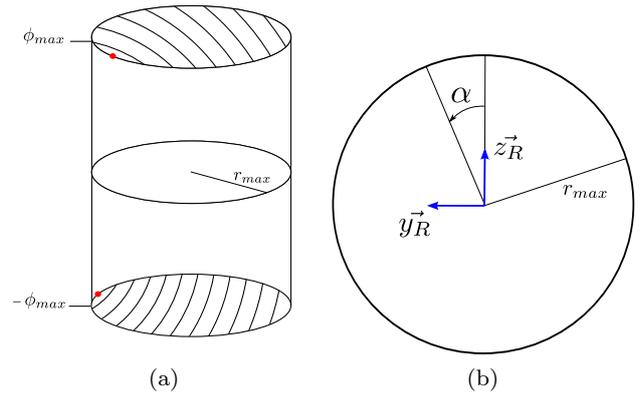

Fig. 4: Representation of the admissible cylindrical set $\mathcal{T} \times \mathcal{R}$ of the problem $(\mathcal{P})$. The contour lines of the criterion $F_k(T_y, T_z, \phi)$ are sketched as functions of T_y, T_z when ϕ takes a constant value corresponding to the bottom ($\phi = -\phi_{max}$) or top ($\phi = \phi_{max}$) side of $\mathcal{T} \times \mathcal{R}$, respectively. The red spots depict the optima of F_k restricted to these sides.

volume (Figure 4a). The height of the cylinder represents the admissible rotations while horizontal sections stand for the feasible translations given a fixed rotation.

Though F_k has no closed form, an approximation of its gradient around a defined translation and rotation $U = (T_y, T_z, \phi)^T$ can be derived by means of successive first order Taylor expansions and the Unscented transform (Julier and Uhlmann 2004). This approximation writes as

$$F_k(U + du) = F_k(U) + \nabla F_k(U)^T du, \quad (11)$$

with $du = (dT_y, dT_z, d\phi)^T$ the infinitesimal motion vector applied around U and $\nabla F_k(U)$ the gradient of F_k evaluated at U, which points to the direction of steepest ascent of F_k around U. A derivation of $\nabla F_k(U)$ is proposed in Appendix 7.

The projected gradient algorithm is then used to solve $(\mathcal{P})$ numerically. It consists in iteratively updating the value of the decision variable $U = (T_y, T_z, \phi)^T$ obtained through the conventional gradient ascent method by projecting it onto the closed convex set $\mathcal{T} \times \mathcal{R}$ by means of the projection operator $\pi_{\mathcal{T} \times \mathcal{R}}(.)$ defined as

$$\pi_{\mathcal{T} \times \mathcal{R}}(U) \triangleq \arg\min_x \{\|U - x\|_2, x \in (\mathcal{T} \times \mathcal{R})\}. \quad (12)$$

This leads to Algorithm 1.

4 Geometrical insights

In this section, the geometry of the maximization problem $(\mathcal{P})$ is depicted.

Algorithm 1: Simplified Projected Gradient

Data:
- Moments of the initial belief at time k: $\hat{x}_{k|k}, P_{k|k}$
- Maximum admissible translation of the head: r_{max}
- Maximum admissible rotation of the head: ϕ_{max}
- Step size: γ
- Number of iterations: M
- Projection operator onto $\mathcal{T} \times \mathcal{R}$: $\pi_{\mathcal{T} \times \mathcal{R}}(T_y, T_z, \phi)$

Outputs: $U^* = (T_y^*, T_z^*, \phi^*)^T \triangleq U_M = (T_{yM}, T_{zM}, \phi_M)^T$

Initialization

1 $U_0 = [T_{y0}, T_{z0}, \phi_0]^T$;

   for $i = 0, \ldots, M-1$ **do**
2 | evaluate $d_i = \nabla F_k(U_i)$, where F_k is defined on the basis of the initial belief at time k;
3 | set $U_{i+1} = \pi_{\mathcal{T} \times \mathcal{R}}(U_i + \gamma d_i)$;
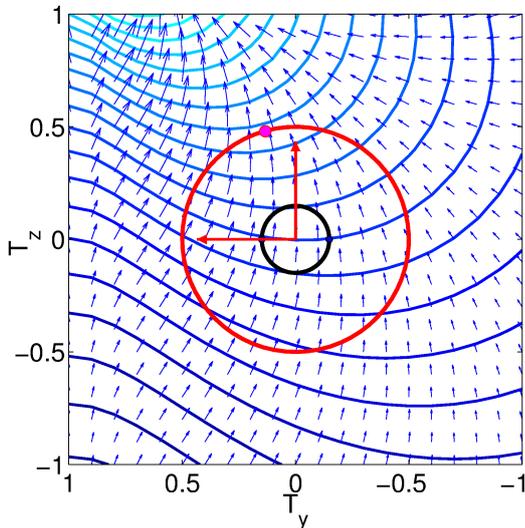4 **end**

Fig. 5: Contour lines and local gradient vectors of the criterion $F_k(T_y, T_z, 0)$ w.r.t. the translation variables T_y, T_z, *i.e.*, when no subsequent rotation is applied to the head ($\phi = \phi_0 = 0$). The red circle delimits the admissible translations. The magenta spot depicts the constrained local maximum.

4.1 Overview

Given a belief $\mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$ on the sensor-to-source position at time k, it is interesting to consider the level sets of the criterion $F_k(T_y, T_z, \phi)$ w.r.t. the translation and rotation variables T_y, T_z, ϕ. The gradient vectors of $F_k(T_y, T_z, \phi)$ are orthogonal to these surfaces and highlight the directions of steepest ascent. Restricting to horizontal sections of the admissible cylindrical set indexed by values ϕ_0 of the rotation variable can ease the analysis. The contour lines of $F(T_y, T_z, \phi_0)$ w.r.t. T_y, T_z

can be observed, as well as the 2-dimensional "local" gradient vectors—which are just obtained by setting the third entry of the genuine 3-dimensional gradient vectors to 0 (Figure 5).

For the instances of the problem ($\mathcal{P}$) considered in Sections 4.2–4.3 below, the optimum solution(s) have been observed to lie on the external surface of $\mathcal{T} \times \mathcal{R}$ in all considered scenarios (this fact has not been proved analytically). So, the contour lines of the criterion F_k constrained to the cylinder surface will also be displayed. To this aim, the following bivariate function is introduced

$$\tilde{F}_k(\alpha, \phi) = [F_k \circ g](\alpha, \phi) \tag{13}$$

with $g : \mathbb{R}^2 \to \mathbb{R}^3$

$$\begin{pmatrix} \alpha \\ \phi \end{pmatrix} \mapsto \begin{pmatrix} r_{max} \sin(\alpha) \\ r_{max} \cos(\alpha) \\ \phi \end{pmatrix},$$

where (α, ϕ) references the position onto the cylinder surface (Figure 4a–4b).

4.2 Iso-entropy contour lines for ITD based exploration

When $\bar{l}(\theta_k)$ in (3) stands for the Woodworth-Schlosberg farfield approximation (4) of the ITD between two antipodal microphones placed on a spherical head, the iso-z_k loci are similar to those depicted in Figure 3.

The contour lines of $F_k(T_y, T_z, \phi_0)$ are plotted on Figures 6a–6c w.r.t. T_y, T_z for various subsequent rotations ϕ_0 of the head, given an initial frame $\mathcal{F}_k$ and a confidence ellipse describing the belief $\mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$, where $\hat{x}_{k|k} = (1, 1.5)^T$. The set of admissible translations is also displayed, as well as the constrained local maximum on the slice of the admissible set defined by ϕ_0.

In Figure 6a, the sensor undergoes a pure translation followed by no rotation. The contour lines of the criterion appear to be distorted—*i.e.*, the gradient of the criterion is subject to important local variations—whenever the translation is either $T = (1, .)^T$ or $T = (., 1.5)^T$. By refering to intuitive arguments from Section 3.2, one can show that for $T = (1, .)^T$ (resp. $T = (., 1.5)^T$), the distorsion is explained by the fact that O_{k+1} lies on the major axis of the confidence ellipse associated to the next predicted state pdf $\mathcal{N}(x_{k+1}; \hat{x}_{k+1|k}, P_{k+1|k})$ (resp. the interaural axis $\overrightarrow{y_{Rk+1}}$ is aligned with the minor axis of this ellipse). For each such restricted value of T, the head must get closer to the source so as to reach a given value of the information criterion, than if a neighboring unrestricted translation were applied.
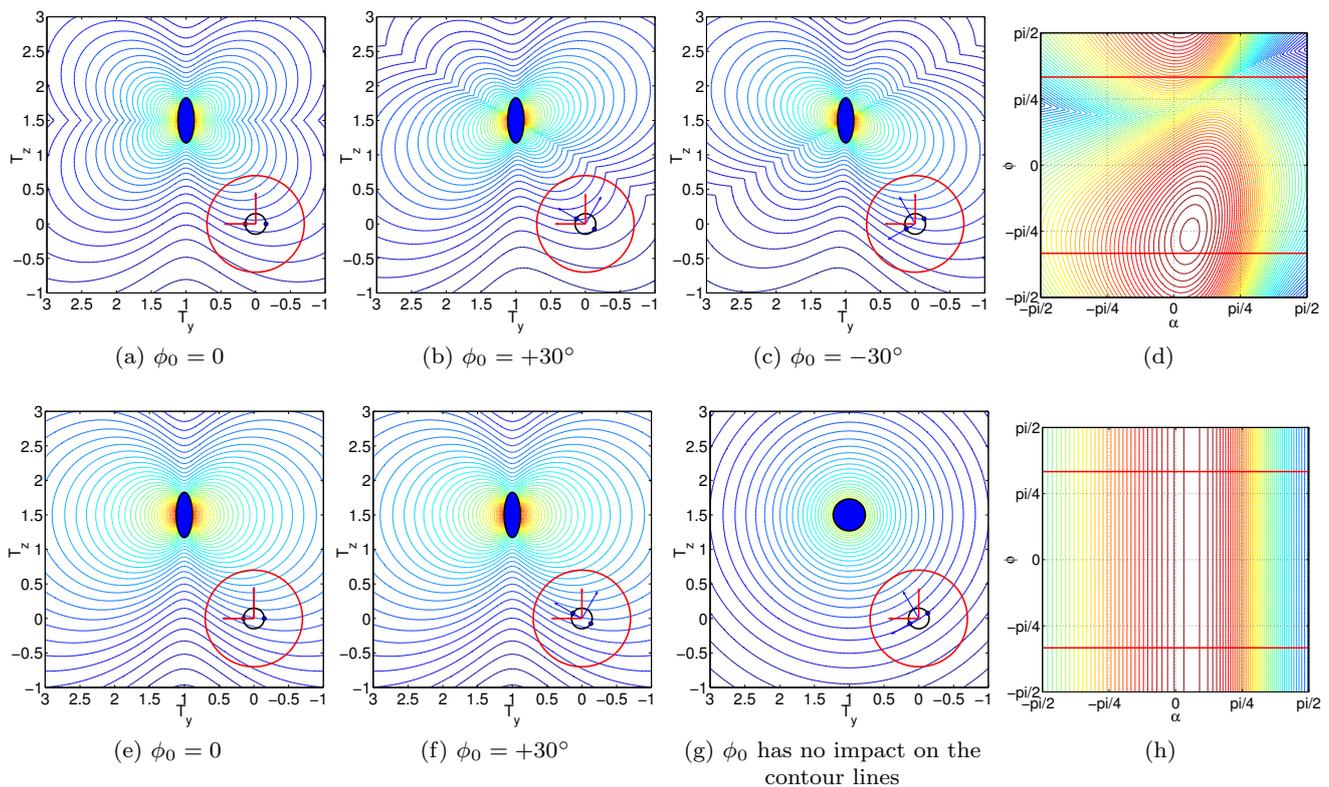
Fig. 6: (a,b,c,e,f,g): Contour lines of the criterion $F_k(T_y, T_z, \phi_0)$ w.r.t. T_y (abscissa, in meters) and T_z (ordinate, in meters). (d,h): Contour lines of $\tilde{F}_k(\alpha, \phi)$ w.r.t. α (abscissa, in radians) and ϕ (ordinate, in radians). In (a,b,c,d) (resp. (e,f,g,h)), the exploration is based on ITD measurements (resp. on ideal azimuth observations). The sensor frame in the initial position $\mathcal{F}_k = (O, \overrightarrow{x_R}, \overrightarrow{y_R}, \overrightarrow{z_R})$ is plotted in red. The initial estimate of the head-to-source position is $\hat{x}_{k|k} = (1, 1.5)^T$. The blue ellipse/circle represents the 99%-probability confidence ellipse associated to the initial belief $\mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$. The red circle delimits the admissible translation $T \in \mathcal{T}$. The blue frame portrays the orientation of $\mathcal{F}_{k+1}$ if a zero translation were applied. The contours are warm (resp. cold) when F_k—or, equivalently, $\tilde{F}_k$—has high (resp. low) values. On Figures (d,h), the horizontal red lines depict the limits of the admissible head rotation, which have been set to $\pm 60°$.

Subsequent rotations of the head by $\phi_0 = +30°$ or $\phi_0 = -30°$ turn Figure 6a into Figure 6b or Figure 6c, respectively. The contour lines are changed, and so is the maximum restricted to the slice defined by ϕ_0. It is more interesting to apply a rotation of $-30°$ than $+30°$, because the obtained optimum for $\phi_0 = -30°$ lies on a contour line with higher value (and thus warmer color). Noticeably, the first distortions explained in the above paragraph for a null rotation remain, while the second ones are just rotated by ϕ_0. Also, as the step size between the index of two consecutive contour lines is constant, and as these contour lines are not regularly spaced, the closer the sensor gets to the source, the higher is the increase in the information criterion F_k.

To get some insight on the maximum value of $F_k(T_y, T_z, \phi)$ on the cylindrical surface of the admissible set, the function $\tilde{F}_k(\alpha, \phi)$ has then been evaluated

for the same initial belief. It appears that its maximum is located on $\phi^* = -48°$ (Figure 6d).

In some cases, e.g., $\hat{x}_{k|k} = (0, 1.5)^T$, the problem $(\mathcal{P})$ has several optimums, see Figures 7a–7b.

4.3 Iso-entropy contour lines for azimuth based exploration

This section considers the following observation model

$$z_k = \theta_k + v_k, \ z_k \in \mathbb{R}, \ v_k \sim \mathcal{N}(0, R_k). \tag{14}$$

Note that observing azimuth measurements contaminated with constant-variance noise is unrealistic in practice. Indeed, when extracting azimuth measurements from the binaural stream, the closer the sound source is to the front axis (resp. to the interaural axis),
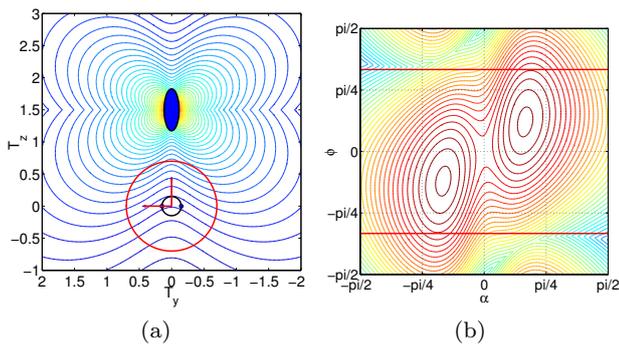
Fig. 7: Contour lines of: (a) $F_k(T_y, T_z, \phi_0)$ w.r.t. T_y, T_z; (b) $\tilde{F}_k(\alpha, \phi)$ w.r.t. α, ϕ when $(\mathcal{P})$ has two solutions. Conventions similar to Figures 6a–6h are used.

the smaller (resp. the bigger) the associated uncertainty is. Nevertheless, this case has been included because it enables a verification of some intuitive features.

The iso-z loci corresponding to equispaced values of the azimuth measurements are equiangular radial lines passing through O. The confidence cones associated to any measured azimuth then have the same width—they are just rotated images of each other. So, given a belief on the source position evenly spread around its genuine location, the assimilation of such an azimuth measurement intuitively brings the same information whether the sensor remains static or whether it moves on a circle centered on the source, regardless of its orientation.

The analysis of the contour lines of $F_k(T_y, T_z, \phi_0)$ w.r.t. T_y, T_z shows that they do not depend on the rotation ϕ_0 (Figure 6e–6f). Consequently, the contour lines of $\tilde{F}_k(\alpha, \phi)$ w.r.t. α, ϕ are vertical (Figure 6h). Nonetheless, the contour lines are still distorted for $T = (1, .)^T$ in (Figure 6e–6f) for the same reasons as those explained in Section 4.2. These distortions vanish when the confidence ellipse associated to the initial belief is circular (Figure 6g), and the contour lines become concentric. In this case, the only way to increase the gained information on the source location is to get closer to it, which is in agreement with the above intuition.

5 Evaluation of the algorithm

The whole three-stage scheme has been implemented on a simulated or real KEMAR binaural head-and-torso-simulator (HATS) from G.R.A.S.®(kemar.us) endowed with omnidirectional planar motion, i.e., with two translational and one rotational degrees of freedom. This section reports the assessment of the obtained audio-motor localization, depending on whether the binaural head undergoes the active motion developed in this paper or other kinds of open-loop movements.

For the sake of simplicity, the binaural head and the robot supporting it move every $T_s = 1$s, then stop in order to acquire binaural signals, perform their short-term analysis (Stage A, Section 2.2) and update the belief on the source position (Stage B, Section 2.3). To drive the exploration, Stage C relies on the Woodworth-Schlosberg measurement equation. The next best position of the robot then comes from the solution of $(\mathcal{P})$ (Section 3.3).

The quality of the short-term azimuth estimation in Stage A critically affects the behavior of the whole binaural active localization. Therefore, in both simulated and live experiments, a non-intermittent white noise signal filtered by a 1kHz bandwidth band-pass filter with 1kHz central frequency has been selected for the sound source, as it endows the azimuth pseudo-likelihood with modes much sharper than with speech sources for instance (Portello et al 2013). Various ways to cope with intermittent sources in Stages A or B have been proposed in (Portello et al 2014b) and (Portello et al 2012), but they have not been implemented here. The movements of the binaural sensor have been limited in translation and rotation by $r \leq r_{max} = 0.1$m and $|\phi| \leq \phi_{max} = 15°$.

5.1 Simulations with audio spatialization

The online rendering of realistic binaural signals caused by a static sound source has first been simulated in an anechoic environment. When the sensor moves, those binaural signals are synthesized by using a database of Head Related Impulse Responses (HRIRs) suited to the used KEMAR HATS. This database as well as a binaural simulator are publicly available at the URLs www.twoears.eu and docs.twoears.eu/en/latest/binsim/.

The sound source is initialized at the position $X = (1, 2)^T$ in the robot frame $\mathcal{F}_0$ at time $k = 0$. To simplify the notation in the legends of the next plots, this frame is denoted as $\mathcal{F}_0 = (O, \overrightarrow{x}_0, \overrightarrow{y}_0, \overrightarrow{z}_0)$.

Various motions of the sensor have been simulated: the proposed active strategy, a translation along the interaural axis, a circular movement such that the front direction of the head stays tangent to its trajectory, and a random movement (Figure 8a). During the five first seconds in all the scenarios, the same rotational movement is applied to the sensor in order to disambiguate front and back, so that at $t = 5$s the Gaussian mixture belief can be better approximated by a single Gaussian pdf. The common progress of the audio-motor localization from initial time $t = 0$s to $t = 5$s is displayed on Figures 8c–8d. Then, each specific movement is applied from time $t = 6$s until the end.

(a)

(b)

(c) Initialisation

(d) Situation at $t = 5$ s

(e) Circular movement at $t = 17$ s
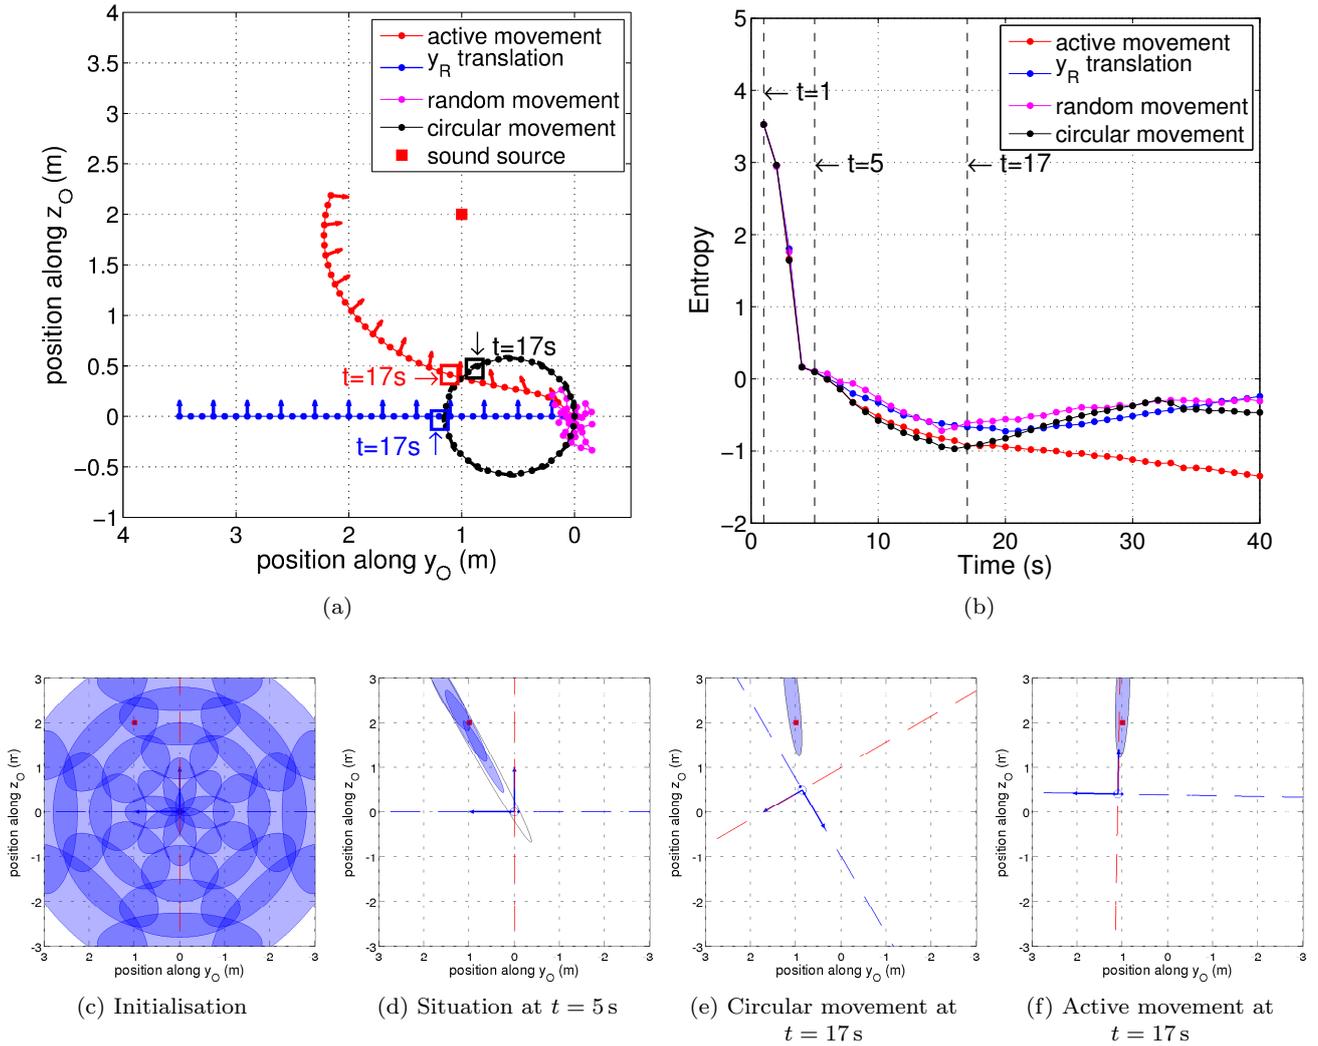
(f) Active movement at $t = 17$ s

Fig. 8: Simulated sound source localization for different scenarios. In the circular movement, the front direction is tangent to the circle. The random path is generated by randomly selecting positions on admissible cylindrical set. (a): Source position and head trajectories in the world frame (*i.e.*, the initial frame $\mathcal{F}_0$). (b): Entropy decrease of the posterior state pdf for the various motion strategies. (c,d,e,f): Interesting snapshots of the localization process showing the binaural head (front direction in dashed red, interaural axis in dashed blue), the source (in red), and the 99%-probability confidence ellipses of the hypotheses constituting the Gaussian mixture belief.

It can be observed that the active motion translates the sensor and rotates its fovea towards the estimated position of the sound source. By computing the Gaussian moment-matched approximation of every state belief $\sum_{i=1}^{I_k} w_k^i \mathcal{N}(x_k; \hat{x}_{k|k}^i, P_{k|k}^i)$, the entropy $h(x_k|z_{1:k})$ has been evaluated for the different strategies (Figure 8b). In terms of localization efficiency, the active motion strategy clearly outperforms the random and translation open-loop movements.

In view of the closeness of the entropies of the passive circular and active motions at each time $t \in [6\,\text{s}, 9\,\text{s}]$ and at $t = 17\,\text{s}$, the confidence ellipses of the respec-

tive beliefs have similar sizes. However, they may have distinct centers and/or orientations, see for instance Figures 8e and 8f.

Between $t = 9\,\text{s}$ and $t = 17\,\text{s}$, the entropies of the posterior state pdfs obtained for the circular motion are minimum. This does not contradict the fundamental property that, between any two consecutive times, the active strategy finds the translation and rotation of the head leading to the maximum decrease of the entropy. In fact, this is an interesting example where the sequence of N one-step-ahead optimum motions does not constitute a N-step-ahead optimum motion.

By Section 3.2, if at $t = 17$ s the head starts from Figure 8e—where its (blue) interaural axis is close to the confidence ellipse—and keeps moving on a circular path tangent to its front axis with no dynamic noise, then only little information can be brought by the measurement. In the simulated experiment, the entropy even increases because the information loss implied by a noisy dynamics cannot be compensated by the little information gain brought by the measurement. If at $t = 17$ s the head starts from Figure 8f—where its (red) front axis intersects the confidence ellipse—then there plausibly exists an admissible translation and rotation which can further decrease the entropy, even if little noise affects the dynamics.

5.2 Live experiments on a binaural robot

A KEMAR HATS has been mounted on a NEOBOTIX MP-L655 nonholonomic mobile robot. In order to make the head omnidirectional, *i.e.*, to endow it with two translational and one rotational degrees of freedom, the neck of the HATS has been equipped with a homemade controllable azimuth degree of freedom (Figure 1). Its software architecture is based on the ROS middleware. Real time components such as binaural audio stream server or three-stage active localization have been synthesized by means of the GenoM3 module generator (Mallet et al 2010). The supervision task which manages the program, plots and saves the results, is performed by a MATLAB® client. The experiments have been conducted in an open-space 15m × 5m × 8m area delimited by dividing walls made of resin, so that reverberation effects were limited.

The results of the audio-motor localization for several motion strategies as well as the genuine position of the source measured by a real-time motion capture system—with ±0.1 mm accuracy—are displayed on Figure 9. A translation along the interaural axis reduces the uncertainty on the distance to the source but cannot disambiguate front from back. A pure rotation (not shown on Figure 9 due to space reasons) resolves the front-back ambiguity but cannot recover the source range. The active motion drives the head in the same way as before. The entropy of the moment-matched approximation of the state posterior pdf is reported on Figure 10. A circular motion is also implemented, leading to a behavior quite similar to Section 5.1.

The whole three-stage framework runs in 5 ms on a i7 quadcore laptop @2.8 GHz with 16 GB RAM. Further videos are available on `http://homepages.laas.fr/danes/AR2016`.

6 Conclusion and prospects

Given a Gaussian belief on the relative position of a sound source with respect to a binaural head, a method has been proposed to determine the admissible planar motion of the sensor which leads to the one-step-ahead best audio-motor localization. It internally relies on a measure of the information brought by the incorporation of TDOA/ITD observations after the sensor has moved. Any other measurement variable could guide the exploration, provided that it is related to the hidden relative position by a closed-form equation similar to (3). Experiments have been conducted on a soundscape rendering simulator and on a mobile binaural robot, when the approach constitutes the "feedback control" component of a three-stage framework to active binaural localization.

Though the "one-step-ahead" statement of the synthesis of the active motion of the sensor is greatly simplified, several issues remain open even in this context. An immediate problem concerns the gap between the Gaussian prior required by the method and the —multimodal—Gaussian mixture (1) provided by the used estimation technique (GMsrUKF) in Section 2.3. This is especially significant at the first localization times, as the combination of the short-term azimuths extracted from the binaural stream and of the sensor motion does not yet enable front-back disambiguation nor range recovery. Several elementary options can be envisaged to get around this problem: (a) keep the most probable hypothesis of the belief provided by the GMsrUKF; (b) turn the genuine multimodal belief into its Gaussian moment-matched approximation; (c) keep the most probable "branch" of the genuine Gaussian mixture—*i.e.*, set of contiguous hypotheses with similar azimuths— and compute its moment-mached approximation; (d) at the early times, apply elementary translation and rotation movements to the head so as to reduce the number of hypotheses in the Gaussian mixture. In this paper, (b) and (d) were jointly used. One more involved alternative is to avoid trading the Gaussian mixture belief for a single Gaussian distribution. As the differential entropy of a Gaussian mixture density cannot be evaluated analytically, two solutions can be envisaged: use an alternative measure of information which can be expressed in closed-form for a Gaussian mixture distribution and supports a rule similar to (8); approximate the differential entropy of a Gaussian mixture wherever needed by a closed-form formula whose accuracy-complexity balance can be handled. These topics are the subject of current research.
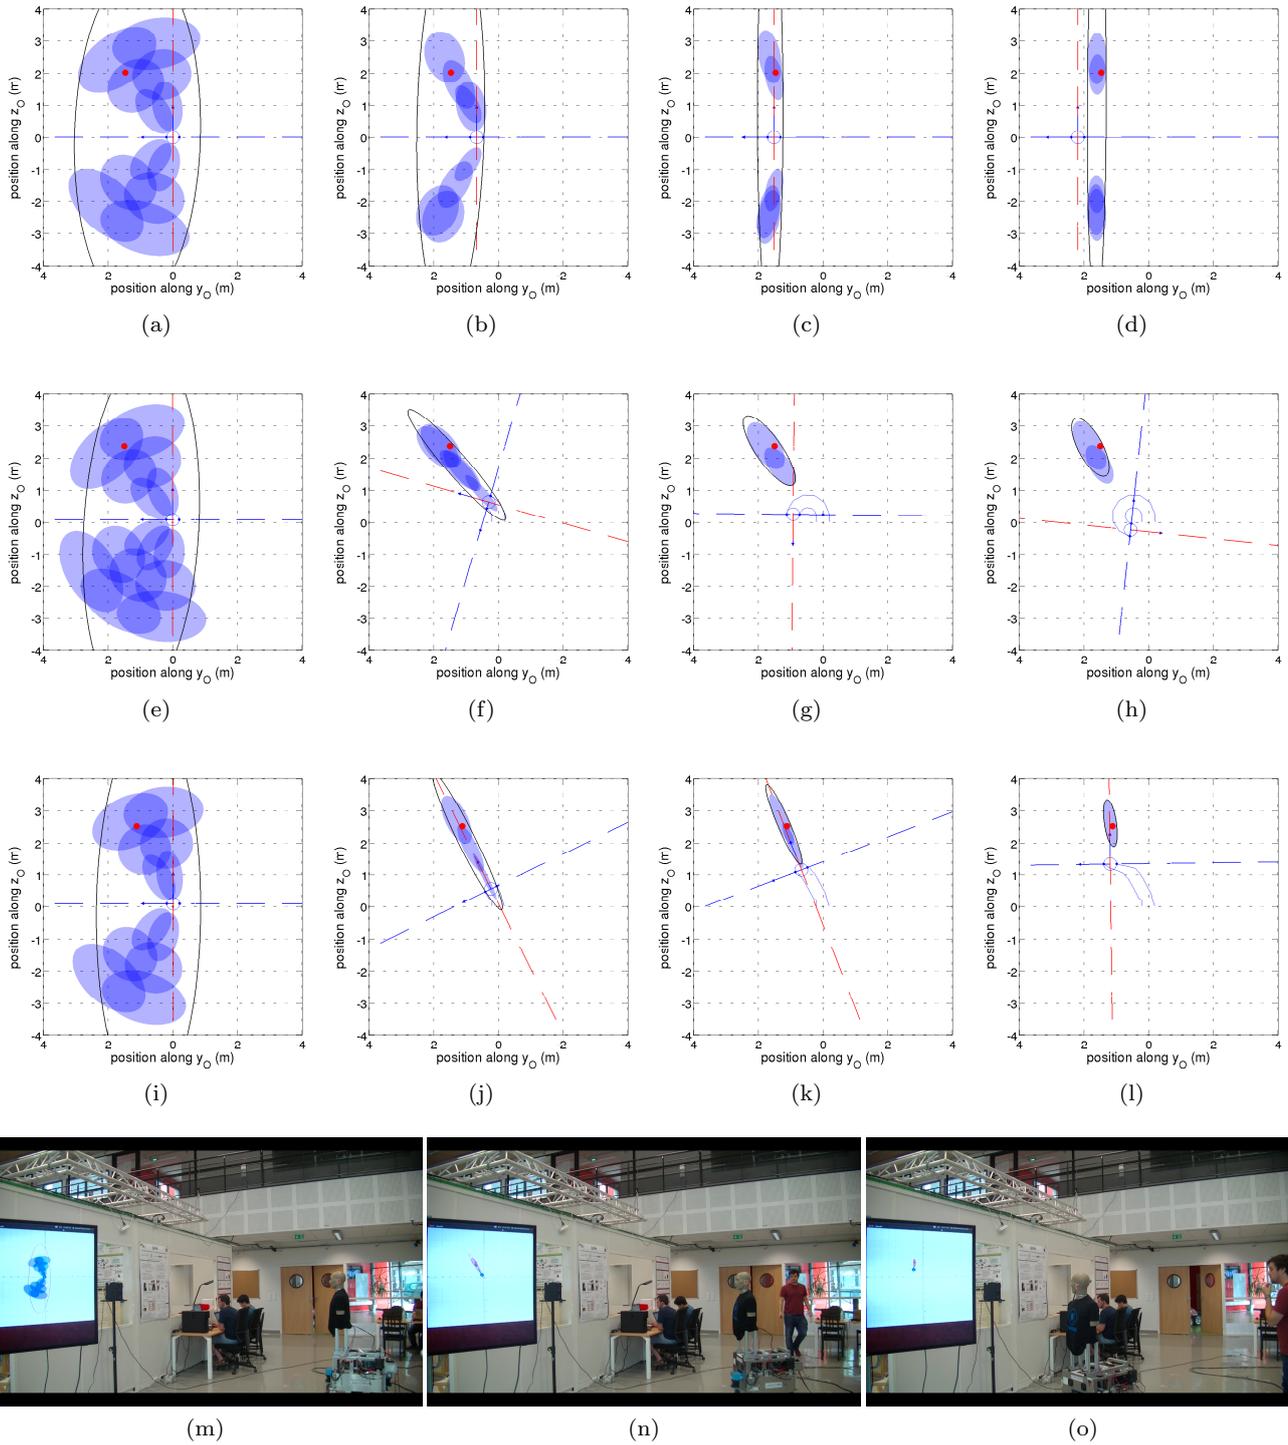
Fig. 9: Single-source localization for different scenarios: (a,b,c,d) translation of the head along the interaural axis; (e,f,g,h) circular movement; (i,j,k,l) active motion. Snapshots (a–l) of the localization process display in the initial frame $\mathcal{F}_0$ the binaural head (front direction in dashed red, interaural axis in dashed blue), the source (in red), and the 99%-probability confidence ellipses of the hypotheses constituting the Gaussian mixture belief. They are displayed at times: (a,e,i) $t = 1\,\mathrm{s}$; (b,f,j) $t = 10\,\mathrm{s}$; (c,g,k) $t = 20\,\mathrm{s}$; (d,h,l) $t = 28\,\mathrm{s}$. Screenshots of the recorded video for the active motion scenario are reported at times: (m) $t = 2\,\mathrm{s}$; (n) $t = 10\,\mathrm{s}$; (o) $t = 34\,\mathrm{s}$.
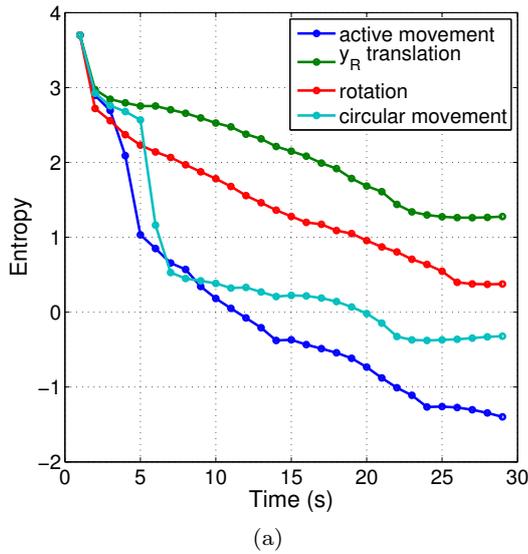
(a)

Fig. 10: Entropy reduction of the posterior state pdf for various motion strategies

Though the proposed strategy does find the translation and rotation optimizing a one-step-ahead criterion, the sequence of N such motions may be outperformed by another sequence of admissible displacements as explained in Section 5.1. This is why current research also focuses on multi-step methods, where the objective is to find a sequence of the robot commands $u^\star = \{u_k, u_{k+1}, \ldots, u_{k+N}\}$ which improves the localization after N-steps. For instance, this optimum N-step sequence may be obtained by expressing the differential entropy of the belief at $k + N$ and minimizing its expected value over the next measurements $z_{k+1}, \ldots, z_{k+N}$, in the vein of Deutsch et al (2004).

A thorough evaluation of Stage A in several kind of acoustic environments is in process. It will be followed by the evaluation of the whole localization framework, including the audio-motor localization Stage B and information-based feedback control Stage C. Finally, the integration of the proposed active localization framework within a comprehensive computational model of human auditory perception—like the one developed in TWO!EARS—requires further investigation. Active localization has been viewed as a sensorimotor function operating on short time ranges, i.e., a low-level "reflexive behavior". So, its interaction with upper-level long-term cognitive processes needs to be refined. Among the important issues are a kind of exploration-exploitation dilemma: when and how must a cognitive process decide between exploring—i.e., parameterizing and triggering an active localization reflexive behavior in order to gather information—and

launching an elaborate reasoning on the basis of its current knowledge?

7 Appendix

Consider the posterior state pdf $p(x_k|z_{1:k})$ of the sensor-to-source position at time k, and $\mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$ the approximate Gaussian belief. This pdf can be mapped into the 1D Gaussian approximation $\mathcal{N}(z_{k+1}; \hat{z}_{k+1|k}, S_{k+1|k})$ of the predicted measurement pdf $p(z_{k+1}|z_{1:k})$, by using the unscented transform. The aim is then to maximize the variance $S_{k+1|k}$ so as to increase the entropy $h(z_{k+1}|z_{1:k})$. This involves the composition of several functions.

First the sigma-points $\{X_i^-\}$ corresponding to $p(x_k|z_{1:k}) = \mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$ are computed from the posterior mean $\hat{x}_{k|k}$ of the state vector at time k and the Cholesky decomposition $P_{k|k} = L_{k|k}L_{k|k}^T$ of the posterior covariance:

$$\{X_i^-\} = \text{Sigma\_points}\left(\hat{x}_{k|k}, L_{k|k}\right) \qquad (15)$$

The sigma-points $\{X_i^+\}$ of the next predicted state pdf $p(x_{k+1}|z_{1:k}) = \mathcal{N}(x_k; \hat{x}_{k+1|k}, P_{k+1|k})$ can be obtained by applying the translation and rotation on each sigma point in the set $\{X_i^-\}$. Note that (2) is defined as a function of (T_y, T_z, ϕ), so that

$$\forall i, \ X_i^+ = \Phi_{X_i^-}(T_y, T_z, \phi). \qquad (16)$$

Then the set of sigma-points $\{Z_i^+\}$ of the predicted measurement pdf $p(z_{k+1}|z_{1:k}) = \mathcal{N}(z_k; \hat{z}_{k+1|k}, S_{k+1|k})$ can be obtained from $\{X_i^+\}$ defined in (16) by:

$$\forall i, \ Z_i^+ = l\left(-\text{atan2}\left(X_i^+(1), X_i^+(2)\right)\right), \qquad (17)$$

with $X_i^+(1)$ and $X_i^+(2)$ the components of X_i^+, and $l(\cdot)$ the measurement equation used to guide the exploration. Finally the mean $\hat{z}_{k+1|k}$ and variance $S_{k+1|k}$ of $p(z_{k+1}|z_{1:k})$ are computed by

$$\hat{z}_{k+1|k} = \sum_i w_m^i Z_i^+ \qquad (18)$$

$$S_{k+1|k} = \sum_i w_c^i \left(Z_i^+ - \hat{z}_{k+1|k}\right)^2, \qquad (19)$$

with $\{w_m^i\}$ and $\{w_c^i\}$ the classic weights of the unscented transform.

The log of the variance $S_{k+1|k}$ comes as a function of the finite translation and rotation, i.e., $\log S_{k+1|k} = F_k(T_y, T_z, \phi)$. However the maximum of this function is not analytically tractable. Its gradient around $U = (T_y, T_z, \phi)$ is then computed as follows.

The first order Taylor expansion of the functions $\Phi_{X_i^-}$, atan2, l, and log, are composed around

U with infinitesimal translations and rotation $du = (dT_y, dT_z, d\phi)^T$:

$$\Phi_{X_i^-}(U + du) = \Phi_{X_i^-}(U) + J\Phi_{X_i^-}(U)\,du$$

$$\text{atan2}(u,v) = \text{atan2}(u_0, v_0) + \nabla^T \text{atan2}(u_0, v_0) \begin{pmatrix} u - u_0 \\ v - v_0 \end{pmatrix}$$

$$l(w) = l(w_0) + l'(w_0)(w - w_0)$$

$$\log(r) = \log(r_0) + \frac{1}{r_0}(r - r_0) \tag{20}$$

with ∇ the gradient operator. $J\Phi_{X_i^-}(U)$ is the Jacobian of $\Phi_{X_i^-}$ at U. Then the result of the composition, noted $Z_i(dT_y, dT_z, d\phi)$, is used to retrieve the mean and the variance with (18) and (19). Finally, the first order Taylor expansion of $F_k(dT_y, dT_z, d\phi)$ is obtained, highlighting the gradient ∇F_k:

$$F_k(U + du) = F_k(U) + \nabla^T F_k(U)\,du. \tag{21}$$

References

Aaronson N, Hartmann W (2014) Testing, correcting, and extending the woodworth model for interaural time difference. The Journal of the Acoustical Society of America 135:817–823

Bourgault F, Makarenko A, Williams S, Grocholsky B, Durrant-Whyte H (2002) Information based adaptive robotic exploration. In: IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, (IROS'2002), Lausanne, Switzerland

Bustamante G, Portello A, Danès P (2015) A three-stage framework to active source localization from a binaural head. In: IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP'2015), Brisbane, Australia

Bustamante G, Danès P, Forgue T, Podlubne A (2016) Towards information-based feedback control for binaural active localization. In: IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (in press ICASSP'2016), Shanghai, China

Cooke M, Lu Y, Lu Y, Horaud R (2007) Active hearing, active speaking. In: Int. Symp. on Auditory and Audiological Research (ISAAR'07), Marienlyst, Helsigør, Denmark

Cover T, Thomas J (1991) Elements of Information Theory. Wiley

Denzler J, Brown C (2002) Information theoretic sensor data selection for active object recognition and state estimation. IEEE Trans on Pattern Analysis and Machine Intelligence 24(2):145–157

Deutsch B, Zobel M, Denzler J, Niemann H (2004) Multi-step entropy based sensor control for visual object tracking. Pattern Recognition 3175:359–366

Feder H, Leonard J, Smith C (1999) Adaptive mobile robot navigation and mapping. The International Journal of Robotics Research 18(7):650–668

Forster C, Pizzoli M, Scaramuzza D (2014) Appearance-based active, monocular, dense reconstruction for micro aerial vehicles. In: Proceedings of Robotics: Science and Systems, Berkeley, USA

Grocholsky B, Makarenko A, Durrant-Whyte H (2003) Information-theoretic coordinated control of multiple sensor platforms. In: IEEE Int. Conf. on Robotics and Automation, (ICRA'03), Taipei, Taiwan

Julian B (2013) Mutual information-based gradient-ascent control for distributed robotics. PhD thesis, Massachusetts Institute of Technology

Julier SJ, Uhlmann JK (2004) Unscented filtering and nonlinear estimation. Proceedings of the IEEE 92(3):401–422, DOI 10.1109/JPROC.2003.823141

Kumon M, Fukushima K, Kunimatsu S, Ishitobi M (2010) Motion planning based on simultaneous perturbation stochastic approximation for mobile auditory robots. In: IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS'2010), Taipei, Taiwan

Le Cadre JP, Laurent-Michel S (1999) Optimizing the receiver maneuvers for bearings-only tracking. Automatica 35(4):591–606

Mallet A, Pasteur C, Herrb M, Lemaignan S, Ingrand F (2010) Genom3: Building middleware-independent robotic components. In: IEEE Int. Conf. on Robotics and Automation, (ICRA'2010), Anchorage, Alaska

Manyika J (1993) An information-theoretic approach to data fusion and sensor management. PhD thesis, University of Oxford

Martinson E, Schultz A (2009) Discovery of sound sources by an autonomous mobile robot. Autonomous Robots 27:221–237

Martinson E, Apker T, Bugajska M (2011) Optimizing a reconfigurable robotic microphone array. In: IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS'2011), San Francisco, California

Nakadai K, Lourens T, Okuno H, Kitano H (2000) Active audition for humanoid. In: Nat. Conf. on Artificial Intelligence (AAAI'2000), Austin, TX

Portello A, Danès P, Argentieri S (2012) Active binaural localization of intermittent moving sources in the presence of false measurements. In: IEEE/RSJ Int. Conf. on Intell. Robots and Systems (IROS'2012)

Portello A, Danès P, Argentieri S, Pledel S (2013) HRTF-based source azimuth estimation and activity detection from a binaural sensor. In: IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS'2013), Tokyo, Japan

Portello A, Bustamante G, Danès P, Mifsud A (2014a) Localization of multiple sources from a binaural head in a known noisy environment. In: IEEE/RSJ Int. Conf. on Intell. Robots and Systems (IROS'2014), Chicago, IL

Portello A, Bustamante G, Danès P, Piat J, Manhès J (2014b) Active localization of an intermittent sound source from a moving binaural sensor. In: Forum Acustium (FA'2014), Krakow, Poland

Ristic B, Arulampalam M (2003) Tracking a manoeuvring target using angle-only measurements: algorithms and performance. Signal Processing 83(6):1223 – 1238

Sasaki Y, Thompson S, Kaneyoshi M, Kagami S (2010) Map-generation and identification of multiple sound sources from robot in motion. In: IEEE/RSJ Intl. Conf. Intell. Robots and Systems (IROS'2010), Taipei, Taiwan, pp 437–443

Sommerlade E, Reid I (2008) Information-theoretic active scene exploration. In: IEEE Conf. on Computer Vision and Pattern Recognition, (CVPR'2008), Anchorage, Alaska

Thrun S, Burgard W, Fox D (2005) Probabilistic Robotics.
    The MIT Press
Vincent E, Sini A, Charpillet F (2015) Audio source local-
    ization by optimal control of a mobile robot. In: IEEE
    Int. Conf. on Acoustics, Speech and Signal Processing
    (ICASSP'2015), Brisbane, Australia