

TOWARDS INFORMATION-BASED FEEDBACK CONTROL FOR BINAURAL ACTIVE LOCALIZATION

Gabriel Bustamante^{*†}

Patrick Danès^{*†}

Thomas Forgue^{*}

Ariel Podlubne^{*}

^{*} CNRS, LAAS, 7 avenue du Colonel Roche, F-31400 Toulouse, France

[†] Univ de Toulouse, UPS, LAAS, F-31400 Toulouse, France

ABSTRACT

This paper takes place within the field of sound source localization by combining the signals sensed by a binaural head with its motor commands. Such so-called “active” schemes are known to overcome limitations occurring in the static context, such as front-back ambiguities or distance non-observability. On the basis of a stochastic filter, which approximates the posterior probability density function of the sensor-to-source situation, a feedback controller of the sensor motion is proposed so as to reduce the associated uncertainty. An information-theoretic analysis of the effect of the sensor motion on the localization uncertainty is first conducted. Then, a gradient ascent scheme is used to drive the head towards the area of minimum uncertainty (maximum information). An evaluation on simulated scenarios, as well as on data coming from real experiments, is included.

Index Terms— Robot audition, active localization, stochastic filtering, information theory, information based control.

1. INTRODUCTION

Auditory robots enable the incorporation of motion into binaural “active” sound source localization, so as to remove front-back ambiguities or recover range [1]. Such schemes can cope with moving and intermittent sound sources [2]. This paper addresses the feedback control of the sensor motion so as to reduce localization uncertainty.

Within the related “exploration problem” in robotics, robots move autonomously so as to maximize their knowledge about the world. Simultaneous localization and mapping (SLAM) has been extended so that robots move in the direction of maximum local information improvement [3]. The control scheme extracts the information about the state variables using the concepts of Shannon entropy and mutual information [4], and, in the end, maximizes a “size” criterion such as the determinant or trace of the inverse of the one-step ahead posterior state covariance matrix [5]. Other information-theoretic controllers have been applied to information retrieval on some targets [6], robot guidance towards

areas of maximum uncertainty [7], control of a robot-mounted camera to optimize depth estimation [8], or sensor parameters selection (e.g., zoom or attitude) for scene analysis [9, 10].

In robot audition, information-based control is sparser. Motion planning was proposed in [11] to improve speech recognition from a monaural robot. Sound localization was improved in [12] by moving microphones deployed in the environment. Recently [13], a robot equipped with a microphone array has been controlled to improve occupancy grid based source localization, by using dynamic programming.

Our single-source localization strategy is organized into three layers [14]. The first stage estimates the source azimuth from the short-term analysis of the binaural stream. The second stage assimilates these azimuths over time and combines them with the motor commands into a stochastic filter, leading to the posterior probability density function (pdf) of the head-to-source relative situation. The paper focuses on the third stage, which aims at controlling the head motion to improve the quality of the fusion performed in the second stage. Simulations and real-life experiments illustrate the approach.

2. PRIOR WORK

This section recalls the azimuth estimation of a single source [15] and the incorporation of the head motion [16] so as to get a Gaussian mixture approximation of the head-to-source pdf. The reading of [14] is strongly recommended. In the sequel, random variables/processes and corresponding samples are written using similar lowercase letters.

2.1. Short-term extraction of directional cues

The left and right microphones are termed R_1 and R_2 . The interaural transfer function is known over an adequate range of source azimuth and frequencies. The frame $\{\mathcal{F} = (O, \vec{e}_x, \vec{e}_y, \vec{e}_z)\}$ is attached to the head, with $R_1\vec{O} = O\vec{R}_2$. R_1, R_2 and the pointwise emitter E lie on a common horizontal plane. The source and sensor noises are modeled as random processes satisfying reasonable hypotheses (Gaussianity, zero-mean, band-limited, “local stationnarity”...). From [15], on the basis of the channel-time-frequency decomposition y_k of the binaural signal on a sliding window

This work was partially supported by EU FET grant TWO!EARS, ICT-618075, www.twoears.eu.

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” $p(y_k|\theta_k)$. It 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, thanks to a notable separation property.

2.2. Fusion of audio information with motor commands

A discrete-time stochastic state space equation is set up, uniting the velocity control vector $u_k \in \mathbb{R}^3$ (2 translations and 1 rotation) of the head to the head-to-source situation $x_k \in \mathbb{R}^2$. A theoretically sound Gaussian mixture square root unscented Kalman filter (GMsrUKF) is defined so as to incorporate the above pseudo likelihood $p(y_k|\theta_k)$ (where θ_k comes as a static function of x_k) [16] and compute a Gaussian mixture approximation of the posterior pdf $p(x_k|y_{1:k}) = \sum_{i=1}^I w_k^i \mathcal{N}(x_k; \hat{x}_{k|k}^i, P_{k|k}^i)$, where $(w_k^i, \hat{x}_{k|k}^i, P_{k|k}^i)$ are the weight, mean and covariance of each hypothesis. Contrarily to several particle filters, a self-initialization as well as posterior covariance consistency is ensured, so that front and back are disambiguated, and both range and azimuth are faithfully recovered.

3. TOWARDS AN INFORMATION-BASED SENSORIMOTOR FEEDBACK

An information measure can be defined from the posterior pdf $p(x_k|y_{1:k})$ at time k , which captures all the information on the head-to-source situation held in the measurements. The one-step ahead control problem is studied, consisting in determining the control vector u_k^* such that the information in $p(x_{k+1}|y_{1:k+1})$, averaged over the (unknown) possible values of the next measurement y_{k+1} , is maximized. Two simplifications make the problem tractable: $p(x_k|y_{1:k})$ is reduced to a single Gaussian pdf $\mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$, e.g., by keeping its most probable hypothesis or by computing its moment-matched approximation; in the definition of u_k^* , the next channel-time-frequency decomposition y_{k+1} is traded for a scalar vector z_{k+1} satisfying a closed-form measurement equation $z_{k+1} = g(x_{k+1}) + v_{k+1}$, with v the measurement noise. Henceforth, the Woodworth-Schlossberg formula for interaural time difference approximation over a spherical head [17] is selected for $g(\cdot)$ and thus guides the exploration. u_k^* is then defined on the basis of $p(x_k|z_{1:k})$ defined as $\mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$ and on the above measurement equation. Once this control signal is applied, the next state posterior pdf $p(x_{k+1}|y_{1:k+1})$ is computed from the GMsrUKF described in §2.2. Then, one defines again $p(x_{k+1}|z_{1:k+1})$ as equal to $p(x_{k+1}|y_{1:k+1})$ and the whole process is repeated to determine u_{k+1}^* .

3.1. Information measures and control input

Let w, x be two random variables with pdfs $p(w), p(x)$. The differential entropy $h(x)$ of x embodies its uncertainty, in that the lower $h(x)$ the higher the information in x . The mutual information $I(w, x) (\geq 0$ by definition) measures the amount of information w contains about x [4]. They are defined by

$$h(x) = - \int p(x) \log p(x) dx; I(w, x) = \iint p(w, x) \log \frac{p(w, x)}{p(w)p(x)} dw dx.$$

If w and x are conditioned on the event that a random variable v takes a given value, then the entropies/information are denoted by $h(w|v)$, $h(x|v)$ and $I(w, x|v)$. The following rule, somewhat similar to [18], holds.

Theorem 1 *Decomposing the negative logarithm of the posterior pdf $p(x_{k+1}|z_{1:k+1})$ as*

$$-\log p(x_{k+1}|z_{1:k+1}) = -\log p(x_{k+1}|z_{1:k}) - \log \left(\frac{p(z_{k+1}|x_{k+1})}{p(z_{k+1}|z_{1:k})} \right),$$

and taking its expectation conditioned on $z_{1:k}$ (which involves the joint pdf $p(x_{k+1}; z_{k+1}|z_{1:k})$), leads to

$$\begin{aligned} \mathbb{E}_{z_{k+1}} \{h(x_{k+1}|z_{1:k+1})\} &= h(x_{k+1}|z_{1:k}) - I(x_{k+1}; z_{k+1}|z_{1:k}) \\ \mathbb{E}_{x_{k+1}} \{h(z_{k+1}|x_{k+1})\} &= h(z_{k+1}|z_{1:k}) - I(x_{k+1}; z_{k+1}|z_{1:k}), \end{aligned}$$

with $h(x_{k+1}|z_{1:k+1})$, $h(x_{k+1}|z_{1:k})$, $h(z_{k+1}|x_{k+1})$, $h(z_{k+1}|z_{1:k})$ the entropies of the next filtered state pdf, the next predicted state pdf, the observation law, the next predicted measurement pdf, and $I(x_{k+1}; z_{k+1}|z_{1:k})$ the mutual information of the next state and measurement conditioned on the sequence of measurements up to current time.

In view of the mutual information positivity, the inequality $\mathbb{E}_{z_{k+1}} \{h(x_{k+1}|z_{1:k+1})\} \leq h(x_{k+1}|z_{1:k})$ holds. Given $p(x_{k+1}|z_{1:k})$, minimizing $\mathbb{E}_{z_{k+1}} \{h(x_{k+1}|z_{1:k+1})\}$ boils down to maximizing $I(x_{k+1}; z_{k+1}|z_{1:k})$.

The entropy of a multivariate normal distribution comes as an increasing affine function of the log-determinant of its covariance matrix [4]. So, $h(z_{k+1}|x_{k+1})$ only depends on the variance of the noise v_{k+1} . In view of the (nonlinear) Kalman equations that could be used to assimilate the measurement z_{k+1} for exploration, $h(x_{k+1}|z_{1:k+1})$ does not depend on z_{k+1} . As the prior state dynamics defines a rigid body motion of the head, the entropy $h(x_{k+1}|z_{1:k})$ of the predicted state is equal to $h(x_k|z_{1:k})$ and does not depend on u_k if the dynamic noise is neglected. So, the following holds.

Theorem 2 *Finding a control input u_k^* which minimizes the entropy $h(x_{k+1}|z_{1:k+1})$ of the next filtered state pdf 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, that is,*

$$\begin{aligned} u_k^* &= \arg \min_{u_k} h(x_{k+1}|z_{1:k+1}) = \arg \max_{u_k} I(x_{k+1}; z_{k+1}|z_{1:k}) \\ &= \arg \max_{u_k} h(z_{k+1}|z_{1:k}). \end{aligned}$$

The Kalman filter equations lead to the approximations $\hat{z}_{k+1|k}$ and $S_{k+1|k}$ of the predicted measurement mean and covariance. Considering that $p(z_{k+1}|z_{1:k})$ is approximated by the Gaussian pdf $\mathcal{N}(z_{k+1}; \hat{z}_{k+1|k}, S_{k+1|k})$, the entropy $h(z_{k+1}|z_{1:k})$ can be rewritten as an increasing affine function of the log-determinant of $S_{k+1|k}$.

3.2. Geometric interpretation

Theorem (2) can be interpreted geometrically. Given a head-to-source situation (Figure 1.a), the 2D Gaussian approximation of the next filtered state pdf resulting from the fusion of the next predicted state pdf with the measurement is represented for various positions: when the sensor is still (Figure 1.b); when after a sensor motion, the interaural axis, supported by \vec{e}_y (resp. the boresight direction, supported by \vec{e}_z), is parallel to the small axis of the confidence ellipse associated to the predicted state pdf (Figure 1.c) (resp. (Figure 1.d)). Importantly, the Woodworth iso-ITD are not uniformly distributed along the azimuths. They are more concentrated along the direction of \vec{e}_z which defines the auditory fovea, while they are sparser along the interaural axis \vec{e}_y .

The variance of the predicted measurement is low when the ellipse is intersected by a few number of iso-ITD (Figure 1.b-c). In this case, the measurement uncertainty due to noise corresponds to a wide spatial sector. Consequently, the measurement update cannot significantly improve the information in the filtered state pdf. The more iso-ITDs intersect the ellipse associated to the predicted state pdf, the higher the variance of the predicted measurement. For instance, when the small axis of this ellipse is parallel to the auditory fovea (Figure 1.d) or when the head gets closer to the source, the measurement uncertainty due to noise corresponds to a narrow cone. Then, the fusion of the predicted state pdf and of the measurement results to a strong increase in the information of the filtered state pdf.

3.3. Feedback control by gradient ascent strategy

As the robot only undergoes rigid body motions, the problem is reduced to find, from the head-to-source situation at time k characterized by $\mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$, the adequate finite translation $\{T_y, T_z\}$ and rotation $\{\phi\}$ which maximizes the variance of the next predicted measurement pdf $S_{k+1|k}$. An expression $S_{k+1|k} = F_k(T_y, T_z, \phi)$ of this variance with respect to the motion variables is then necessary. In practice this expression has no closed form. An analytic derivation of an approximation of its gradient around $\vec{D}_0 = (0, 0, 0)^T$ is proposed so as to point out the direction of its steepest ascent. Using successive first order Taylor expansions and the unscented transform, an analytic approximation

$$F_k(\vec{D}_0 + \vec{du}) = F_k(\vec{D}_0) + \nabla F_k(\vec{D}_0)^T \vec{du}$$

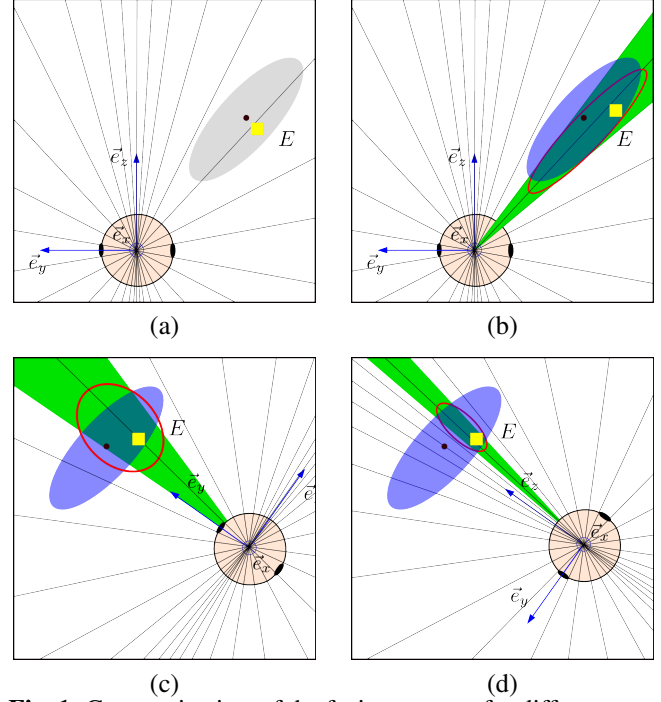


Fig. 1. Geometric view of the fusion process for different motion strategies. The blue frame $\mathcal{F} : (O, \vec{e}_x, \vec{e}_y, \vec{e}_z)$ is attached to the binaural head (represented with its “ears”). The sound source real position is pointed with a yellow square. The grey confidence ellipse is related to the current posterior head-to-source pdf. The measurement space is materialized by the Woodworth iso-ITDs. The blue ellipses are associated with the next predicted state pdf (after applying the velocity command to the head, which is 0 in b). The green sector/cone describes the spatial uncertainty due to measurement noise. The red ellipse sketches the confidence ellipsoid associated to the next filtered state pdf, after the incorporation of the Woodworth ITD for the source position.

can be written when an infinitesimal motion vector $\vec{du} = (dT_y, dT_z, d\phi)^T$ is applied around \vec{D}_0 . Therein, $F_k(\vec{D}_0)$ is the variance of the predicted measurement without any displacement, and $d_k = \nabla F_k(\vec{D}_0)$ is the gradient of F_k evaluated at \vec{D}_0 , i.e., the direction of steepest ascent of the next predicted measurement variance. Note that the function F_{k+1} is not the same as F_k due to the incorporation of the observation y_{k+1} . In other words, the strategy does not consist in iteratively maximizing a function by the gradient method.

4. SIMULATIONS AND REAL-LIFE EXPERIMENTS

4.1. Evaluation of the method

Simulations have been conducted to test the validity of the method (Figure 2). The GMsrUKF coupled with the Woodworth-Schlosberg measurement equation for exploratory control have been implemented in MATLAB. The filter fre-

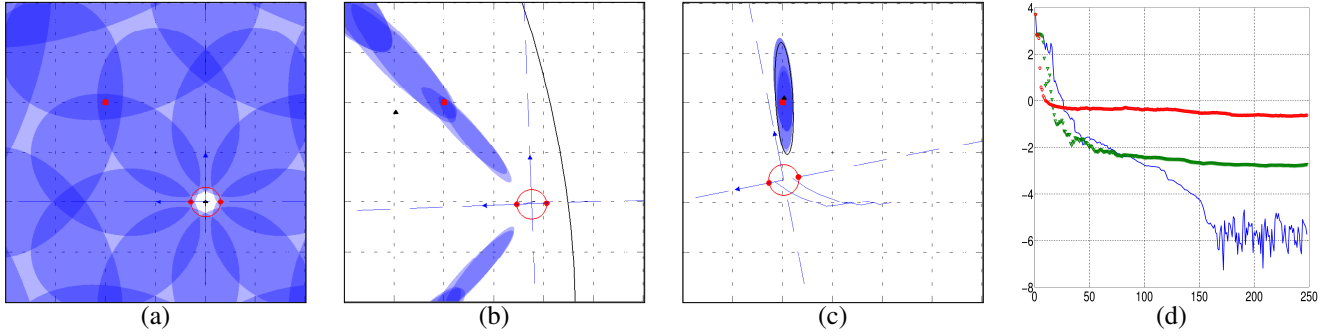


Fig. 2. Audio-motor localization of a sound source by moving the sensor towards the gradient direction of F_k . (a): Self-initialization. (b): Front-back ambiguity with no motion. (c): Information-based motion towards the sound source. (d) Entropy of the moment-matched approximation of the posterior state pdf along time for different control strategies: pure rotation (red), circular trajectory around the sound source (green) and information based strategy (blue).

quency is 20Hz and a new control input is updated every 100ms . The ground truth position of the sensor is depicted with a blue frame equipped with head and ears while the sound source position is plotted with a red square. Blue ellipses represent 99% confidence regions associated with the various hypotheses of Gaussian mixture approximation of the state posterior pdf. Several scenarios have been considered so as to compare the efficiency of the command strategy (Figure 2.d). During the transient mode, the proposed strategy is not better than purely rotational or circular open-loop motions, because it is based on a single Gaussian pdf which does not capture all the hypotheses of the Gaussian sum approximation of the posterior state pdf. When only one hypothesis remains, the decrease is more significant. The plotted final entropy of the moment-matched approximation of the genuine posterior state pdf reaches a steady state value which is much lower than the steady state value of this entropy for other motions.

4.2. Real life experiment

The live experiments have been conducted on the robot Jido which is a MP-L655 platform from Neobotix®. The KEMAR binaural head and torso from G.R.A.S.® is embedded on it. The head has been equipped with a homemade controllable azimuth dof on its neck (Figure 3). The software architecture is based on the ROS middleware, and involves the GenoM3 module generator [19]. The speaker emits a white noise filtered with 1 kHz bandwidth band-pass filter whose central frequency is 1 kHz. Due to space reasons, active localization results as well as technical complements on the gradient derivation, are shown on <http://homepages.laas.fr/danes/ICASSP2016>.

5. CONCLUSION AND PROSPECTS

This paper has presented a feedback control scheme for active binaural localization. The method has been assessed in simulations as well as on an embedded auditory platform. An on-

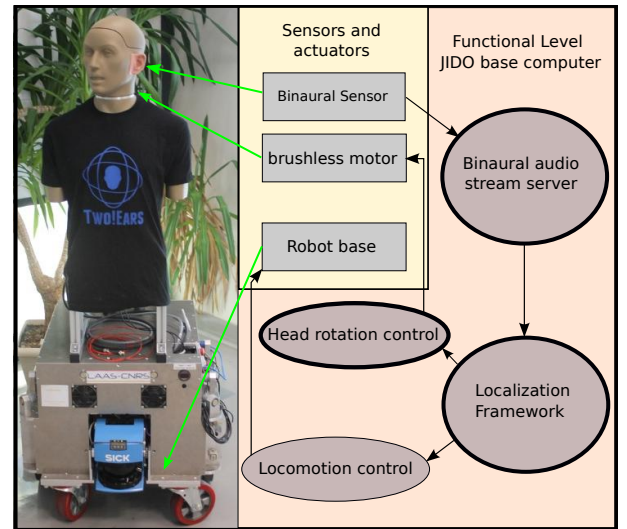


Fig. 3. Audition robot architecture. All the curved shapes are ROS or GenoM3/ROS nodes. The binaural audio stream server publishes binaural data from the KEMAR head. The localization framework implements the three-step strategy. It sends commands the head rotation control node and to the robot locomotion control node.

going prospect is to implement a complete gradient or Newton algorithm on each F_k so as to find the admissible optimal *finite* translation and rotation. The nature of local optimums will be also investigated.

Multi-step method will then be considered. For instance, a criterion based on the expectation of the differential entropy over several steps could be used to guide the motion, in the vein of [20]. The guidance will thus be viewed as a multi-step optimization problem where the objective is to find a sequence of the robot commands $u^* = \{u_k, u_{k+1}, \dots, u_{k+N}\}$ so as to improve the localization over N -steps. The resolution may rely on Partial Observable Markov Decision Processes incorporating reward functions based on information analysis [21] or on other optimal control techniques.

6. REFERENCES

- [1] K. Nakadai, T. Lourens, H.G. Okuno, and H. Kitano, "Active audition for humanoid," in *Nat. Conf. on Artificial Intelligence (AAAI'2000)*, Austin, TX, 2000.
- [2] A. Portello, P. Danès, and S. Argentieri, "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)*, 2012.
- [3] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*, The MIT Press, 2005.
- [4] T. Cover and J. Thomas, *Elements of Information Theory*, Wiley, 1991.
- [5] F. Bourgault, A.A. Makarenko, S.B. Williams, B. Grocholsky, and H.F. Durrant-Whyte, "Information based adaptive robotic exploration," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, (IROS'2002)*, Lausanne, Switzerland, 2002.
- [6] B. Grocholsky, A. Makarenko, and H. Durrant-Whyte, "Information-theoretic coordinated control of multiple sensor platforms," in *IEEE Int. Conf. on Robotics and Automation, (ICRA'03)*, Taipei, Taiwan, 2003.
- [7] B.J. Julian, *Mutual Information-based Gradient-ascent Control for Distributed Robotics*, Ph.D. thesis, Massachusetts Institute of Technology, 2013.
- [8] Christian Forster, Matia Pizzoli, and Davide Scaramuzza, "Appearance-based active, monocular, dense reconstruction for micro aerial vehicles," in *Proceedings of Robotics: Science and Systems*, Berkeley, USA, July 2014.
- [9] J. Denzler and C.M. Brown, "Information theoretic sensor data selection for active object recognition and state estimation," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 24, no. 2, pp. 145–157, 2002.
- [10] E. Sommerlade and I. Reid, "Information-theoretic active scene exploration," in *IEEE Conf. on Computer Vision and Pattern Recognition, (CVPR'2008)*, Anchorage, Alaska, 2008.
- [11] M. Kumon, K. Fukushima, S. Kunitatsu, and M. Ishitobi, "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, 2010.
- [12] E. Martinson, T. Apker, and M. Bugajska, "Optimizing a reconfigurable robotic microphone array," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS'2011)*, San Francisco, California, 2011.
- [13] E. Vincent, A. Sini, and F. Charpillat, "Audio source localization by optimal control of a mobile robot," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP'2015)*, Brisbane, Australia, 2015.
- [14] G. Bustamante, A. Portello, and P. Danès, "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, 2015.
- [15] A. Portello, P. Danès, S. Argentieri, and S. Pledel, "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, 2013.
- [16] A. Portello, G. Bustamante, P. Danès, J. Piat, and J. Manhès, "Active localization of an intermittent sound source from a moving binaural sensor," in *Forum Acusticum (FA'2014)*, Krakow, Poland, 2014.
- [17] N.L. Aaronson and W. Hartmann, "Testing, correcting, and extending the woodworth model for interaural time difference," *The Journal of the Acoustical Society of America*, vol. 135, pp. 817–823, 2014.
- [18] J.M. Manyika, *An Information-Theoretic Approach to Data Fusion and Sensor Management*, Ph.D. thesis, University of Oxford, 1993.
- [19] A. Mallet, C. Pasteur, M. Herrb, S. Lemaignan, and F. Ingrand, "Genom3: Building middleware-independent robotic components," in *IEEE Int. Conf. on Robotics and Automation, (ICRA'2010)*, Anchorage, Alaska, 2010.
- [20] B. Deutsch, M. Zobel, J. Denzler, and H. Niemann, "Multi-step entropy based sensor control for visual object tracking," *Pattern Recognition*, vol. 3175, pp. 359–366, 2004.
- [21] M. Araya-López, O. Buffet, V. Thomas, and F. Charpillat, "A POMDP extension with belief-dependent rewards," in *Advances in Neural Information Processing Systems 23*, pp. 64–72. Curran Associates, Inc., 2010.