

A BINAURAL SOUND SOURCE LOCALIZATION METHOD USING AUDITIVE CUES AND VISION

Karim Youssef, Sylvain Argentieri, and Jean-Luc Zarader

Abstract—A fundamental task for a robotic audition system is sound source localization. This paper addresses the localization problem in a robotic humanoid context, providing a novel learning algorithm that uses binaural cues to determine the sound source's position. Sound signals are extracted from a humanoid robot's ears. Binaural cues are then computed to provide inputs for a neural network. The neural network uses pixel coordinates of a sound source in a camera image as outputs. This learning approach provides good localization performances as it reaches very small errors for azimuth and elevation angles estimates.

Keywords — Binaural audition, sound localization, sound processing, learning.

I. INTRODUCTION

Robots and intelligent systems are becoming more and more reliable as partners in the humans' everyday life. Nowadays, it has become possible to envision machines in social interaction. For that purpose, sound is of particular interest. Indeed, a sound signal holds various information: sound sources identities, their spatial locations and the contents of the emitted sounds. This brings to the fore multiple problems that an artificial audition system has to deal with, like Voice Activity Detection (VAD), speaker and speech recognition, sound source separation and localization.

Sound source localization has been widely studied in the last few decades. Most of the previously built systems use microphone arrays together with techniques like beamforming [14]. But microphone-array based systems are often computationally expensive, which makes it important to resort to less complex methods. In this context, binaural audition has emerged as an interesting, low-complexity and biologically-inspired sound processing domain. It is based on the use of the signals captured by only two microphones to reach human-like auditive capabilities. Indeed, binaural processing has been used in multiple applications, like sound source localization [12], [4], speech enhancement [2], and voice activity detection [2]. We have also proposed in [15] a binaural speaker recognition system, which offered better performances than conventional monaural systems, but was sensitive to the speaker position.

Binaural sound source localization systems have been conceived in multiple aspects. Some of them proposed studies that try to model the human head and to link the auditive cues to its geometry and to the sound source direction, see [4] or [5]. Thus, inverting these models could then allow to

deduce the sound source position. But when such systems are used in experimental conditions, the models fail and do not comply with reality [8]. On the contrary, learning-based approaches might be better adapted and more robust to such problems [7], [9]. In this context, this paper presents a novel sound source localization system. It provides a new way of coupling vision and sound in a learning-based approach that provides effective localization capabilities. The system uses a neural network to learn and exploit the relationship between the visualized positions of the sound source and the auditory cues extracted from two ears. This approach is inspired by the impressive human capabilities where vision plays a fundamental role in auditory scene analysis. Indeed, some recent works hypothesize some visually guided auditory adaptation processes for seer people [6]. In our approach, vision provides a tool to represent the sound source's position in the scene, and the localization, expressed in pixel positions inside an image, relies on the information provided by the auditive cues. The approach is less complex and has a better spatial resolution than related works, see [11], [7] and [9]. For now, the work disregards the hypothesis of multiple sound sources, including noises and reverberations. The feasibility of the approach, applied on white noise signals is tested and discussed in this paper, for a generalization to more realistic environments in following works, including speech and additive unwanted signals. Contrarily to many studies that only address the azimuth estimation like [7], [11], this work aims at estimating both azimuth and elevation at the same time.

The paper is organized as follows: the azimuth and elevation estimation methods are presented in the next section. Simulation and experimental datasets and tests results are presented and discussed in Section III. Finally, a conclusion ends the paper.

II. AZIMUTH AND ELEVATION ESTIMATION

A. Azimuth estimation

In all the following, a neural network is used in order to estimate the sound source position. Its inputs are code vectors composed of interaural cues: Interaural Level Differences (ILD), Interaural Phase Differences (IPD) and Interaural Time Differences (ITD). These cues are precisely described first. The neural network itself, together with the learning algorithm is depicted later, and its outputs are last introduced.

1) *Network inputs and auditory cues extraction:* As it can be seen in Figure 1, signals from both robot ears are exploited to compute the interaural auditory cues. The human cochlear filtering is artificially reproduced by a set of 20

K. Youssef, S. Argentieri and J.-L. Zarader are with UPMC Univ. Paris 06 and ISIR (CNRS UMR 7222), F-75005, Paris, FRANCE
lastname@isir.upmc.fr




Fig. 1. Auditory cues extraction diagram.

gammatone filters defined in [10]. Their central frequencies $f_c(i)$ range from 100Hz to about half of the sampling frequency $f_s = 44100$ Hz. This process leads to 20 signals per ear, the interaural cues are then extracted from these 20 signals through the following methodology.

a) *Interaural Level Difference*: The ILD is a frequency-dependent cue that reflects the difference in powers of the signals reaching the two ears. An ILD for each gammatone filter's frequency range can be extracted according to:

$$\text{ILD}(f_c(i)) = 20 \log_{10} \frac{E_l(f_c(i))}{E_r(f_c(i))}, \quad (1)$$

where $E_l(f_c(i))$ and $E_r(f_c(i))$ respectively represent the left and right cochlear filter output powers corresponding to the i^{th} gammatone response centered at frequency $f_c(i)$, $i \in [1, 20]$.

b) *Interaural Phase Difference*: IPD refers to the difference in the phases of waves reaching the ears. It is obtained with:

$$\text{IPD}(f_c(i)) = 2\pi f_c(i) \tau_{lr}(f_c(i)), \quad \text{with} \quad (2)$$

$$\tau_{lr}(f_c(i)) = k/f_s \text{ and } k = \arg \max_n (R_{lr}^{(i)}[n]),$$

where $R_{lr}^{(i)}[n] = \frac{1}{N} \sum_{m=0}^{N-n-1} l_i[m+n] r_i[m]$ is the biased estimate of the cross-correlation function between the two signals $l_i[n]$ and $r_i[n]$ originating from the i^{th} left and right gammatone filters respectively.

c) *Interaural Time Difference*: ITD reflects the difference between the lengths of the paths traveled by the sound wave before reaching the ears.

$$\text{ITD} = \frac{1}{2\pi} \mathbf{f}^+ \text{IPD}(\mathbf{f}), \quad (3)$$

where $(\cdot)^+$ denotes the Moore-Penrose pseudo-inverse, $\mathbf{f} = (f_c(1), f_c(2), \dots, f_c(N_{\text{filter}}))^T$ and $\text{IPD}(\mathbf{f}) = (\text{IPD}(f_c(1)), \dots, \text{IPD}(f_c(N_{\text{filter}})))^T$. Consequently, the ITD value is obtained by a least square operation performed on the IPD.

2) *Code vector composition*: ILDs and IPDs are known to be not salient for low and high frequencies respectively (Duplex theory by Lord Rayleigh). An observation of their patterns as a function of frequency led to taking them into consideration for frequencies higher than 1.5kHz, and lower than 3kHz respectively, as they don't carry much information outside these frequency intervals. So, the final neural network's input code vectors are composed of 13 ILDs, 12 IPDs and a single ITD value each, which makes a total input dimension of 26.

3) *Network constitution and learning algorithm*: The neural network used in this study is a feed-forward multi-layer perceptron (MLP) with one hidden layer composed of 15 cells. Since the input code vectors contain data of different types (amplitudes, phases and times), a regular complete connections neural network—i.e. a network where each hidden cell is connected to all input cells—, is not physically adapted to these inputs. But note that a more classical full-connected network could also be used. A hidden cell should not be connected to two inputs of two different physical types. Therefore, one hidden cell is dedicated to the ITD, 7 are dedicated to the ILDs and 6 to the IPDs. And the connections between the hidden cells and the outputs are kept unmodified. The training uses the full gradient backpropagation algorithm. Cross-validation steps are performed periodically, and the training is stopped when the network's performances stop improving. The amounts of training, cross-validation and testing data are specified later.

4) *Network outputs and sound source representation*: The proposed sound source is a loudspeaker carrying three colored markers. A camera mounted on the robot's head takes movies of the moving sound source and an image processing system analyses the captured images and evaluates the line and column indices of each marker's center. For the azimuth estimation, and since column variations reflect variations in the horizontal plane, the outputs of the network are only the three markers' column indices in the image.

B. Elevation estimation

Tests performed with interaural cues on elevation estimation do not show satisfactory results. Indeed, interaural cues contain powerful information about the azimuth, and very weak information about the elevation [13]. More generally, it is known that the elevation can be accessed from monaural cues. Indeed, the human pinna shape is at the origin of interferences with the waves directly entering the auditory canal, causing constructive and/or destructive reflections at specific frequencies depending of the sound source location. This phenomenon produces spectral peaks and notches which are supposed to be used by humans when evaluating the elevation of a sound source [3]. In this field, one can cite [13], where a method using spectral cues for the elevation estimation is presented, with a robot having two logarithmic-shaped reflectors as pinnas. But spectral peaks and notches are known to be very hard to extract. As a solution, we propose here to compute for each source position the energies coming from the 2 cochlear filter-banks. These energies are expected to capture the aforementioned reflections translated by high and low energies in specific spectral areas, and thus to better the elevation estimation performances of the proposed approach. A regular neural network is now used to estimate the three line coordinates of the three markers in the image. The input of this network is not made of interaural cues. It consists in the 40 energy values corresponding to the 2×20 gammatone filters, and the outputs are the three markers' line coordinates.

III. SIMULATIONS AND EXPERIMENTS

In order to evaluate the proposed approach, simulated and experimental databases have been elaborated. In both cases, the sound source emits a white discrete Gaussian noise (useful here as its spectrum spreads over a wide frequency band). The cues are extracted on the basis of 1024-points time windows lasting 23ms with a sampling frequency of $f_s = 44.1\text{kHz}$.

In both cases, the learning of the neural networks is done with 60% of the total amount of data, the cross-validation uses 20% and the remaining 20% are used for the testing. In the testing phase, the networks provide estimations of the three markers' line and column indices based on the perceived auditory inputs. The mean estimation errors of the system are defined as the mean Euclidean distances between the estimated outputs and the real ones.

A. Simulations

This subsection presents an artificially generated database and the resulting localization performances. A virtual robot is placed in an environment where a sound source is moving. For each source position, the left and right ear signals are computed and a virtual camera placed on the robot's head provides the source position in an image.

1) *Database generation*: The simulated input database is based on left and right ear signals generated from the source noise signal, convolved with impulse responses known as Head Related Impulse Responses (HRIRs) for different spatial positions. The CIPIC database provides these left and right impulse responses for various azimuths and elevations [1]. To obtain the outputs, a camera model has been simulated. It performs a 3D to 2D transformation and projects the visual markers into an image plane. Thus, their horizontal and vertical positions pix_x and pix_y in pixels are obtained. In the testing phase, the networks provide estimations of the line and column indices, $\hat{\text{pix}}_y$ and $\hat{\text{pix}}_x$ respectively. Then, an inverse of the camera model gives the corresponding estimated elevation and azimuth angles $\hat{\phi}$ and $\hat{\theta}$. The estimation errors are defined as $\epsilon_\phi = |\phi - \hat{\phi}|$ and $\epsilon_\theta = |\theta - \hat{\theta}|$ respectively.

2) *Localization results*: During the learning and testing steps, the database is restricted to angles between -45° and 45° with a 1° step for both azimuth and elevation for a total number of 8281 examples. This resolution has been reached with a spline interpolation performed on the azimuth and/or elevation impulse responses. The resulting estimations are shown in Figure 2. As expected, they show a high accuracy in the azimuth and elevation estimation, having mean errors of only 1.22° and 2.06° , and mean standard deviations of 1.53° and 1.69° respectively. So the approach is able to provide a correct estimation, but takes benefit from the precise knowledge of the sound angular position so as to compare it with the estimated one. This is rarely the case in an experimental setup, thus justifying the idea to work directly in an image.




Fig. 2. Simulation: estimation results, predicted angles as a function of the real angles. (a) azimuth angle, (b) elevation angle.

B. Experiments

This subsection presents experimental results obtained with real binaural signals recorded by using a dummy head and images provided by a camera mounted on top of it.

1) *Database*: The experimental database has been recorded in an acoustically prepared room equipped with 3D-pyramidal-pattern sound-absorbing studio foams placed on the roof and on the walls. A KU100 dummy head from Neumann is employed. It has two microphone capsules built inside two human-like ears, thus reproducing the effects of the human head and outer ears on a sound signal, before reaching the inner ear. The two microphone outputs are synchronously acquired by a National Instruments PCI acquisition board through 24 bits delta-sigma converters operating at a sampling frequency $f_s = 44.1\text{kHz}$. A camera from Baumer is placed on top of the head, and provides 44 photos per second with a 640×480 resolution. This frame rate is selected to easily synchronize each image with an approximately 23ms sound frame. A small portable round loudspeaker with a frequency response ranging from 200Hz to 16kHz is used to emit a white Gaussian noise. 3 colored patches are stuck in front of it, and an image processing algorithm gives the coordinates of the centers of the patches. During a recording, a person holds the loudspeaker emitting the noise in the camera field of view, and moves it in different directions (left, right, up, or down).

2) *Localization results*: As already pointed out, the exact relative angular location of each marker with respect to the head does not need to be known when working directly in the image. So, the experiments will be assessed by comparing the actual pixel coordinates pix_y and pix_x to the predictions $\hat{\text{pix}}_y$ and $\hat{\text{pix}}_x$ produced by the neural networks thanks to the auditive cues. Having three points, the mean real and estimated coordinates are compared, so as to compare the real and estimated loudspeaker centers. Note that the results presented in this section are obtained on a 12s-

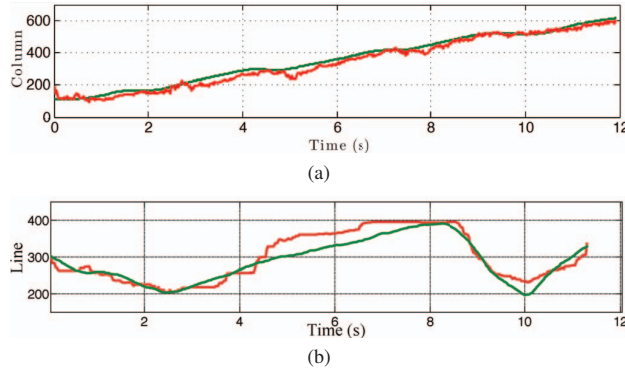


Fig. 3. Experiments: estimation results, predicted dimensions as a function of time. (a) columns, (b) lines.




Fig. 4. Experiments: a truncated view of an image taken by the camera. It shows a predicted trajectory made by the sound source (blue) and the corresponding real trajectory (red).

recording during which the sound source moves in the image. Estimations are reported in Figure 3. They show that the predicted pixel coordinates follow the real ones quite well, while the column estimation results are better than the line estimation results, which is inline with the simulation results showing a better performance in the azimuth case. Figure 4 shows a comparison between a predicted trajectory and a real one. It can be seen that the system localizes the target source quite accurately. The observed differences between the two trajectories are mainly caused by the line coordinates having bigger estimation errors than the columns.

C. Discussion

The tests made on the simulated and the experimental databases lead to the same conclusions: the system is able to efficiently estimate the position of the sound source, with better performances in the azimuth estimation than in the elevation estimation. Compared to related works, this system has a higher resolution and is less complex. For example, one can cite [11] where a parametric model computing ILDs and IPDs as a function of the azimuth is used, and these cues are inverted to deduce the azimuth. But the estimation errors are higher than those obtained in our study, and the resolution is weaker (5 degrees in azimuth). Also in [7] and [9], the binaural systems rely on probabilistic approaches needing large databases and computational capabilities and times, and have position resolutions of 5° .

IV. CONCLUSION

A sound source localization system has been presented. It deals with the localization problem in a new learning fashion using cues extracted from both human-like ears of a humanoid robot and visual information from a camera placed on its head. While the interaural cues provided very satisfactory results for azimuth estimation, output energies from a set of cochlear filters allowed to efficiently determine the source's elevation. Nevertheless, the described work provides an efficient tool in adequate acoustic conditions. Ongoing works are now aiming at generalizing the tests to more complex situations, with noises and reverberations, together with human voice sound signals.

ACKNOWLEDGMENT

This work was conducted within the French/Japan BINAHR (BINAural Active Audition for Humanoid Robots) project under Contract n° ANR-09-BLAN-0370-02 funded by the French National Research Agency.

REFERENCES

- [1] V. Algazi, R. Duda, R. Morisson, and D. Thompson. The cipc hrtf database. *Proceedings of the 2001 IEEE Workshop on Applications of Signal Processing to audio and Acoustics*, pages pp. 99–102, 2001.
- [2] R. Brueckmann, A. Scheidig, and H.-M. Gross. Adaptive noise reduction and voice activity detection for improved verbal human-robot interaction using binaural data. *IEEE International Conference on Robotics and Automation*, April 2007.
- [3] J. Garas. Adaptive 3d sound systems. *Kluwer*, 2000.
- [4] A. A. Handzel and P. Krishnaprasad. Biomimetic sound-source localization. *IEEE Sensors Journal*, 2:607–616, 2002.
- [5] M. Jeub, M. Dörbecker, and P. Vary. A semi-analytical model for the binaural coherence of noise fields. *IEEE Signal Processing Letters*, 18(3), March 2011.
- [6] J. Lewald and S. Getzmann. Horizontal and vertical effects of eye-position on sound localization. *Hearing Research*, 213(1-2):99–106, Mar 2006.
- [7] T. May, S. van de Par, and A. Kohlrausch. A probabilistic model for robust localization based on a binaural auditory front-end. *IEEE Transactions on Audio, Speech and Language Processing*, 19(1), 2011.
- [8] K. Nakadai, D. Matsuura, H. G. Okuno, and H. Kitano. Applying scattering theory to robot audition system: Robust sound source localization and extraction. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2003.
- [9] J. Nix and V. Hohmann. Sound source localization in real sound fields based on empirical statistics of interaural parameters. *Journal of the Acoustic Society of America*, 119(1), 2006.
- [10] R. Patterson, K. Robinson, J. Holdsworth, D. McKeown, C. Zhang, and M. Allerhand. Complex sounds and auditory images. In *International Symposium on Hearing, Auditory physiology and perception*, pages 429–446, 1992.
- [11] M. Raspaud, H. Viste, and G. Evangelista. Binaural source localization by joint estimation of ild and itd. *IEEE Transactions on Audio, Speech and Language Processing*, 18(1), 2010.
- [12] T. Rodemann, G. Ince, F. Joubin, and C. Goerick. Using binaural and spectral cues for azimuth and elevation localization. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, September 2008.
- [13] T. Shimoda, T. Nakashima, M. Kumon, R. Kohzawa, Z. Iwai, and M. Iwai. Spectral cues for robust sound localization with pinnae. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2006.
- [14] J.-M. Valin, F. Michaud, and J. Rouat. Robust 3d localization and tracking of sound sources using beamforming and particle filtering. *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP Proceedings*, 2006.
- [15] K. Youssef, S. Argentieri, and J.-L. Zarader. From monaural to binaural speaker recognition for humanoid robots. In *IEEE-RAS International Conference on Humanoid Robots*, pages 580 – 586, Dec. 2010.