

ABSOLUTE GEOMETRY CALIBRATION OF DISTRIBUTED MICROPHONE ARRAYS IN AN AUDIO-VISUAL SENSOR NETWORK

Florian Jacob, Reinhold Haeb-Umbach

Department of Communications Engineering, University of Paderborn, Germany

{jacob, haeb}@nt.uni-paderborn.de

ABSTRACT

Joint audio-visual speaker tracking requires that the locations of microphones and cameras are known and that they are given in a common coordinate system. Sensor self-localization algorithms, however, are usually separately developed for either the acoustic or the visual modality and return their positions in a modality specific coordinate system, often with an unknown rotation, scaling and translation between the two. In this paper we propose two techniques to determine the positions of acoustic sensors in a common coordinate system, based on audio-visual correlates, i.e., events that are localized by both, microphones and cameras separately. The first approach maps the output of an acoustic self-calibration algorithm by estimating rotation, scale and translation to the visual coordinate system, while the second solves a joint system of equations with acoustic and visual directions of arrival as input. The evaluation of the two strategies reveals that joint calibration outperforms the mapping approach and achieves an overall calibration error of 0.20 m even in reverberant environments.

Index Terms— coordinate mapping, absolute geometry calibration

1. INTRODUCTION

Advanced teleconferencing systems, smart rooms or surveillance and monitoring systems are example applications of distributed audio-visual sensor networks. For many tasks, such as automatic camera steering, events or objects of interest have to be localized either acoustically, visually or jointly, which in turn requires that the positions of the sensors need to be known. While the sensor positions can be determined manually, it is more convenient to do so automatically, in particular if they can change over time, e.g., if a smartphone, which is part of the network, is carried by a moving person.

Automatic geometry calibration of sensors is typically realized by localizing and tracking an object and subsequently determining the position of the sensors, such that the measurements of the object's positions are most plausible.

Visual calibration algorithms work on features extracted from the camera images. They can be divided into two categories [1]. The first one tries to extract these features from easily recognizable objects [2], whereas the second group extracts features from an arbitrary scene to compare the field of view of the individual cameras [3].

For acoustic sensor nodes time of flight (ToF) based algorithms, which employ special calibration hardware and signals to achieve high positioning accuracies [4, 5], have been proposed. However, a tight clock synchronization between transmitter and receiver is required, whereas [6] relaxed this limitation by estimating the differences in the sampling phase and the sensor positions jointly. If the calibration is based on time difference of arrival (TDoA) measurements, loudspeaker and microphones need no longer be synchronized, and a human speaker can be used as sound source. However, a clock synchronisation of the A/D converters of the distributed microphones is still required. Even this requirement becomes obsolete if direction of arrival (DoA) based techniques are employed [7, 8]. TDoA and DoA based calibration if carried out with artificial calibration signals with appropriate correlation properties, will typically achieve higher accuracy compared to speech signal based approaches [9]. Calibration based on natural speech is preferable from a usability point of view, as it can be carried out in the background unnoticed by the users of the audio-visual sensor network.

Most geometry calibration techniques are unable to report the sensor positions in absolute coordinates. They return their estimates in a modality specific coordinate system, resulting in an unknown rotation, translation and scaling between the coordinate axes of the acoustic and visual sensor network. The scaling ambiguity can be fixed if ToA or TDoA measurements are employed [10, 11]. If the calibration is based solely on DoA measurements, the scale ambiguity still remains, regardless of the modality used [11, 12].

If the sensor positions of one modality are known, the displacement between the coordinate systems can be resolved by exploiting audio-visual correlates, i.e., events or objects that can be localized both acoustically and visually [13, 14]. In this paper we build upon this idea and present two strategies to localize the acoustic sensors in a joint audio-visual coordinate system. Both the acoustic and visual localization is

This work has been supported by Deutsche Forschungsgemeinschaft (DFG) under contract no. Ha3455/7-2.

solely based on DoA measurements as they impose the least synchronisation requirements as detailed above.

The first approach uses the existing acoustic sensor calibration techniques from [7]. Based on the relative geometry estimates the speaker trajectory can be recovered with the intersection based approach from [15], while simultaneously the speaker trajectory is estimated by the visual sensor network. By computing the optimal mapping between the acoustic and visual trajectory we are able to reveal rotation, translation and scale between both modalities. The second approach exploits the fact that the sensors of both modalities deliver DoA estimates. Thus, acoustic and visual measurements can be cast in a single system of equations to determine the acoustic sensor positions, while the known visual sensor positions serve as anchor positions to eliminate the scale ambiguity. A key component of both DoA based calibration methods is the random sample consensus (RANSAC) outlier rejection algorithm [16], which diminishes the impact of poor DoA estimates on the localization performance. In case of the joint calibration, this scheme will not only reject acoustic DoA outliers, it will reject visual DoA outliers as well.

In the next section we introduce the first approach based on a coordinate mapping, whereas Sec. 3 describes the joint calibration approach. The performance of both algorithms is evaluated in Sec. 4, before Sec. 5 concludes this paper.

2. COORDINATE MAPPING

Our goal is the estimation of the coordinates of I acoustic sensors, where the coordinate system is defined by the known positions of K visual sensors. The location of the k -th visual sensor node is described in 2D by the position vector \mathbf{c}_k and orientation γ_k . Now, consider a moving speaker located at position \mathbf{e}_t at time t , who is seen by the visual sensors at DoAs $\delta_{k,t}$, $k=1, \dots, K$. A position estimate \mathbf{e}_t is obtained from the DoAs by the intersection based technique presented in [15].

The acoustic DoA estimates $\varphi_{i,t}$, $i=1, \dots, I$; $t=1, \dots, T$, captured from the same speaker trajectory are used to determine estimates $\tilde{\mathbf{m}}_i$, $i=1, \dots, I$, of the acoustic sensor positions and estimates Θ_i of the orientations, using the calibration algorithm from [7]. This algorithm can only provide a relative geometry with an unknown scale factor. Therefore, only relative speaker position estimates $\tilde{\mathbf{e}}_t$ are obtained, using the same intersection based method as above.

Since the acoustic event locations $\tilde{\mathbf{e}}_t$ are described in a different coordinate system as the visual estimates \mathbf{e}_t , there arises the following coordinate mapping problem:

$$\mathbf{e}_t = s\mathbf{R}\tilde{\mathbf{e}}_t + \mathbf{d}; \quad t = 1, \dots, T, \quad (1)$$

where s models the unknown scale factor and \mathbf{R} and \mathbf{d} the rotation and translation between the acoustic and the visual coordinate system. Mapping a set of points from one coordinate system to another is known as Rigid Body Transforma-

tion (RBT). In contrast to the widespread approach from [17] to compute the RBT parameters (scale, rotation and translation) via a Singular Value Decomposition (SVD) we suggest a computation in the Discrete Fourier transform (DFT) domain, which turned out to be computationally more efficient.

Hence, we introduce a complex representation of the estimated speaker positions as $u_t = \tilde{e}_{1,t} + j\tilde{e}_{2,t}$ and $v_t = e_{1,t} + je_{2,t}$ respectively, where $\tilde{\mathbf{e}}_t = (\tilde{e}_{1,t}, \tilde{e}_{2,t})$ and $\mathbf{e}_t = (e_{1,t}, e_{2,t})$ are the two-dimensional speaker positions in the acoustic and visual coordinate system, respectively. Thus, the mapping problem of Eq. (1) is expressed as

$$v_t = \alpha u_t + \beta; \quad \alpha, \beta \in \mathbb{C}. \quad (2)$$

The absolute value and the phase of α correspond to scale and orientation, while β corresponds to the translation. Arranging all observations into vectors $\mathbf{v} = [v_1, \dots, v_T]^T$ and $\mathbf{u} = [u_1, \dots, u_T]^T$ the least squares estimate for the RBT parameters in the complex space is given by

$$\langle \alpha^*, \beta^* \rangle = \underset{\alpha, \beta}{\operatorname{argmin}} (\alpha \mathbf{u} + \beta \mathbf{1} - \mathbf{v})^H (\alpha \mathbf{u} + \beta \mathbf{1} - \mathbf{v}), \quad (3)$$

where $\mathbf{1}$ denotes an T -element vector of ones and $(\cdot)^H$ the complex conjugate transpose of a vector.

Let \mathbf{x} and \mathbf{y} denote the DFTs of \mathbf{u} and \mathbf{v} . The optimization problem of Eq. (3) is expressed in the DFT domain as

$$\langle \alpha^*, \beta^* \rangle = \underset{\alpha, \beta}{\operatorname{argmin}} (\alpha \mathbf{x} + \beta \mathbf{z} - \mathbf{y})^H (\alpha \mathbf{x} + \beta \mathbf{z} - \mathbf{y}), \quad (4)$$

where $\mathbf{z} = [1, 0, \dots, 0]^T$ is a vector of length T . Due to the orthogonality properties of the DFT the joint optimization is decoupled into two separate optimizations:

$$\alpha^* = \underset{\alpha}{\operatorname{argmin}} (\alpha \mathbf{x}_{2:T} - \mathbf{y}_{2:T})^H (\alpha \mathbf{x}_{2:T} - \mathbf{y}_{2:T}) \quad \text{and} \quad (5)$$

$$\beta^* = \underset{\beta}{\operatorname{argmin}} (\alpha^* x_1 + \beta - y_1)^H (\alpha^* x_1 + \beta - y_1), \quad (6)$$

where the first bin of the DFTs is denoted by $(\cdot)_1$ and all other bins by $(\cdot)_{2:T}$. Since Eq. (5) and Eq. (6) are general least squares problems, the solution is found to be

$$\alpha^* = \mathbf{x}_{2:T}^H \mathbf{y}_{2:T} / (\mathbf{x}_{2:T}^H \mathbf{x}_{2:T}) \quad \text{and} \quad \beta^* = y_1 - \alpha^* x_1. \quad (7)$$

The RBT parameters can be retrieved as follows:

$$s = |\alpha|, \quad \mathbf{R} = \begin{bmatrix} \Re\{\frac{\alpha}{s}\} & -\Im\{\frac{\alpha}{s}\} \\ \Im\{\frac{\alpha}{s}\} & \Re\{\frac{\alpha}{s}\} \end{bmatrix} \quad \text{and} \quad \mathbf{d} = \frac{1}{N} \begin{bmatrix} \Re\{\beta\} \\ \Im\{\beta\} \end{bmatrix}, \quad (8)$$

where \Re and \Im denote real and imaginary part, respectively. If this transformation is applied to the relative acoustic sensor position estimate $\tilde{\mathbf{m}}_i$ according to Eq. (1), the absolute acoustic sensor positions \mathbf{m}_i in the visual coordinate system are obtained.

To summarize, the calibration algorithm to recover the acoustic sensor positions in the visual coordinate system consists of three steps. First, run the relative acoustic calibration

algorithm and estimate the speaker trajectory. At the same time, track the speaker in the visual domain. Secondly, compute the DFTs of both trajectories, evaluate Eq. (7) and compute the RBT parameters by Eq. (8). Finally, use the RBT parameters to transform the acoustic sensor position estimates from the first step into the visual coordinate system.

The DFT based RBT parameter estimation delivers the same results as the conventional SVD based technique [17], but our FFT based implementation is twice as fast as the SVD.

3. JOINT CALIBRATION

Since the acoustic and the visual sensors deliver DoA estimates, we propose to extend the calibration algorithm that was used for the acoustic sensors only in step one of the algorithm presented in the last section, to both modalities and jointly calibrate the audio-visual network. Due to the known positions of the visual sensors, the scale ambiguity vanishes.

In the local coordinate system of the i -th acoustic sensor a DoA measurement can be modelled as a unit length vector

$$\mathbf{f}_{i,t} = [\cos(\varphi_{i,t}) \quad \sin(\varphi_{i,t})]^T, \quad (9)$$

pointing from the sensor position to the event location. This measurement vector will be compared with a prediction vector

$$\hat{\mathbf{f}}_{i,t} = \begin{bmatrix} \cos(\hat{\varphi}_{i,t} - \Theta_i) \\ \sin(\hat{\varphi}_{i,t} - \Theta_i) \end{bmatrix}, \quad (10)$$

where $\hat{\varphi}_{i,t} = \arg\{\mathbf{e}_t - \mathbf{m}_i\}$, see Fig. 1. Following our previous publication [7] this prediction can be formulated as a function of the geometry parameters as follows:

$$\hat{\mathbf{f}}_{i,t} = \begin{bmatrix} \cos(\Theta_i) & \sin(\Theta_i) \\ -\sin(\Theta_i) & \cos(\Theta_i) \end{bmatrix} \frac{\mathbf{e}_t - \mathbf{m}_i}{\|\mathbf{e}_t - \mathbf{m}_i\|}. \quad (11)$$

By introducing the abbreviation

$$f = \sum_{i=1}^I \sum_{t=1}^T \|\mathbf{f}_{i,t}^T \hat{\mathbf{f}}_{i,t}\|^2 \quad (12)$$

and arranging the sensor positions, sensor orientations and events into matrices $\mathbf{M} = [\mathbf{m}_1, \dots, \mathbf{m}_I]$, $\Theta = [\Theta_1, \dots, \Theta_I]$




Fig. 1. Geometric relation between acoustic sensor and event location.

and $\mathbf{E} = [\mathbf{e}_1, \dots, \mathbf{e}_T]$ respectively, the geometry can be recovered by

$$\langle \mathbf{M}^*, \Theta^*, \mathbf{E}^* \rangle = \underset{\mathbf{M}, \Theta, \mathbf{E}}{\operatorname{argmax}} \{f\}. \quad (13)$$

The maximization problem of Eq. (13) can easily be transformed into a root-finding problem, since $\mathbf{f}_{i,t}$ and $\hat{\mathbf{f}}_{i,t}$ are unit length vectors. Subsequently, the minimization is carried out by Newton's method.

The formulation for the estimated and predicted DoA vectors hold for the visual sensors, too. Thus we define

$$\mathbf{g}_{k,t} = [\cos(\delta_{k,t}) \quad \sin(\delta_{k,t})]^T \text{ and} \quad (14)$$

$$\hat{\mathbf{g}}_{k,t} = \begin{bmatrix} \cos(\gamma_k) & \sin(\gamma_k) \\ -\sin(\gamma_k) & \cos(\gamma_k) \end{bmatrix} \frac{\mathbf{e}_t - \mathbf{c}_k}{\|\mathbf{e}_t - \mathbf{c}_k\|}, \quad (15)$$

with the only difference that the visual sensor positions \mathbf{c}_k and the corresponding orientations γ_k are known. Hence, the visual DoA measurements form additional constraints for the optimization of Eq. (13), and we incorporate them to obtain a formulation which allows a joint audio-visual calibration:

$$\langle \mathbf{M}^*, \Theta^*, \mathbf{E}^* \rangle = \underset{\mathbf{M}, \Theta, \mathbf{E}}{\operatorname{argmax}} \left\{ f + \sum_{k=1}^K \sum_{t=1}^T \|\mathbf{g}_{k,t}^T \hat{\mathbf{g}}_{k,t}\|^2 \right\}. \quad (16)$$

The optimization of Eq. (16) is again turned into a root-finding problem in order to apply Newton's method, where the visual measurements provide the required constraints to obtain an absolute sensor position estimate in the coordinate system defined by the visual sensors.

In the noise free case, with perfect DoA measurements, the sensor positions and orientations can perfectly be recovered, but imperfect acoustic or visual DoA estimates caused by reverberation or false detections can prevent a successful optimization. Our earlier investigations presented in [7] showed, that this issue can successfully be addressed by the RANSAC [16]. Since the application of the RANSAC is straightforward we highlight only the relevant parts. The procedure can be summarized as follows:

1. Randomly select the minimal number of observations necessary to solve Eq. (16), e.g. $T > \frac{3I}{I+K-2}$.
2. Determine sensor positions and orientations based on the selected observations by solving Eq. (16).
3. Compute the intersection of all DoA axes for each event. The hypothesized event location is the mean of all intersections. A DoA measured by a sensor becomes part of the candidate set \mathcal{C} , if the average distance of all its intersection points to the hypothesized event location is smaller than a threshold.
4. If the number of elements in \mathcal{C} is larger than the consensus set, estimate the sensor positions and orientations based on \mathcal{C} . It becomes the new consensus if its error is smaller than the error of the current consensus.
5. If the number of elements in \mathcal{C} is smaller than consensus set, choose a new initial set or stop the algorithm as soon as the maximum number of iterations is reached.

As a modification of this standard approach, we used the updated consensus set of step 4 as the input for the second step.

4. SIMULATION RESULTS

In order to evaluate the performance of both calibration strategies we used the following simulation framework. We simulated 3 random speaker trajectories, where the speaker stops at approximately 140 positions for 5 seconds before he moves on. The sensors are located in a room of size $6.2\text{ m} \times 7.2\text{ m}$. 4 simulated cameras and 4 simulated five-element circular microphone arrays (radius 5 cm) are located sufficiently far apart from the walls, where the cameras were oriented towards the center of the room. The microphone signals are generated by the Image Method [18], for reverberation times from 0 ms up to 500 ms. Acoustic DoA estimates are obtained by correlating the filter impulse responses of a filter-and-sum beamformer, which continuously adapts to the moving source [19].

Rather than working on a true camera signal, visual DoA estimates are simulated as follows. We employ Hidden Markov Models (HMMs) to describe the errors in the DoA estimation. A limited field of view of the camera is taken into account by dropping all angles outside a window of $\pm 30^\circ$ relative to the camera orientation. This effect is modelled by two separate HMMs. The first HMM is for the case that a speaker is inside the visible region of the camera. Here we distinguish the states 'detection', 'missed detection' and 'false detection'. The second HMM models the case that no speaker is inside the visible region. It incorporates the states 'false detection' and 'no detection'. The transition probabilities of these models and the variance of the error distribution have been learned by computing histograms of oriented gradients (HOG) and applying a support vector machine (SVM) to identify the head and shoulder region of the speaker on the AV16.3 audio-visual corpus [20], using the annotated sequences *seq01-1p-0000* and *seq15-1p-0100*.

In order to perform a fair comparison between the approaches presented in Sec. 2 and Sec. 3 the estimation of the RBT parameters is embedded into a RANSAC framework, too, since we have shown in [13] that the RANSAC can boost the performance of the estimation of the RBT parameters.

Since the RANSAC is a random process, we average over multiple runs. A sensor configuration is characterized by the positions and orientations. Fig. 2 compares the mean positioning error (MPE) of the coordinate mapping based calibration (RBT) and the joint calibration strategy (Joint). It can be observed that the joint calibration clearly outperforms the RBT approach, in particular at low reverberation times. Obviously, it is advantageous to avoid premature decisions on acoustic source and sensor positions until the visual information is accounted for, as it is done in the joint calibration approach.

The coordinate mapping approach has limited capabilities to determine a precise scale factor, and errors in the scale factor dominate its performance. In order to isolate scale fac-

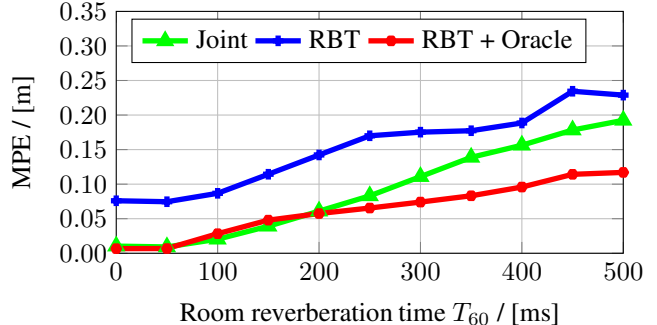


Fig. 2. Comparison of mean positioning error (MPE) for joint audio-visual calibration (joint), calibration by coordinate mapping (RBT) and coordinate mapping with an oracle information (RBT + oracle).

tor estimation errors from orientation and translation errors we performed an oracle experiment, where the scaling is assumed to be known. Indeed, the performance is now similar to that of the joint approach for low reverberation times and superior in a highly reverberant environment. The sensor orientation error of both approaches is approximately the same and smaller than 2° for all reverberation times.

To achieve precise calibration results a suitable spatial event configuration is more important than the total number of available events. Thus, we selected 15 events with an appropriate configuration of one exemplar trajectory and perform a joint calibration. The results of Tab. 1 show that a similar performance as in the previous experiment, which used the complete trajectory, is possible.

T_{60} / ms	0	100	200	300	400	500
MPE / m	0.01	0.02	0.06	0.14	0.13	0.23

Table 1. Joint calibration using 15 events with appropriate spatial configuration.

5. CONCLUSIONS

We have described two different strategies to obtain an absolute calibration of an acoustic sensor network if it is combined with a visual sensor network, whose sensor positions are known. By using one of the two strategies, the scaling problem identified in earlier publications [3, 7, 11] can be solved. The first approach, which relies on the mapping of an acoustic to a visual speaker trajectory, works with arbitrary acoustic calibration strategies and is therefore very flexible. However, the performance is limited due to the scale estimation errors. The second approach, which is based on the solution of a system of nonlinear equations employing acoustic and visual DoA measurements, is computationally more complex. It outperformed the first approach for all reverberation times and delivered a calibration error smaller than 0.20 m and 2° even in reverberant environments.

6. REFERENCES

- [1] M. Brückner, F. Bajramovic, and J. Denzler, "Self-Calibration of Camera Networks: Active and Passive Methods," in *Emerging Topics in Computer Vision and its Applications*, pp. 447–469. Dec. 2011.
- [2] X. Chen, J. Davis, and P. Slusallek, "Wide Area Camera Calibration Using Virtual Calibration Objects," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2000, pp. 2520–2527.
- [3] M. Brückner and J. Denzler, "Active Self-calibration of Multi-camera Systems," in *DAGM Symposium on Pattern Recognition*, Sept. 2010, pp. 31–40.
- [4] J. M. Sachar, H. F. Silverman, and W. R. Patterson, "Microphone position and gain calibration for a large-aperture microphone array," *IEEE Transactions on Speech and Audio Processing*, pp. 42–52, Jan. 2005.
- [5] M. Crocco, A. Del Bue, M. Bustreo, and V. Murino, "A Closed Form Solution to the Microphone Position Self-Calibration Problem," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, Mar. 2012, pp. 2597–2600.
- [6] N. D. Gaubitch, W. B. Kleijn, and R. Heusdens, "Auto-localization in ad-hoc microphone arrays," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, May 2013, pp. 106–110.
- [7] F. Jacob, J. Schmalenstroeer, and R. Haeb-Umbach, "Microphone Array Position Self-Calibration from Reverberant Speech Input," in *International Workshop on Acoustic Signal Enhancement*, Sept. 2012.
- [8] F. Jacob, J. Schmalenstroeer, and R. Haeb-Umbach, "DOA-based microphone array position self-calibration using circular statistics," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, May 2013, pp. 116–120.
- [9] S.D. Valente, M. Tagliasacchi, F. Antonacci, P. Bestagini, A. Sarti, and S. Tubaro, "Geometric calibration of distributed microphone arrays from acoustic source correspondences," in *IEEE International Workshop on Multimedia Signal Processing*, Oct. 2010, pp. 13–18.
- [10] M. H. Hennecke and G. A. Fink, "Towards Acoustic Self-Localization of Ad Hoc Smartphone Arrays," in *Hands-Free Speech Communication and Microphone Arrays*, May 2011, pp. 127–132.
- [11] J. Schmalenstroeer, F. Jacob, R. Haeb-Umbach, M. Hennecke, and Gernot A. Fink, "Unsupervised Geometry Calibration of Acoustic Sensor Networks Using Source Correspondences," in *Interspeech*, Aug. 2011, pp. 597–600.
- [12] M. Brückner, F. Bajramovic, and J. Denzler, "Geometric and probabilistic image dissimilarity measures for common field of view detection," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2009, pp. 2052–2057.
- [13] F. Jacob and R. Haeb-Umbach, "Coordinate Mapping Between an Acoustic and Visual Sensor Network in the Shape Domain for a Joint Self-Calibrating Speaker Tracking," in *11. ITG Fachtagung Sprachkommunikation*, Sept. 2014.
- [14] A. Plinge and Gernot A. Fink, "Geometry Calibration of Distributed Microphone Arrays exploiting audio-visual Correspondences," in *European signal processing conference*, Sept. 2014.
- [15] M.S. Brandstein, J.E. Adcock, and H.F. Silverman, "A closed-form method for finding source locations from microphone-array time-decay estimates," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, May 1995, pp. 3019–3022.
- [16] M. Fischler and R. Bolles, "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography," *Communications of the ACM*, pp. 381–395, June 1981.
- [17] J. H. Challis, "A procedure for determining rigid body transformation parameters," *Journal of Biomechanics*, 1995.
- [18] J. B. Allen and D. A. Berkley, "Image method for efficiently simulating small-room acoustics," *Journal Acoustical Society of America*, 1979.
- [19] E. Warsitz and R. Haeb-Umbach, "Acoustic filter-and-sum beamforming by adaptive principal component analysis," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Mar. 2005, pp. 797–800.
- [20] G. Lathoud, J.-M. Odobez, and D. Gatica-Perez, "AV16.3: an Audio-Visual Corpus for Speaker Localization and Tracking," in *Workshop for Machine Learning for Multimodal Interaction*, July 2005.