forum acusticum 2026
8-12 September · Graz, Austria

DIRECTIONAL MIXTURE MODELS OF CONTINUOUS PLANE-WAVE DISTRIBUTIONS FOR SOUND FIELD ESTIMATION

Matthias Blochberger[1], Toon van Waterschoot[1]

[1]*Department of Electrical Engineering (ESAT), STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, KU Leuven, Belgium*

Abstract

Plane-wave-based sound field estimation commonly relies on discretising the plane-wave domain into a finite set of propagation directions, which can lead to high-dimensional models and unfavourable conditioning in sparse inverse problems. In this work, we investigate continuous directional latent models by parameterising the plane-wave distribution on the unit sphere with compact mixtures. We consider two component families within the same framework: von Mises–Fisher (vMF) components, which provide an intrinsic isotropic model on the sphere, and tangent-space Gaussian (TG) components, which provide explicit covariance control and analytically convenient local approximations. Inserting these models into the Herglotz plane-wave representation yields compact sound-field models whose atoms are tied to continuous directional distributions rather than to a fixed propagation grid. The unknown amplitudes and directional parameters are estimated from sparse pressure measurements using a variable-projection style procedure: amplitudes are refit by ridge least squares for fixed directional parameters, while directions and spreads or covariances are optimised iteratively. The comparison is motivated by the trade-off between intrinsic spherical fidelity for vMF mixtures and local anisotropic flexibility for TG. Numerical experiments in simulated room-acoustic fields compare reconstruction accuracy and parameter efficiency against ridge and sparse fixed-grid plane-wave estimation. The results show improved reconstruction at matched intrinsic parameter counts in several settings for both methods, with vMF mixtures giving the most consistent high-frequency gains, at the cost of iterative nonlinear fitting.

Corresponding author: *matthias.blochberger@kuleuven.be*

Keywords: *Helmholtz equation, Gaussian mixture model, plane-wave decomposition, sound field estimation, sparse approximation*

1. Introduction

Sound-field estimation in source-free regions is commonly approached by expanding the pressure field in functions that satisfy the Helmholtz equation [1], such as spherical harmonics (SHs) [2], equivalent sources [3], or plane waves [4]. SH methods provide a continuous-region representation but are tied to truncation order and array geometry, while equivalent-source and Helmholtz-equation least-squares (LS) formulations represent the field through physically motivated basis functions and regularized inverse problems. Fixed-grid plane-wave decomposition (PWD) is particularly attractive because it leads to a linear coefficient-estimation problem, but it requires discretizing the sphere of propagation directions. This discretization can introduce basis mismatch, whereas very dense grids increase dictionary coherence and computational cost.

Sparse reconstruction and compressive-sensing methods [5], [6] address the limited-measurement setting by exploiting sparsity in a chosen spatial basis, but their performance depends strongly on basis construction and regularization. Recent work has therefore explored spatial bases tailored to source geometry, sparse Bayesian formulations, and adaptive direction refinement to reduce plane-wave basis mismatch. Our approach follows the same motivation of improving parameter efficiency, but instead of selecting or refining individual fixed-grid directions, it parameterizes the continuous plane-wave distribution by a small number of adaptive directional mixture components.

A complementary line of work uses kernel [4], Gaussian-process [7], or neural models with physical constraints [8]. Helmholtz-constrained kernel ridge regression and Gaussian-process sound-field reconstruction avoid generic interpolation kernels by using function spaces compatible with acoustic wave propagation, while recent physics-informed neural networks impose wave or Helmholtz constraints through training objectives or boundary-integral formulations. In contrast, the present method remains explicitly parametric: each atom is generated by inserting a directional distribution into the plane-wave representation, so the reconstructed field is Helmholtz-compatible by construction and the model dimension is directly controlled by the number of mixture components.

In this work, we replace the fixed directional grid by a small number of adaptive continuous directional components. We consider spherically isotropic von Mises–Fisher (vMF) distributions as an intrinsic component family, and tangent-space Gaussian (TG) distributions as a local anisotropic alternative. The aim is improved reconstruction accuracy at a fixed intrinsic parameter count. The main contributions are: (i) the formulation of continuous directional mixture models based on vMF and TG atoms within the plane-wave representation, (ii) a variable-projection style estimation procedure for fitting the resulting nonlinear models, and (iii) an empirical comparison against ridge and sparse fixed-grid PWD in randomized image-source room simulations.

2. Method

Let $p : \Omega \subset \mathbb{R}^3 \to \mathbb{C}$ be the sound field at $k = \frac{\omega}{c}$, where Ω is the source-free region of interest within a room. Standard plane-wave-based methods solving the Helmholtz equation

$$(\Delta + k^2)p(\mathbf{x}) = 0, \quad x \in \Omega, \tag{1}$$

rely on the discretization of

$$p(\mathbf{x}) = \int_{\mathbb{S}^2} a(\mathbf{s}) e^{ik\mathbf{s}^\top \mathbf{x}} \, \mathrm{d}\mathbf{s} \tag{2}$$

with the plane wave distribution $a(\mathbf{s})$, into a finite number of plane-wave directions,

$$p(\mathbf{x}) \approx \sum_{n=1}^{Q} a_n e^{ik\mathbf{s}_n^\top \mathbf{x}}, \tag{3}$$

where a_n are the plane wave coefficients and $\mathbf{s}_n$ are the plane wave directions.

Here, we formulate a model that is based on estimation of continuous directional distribution $a(\mathbf{s})$ approximated by a mixture

$$a(\mathbf{s}) \approx \sum_{n=1}^{N} \alpha_n \rho_n(\mathbf{s}; \boldsymbol{\theta}_n), \tag{4}$$

where α_n are the mixture weights, $\boldsymbol{\theta}_n$ are the parameters of the mixture components, and ρ_n is a directional atom. The field can then be approximated as

$$p(\mathbf{x}) \approx \sum_{n=1}^{N} \alpha_n \psi_n(\mathbf{x}), \tag{5}$$

with

$$\psi_n(\mathbf{x}) = \int_{\mathbb{S}^2} \rho_n(\mathbf{s}; \boldsymbol{\theta}_n) e^{ik\mathbf{s}^\top \mathbf{x}} \, \mathrm{d}\mathbf{s}. \tag{6}$$

Subsequently, we define two comparable types of atoms describing directional components.

2.1. Von Mises-Fisher Distribution

Let us define the vMF as the intrinsic spherical component family on $\mathbb{S}^2$. It is defined as

$$\rho_n^{\mathrm{vMF}}(\mathbf{s}; \boldsymbol{\theta}_n) = c(\kappa_n) e^{\kappa_n \boldsymbol{\mu}_n^\top \mathbf{s}}, \quad \mathbf{s} \in \mathbb{S}^2, \tag{7}$$

with $\boldsymbol{\theta}_n = (\boldsymbol{\mu}_n, \kappa_n)$, where the normalization constant is

$$c(\kappa_n) = \frac{\kappa_n}{4\pi \sinh \kappa_n}, \quad \kappa_n \geq 0. \tag{8}$$

By continuity, $\kappa_n = 0$ corresponds to the uniform distribution on $\mathbb{S}^2$.

The field atom induced by one vMF component is

$$\psi_n^{\mathrm{vMF}}(\mathbf{x}) = c(\kappa_n) \int_{\mathbb{S}^2} e^{(\kappa_n \boldsymbol{\mu}_n + ik\mathbf{x})^\top \mathbf{s}} \, \mathrm{d}\mathbf{s} \tag{9}$$

Using standard spherical exponential identity, this evaluates to

$$\psi_n^{\mathrm{vMF}}(\mathbf{x}) = \frac{\kappa_n}{\sinh \kappa_n} \frac{\sinh \gamma_n(\mathbf{x})}{\gamma_n(\mathbf{x})}, \tag{10}$$

where

$$\gamma_n(\mathbf{x}) = \sqrt{\kappa_n^2 - k^2 \|\mathbf{x}\|^2 + 2i\kappa_n k \boldsymbol{\mu}_n^\top \mathbf{x}}. \tag{11}$$

Hence the full vMF mixture model becomes

$$p(\mathbf{x}) \approx \sum_{n=1}^{N} \alpha_n \frac{\kappa_n}{\sinh \kappa_n} \frac{\sinh \gamma_n(\mathbf{x})}{\gamma_n(\mathbf{x})}. \tag{12}$$

The vMF function family provides natural canonical mixture atoms because it is defined intrinsically on $\mathbb{S}^2$, is rotationally symmetric around its mean direction $\boldsymbol{\mu}_n$ and is controlled by a single concentration parameter κ_n.

2.2. Tangent-Space Gaussian Distribution

Let the center direction be $\boldsymbol{\mu}_n \in \mathbb{S}^2$ and let the tangent plane at that direction be denoted as $T_{\boldsymbol{\mu}_n}\mathbb{S}^2$. We define an orthonormal basis of the tangent plane

$$\begin{aligned} \mathbf{U}_n &= [\mathbf{u}_{n,1}, \mathbf{u}_{n,2}] \in \mathbb{R}^{3\times 2}, \\ \mathbf{U}_n^\top \mathbf{U}_n &= I_2, \\ \mathbf{U}_n^\top \boldsymbol{\mu}_n &= 0, \end{aligned} \tag{13}$$

where the choice of basis vectors $\mathbf{u}_{n,1}$ and $\mathbf{u}_{n,2}$ need to be consistently constructed from $\boldsymbol{\mu}_n$. Then let there be the local tangent coordinates $\boldsymbol{\xi} \in \mathbb{R}^2$. Nearby directions on the sphere can be approximated by

$$\mathbf{s}_n(\boldsymbol{\xi}) \approx \boldsymbol{\mu}_n + \mathbf{U}_n \boldsymbol{\xi} - \frac{1}{2}\|\boldsymbol{\xi}\|^2 \boldsymbol{\mu}_n. \tag{14}$$

Define

$$\rho_n^{\mathrm{TG}}(\boldsymbol{\xi}; \boldsymbol{\theta}_n) = \frac{1}{2\pi \det(\boldsymbol{\Sigma}_n)^{\frac{1}{2}}} e^{-\frac{1}{2}\boldsymbol{\xi}^\top \boldsymbol{\Sigma}_n^{-1} \boldsymbol{\xi}}, \tag{15}$$

where $\boldsymbol{\theta}_n = (\boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n)$ are the parameters of the TG component. This defines a local tangent-space density, which is mapped approximately to nearby directions on the sphere through Eq. (14).

The forward model using the tangent-space coordinates

$$\psi_n^{\mathrm{TG}}(\mathbf{x}) \approx \int_{\mathbb{R}^2} \rho_n^{\mathrm{TG}}(\boldsymbol{\xi};\boldsymbol{\theta}_n) e^{jk\mathbf{s}_n(\boldsymbol{\xi})^\top \mathbf{x}} J_n(\boldsymbol{\xi})\,\mathrm{d}\boldsymbol{\xi}, \quad (16)$$

where $J_n(\boldsymbol{\xi})$ is the Jacobian of the transformation from tangent coordinates to the sphere. We assume $J_n(\boldsymbol{\xi}) \approx 1$ for reasonably narrow Gaussians. Using the approximation Eq. (14), we define the coordinates

$$\begin{aligned} z_n(\mathbf{x}) &:= \boldsymbol{\mu}_n^\top \mathbf{x} \in \mathbb{R}, \\ \mathbf{b}_n(\mathbf{x}) &:= \mathbf{U}_n^\top \mathbf{x} \in \mathbb{R}^2. \end{aligned} \quad (17)$$

Then we can write

$$\psi_n(\mathbf{x}) \approx e^{jkz_n(\mathbf{x})} \frac{1}{2\pi \det(\boldsymbol{\Sigma}_n)^{\frac{1}{2}}} \int_{\mathbb{R}^2} e^{-\frac{1}{2}q}\,\mathrm{d}\boldsymbol{\xi}. \quad (18)$$

with

$$q = \boldsymbol{\xi}^\top(\boldsymbol{\Sigma}_n^{-1} + ikz_n(\mathbf{x})I_2)\boldsymbol{\xi} - ik\boldsymbol{\xi}^\top \mathbf{b}_n(\mathbf{x}). \quad (19)$$

The expression Eq. (18) is a Gaussian integral and can be solved in closed form. Let

$$A_n(\mathbf{x}) := \boldsymbol{\Sigma}_n^{-1} + ikz_n(\mathbf{x})I_2 \quad (20)$$

then the Gaussian integral yields

$$\psi_n^{\mathrm{TG}}(\mathbf{x}) \approx e^{ikz_n(\mathbf{x})} \frac{e^{-\frac{k^2}{2}\mathbf{b}_n(\mathbf{x})^\top A_n(\mathbf{x})^{-1}\mathbf{b}_n(\mathbf{x})}}{\det(I_2 + ikz_n(\mathbf{x})\boldsymbol{\Sigma}_n)^{\frac{1}{2}}}. \quad (21)$$

This expression contains the wave $e^{ikz_n(\mathbf{x})}$, a beam-like term $e^{-\frac{k^2}{2}\mathbf{b}_n(\mathbf{x})^\top A_n(\mathbf{x})^{-1}\mathbf{b}_n(\mathbf{x})}$ and a normalization term $\det\left(I_2 + ikz_{n(\mathbf{x})}\boldsymbol{\Sigma}_n\right)^{\frac{1}{2}}$. The latter modulates both amplitude and phase.

Unlike the exact intrinsic and isotropic vMF atom, the TG atom is a local tangent-space approximation that introduces anisotropic covariance control at the cost of global spherical fidelity.

2.3. Parameter Count and Model Complexity

For the comparisons below, we count intrinsic real-valued parameters. A vMF component has one complex amplitude, one direction on $\mathbb{S}^2$, and one concentration parameter. Thus, $P_{\mathrm{vMF}} = 5N$, with two real parameters for the amplitude, two for the direction, and one for the concentration. A TG component has one complex amplitude, one direction on $\mathbb{S}^2$, and a symmetric positive definite 2×2 covariance matrix in the tangent plane, giving $P_{\mathrm{TG}} = 7N$. A fixed-grid PWD model with Q fixed directions estimates Q complex coefficients and therefore has $P_{\mathrm{PWD}} = 2Q$.

For sparse PWD, the full optimization variable still has $2Q$ real parameters, but the active parameter count $2Q_{\mathrm{active}}$ is also reported because the sparsity regularizer sets many coefficients to zero.

2.4. Relation to Spherical Harmonics

The proposed mixtures can also be interpreted in the spherical-harmonic domain. If each directional atom has coefficients $\beta_{n,l,m}$, then the mixture distribution has coefficients

$$a_{l,m} = \sum_{n=1}^{N} \alpha_n \beta_{n,l,m}. \quad (22)$$

Thus, the model restricts the general plane-wave distribution to an adaptive low-dimensional nonlinear manifold. This contrasts with SH truncation, where the basis functions are fixed and the model order is controlled by the maximum harmonic degree. In this sense, the proposed mixture is adaptive in the directional domain rather than fixed by harmonic order.

2.5. Estimation

Assume that the pressure is observed at positions $\mathbf{x}_m$, $m = 1, \ldots, M$, with

$$y_m = p(\mathbf{x}_m) + \eta_m. \quad (23)$$

For fixed nonlinear component parameters $\boldsymbol{\Theta} = \{\boldsymbol{\theta}_n\}_{\{n=1\}}^{N}$, the mixture model is linear in the complex amplitudes. We define the atom matrix

$$\boldsymbol{\Psi}_{\boldsymbol{\Theta}}[m,n] = \psi_n(\mathbf{x}_m;\boldsymbol{\theta}_n), \quad (24)$$

so that the predicted measurement vector is

$$\hat{\mathbf{y}} = \boldsymbol{\Psi}_{\boldsymbol{\Theta}}\boldsymbol{\alpha}. \quad (25)$$

For fixed nonlinear parameters, the amplitudes are obtained by ridge LS,

$$\hat{\boldsymbol{\alpha}}(\boldsymbol{\Theta}) = \arg\min_{\boldsymbol{\alpha}} \|\boldsymbol{\Psi}_{\boldsymbol{\Theta}}\boldsymbol{\alpha} - \mathbf{y}\|_2^2 + \lambda_\alpha \|\boldsymbol{\alpha}\|_2^2. \quad (26)$$

This has the closed-form normal-equation solution

$$\hat{\boldsymbol{\alpha}}(\boldsymbol{\Theta}) = \left(\boldsymbol{\Psi}_{\boldsymbol{\Theta}}^H \boldsymbol{\Psi}_{\boldsymbol{\Theta}} + \lambda_\alpha I\right)^{-1} \boldsymbol{\Psi}_{\boldsymbol{\Theta}}^H \mathbf{y}. \quad (27)$$

The nonlinear parameters are then updated by minimizing the projected objective

$$\mathcal{L}(\boldsymbol{\Theta}) = \|\boldsymbol{\Psi}_{\boldsymbol{\Theta}}\hat{\boldsymbol{\alpha}}(\boldsymbol{\Theta}) - \mathbf{y}\|_2^2 + \lambda_\alpha \|\hat{\boldsymbol{\alpha}}(\boldsymbol{\Theta})\|_2^2 + \mathcal{R}(\boldsymbol{\Theta}) \quad (28)$$

using automatic differentiation. In the implementation used here, the LS amplitude update is treated as a projection step and the subsequent gradient update is applied to the nonlinear parameters with the projected amplitudes considered fixed. This follows the variable-projection principle for separable nonlinear least-squares problems [9]. For vMF and for the local TG approximation, the resulting atom expressions are available in closed form; this allows the atom matrix and its gradients to be evaluated directly by automatic differentiation, without numerical quadrature over propagation directions during fitting.

The implemented procedure initializes component directions and spreads/covariances, fits the amplitudes by ridge LS, evaluates the projected objective, updates only the nonlinear parameters by Adam, and refits the amplitudes after the update. This cycle is repeated for a fixed number of iterations. Thus, ridge PWD requires one linear LS solve, sparse PWD is solved iteratively by the fast iterative shrinkage-thresholding algorithm (FISTA) [10], and the proposed mixtures require an iterative nonlinear variable-projection procedure.

3. Experiments

We conduct a numerical simulation study to evaluate the effectiveness of the proposed continuous mixture models and compare them to two fixed-grid PWD baselines.

3.1. Experimental Setup

We evaluate the models in simulated room-acoustic fields generated with a randomized image-source method [11]. Two scenarios are used. The sparse scenario contains few physical sources and weak late reverberation, so that the local field is expected to contain a small number of dominant directional arrivals. The dense scenario uses more sources and stronger reverberation, producing a less sparse plane-wave distribution.
The receiver region is a compact source-free cubical volume with a side length of 0.7 m centered in a $6 \times 5 \times 3$ m room. Sparse pressure measurements are sampled at $M = 96$ microphone positions, while reconstruction is evaluated at $M_{\mathrm{eval}} = 2048$ separate positions in the same source-free region. The frequency range is 100–1000 Hz. For all plane-wave models, the spatial coordinates are centered at the receiver region.
The sparse scenario places 2 sources and has a reverberation time of 0.15 s, while the dense scenario places 5 sources and a reverberation time of 0.8 s. The wall reflection coefficient is set from the target reverberation time using Sabine's formula, yielding approximate sparse and dense reverberation settings [12]. We simulate 20 random realizations for both the sparse and dense scenario. For each realization, the measurement and evaluation positions are sampled uniformly at random within the source-free region, and the source positions are sampled uniformly at random outside the measurement region. For each room realization, each physical source was assigned a fixed complex source weight. The source magnitude was sampled independently from a uniform distribution on $[0.5, 1.5]$, and the phase was sampled independently from a uniform distribution on $[0, 2\pi)$. These weights were kept fixed across the frequency sweep for that realization. Complex white Gaussian noise was added to the measured pressures at a signal-to-noise ratio (SNR) of 20 dB, with independent real and imaginary parts of equal variance.

3.2. Compared Methods

We compare four model classes. Ridge PWD uses a fixed set of Q directions and estimates complex coefficients by one ridge LS solve. Sparse PWD uses the same fixed directional grid but adds a complex sparsity penalty and is solved iteratively using FISTA [10]. The vMF mixture estimates adaptive directions and isotropic angular concentration. The TG mixture estimates adaptive directions and local anisotropic covariance matrices.
The main comparisons are performed at matched intrinsic real-parameter counts. Thus, the orders N and Q are chosen such that $5N$, $7N$, and $2Q$ are comparable for vMF, TG, and fixed PWD, respectively. Sparse PWD is additionally reported as a function of active parameters where the reported curve is an oracle envelope over the tested grid sizes and sparsity penalties: for each realization and parameter count, the lowest evaluation normalized mean-squared error

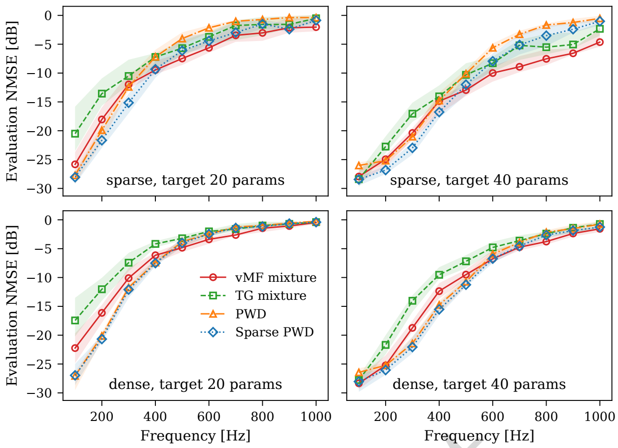

Figure 1: Median NMSE and 25-75 quantile spreads at the evaluation positions over the frequency range $[100, 1000]$ Hz. Sparse PWD is shown as an oracle over the tested grid sizes Q and sparsity penalties. The vMF mixture improves performance mainly at higher frequencies, particularly in the sparse scenario, while fixed-grid PWD remains competitive at lower frequencies. The TG mixture is less consistent, with improvements mainly at higher frequencies and weaker performance at lower frequencies.

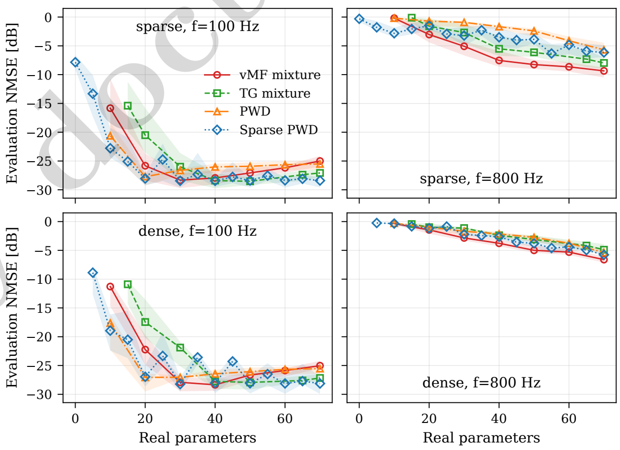

Figure 2: Median NMSE and 25-75 quantile spreads at the evaluation positions over the parameter-count range $[10, 70]$. Sparse PWD is shown as an oracle over the tested grid sizes Q and sparsity penalties. The continuous mixtures are most beneficial at higher frequency, with vMF providing the most consistent improvement.

(NMSE) among the tested sparse PWD settings is selected before aggregating over realizations.
Table 1 lists the solver setup for each method.

3.3. Evaluation Metric

The primary metric is the NMSE at evaluation positions $j = 1, \ldots, M_{\mathrm{eval}}$,

$$\mathrm{NMSE}_{\mathrm{eval}} = 10 \log_{10} \left(\frac{\sum_j \left| \hat{p}(\mathbf{x}_j) - p(\mathbf{x}_j) \right|^2}{\sum_j \left| p(\mathbf{x}_j) \right|^2} \right) \tag{29}$$

reported in dB.

Table 1: Optimization settings used in the experiments. For vMF and TG, the two step sizes denote direction and spread/covariance updates, respectively. Gradient clipping with norm 1 was used for both continuous mixture models. VarPro denotes "variable-projection".

Method	Solver	Iterations	Ridge / Penalty	Step sizes
Ridge PWD	LS	1	10^{-4}	–
Sparse PWD	FISTA	2000	10 values, 10^{-7}–10^{-2}	–
vMF	VarPro + Adam	500	10^{-8}	10^{-3} / 10^{-2}
TG	VarPro + Adam	500	10^{-8}	10^{-3} / $5 \cdot 10^{-3}$

3.4. Results

The results show a frequency-dependent behaviour. Figure 1 shows the median NMSE at the evaluation positions over the frequency range $[100, 1000]$ Hz at two fixed parameter counts $\{20, 40\}$ and for the sparse and dense acoustic scenario. At higher frequencies, vMF mixtures are more parameter-efficient than fixed-grid PWD, especially in the sparse scenario. At lower frequencies, fixed-grid PWD remains competitive or performs better. The TG model is less consistently beneficial; it is mainly competitive in sparse high-frequency fields. A plausible explanation is that the extra anisotropic covariance degrees of freedom may reduce the number of available directional components at a fixed parameter count and make the nonlinear optimization problem more difficult.

Figure 2 shows the median NMSE at the evaluation positions over parameter count for fixed frequencies $\{100, 800\}$ Hz for the sparse and dense scenario. This figure highlights the advantage of the mixture-based methods at higher frequencies: The vMF approach shows the best high-frequency improvement in the sparse scenario. In the dense scenario, TG improves for some higher-parameter settings, suggesting that local anisotropy may be useful in more complex directional fields, although this effect is less consistent than the high-frequency vMF gain.

4. Conclusion

We studied continuous adaptive directional mixture models for plane-wave sound-field estimation. In randomized image-source room simulations, vMF mixtures improved reconstruction at matched intrinsic parameter counts primarily at higher frequencies, where adaptive directional atoms are beneficial. TG mixtures were less robust and only became competitive in sparse high-frequency scenarios, suggesting that their additional anisotropic covariance parameters and harder optimization can outweigh their modeling flexibility. Compared with ridge PWD, the proposed models trade the simplicity of a one-shot LS solve for an iterative nonlinear fitting procedure and a more compact adaptive representation. The observed frequency dependence suggests that hybrid models or frequency-dependent model selection, for example combining fixed-grid PWD at low frequencies with adaptive vMF mixtures at higher frequencies, may be a useful direction for future work. Future work should also consider more robust optimization, principled model selection, and validation on measured room data.

5. Acknowledgments

This research work was carried out at the ESAT Laboratory of KU Leuven, and has received funding from KU Leuven internal funds (C14/21/075), Research Foundation Flanders (FWO projects S005525N and G0A0424N), the Flemish Government (AI Research Program), and the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 956369.

References

[1] E. G. Williams, *Fourier Acoustics: Sound Radiation and Nearfield Acoustical Holography*. Academic Press, 1999.

[2] B. Rafaely, *Fundamentals of Spherical Array Processing*, vol. 16. in Springer Topics in Signal Processing, vol. 16. Cham: Springer International Publishing, 2019. doi: 10.1007/978-3-319-99561-8.

[3] N. Antonello, E. De Sena, M. Moonen, P. A. Naylor, and T. van Waterschoot, "Room Impulse Response Interpolation Using a Sparse Spatio-Temporal Representation of the Sound Field," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 25, no. 10, pp. 1929–1941, Oct. 2017, doi: 10.1109/TASLP.2017.2730284.

[4] N. Ueno, S. Koyama, and H. Saruwatari, "Kernel Ridge Regression with Constraint of Helmholtz Equation for Sound Field Interpolation," in *Proc. 2018 Int. Workshop Acoustic Signal Enhancement (IWAENC '18)*, Tokyo: IEEE, Sep. 2018, pp. 1–440. doi: 10.1109/IWAENC.2018.8521334.

[5] G. Chardon, L. Daudet, A. Peillot, F. Ollivier, N. Bertin, and R. Gribonval, "Near-Field Acoustic Holography Using Sparse Regularization and Compressive Sampling Principles," *J. Acoust. Soc. Am.*, vol. 132, no. 3, pp. 1521–1534, Sep. 2012, doi: 10.1121/1.4740476.

[6] J. Xu, K. Chen, L. Wang, and J. Zhang, "Sparse Plane Wave Approximation of Acoustic Modes to Address Basis Mismatch," *Applied Sciences*, vol. 12, no. 2, p. 837, Jan. 2022, doi: 10.3390/app12020837.

[7] D. Caviedes-Nozal, N. A. B. Riis, F. M. Heuchel, J. Brunskog, P. Gerstoft, and E. Fernandez-Grande, "Gaussian Processes for Sound

Field Reconstruction," *J. Acoust. Soc. Am.*, vol. 149, no. 2, pp. 1107–1119, Feb. 2021, doi: 10.1121/10.0003497.

[8] S. Damiano and T. Van Waterschoot, "Sound Field Reconstruction Using Physics-Informed Boundary Integral Networks," in *Proc. 33rd European Signal Process. Conf. (EUSIPCO '25)*, Sep. 2025, pp. 76–80. doi: 10.23919/EUSIPCO63237.2025.11226620.

[9] G. H. Golub and V. Pereyra, "The Differentiation of Pseudo-Inverses and Nonlinear Least Squares Problems Whose Variables Separate," *SIAM J. Numer. Anal.*, vol. 10, no. 2, pp. 413–432, Apr. 1973, doi: 10.1137/0710036.

[10] A. Beck and M. Teboulle, "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems," *SIAM J. Imaging Sci.*, vol. 2, no. 1, pp. 183–202, Jan. 2009, doi: 10.1137/080716542.

[11] E. De Sena, N. Antonello, M. Moonen, and T. van Waterschoot, "On the Modeling of Rectangular Geometries in Room Acoustic Simulations," *IEEE Trans. Audio Speech Lang. Process.*, vol. 23, no. 4, pp. 774–786, 2015, doi: 10.1109/TASLP.2015.2405476.

[12] H. Kuttruff, *Room Acoustics*, 5. ed. Boca Raton, Fla.: CRC Press [u.a.], 2009.