



# A Time–Space Beamformer and Deconvolution for Rotating Acoustic Sources

Daniel Ernst<sup>1</sup>, Marius Lehmann<sup>2</sup>

<sup>1</sup>German Aerospace Center (DLR)

Bunsenstr. 10, 37073 Göttingen, Germany

<sup>2</sup>ebm-papst Mulfingen GmbH & Co. KG

## Abstract

We present a time-domain beamforming framework for rotating acoustic source fields and an associated sparse deconvolution method. The proposed *Time–Space Beamformer* (TSB) reformulates classical frequency-domain delay-and-sum processing as a four-stage pipeline: (i) stationary-grid frequency-domain beamforming on the microphone pressure spectra, (ii) per-focus inverse FFT to time-domain beamformer signals, (iii) spatial interpolation onto a rotating focus grid with explicit emission-time correction, and (iv) Welch spectral estimation at each rotating focus point. The output is a non-negative beamforming vector  $\mathbf{b}(f_k) \in \mathbb{R}_+^N$  per analysis frequency. Building on this output, we construct a frequency-dependent synthetic response matrix by passing unit-power point sources at every focus point through the identical TSB pipeline and solve a non-negative, mixed  $\ell_1/\ell_2$  regularised inverse problem (FISTA with a closed-form proximal operator). The method is validated on a fully reproducible synthetic rotating-source benchmark inspired by the Microphone Array Benchmark b11. TSB collapses the trajectory smear of stationary-grid beamforming to focused lobes at the rotating-frame source positions, and the additional sparse deconvolution reduces the localization error to grid-scale values over most of the analysis band while improving amplitude consistency. A sensitivity study quantifies robustness against sensor noise, rotor-speed and sound-speed mismatch, source count, and the two regularisation weights.

## 1 Introduction

Acoustic source mapping of rotating components (fans, propellers, electric machine rotors) remains challenging because the source motion is coupled to the propagation delay and the array sees the source from a continuously changing direction. Classical beamforming and deconvolution methods such as DAMAS [1], CLEAN-SC [2], and functional beamforming [3] were

originally designed for stationary scenes, and dedicated treatments for rotating sources rely on angular re-gridding, focusing in the rotating frame, or time-domain resampling strategies [4, 5].

Motivated by these approaches, we formulate a simple, implementation-oriented algorithm that separates the temporal and spatial operations explicitly. This decoupling has two practical benefits: it allows direct re-use of existing frequency-domain beamforming infrastructure (Steps 1 and 4), and it provides a natural, modular forward operator for sparse deconvolution (each column of the system matrix is one TSB evaluation of a synthetic rotating point source).

The contributions of this paper are: (i) a four-step time–space beamformer (TSB) for rotating sources with an explicit emission-time correction that removes the systematic angular bias introduced by finite propagation delay; (ii) a frequency-wise sparse deconvolution scheme whose forward operator is built by the identical TSB pipeline, minimising model mismatch between data and operator; and (iii) a fully reproducible synthetic benchmark with exact ground truth and a comprehensive sensitivity analysis quantifying the dominant error drivers.

## 2 Time–Space Beamformer (TSB)

We consider an array of  $M$  microphones recording  $T$  time samples at sampling rate  $f_s$ , and  $N$  focus points on a regular Cartesian grid in the source plane at axial distance  $z_s$  from the array. Let  $\mathbf{p}(t) \in \mathbb{R}^M$  be the microphone pressure vector and  $\mathbf{P}(f_k) = \text{FFT}\{\mathbf{p}(t)\}$  its discrete Fourier transform. We use free-field monopole transfer functions

$$g_{n,m}(f) = \frac{e^{-j2\pi f r_{n,m}/c}}{4\pi r_{n,m}}, \quad n = 1, \dots, N, \quad m = 1, \dots, M, \quad (1)$$

with  $r_{n,m}$  the distance from focus point  $n$  to microphone  $m$  and  $c$  the speed of sound, and unit-gain steering weights  $\mathbf{w}_n(f) = \mathbf{g}_n(f)/(\mathbf{g}_n^H \mathbf{g}_n)$  so that  $\mathbf{w}_n^H \mathbf{g}_n = 1$  (Sarradj formulation III). This particular choice is not intrinsic to the TSB algorithm: other steering weighting strategies can be substituted without changing the four-stage processing chain. Likewise, more elaborate Green’s functions, for example including boundary reflections or flow effects, can replace the free-field monopole model without modifying the TSB concept itself.

The proposed TSB computes rotating-source spectra in four steps.

**Step 1: Frequency-domain beamforming on a stationary grid.** For each frequency  $f_k$  in the analysis band  $[f_{lo}, f_{hi}]$ , beamform the microphone spectra directly (without cross-spectral matrix construction),

$$\tilde{b}_n(f_k) = \mathbf{w}_n^H(f_k) \mathbf{P}(f_k), \quad (2)$$

on the fixed (non-rotating) focus grid. Frequency bins outside the analysis band are set to zero. This acts as a band-pass that rejects out-of-band sensor noise and skips the steering product for unused bins, reducing computational cost.

**Step 2: Inverse FFT to beamformer time signals.** For each focus point, transform the (Hermitian-symmetric) frequency sequence back to a real time signal,

$$\tilde{b}_n(t_m) = \mathcal{F}^{-1}\{\tilde{b}_n(f_k)\}, \quad m = 1, \dots, T. \quad (3)$$

**Step 3: Spatial interpolation onto the rotating focus grid with emission-time correction.** The rotating focus grid co-rotates with the rotor: a focus point fixed in the rotating frame oc-

copies the stationary-frame position obtained by rotating its rotating-frame coordinates by the instantaneous rotor angle. To account for the finite propagation delay between source and array, the effective rotation angle applied to focus point  $n$  is

$$\theta_n(t_m) = \omega(t_m + \tau_n), \quad \tau_n = \frac{1}{M} \sum_{m=1}^M \frac{r_{n,m}}{c}, \quad (4)$$

where  $\omega$  is the angular rotor velocity and  $\tau_n$  is the mean propagation delay from focus point  $n$  to the array. This correction shifts the interpolation phase so that the stationary-grid map is sampled at the position the source occupied when the energy currently arriving at the microphones was *emitted*, rather than at the instantaneous rotor position at the time of arrival. The rotating-frame value  $b_n(t_m)$  is then obtained by bilinear interpolation of  $\tilde{b}(t_m)$  on the regular stationary grid at the query position determined by  $\theta_n(t_m)$ . When the corresponding ROSI time-domain weighting is used, this rotating-frame signal is equivalent to the output of the Rotating Source Identifier (ROSI) algorithm of Sijtsma et al. [6].

**Step 4: Welch spectral estimation on rotating-grid signals.** Apply Welch averaging (Hann window, 50 % overlap) to each  $b_n(t_m)$  to estimate the rotating-source power spectrum. The TSB output is the non-negative vector

$$\mathbf{b}(f_k) = [b_1(f_k), \dots, b_N(f_k)]^\top \in \mathbb{R}_+^N, \quad (5)$$

for every analysis frequency  $f_k$ , in which a rotating point source appears as a focused lobe at its (stationary) rotating-frame position.

### 3 Sparse Deconvolution Algorithm

Because Welch power spectra are non-negative and approximately linear in the source powers when the sources are mutually incoherent, the TSB output at each frequency can be modelled as

$$\mathbf{b}(f_k) \approx \mathbf{A}(f_k) \mathbf{x}(f_k), \quad \mathbf{b}(f_k) \geq 0, \mathbf{x}(f_k) \geq 0, \quad (6)$$

where  $\mathbf{x}(f_k) \in \mathbb{R}_+^N$  is the unknown source-power map in the rotating frame and  $\mathbf{A}(f_k) \in \mathbb{R}_+^{N \times N}$  is the frequency-dependent point-spread (response) matrix.

**Step 1: Compute the TSB result.** Obtain the measured beamforming vector  $\mathbf{b}(f_k)$  for each  $f_k$  from the four-step TSB pipeline of Sec. 2.

**Step 2: Build the synthetic forward matrix.** For each focus point  $j \in \{1, \dots, N\}$ , place a synthetic point source at  $j$ , rotating synchronously with the rotor and emitting the same deterministic, unit-power broadband excitation covering the analysis band, push the resulting microphone signals through the identical TSB pipeline, and store the output power-spectrum vector as column  $j$  of  $\mathbf{A}(f_k)$ . Reusing one excitation for every column ensures that  $\mathbf{A}(f_k)$  encodes the spatial point-spread alone and is independent of the excitation randomness. Because the steering tensor and the rotating-grid interpolation indices depend only on the geometry, rotor speed, and sound speed, they are computed once and reused across all  $N$  columns, which makes the  $\mathcal{O}(N)$  forward simulations tractable.

**Algorithm 1** Non-negative Elastic-Net Deconvolution via FISTA

**Require:** Forward matrix  $\mathbf{A} \in \mathbb{R}_+^{N \times N}$ , TSB map  $\mathbf{b} \in \mathbb{R}_+^N$ , weights  $\alpha_1 > 0$ ,  $\alpha_2 \geq 0$ , max. number of iterations  $K_{\max}$ , convergence threshold  $\tau$

**Ensure:** Non-negative source map  $\mathbf{x}^* \in \mathbb{R}_+^N$

- 1: **Precompute**  $\tilde{\mathbf{A}} \leftarrow \mathbf{A}^\top \mathbf{A}$ ,  $\tilde{\mathbf{b}} \leftarrow \mathbf{A}^\top \mathbf{b}$
- 2: **Estimate Lipschitz constant** via 30 power iterations of  $\tilde{\mathbf{A}}$ : initialise unit random vector  $\mathbf{v}$ , iterate  $\mathbf{v} \leftarrow \tilde{\mathbf{A}}\mathbf{v}/\|\tilde{\mathbf{A}}\mathbf{v}\|$ , set  $L \leftarrow \mathbf{v}^\top \tilde{\mathbf{A}}\mathbf{v}$
- 3:  $\eta \leftarrow \frac{1}{2L}$ ,  $d \leftarrow 1 + 2\eta\alpha_2$
- 4:  $\mathbf{x}^{(0)} \leftarrow \mathbf{0}$ ,  $\mathbf{z}^{(1)} \leftarrow \mathbf{0}$ ,  $t_1 \leftarrow 1$
- 5: **for**  $k = 1, 2, \dots, K_{\max}$  **do**
- 6:      $\mathbf{g} \leftarrow 2 \left( \tilde{\mathbf{A}}\mathbf{z}^{(k)} - \tilde{\mathbf{b}} \right)$  ▷ gradient of  $\|\mathbf{A}\mathbf{z}^{(k)} - \mathbf{b}\|_2^2$
- 7:      $\mathbf{x}^{(k)} \leftarrow \max \left( \mathbf{0}, \frac{\mathbf{z}^{(k)} - \eta\mathbf{g} - \eta\alpha_1}{d} \right)$  ▷ proximal map, element-wise
- 8:      $t_{k+1} \leftarrow \frac{1 + \sqrt{1 + 4t_k^2}}{2}$  ▷ FISTA momentum parameter
- 9:      $\mathbf{z}^{(k+1)} \leftarrow \mathbf{x}^{(k)} + \frac{t_k - 1}{t_{k+1}} \left( \mathbf{x}^{(k)} - \mathbf{x}^{(k-1)} \right)$  ▷ Nesterov extrapolation
- 10:     **if**  $\frac{\|\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}\|}{\|\mathbf{x}^{(k)}\| + \varepsilon} < \tau$  **then break**
- 11:     **end if**
- 12: **end for**
- 13: **return**  $\mathbf{x}^{(k)}$

**Step 3: Solve the non-negative mixed  $\ell_1/\ell_2$  inverse problem.** For each frequency, estimate

$$\mathbf{x}^*(f_k) = \arg \min_{\mathbf{x} \geq 0} \|\mathbf{A}(f_k)\mathbf{x} - \mathbf{b}(f_k)\|_2^2 + \alpha_1 \|\mathbf{x}\|_1 + \alpha_2 \|\mathbf{x}\|_2^2, \quad (7)$$

where  $\alpha_1 > 0$  promotes sparsity and  $\alpha_2 \geq 0$  stabilises amplitudes in the presence of correlated columns in  $\mathbf{A}(f_k)$ . Because  $\mathbf{x} \geq 0$  the  $\ell_1$  term reduces to  $\sum_i x_i$ , and the proximal operator of the regulariser combined with the non-negativity constraint is closed-form,

$$\text{prox}_\eta(\mathbf{z}) = \max \left( \mathbf{0}, \frac{\mathbf{z} - \eta\alpha_1}{1 + 2\eta\alpha_2} \right). \quad (8)$$

We solve Eq. (7) by accelerated projected proximal gradient (FISTA) [7], following related aeroacoustic applications of FISTA-based sparse reconstruction [8], with step size  $\eta = 1/(2\lambda_{\max}(\mathbf{A}^\top \mathbf{A}))$ , where  $\lambda_{\max}$  is estimated by 30 power iterations of  $\mathbf{A}^\top \mathbf{A}$ . The weights are scaled adaptively to the data,  $\alpha_1 = \beta_1 \max(\mathbf{A}^\top \mathbf{b})$  and  $\alpha_2 = \beta_2 \alpha_1$ , where  $\beta_1$  and  $\beta_2$  are dimensionless hyperparameters that transfer across frequencies and scenes. The complete procedure is stated as Algorithm 1.

In the application study (Section 4) the algorithm is run with  $K_{\max} = 800$ ,  $\tau = 10^{-9}$ ,  $\varepsilon = 10^{-30}$ ,

$\beta_1 = 0.1$ , and  $\beta_2 = 10^{-3}$ .

## 4 Application to a Synthetic Rotating-Source Benchmark

Because publicly accessible measured rotating-source data with exact ground truth is limited, the method is validated on a fully reproducible synthetic test case inspired by the *Microphone Array Benchmark b11* scenario (compact rotating point sources) while giving exact control over the ground truth.

**Setup.** A 32-microphone, four-arm spiral array of 0.5 m radius observes a source plane at distance  $z_s = 0.70$  m (Fig. 1). The rotor turns at 1500 rpm. Two compact point sources, located at radii 0.10 m and 0.18 m and rotating synchronously with the rotor, emit independent band-limited random noise covering 50 Hz to 9000 Hz with relative amplitudes 1.0 and 0.8. Microphone signals are generated by a time-varying-delay propagation model,  $p_m(t) = s(t - r_m(t)/c)/(4\pi r_m(t))$ , evaluated by linear interpolation in the retarded sample index, and corrupted with additive white sensor noise at 20 dB broadband SNR. Signals are sampled at  $f_s = 25\,600$  Hz; the measurement record is  $T = 160$  ms (4096 samples, 4 rotor revolutions) and the Welch estimator uses 512-sample segments (Hann window, 50 % overlap), yielding  $K = 15$  averaged segments and a frequency resolution of  $f_s/512 = 50$  Hz. The rotating focus grid is a  $40 \times 40$  Cartesian grid (0.6 m side, 15 mm spacing) centred on the rotor axis, giving a forward matrix  $\mathbf{A}(f_k) \in \mathbb{R}_+^{1600 \times 1600}$  built once and cached.

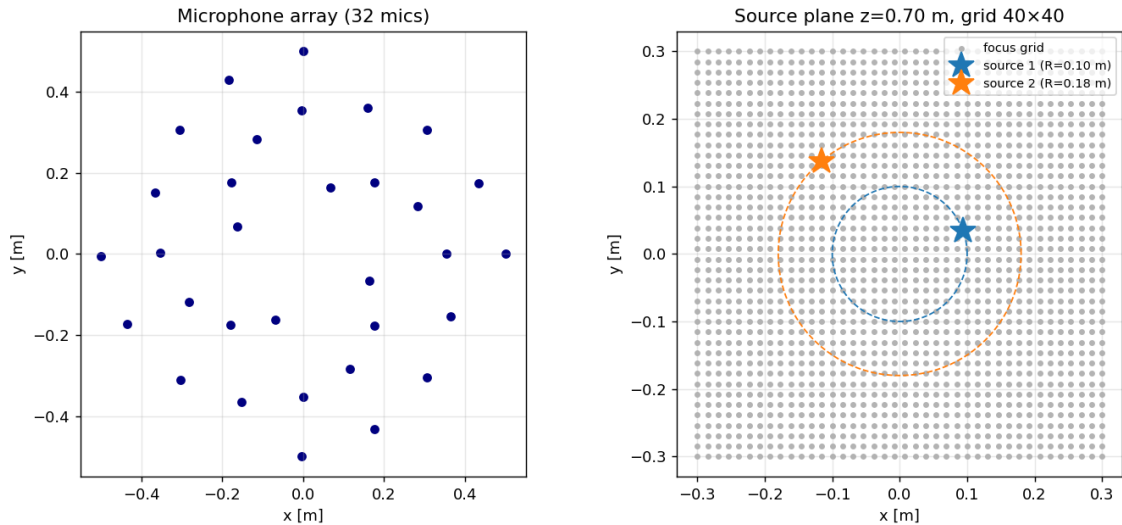


Figure 1: Synthetic test case. Left: spiral microphone array. Right: rotating focus grid in the source plane, the two source trajectories (dashed circles), and the source positions in the rotating frame (stars).

**Results.** Figure 2 shows the source maps at three representative analysis frequencies, 500 Hz, 1000 Hz and 2000 Hz. TSB concentrates the energy near the true rotating-frame positions with a main-lobe whose width is governed by the array aperture. The additional sparse deconvolution generally collapses the response to compact peaks near the ground-truth locations, with

the lowest-frequency case remaining the most challenging because of the broad point-spread function.

Table 1: Quantitative comparison at the five analysis frequencies (mean over the two rotating sources).

Frequency	Method	Loc. error [mm]
500 Hz	TSB	90.1
	TSB + deconv	117.0
1000 Hz	TSB	90.0
	TSB + deconv	13.8
2000 Hz	TSB	121.4
	TSB + deconv	7.3
4000 Hz	TSB	7.3
	TSB + deconv	7.3
8000 Hz	TSB	7.3
	TSB + deconv	7.3

## 5 Regularisation Strategy Comparison

The elastic-net penalty used in Section 4 combines an  $\ell_1$  term that promotes sparsity with a small  $\ell_2$  term that stabilises the solution. To justify this choice we compare four inversion strategies applied to the same TSB output:

- **NNLS** — non-negative least squares ( $\alpha_1 \approx 0$ ,  $\alpha_2 \approx 0$ ); the only constraint is  $\mathbf{x} \geq \mathbf{0}$ .
- **Tikhonov** ( $\ell_2$  only) —  $\alpha_1 = 0$ ,  $\alpha_2 = 0.10 \cdot \max(\mathbf{A}^\top \mathbf{b})$ .
- **LASSO** ( $\ell_1$  only) —  $\alpha_1 = 0.10 \cdot \max(\mathbf{A}^\top \mathbf{b})$ ,  $\alpha_2 = 0$ .
- **Elastic net** ( $\ell_1 + \ell_2$ ) —  $\alpha_1 = 0.10 \cdot \max(\mathbf{A}^\top \mathbf{b})$ ,  $\alpha_2 = 10^{-3} \alpha_1$  (default, as in Section 4).

Figure 4 shows the resulting source maps on the  $40 \times 40$  focus grid at 1000 Hz and 4000 Hz. NNLS removes negative-valued artefacts but retains the broad sidelobe pattern of the raw TSB map, since without any penalty the inverse problem is severely underdetermined. Tikhonov regularisation narrows the sidelobes but tends to produce blurred, spatially spread solutions that overestimate source widths. LASSO achieves compact peaks via the sparsity-inducing  $\ell_1$  penalty; however, without the stabilising  $\ell_2$  term, the solution can exhibit amplitude bias or split a single source into adjacent active grid points, especially at lower frequencies where the beamformer point-spread function is broader. The elastic net combines the compactness of LASSO with the stability of Tikhonov regularisation: the  $\ell_2$  component prevents degenerate sparse solutions and yields well-localised, amplitude-consistent peaks at all frequencies.




Figure 2: Source maps (30 dB dynamic range) at  $\{500, 1000, 2000\}$  Hz (rows from top to bottom). Open circles mark the true source positions in the rotating frame. Left column: TSB; right column: TSB + sparse deconvolution.

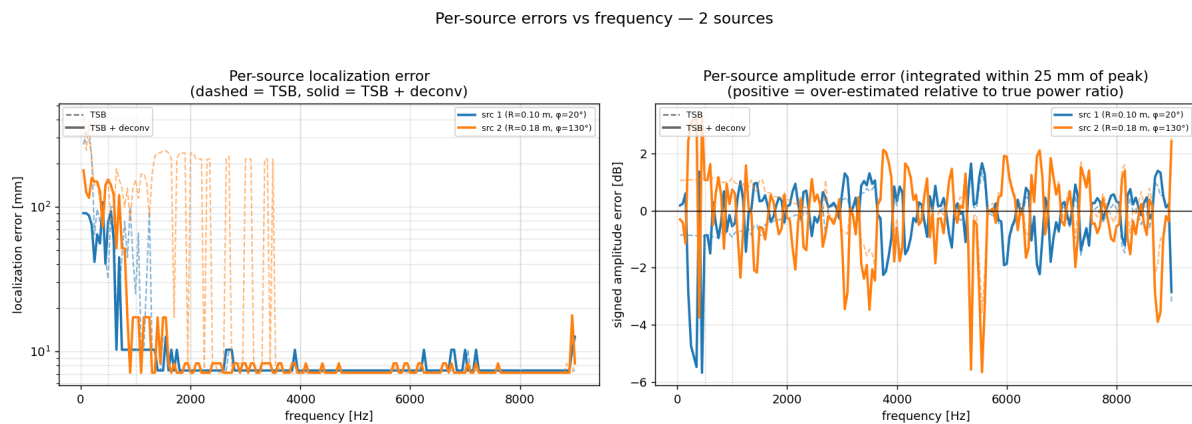


Figure 3: Per-source localization and amplitude error for the two rotating sources at the representative analysis frequencies.

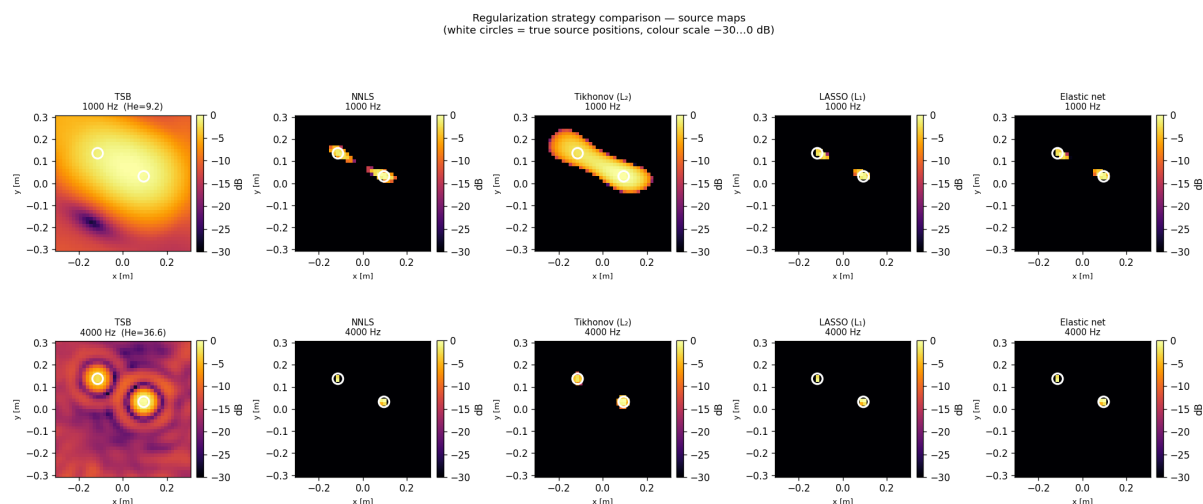


Figure 4: Source maps (30 dB dynamic range,  $40 \times 40$  focus grid) at 1000 Hz (top row) and 4000 Hz (bottom row) for the raw TSB output and four deconvolution strategies: NNLS, Tikhonov ( $\ell_2$  only), LASSO ( $\ell_1$  only), and elastic net ( $\ell_1 + \ell_2$ ). White circles mark the true source positions. The elastic net combines the compact peaks of the  $\ell_1$  penalty with the amplitude stability of the  $\ell_2$  term.

## 6 Effect of Measurement Duration

The Welch spectral estimator averages  $K$  periodogram segments. For a fixed segment length ( $N_{\text{seg}} = 512$  samples at  $f_s = 25\,600$  Hz, 50% overlap),  $K$  grows roughly linearly with the total measurement duration  $T$ ; longer records also correspond to more rotor revolutions, improving the angular coverage of the rotating focus grid. Figure 5 shows the localization error of TSB + sparse deconvolution as  $T$  is swept from 512 to 16384 samples (0.5–16 rotor revolutions,  $K = 1$ –63 Welch averages) on a  $13 \times 13$  focus grid at 20 dB SNR.

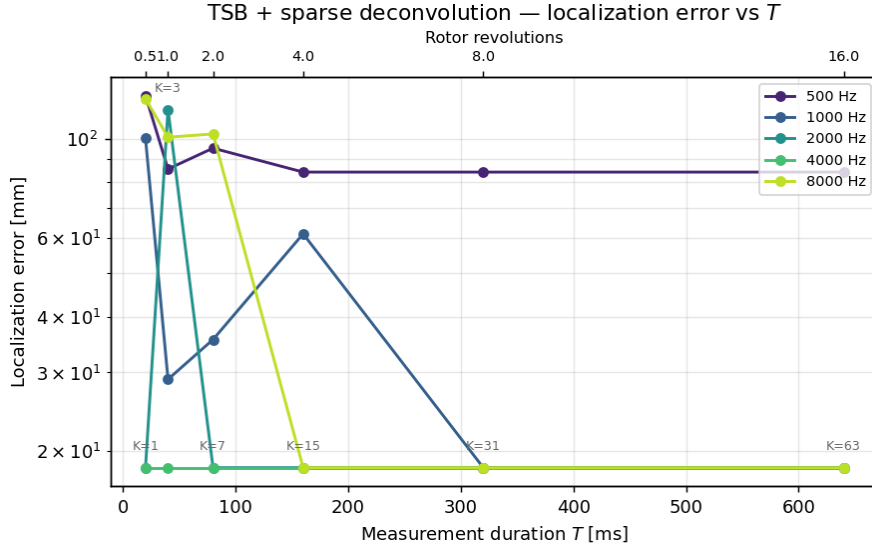


Figure 5: Localization error of TSB + sparse deconvolution versus measurement duration  $T$  at five analysis frequencies (colour-coded). The upper x-axis gives the number of rotor revolutions;  $K$  denotes the number of Welch spectral averages. The error converges to sub-grid values for  $K \geq 7$  ( $T \geq 2048$  samples,  $\approx 2$  revolutions). Results are for a  $13 \times 13$  focus grid at 1500 rpm and 20 dB SNR.

The localization error converges rapidly with  $K$ : already at  $T = 2048$  samples (2 revolutions,  $K = 7$ ) it reaches sub-grid values at all five analysis frequencies. The nominal  $T = 4096$  ( $K = 15$ ) used throughout the application study lies well within the converged plateau, confirming that a few rotor revolutions suffice for reliable deconvolution when the signal-to-noise ratio is adequate.

## 7 Sensitivity Analysis

The robustness of the TSB + sparse deconvolution result is studied at 2000 Hz by independently perturbing six factors. Two figures of merit are reported per panel (Fig. 6): the mean localization error, and the *amplitude error* in decibels, defined as the mean absolute level deviation between the integrated source amplitudes and the reference values,

$$e_{\text{amp}} = \frac{1}{N_s} \sum_{i=1}^{N_s} |10 \log_{10}(a_i) - 10 \log_{10}(a_{\text{ref},i})| \text{ [dB]}, \quad (9)$$

where  $a_i$  is the integrated source amplitude for source  $i$ , defined as the sum of all deconvolved map values within a 25 mm radius of the greedy-extracted peak nearest to the true source position,  $a_{\text{ref},i}$  is the corresponding reference amplitude obtained from the same pipeline at the nominal operating point (so that the unknown overall point-spread gain and integration area cancel), and  $N_s$  is the number of sources. A value of 0 dB indicates perfect amplitude recovery; larger values reflect bias caused by noise, mismatch, or mis-tuned regularisation. The localiza-

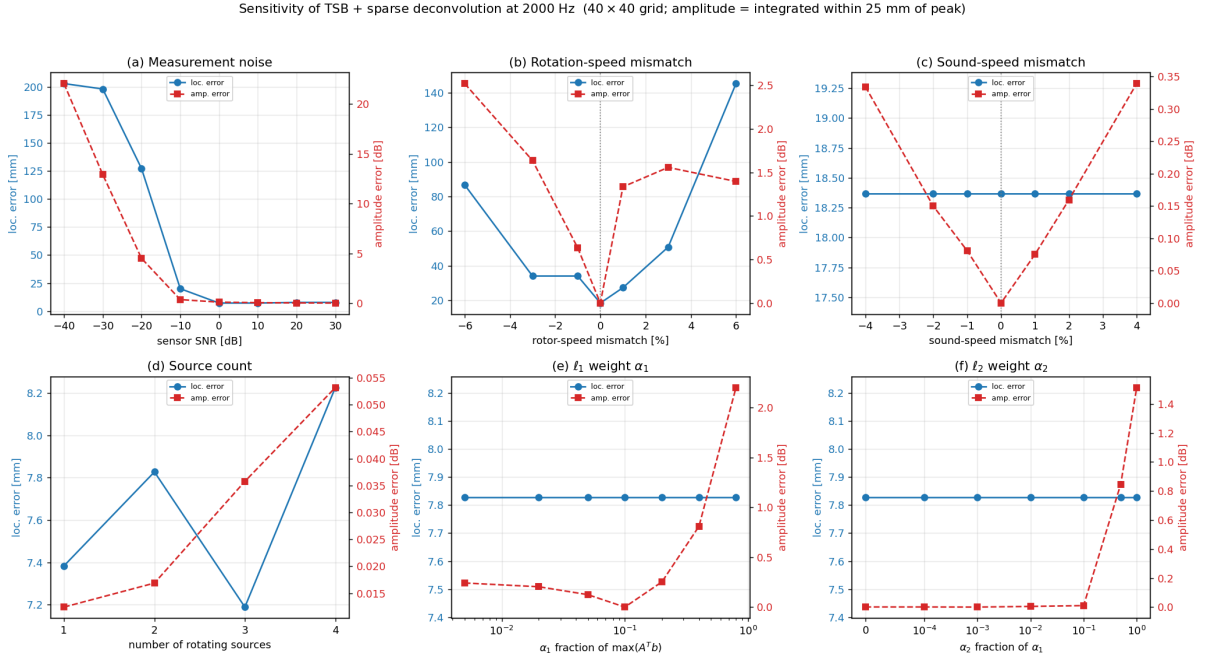


Figure 6: Sensitivity of TSB + sparse deconvolution at 2000 Hz (40 × 40 focus grid; mismatch panels (b)–(c) use a 13 × 13 grid). Each panel shows the localization error (blue, left axis) and the amplitude error  $e_{\text{amp}}$  (Eq. (9); red, right axis; amplitude integrated within 25 mm of the nearest extracted peak) versus (a) measurement noise, (b) rotor-speed mismatch, (c) sound-speed mismatch, (d) number of sources, (e) the  $\ell_1$  weight  $\alpha_1$ , and (f) the  $\ell_2$  weight  $\alpha_2$ .

tion error is grid-quantised and therefore a coarse indicator, while the amplitude error resolves perturbations well below one grid cell. The sweeps for SNR, source count, and the regularisation weights reuse the 40 × 40 forward matrix already built for Sec. 4; the rotor-speed and sound-speed mismatch sweeps use a reduced 13 × 13 grid for computational efficiency, since each mismatch condition requires rebuilding  $\mathbf{A}(f_k)$ .

**(a) Measurement noise.** The sensor SNR is swept from −40 dB to 30 dB. The result is essentially unaffected for  $\text{SNR} \gtrsim -10$  dB: because the sources are band-limited while the sensor noise is broadband, and because the 32-microphone array together with Welch averaging provides a large processing gain, the effective in-band SNR at the focus is roughly 20 dB higher than the nominal broadband value. The localization error, being grid-quantised, stays flat over this whole range, so an overly narrow or favourable SNR sweep can falsely suggest no noise dependence at all; the amplitude error reveals the true behaviour, rising sharply only for strongly negative SNR where the noise floor finally overwhelms the in-band source.

**(b) Rotor-speed mismatch.** The rotor speed assumed in processing is detuned by  $\pm 6\%$  relative to the true speed. Both metrics exhibit a sharp minimum at zero mismatch: an incorrect rotation rate de-synchronises the rotating-grid interpolation and rapidly degrades both the focus and the recovered amplitudes, confirming that accurate tachometer information is essential.

**(c) Sound-speed mismatch.** A  $\pm 4\%$  error in the assumed speed of sound (e.g. from temper-

ature uncertainty) produces a comparatively mild, gradual amplitude degradation, because both the measured-data steering and the forward operator use the same (mismatched) value, so part of the model error cancels.

**(d) Source count.** Increasing the number of simultaneous rotating sources from one to four keeps both the localization and amplitude errors low for the deconvolved result, demonstrating that the sparse inversion separates multiple co-rotating emitters.

**(e) Sparsity weight  $\alpha_1$ .** Sweeping  $\beta_1$  over two decades shows the expected trade-off: too small an  $\alpha_1$  leaves many spurious active grid points and inflates the amplitude error, while an excessive  $\alpha_1$  over-sparsifies and biases the peak locations. A broad plateau around  $\beta_1 \approx 0.1$  yields both a sparse map and minimal error.

**(f) Stabilisation weight  $\alpha_2$ .** The result is largely insensitive to  $\alpha_2$  over several decades; a small positive value improves robustness without affecting accuracy, while a very large  $\alpha_2$  eventually over-smooths the map and increases the amplitude error.

## 8 Discussion

Compared to direct rotating-frame beamforming, TSB isolates the motion handling to a single interpolation stage and preserves compatibility with common frequency-domain preprocessing. The emission-time correction in Step 3 (Eq. (4)) is critical: without it the stationary-grid map is sampled at the wrong angular position, introducing a systematic localization bias that grows with the propagation delay and the rotor speed. The extended deconvolution inherits the modularity of TSB: because the forward operator is generated by the identical TSB chain, the measured data and the operator are produced by the same numerical pipeline, which minimises model mismatch.

The synthetic study confirms the expected benefits — improved localization, amplitude consistency, and separation of closely spaced rotating sources — and quantifies the dominant error driver, namely the rotor-speed accuracy. The mild sensitivity to sound-speed and the wide plateau in  $\alpha_1$  indicate that the method is practical to tune. The regularisation comparison (Section 5) confirms that the elastic-net combination of  $\ell_1$  and  $\ell_2$  penalties outperforms either penalty alone: pure NNLS or Tikhonov regularisation fails to concentrate the solution at the true source locations, while LASSO alone can be unstable at lower frequencies where the point-spread function is broad. The measurement-duration study (Section 6) shows that the deconvolved localization error converges rapidly with the number of Welch averages  $K$ , confirming that practical measurement durations of a few rotor revolutions are sufficient. The linear non-negative forward model (Eq. (6)) relies on the source-incoherence assumption underlying Welch power averaging; strongly coherent rotating sources would couple between columns of  $\mathbf{A}(f_k)$  and require additional treatment, which we leave to future work.

The principal cost of the method is the construction of  $\mathbf{A}(f_k)$ , which requires  $\mathcal{O}(N)$  TSB evaluations per setup: approximately 90 min for the  $40 \times 40$  grid on modern hardware. Caching the signal-independent steering tensor and interpolation map, restricting the steering product to in-band bins, using single-precision complex arithmetic, and evaluating the focus points as batched matrix products together reduce this cost by more than an order of magnitude and make larger focus grids feasible.

## 9 Conclusion

We have presented a Time–Space Beamformer (TSB) for rotating acoustic sources together with an extended sparse deconvolution algorithm. TSB consists of stationary-grid frequency-domain beamforming, per-focus inverse FFT, rotating-grid interpolation with explicit emission-time correction, and Welch spectral estimation. The resulting frequency-wise non-negative vector output naturally supports a synthetic-response matrix formulation and a non-negative mixed  $\ell_1/\ell_2$ -regularised inversion. On a reproducible synthetic rotating-source benchmark the combination reduces the localization error from the broad trajectory smear of conventional stationary-grid beamforming to grid-scale values over most of the analysis band while improving amplitude consistency. A comparison of four regularisation strategies confirms that the elastic-net penalty ( $\ell_1 + \ell_2$ ) yields the most compact and amplitude-consistent source maps. A measurement-duration study shows that the deconvolved localization error converges to sub-grid values for as few as  $K = 7$  Welch averages ( $\approx 2$  rotor revolutions). A sensitivity study using both localization and amplitude error (Eq. (9)) shows that the method is robust to measurement noise down to low SNR and to moderate sound-speed and regularisation changes, while accurate rotor-speed information is the critical requirement.

## References

- [1] T. F. Brooks and W. M. Humphreys, “A deconvolution approach for the mapping of acoustic sources (DAMAS) determined from phased microphone arrays,” *Journal of Sound and Vibration*, vol. 294, no. 4–5, pp. 856–879, 2006.
- [2] P. Sijtsma, “CLEAN based on spatial source coherence,” *International Journal of Aeroacoustics*, vol. 6, no. 4, pp. 357–374, 2007.
- [3] R. P. Dougherty, “Functional beamforming,” *AIAA Journal*, vol. 52, no. 2, pp. 313–319, 2014.
- [4] G. Herold and E. Sarradj, “Microphone array method for the characterization of rotating sound sources in axial fans,” *Noise Control Engineering Journal*, vol. 65, no. 6, pp. 546–551, 2017.
- [5] S. Oertwig, B. Maier, and E. Sarradj, “A beamforming approach for rotating sound sources using virtual rotating microphone arrays,” in *Proceedings of the 8th Berlin Beamforming Conference (BeBeC)*, Berlin, Germany, 2020.
- [6] P. Sijtsma, S. Oerlemans, and H. Holthusen, “Location of rotating sources by phased array measurements,” in *Proceedings of the 7th AIAA/CEAS Aeroacoustics Conference*, Maastricht, The Netherlands, AIAA Paper 2001-2167, 2001.
- [7] A. Beck and M. Teboulle, “A fast iterative shrinkage-thresholding algorithm for linear inverse problems,” *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.
- [8] H.-G. Raumer and C. Spehr, “Wavenumber spectrum determination for aeroacoustic applications using FISTA,” in *Proceedings of Inter-Noise 2022*, Glasgow, UK, pp. 1–9, 2022.