



Sparse hyperparametric Itakura-Saito nonnegative matrix factorization via bi-level optimization

Laura Selicato ^{a,b}, Flavia Esposito ^{b,*}, Andersen Ang ^c,
Nicoletta Del Buono ^b, Rafał Zdunek ^d

^a Water Research Institute - National Research Council (IRSA-CNR), Italy

^b Department of Mathematics, University of Bari Aldo Moro, Italy

^c School of Electronics and Computer Science, University of Southampton, UK

^d Faculty of Electronics, Photonics, and Microsystems, Wrocław University of Science and Technology, Poland

ARTICLE INFO

Keywords:

Nonnegative matrix factorization
Itakura-Saito divergence
Hyperparameter optimization
Sparsity
Bi-level optimization
Dynamical system

ABSTRACT

The selection of penalty hyperparameters is a critical aspect in Nonnegative Matrix Factorization (NMF), since these values control the trade-off between reconstruction accuracy and adherence to desired constraints. In this work, we focus on an NMF problem involving the Itakura-Saito (IS) divergence, which is particularly effective for extracting low spectral density components from spectrograms of mixed signals, and benefits from the introduction of sparsity constraints. We propose a new algorithm called SHINBO, which introduces a bi-level optimization framework to automatically and adaptively tune the row-dependent penalty hyperparameters, enhancing the ability of IS-NMF to isolate sparse, periodic signals in noisy environments. Experimental results demonstrate that SHINBO achieves accurate spectral decompositions and demonstrates superior performance in both synthetic and real-world applications. In the latter case, SHINBO is particularly useful for noninvasive vibration-based fault detection in rolling bearings, where the desired signal components often reside in high-frequency subbands but are obscured by stronger, spectrally broader noise. By addressing the critical issue of hyperparameter selection, SHINBO improves the state-of-the-art in signal recovery for complex, noise-dominated environments.

1. Introduction

Nonnegative Matrix Factorization (NMF) is a dimensionality reduction technique that approximates a nonnegative data matrix as the product of two (lower-dimensional) nonnegative matrices. A key challenge in NMF is to determine penalty coefficients when additional constraints, such as sparsity or smoothness, are imposed [1]. These coefficients control the trade-off between reconstruction accuracy and constraint adherence, but their optimal values are highly dependent on both the dataset and the specific application, making the selection process non-trivial. For example, Corsetti and Fokoué [2] proposes a variant of NMF that incorporates data-dependent penalties and introduces auxiliary constraints to enhance performance in tasks such as face recognition. Additionally, Lantéri et al. [3] presents multiplicative algorithms for NMF that enforce non-negativity and flow preservation constraints while introducing regularizations to ensure smoothness or sparsity. Finally, Fabregat et al. [4] adapts a minimization scheme for functions with non-differentiable constraints, known as PALM, to solve NMF problems, yielding solutions that can be both smooth and

* Corresponding author.

E-mail addresses: laura.selicato@uniba.it (L. Selicato), flavia.esposito@uniba.it (F. Esposito), andersen.ang@soton.ac.uk (A. Ang), nicoletta.delbuono@uniba.it (N. Del Buono), rafal.zdunek@pwr.edu.pl (R. Zdunek).

<https://doi.org/10.1016/j.cam.2025.117316>

Received 24 February 2025; Received in revised form 9 December 2025

Available online 30 December 2025

0377-0427/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

sparse-two highly desirable properties. In this work, we rely on previous studies [5,6] in which the penalized problem is reformulated in a general form, and a strategy is proposed to automatically adjust the penalty coefficient.

Formally, let $X \in \mathbb{R}_+^{m \times n}$ with $m, n \in \mathbb{N}$ be a data matrix, NMF aims to approximate it as the product of $W \in \mathbb{R}_+^{m \times r}$ and $H \in \mathbb{R}_+^{r \times n}$ with $r \leq \min\{m, n\}$, so that $X \approx WH$. In our problem, we want to solve

$$(W^*, H^*) \in \arg \min_{W \geq 0, H \geq 0} D_\beta(X, WH) + \mathcal{P}(\text{Diag}(\lambda)H) \tag{1}$$

with the objective function $D_\beta(\cdot, \cdot)$ being the β -divergence [7,8], assessing the quality of the reconstruction WH in fitting X ¹. The function $\mathcal{P} : \mathbb{R}^{r \times n} \rightarrow \mathbb{R}$ is a penalty term on H that enforces a particular constraint. The vector $\lambda \in \mathbb{R}_+^r$ in (1) contains nonnegative penalty hyperparameters $[\lambda_1, \dots, \lambda_r]$ associated to each row of H .

In this study, we focus on a special case of the β -divergence within NMF. The β -divergence offers a tunable objective function that adapts to the statistics of the data: by varying β it interpolates between the Euclidean distance ($\beta = 2$), the Kullback-Leibler divergence ($\beta = 1$), and-in the limit $\beta \rightarrow 0$ -the Itakura-Saito (IS) divergence. The IS divergence is scale-invariant and normalizes the reconstruction error by the signal magnitude, so weak spectral components (i.e. frequency components with low spectral density) contribute proportionally to the objective rather than being dominated by high-energy regions. As a result, IS-NMF is well suited to preserving fine, low-energy structures (e.g., weak higher-order harmonics, background noise, or faint events) that are often missed by Euclidean/KL objectives. Moreover, IS connects naturally to a multiplicative Gamma-noise model for power spectra, which aligns with the heavy-tailed, non-Gaussian behavior commonly encountered in audio and vibration spectrograms.

Our target application is non-invasive, vibration-based fault detection in rolling bearings [9–13]. The desired signal of interest (SOI) is periodic and impulsive, yet it typically resides in a high-frequency subband and is masked by much stronger, spectrally broad disturbances. We validate our approach on laboratory test-rig measurements of faulty bearings [14]. Although time-domain perturbations are often modeled as additive i.i.d. Gaussian noise, after time-frequency transformation, the IS divergence provides a statistically well-matched criterion for factorizing spectrograms with multiplicative fluctuations.

Related work has applied β -NMF to source separation in audio [15] and established identifiability via *minimum-volume* (MinVol) regularization. The “sufficiently scattered” condition [16] relaxes separability assumptions, and subsequent studies showed that minimizing the convex-hull volume of basis columns can recover factors even in rank-deficient or nonseparable regimes [15], thereby linking MinVol-NMF to practical uniqueness guarantees.

Methodologically, we apply NMF to the spectrogram of the measured signal. Let H denote the temporal activations. We expect the SOI to be captured by one row of H that exhibits sparse, periodic, spike-like activity, whereas rows associated with broadband disturbances are comparatively dense. To bias the factorization toward the SOI, we impose a row-sparsity-promoting penalty on H via the term $\mathcal{P}(\text{Diag}(\lambda)H)$ in (1). We further adopt a bi-level scheme that adapts the hyperparameters λ from the data, aligning the penalty strength with the latent structure of the estimated components.

Contribution. The contribution of this work is twofold.

1. **New model.** In this work, we present a new model for minimizing the Itakura-Saito divergence (Problem (1) with $\beta = 0$) while penalizing rows of H . In particular
 - The penalty hyperparameter is not known in advance. We treat this parameter as an optimization variable, embedding it in a bi-level optimization framework and solving it by a bi-level optimization method.
 - The penalty hyperparameter is row-dependent. Note that Problem (1) is different from standard penalized NMF [1], which applies the same penalty coefficient to all rows in the matrix H . In Problem (1), each row of H is penalized by its own penalty parameter.
2. **New algorithm.** For Problem (1), we present a new multiplicative update (see Eq. (12)), and a way to automatically tune the penalty hyperparameter based on a bi-level strategy (see details in Section 3).

Paper Organization. We introduce the problem and overall algorithm framework in Section 2. In Section 3 we first review the bi-level optimization, then we discuss the details of the bi-level approach proposed in this work for solving the Problem 1 in Section 4. Experimental results on synthetic and real datasets are presented in Section 5. We conclude the paper in Section 6, giving an outline of possible future directions.

Notation. The symbol $\mathbf{0}$ denotes the zero matrix, the symbol $E_{a \times b}$ is the all-in matrix size $a \times b$ with $a, b \in \mathbb{N}$. The notation v denotes a column vector and the notation \underline{v} means that v is a row vector. On matrix, A^T is the transpose of A , and $A^2 = AA$. The symbol $A \odot B$ refers to the Hadamard (element-wise) product between A and B of conformal dimensions, and the symbol $A \oslash B$ with $B \neq \mathbf{0}$ refers to the Hadamard division, and we denote $A^{\odot k}$ as the Hadamard power- k of A . Given $n \in \mathbb{N}$, we denote $[n] := \{1, 2, \dots, n\}$. The

¹ We remind that the β -divergence for matrices can be computed element-wise as $D_\beta(A, B) = \sum_i \sum_j d_\beta(a_{ij}, b_{ij})$, where the function d_β for each $x, y \in \mathbb{R}$ is defined as

$$d_\beta(x, y) = \begin{cases} \frac{1}{\beta(\beta-1)}(x^\beta + (\beta-1)y^\beta - \beta xy^{\beta-1}) & \beta \in \mathbb{R} \setminus \{0, 1\}; \\ x \log\left(\frac{x}{y}\right) - x + y & \beta = 1; \\ \frac{x}{y} - \log\left(\frac{x}{y}\right) - 1 & \beta = 0. \end{cases}$$

symbols $k, t \in \mathbb{N}$ indicate iteration counters. The symbol \mathbf{A}^k refers to the variable \mathbf{A} at the iteration k , A_{ij} or a_{ij} is the (i, j) th element of \mathbf{A} . Lastly $\mathbf{A}_{i \cdot}$ and $\mathbf{A}_{\cdot j}$ are the i th and j th row and column of \mathbf{A} , respectively.

2. The overall optimization framework of SHINBO

In this section, we discuss the focus of the paper, Problem (1) and the overall framework of the proposed algorithm.

The optimization problem. We focus on Problem (1) with $\beta = 0$. As \mathcal{P} we chose a particular penalty function, effective in increasing sparsity, that is the diversity measure J [6,17]. In a particular case, if the matrix is nonnegative, the diversity measure can be written as

$$J(\mathbf{A}) = \sum_{i=1}^n \|\mathbf{A}_{i \cdot}\|_1^2 = \text{Tr}(\mathbf{A}\mathbf{E}\mathbf{A}^\top), \tag{2}$$

where \mathbf{E} is the matrix of 1s. By properties of trace, Problem (1) becomes

$$(\mathbf{W}^*, \mathbf{H}^*, \lambda^*) \in \arg \min_{\mathbf{W} \geq 0, \mathbf{H} \geq 0, \lambda \geq 0} D_0(\mathbf{X}, \mathbf{W}\mathbf{H}) + \text{Tr}(\text{Diag}(\lambda)^2 \mathbf{H}\mathbf{E}\mathbf{H}^\top). \tag{3}$$

We remark that (3) is a nonconvex minimization problem, in which finding global minima is NP-hard, therefore we are interested in finding local minima for the triple $(\mathbf{W}^*, \mathbf{H}^*, \lambda^*)$.

The proposed optimization algorithm SHINBO. We propose an algorithm, called SHINBO, to find a local minima for Problem (3) as follows.

Algorithm 1: SHINBO.

```

1 Input:  $\mathbf{X} \in \mathbb{R}_+^{m \times n}$  and factorization rank  $r$ .
2 Initialize  $\mathbf{W}^0 \in \mathbb{R}_+^{m \times r}$ ,  $\mathbf{H}^0 \in \mathbb{R}_+^{r \times n}$  and  $\lambda^0 \in \mathbb{R}_+^r$ 
3 for  $k = 1, 2, \dots$  do
4      $\mathbf{W}^k = \text{update}(\mathbf{X}, \mathbf{W}^{k-1}, \mathbf{H}^{k-1})$  % classic MU-update
5     for  $l \in [r]$  do
6          $\underline{h}_l^{k-1,0} = \underline{h}_l^{k-1}, \lambda_l^{k-1,0} = \lambda_l^{k-1}$  % initialization
7         for  $t \in [T]$  do
8              $\underline{h}_l^{k,t} = \text{update}(\underline{h}_l^{k,t-1}, \mathbf{W}^k, \mathbf{X}, \lambda_l^k)$  as in (12)
9              $\frac{\partial \mathcal{R}}{\partial \lambda_l}$  as in (10) % hypergradient
10        end
11         $\lambda^{k,T} = \text{update}(\lambda^{k,T}, \nabla_{\lambda} \mathcal{R}(\lambda))$  % projected gradient update
12    end
13 end
14 Return  $\mathbf{W}, \mathbf{H}, \lambda$  at the last iteration.

```

Note that SHINBO is composed by two main parts: one devoted to the update of \mathbf{W} (reviewed in the following) and the other one based on bi-level strategy to optimize simultaneously on \mathbf{H} and λ .

Update on \mathbf{W} . The update of \mathbf{W} can be done simply by the following multiplicative update [18] as $\mathbf{W} = \mathbf{W} \odot (\mathbf{X}\mathbf{H}^\top) \oslash (\mathbf{W}\mathbf{H}\mathbf{H}^\top)$.

The following section introduces the bi-level update on \mathbf{H} and λ by a bi-level method, which is the main contribution of the paper.

3. Bi-level optimization for the subproblem

In this section, we discuss the steps for updating \mathbf{H} and λ in Algorithm 1. Given a fixed \mathbf{W}^k , we have the following optimization subproblem

$$(\mathbf{H}^*, \lambda^*) \in \arg \min_{\mathbf{H} \geq 0, \lambda \geq 0} D_0(\mathbf{X}, \mathbf{W}\mathbf{H}) + \text{Tr}(\text{Diag}(\lambda)^2 \mathbf{H}\mathbf{E}\mathbf{H}^\top). \tag{4}$$

The core idea of this paper is to solve Problem (4) as a bi-level problem, in which we incorporate the problem of tuning hyperparameter λ simultaneously into the update of \mathbf{H} . To do so, first we review the general theory of bi-level optimization and its application to Problem (4).

Section organization and overview of the approach. Under a fixed and given \mathbf{W}^k , the goal is to obtain an updated version of \mathbf{H} and λ that approximately solves Problem (4). The bi-level approach has the following steps:

1. We first rewrite the constrained optimization problem (4) by a bi-level optimization problem for $(\underline{h}_l, \lambda_l)$, with $l \in [r]$, see (5).
2. The inner problem (IP) in Problem (5) is then approximated by the solution of a dynamical system. See Problem IVP- Φ .

3. We then solve Problem IVP- Φ to obtain a solution for \underline{h}_l , and also the hypergradient at the last time point T . See (hypergrad).
4. Lastly, we use the hypergradient to obtain the solution λ by a gradient descent approach. See (6).

We now proceed to discuss each of the steps below.

1. Bi-level formulation. In Algorithm 1, the update of H and λ is performed in r steps, where each step is aimed to update the l th component $(\underline{h}_l, \lambda_l)$. This is achieved by solving the following bi-level problem

$$\min_{\lambda_l \geq 0} \left\{ \begin{array}{l} r(\lambda_l) = \inf_{\underline{h}_l(\lambda_l)} D_2(\mathbf{X}, \mathbf{R} + \mathbf{w}_l \underline{h}_l(\lambda_l)) \quad (\text{OP}) \\ \text{s.t. } \underline{h}_l(\lambda_l) \in \arg \min_{\mathbf{u} \in \mathbb{R}_+^n} D_0(\mathbf{X}, \mathbf{R} + \mathbf{w}_l \mathbf{u}) + \lambda_l^2 \|\mathbf{u}\|_1^2 \quad (\text{IP}) \end{array} \right\}, \tag{5}$$

where the matrix \mathbf{R} is the residual obtained isolating in \mathbf{WH} the l th component of $\mathbf{w}_l \in \mathbb{R}^m$ (column of \mathbf{W}) and $\underline{h}_l \in \mathbb{R}^n$ (row of \mathbf{H}), which is

$$\mathbf{R} = \mathbf{X} - \sum_{j \neq l} \mathbf{w}_j \underline{h}_j.$$

The function $r : \mathbb{R} \rightarrow \mathbb{R}$ is called the Response function of the outer problem related to \underline{h}_l . In the outer problem, the objective D_2 is the β -divergence with $\beta = 2$, which is the Frobenius norm. The inner problem is represented by the β -divergence with $\beta = 0$, which is the Itakura-Saito divergence D_0 , regularized by the row-wise diversity measure defined in (2).

Remark 1. Note that the squared ℓ_1 -norm in the inner problem on \mathbf{u} can be seen as a non-smooth regularization term and therefore proximal gradient descent can be applied on \mathbf{u} (see [19] and [20, Lemma 6.70]), however, their approach is applied to convex objective function, where here in (5) the objective function (the IS-divergence) is possibly nonconvex [18] thus proximal gradient descent do not have convergence guarantee.

2. Dynamical system approach on H. An approach to solve the bi-level Problem (5) over H is to replace the inner problem with a dynamical system [21,22] and compute an approximation solution.

We now omit the k index in $\underline{h}_l^{k,t}$, due to the fact that the update focuses on the iteration over t under a constant k .

Given \underline{h}_l^0 (which depends implicitly on λ_l), we build a dynamical system (IVP- Φ) in the form of a discrete initial value problem as

$$\begin{cases} \underline{h}_l^t = \Phi_t(\underline{h}_l^{t-1}, \lambda_l), & t \in [T], \\ \underline{h}_l^0 = \Phi_0(\lambda_l) \end{cases} \quad (\text{IVP-}\Phi)$$

where $\Phi_t : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}^n$ is a smooth map for $t \in [T]$.

Bi-level optimization uses the IVP- Φ to approximate the solution of the Problem (5). We do so by solving the following minimization problem

$$\lambda_l^* = \arg \min_{\lambda_l} \begin{array}{l} r(\lambda_l) \\ \text{s.t. } \underline{h}_l^t = \Phi_t(\underline{h}_l^{t-1}, \lambda_l) \text{ for } t \in [T], \end{array} \tag{6}$$

in which we approximate the solution of the inner problem with the solution of the dynamical system. This is possible because problem (6) satisfies the existence and convergence theorems as proved in [5].

As a preview, we will derive $\Phi_t(\underline{h}_l^{t-1}, \lambda_l)$ in Section 4.

3. Hypergradient. To find the solution λ^* for Problem (4), we solve the problem formed by joining all the Response functions $r(\lambda_l)$ in (5) as

$$\arg \min_{\lambda \in \mathbb{R}_+^r} \left\{ \mathcal{R}(\lambda) := \sum_j r(\lambda_l) \right\}. \tag{7}$$

We would like to compute the hypergradient, i.e., the gradient of \mathcal{R} with respect to λ , in order to use a gradient descent approach on λ . The hypergradient $\nabla_\lambda \mathcal{R}$, by using chain rule, is

$$\frac{\partial \mathcal{R}}{\partial \lambda_l} = \frac{\partial r}{\partial \lambda_l} + \frac{\partial r}{\partial \underline{h}_l^T} \cdot \frac{d \underline{h}_l^T}{d \lambda_l}, \quad l \in [r]. \tag{hypergrad}$$

Note that \underline{h}_l^T denotes the row-vector \underline{h}_l at the time T .

It is well known that the computation of the hypergradient can be done using Reverse-Mode Differentiation (RMD) or Forward-Mode Differentiation (FMD). Since RMD requires storing specific variables across all iterations and indices in memory, in this work, we use FMD, making it more suitable for scenarios where the total quantity of interest is small. For details we refer the reader to [5,21].

Forward-Mode. FMD computes the differentiation in (hypergrad) using the chain rule. Function Φ_t for $t \in [T]$ depends on λ_l explicit and on \underline{h}_l^{t-1} implicitly, then we have the derivative

$$\frac{d \underline{h}_l^t}{d \lambda_l} = \frac{\partial \Phi_t(\underline{h}_l^{t-1}, \lambda_l)}{\partial \underline{h}_l^{t-1}} \cdot \frac{d \underline{h}_l^{t-1}}{d \lambda_l} + \frac{\partial \Phi_t(\underline{h}_l^{t-1}, \lambda_l)}{\partial \lambda_l}.$$

Let $s^t = \frac{d\mathbf{h}_l^t}{d\lambda_l}$, then each FMD iterate behaves as

$$\begin{cases} s^t = \mathbf{A}_t s^{t-1} + \mathbf{b}_t, & t \in [T] \\ s^0 = \mathbf{b}_0 \end{cases} \tag{8}$$

where $\mathbf{A}_t = \frac{\partial \Phi_t(\mathbf{h}_l^{t-1}, \lambda_l)}{\partial \mathbf{h}_l^{t-1}} \in \mathbb{R}^{n \times n}$ and $\mathbf{b}_t = \frac{\partial \Phi_t(\mathbf{h}_l^{t-1}, \lambda_l)}{\partial \lambda_l} \in \mathbb{R}^n$. Now (hypergrad) can be expressed as

$$\frac{\partial \mathcal{R}}{\partial \lambda_l} = \langle \underline{\mathbf{g}}^T, s^T \rangle \in \mathbb{R} \quad \text{where} \quad \underline{\mathbf{g}}^T = \frac{\partial r}{\partial \mathbf{h}_l} \in \mathbb{R}^n. \tag{9}$$

We note that $\underline{\mathbf{g}}^T$ is denoting the row-vector $\underline{\mathbf{g}}$ at the time T . Lastly, the solution of Problem (8) solves the following equation

$$\frac{\partial \mathcal{R}}{\partial \lambda_l} = \frac{\partial r}{\partial \mathbf{h}_l^T} \left(\mathbf{b}_T + \sum_{t=0}^{T-1} \left(\prod_{s=t+1}^T \mathbf{A}_s \right) \mathbf{b}_t \right). \tag{10}$$

4. Update of λ . Given all the components $\mathbf{h}_1^T, \mathbf{h}_2^T, \dots, \mathbf{h}_r^T$ at the last time point of the dynamic system and the hypergradient (discussed previously), we update λ with a projected gradient descent approach as $\lambda = [\lambda - \alpha \nabla_{\lambda} \mathcal{R}(\lambda)]_+$, for a pre-defined stepsize $\alpha > 0$.

4. Derivation of the algorithm SHINBO

We now introduce the overall bi-level optimization approach, discussed in the previous section, to the Subproblem (4). We use the method of partial Lagrangian multiplier, which is applied only on \mathbf{H} . First, let $\Psi \in \mathbb{R}_+^{r \times n}$ be the matrix of Lagrangian multipliers associated to the nonnegative constraints of \mathbf{H} , then the expression of the (partial) Lagrangian of Subproblem (4) is

$$\mathcal{L}(\mathbf{H}) = D_0(\mathbf{X}, \mathbf{W}\mathbf{H}) + \text{Tr}(\text{Diag}(\lambda)^2 \mathbf{H}\mathbf{E}\mathbf{H}^T) + \text{Tr}(\Psi \mathbf{H}^T).$$

Recall that $\mathbf{M}^{\odot 2}$ denotes the Hadarmard power of \mathbf{M} , then the partial derivative of \mathcal{L} with respect to \mathbf{H} is

$$\frac{\partial \mathcal{L}}{\partial \mathbf{H}} = \mathbf{W}^T \left((\mathbf{W}\mathbf{H})^{\odot -2} (\mathbf{W}\mathbf{H} - \mathbf{X}) \right) + 2\text{Diag}(\lambda)^2 \mathbf{H}\mathbf{E} + \Psi.$$

Now denote H_{ij} the (i, j) th element of the matrix \mathbf{H} , and recall that \odot is the Hadarmard product, then by the complementary slackness $\Psi_{ij} H_{ij} = 0$ in the KKT conditions, we get

$$\left[\mathbf{W}^T \left((\mathbf{W}\mathbf{H})^{\odot -2} \odot (\mathbf{W}\mathbf{H} - \mathbf{X}) \right) \right]_{ij} H_{ij} + 2 \left[\text{Diag}(\lambda)^2 \mathbf{H}\mathbf{E} \right]_{ij} H_{ij} + \Psi_{ij} H_{ij} = 0.$$

These equations lead to the multiplicative update

$$H_{ij} = H_{ij} \left[\mathbf{W}^T (\mathbf{W}\mathbf{H})^{\odot -2} \mathbf{X} \right]_{ij} / \left[\mathbf{W}^T (\mathbf{W}\mathbf{H})^{\odot -1} + 2\text{Diag}(\lambda)^2 \mathbf{H}\mathbf{E} \right]_{ij}, \tag{11}$$

where the division is performed element-wise as the Hadamard division \oslash . By fixing $l \in \{1, \dots, r\}$, the update in row-wise format for the l th row of \mathbf{H} can be rewritten equivalently as

$$\underline{\mathbf{h}}_l^t = \underline{\mathbf{h}}_l^{t-1} \oslash \frac{\left(\mathbf{W}^T (\mathbf{W}\mathbf{H})^{\odot -2} \mathbf{X} \right)_{l:}^{t-1}}{\left(\mathbf{W}^T \odot (\mathbf{W}\mathbf{H})^{\odot -1} \right)_{l:} + 2(\lambda_l^{t-1})^2 \|\underline{\mathbf{h}}_l^{t-1}\|_1 \mathbf{E}_{l:}}. \tag{12}$$

We remark that Eq. (12) is one of the main contribution of this paper.

Remark 2. The key idea in this methodology is that $\Phi_t(\mathbf{h}_l^{t-1}, \lambda_l)$ in the dynamical system IVP- Φ is the right hand side of the update (12).

4.1. The implementation of the bi-level approach in SHINBO

Following the procedure and the discussion in Section 3, we consider $\Phi_t(\mathbf{h}_l^{t-1}, \lambda_l)$ in the dynamical system IVP- Φ (represented by the multiplicative update (12)) and we compute $\mathbf{A}^t = \frac{\partial \Phi_t}{\partial \mathbf{h}_l^{t-1}}$ and $\mathbf{b}^t = \frac{\partial \Phi_t}{\partial \lambda_l}$ for $t \in [T]$ required for the FMD.

- The computation of $\mathbf{A}^t = \frac{\partial \Phi_t}{\partial \mathbf{h}_l^{t-1}}$ gives a diagonal matrix

$$\mathbf{A}_{jj}^t = \frac{N_{lj}}{D_{lj}} - h_{lj} \left(\frac{2 \sum_i^n w_{il}^2 \frac{x_{ij}}{(\mathbf{W}\mathbf{H})_{ij}^3} D_{lj} - N_{lj} \sum_i^n \frac{w_{il}^2}{(\mathbf{W}\mathbf{H})_{lj}^2} + 2\lambda_l^2 N_{lj}}{D_{lj}^2} \right),$$

where the derivative is computed with respect to the (l, j) th element of \mathbf{H} , and

$$N_{lj} = (\mathbf{W}^T (\mathbf{W}\mathbf{H})^{\odot -2} \mathbf{X})_{lj}, \quad D_{lj} = (\mathbf{W}^T (\mathbf{W}\mathbf{H})^{\odot -1} + 2\text{Diag}(\lambda)^2 \mathbf{H}\mathbf{E})_{lj}.$$

Table 1
The method compared in the experiment.

Algorithm	λ	λ setting
MU [18]	0	constant
Penalized MU	0.1	constant for each column
Penalized MU	0.5	constant for each column
SHINBO	not required	automatically optimized, different penalization for each column

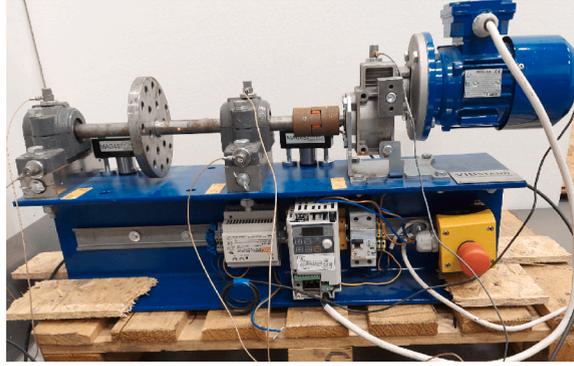


Fig. 1. The test rig used in the experiment [14].

- The vector \mathbf{b}_i is given by $\frac{\partial \Phi_i}{\partial \lambda_i} = -4\lambda_i h_{ij}^2 \frac{N_{ij}}{D_{ij}^2}$.

Finally, having chosen as an outer problem the Frobenius norm, $\underline{\mathbf{g}}^T$ in (9) for computing (hypergrad) is $\underline{\mathbf{g}}^T = -2\mathbf{w}_i^T(\mathbf{X} - \mathbf{R} - \mathbf{w}_i \mathbf{h}_i)$.

Computational cost. The computational cost of the proposed algorithm depends on the matrix-vector multiplications, as is usual in NMF algorithms, and on the bi-level framework, which requires the use of internal loops. The total computational cost of SHINBO is $\mathcal{O}(KTmnr)$, where K is the number of iterations, and T is the number of internal iterations. Although the proposed method appears to suffer from a high computational cost in construction, the gain in efficiency and accuracy in hyperparameter estimation justifies the high cost. Indeed, solving the NMF problem by running a well-known grid or random search methods requires more effort than SHINBO, which automatically selects the best hyperparameter vector for each row of \mathbf{H} .

5. Experimental results

In this section, we present the numerical results of comparing the proposed algorithm SHINBO with the multiplicative update (MU) algorithm in [18]. We also compare SHINBO with the penalized MU, which is the update (12) under a fixed penalty hyperparameter λ with $\lambda_1 = \lambda_2 = \dots = \lambda_r$ for every row of \mathbf{H} . We summarize their difference in Table 1.

Datasets. We evaluate the algorithms on two datasets: one synthetically generated and the other obtained from real vibrational signals measured under laboratory conditions from damaged rolling bearing elements.

- The synthetic dataset is generated starting from a full-rank decomposition $\mathbf{X} \approx \mathbf{W}\mathbf{H}$, where the factor matrix \mathbf{W} contains 10% nonzeros in the mr entries and the factor matrix \mathbf{H} contains 70% nonzeros in the rn entries. In this dataset, we have $(m, n, r) = (100, 70, 3)$.
- The vibrational signals were measured on the test rig presented in Fig. 1 (taken from the previous study [14]). The platform is equipped with an electric motor, gearbox, couplings, and two bearings. One of the latter was deliberately damaged. Diagnostic signals were measured with the accelerometer (KISTLER Model 8702B500), stacked horizontally to the shaft bearing. The 40-s-long vibration signal was recorded with a sampling rate of 50 kHz. For easier visualization, one-second excerpt was selected and then transformed to a spectrogram using a 128-window length, 100-sample overlapping, and 512 frequency bins. The raw observed signal and its spectrogram are shown in Fig. 2. Assuming $r = 4$ (four main components of the mixed signal), we have $(m, n, r) = (257, 1782, 4)$.

5.1. Experimental setup

Initialization. For synthetic dataset, to ensure a fair comparison, all algorithms were initialized using the same initial factor matrices, which were generated from an unpenalized MU-based NMF warm start, that is initialized via nonnegative double singular value

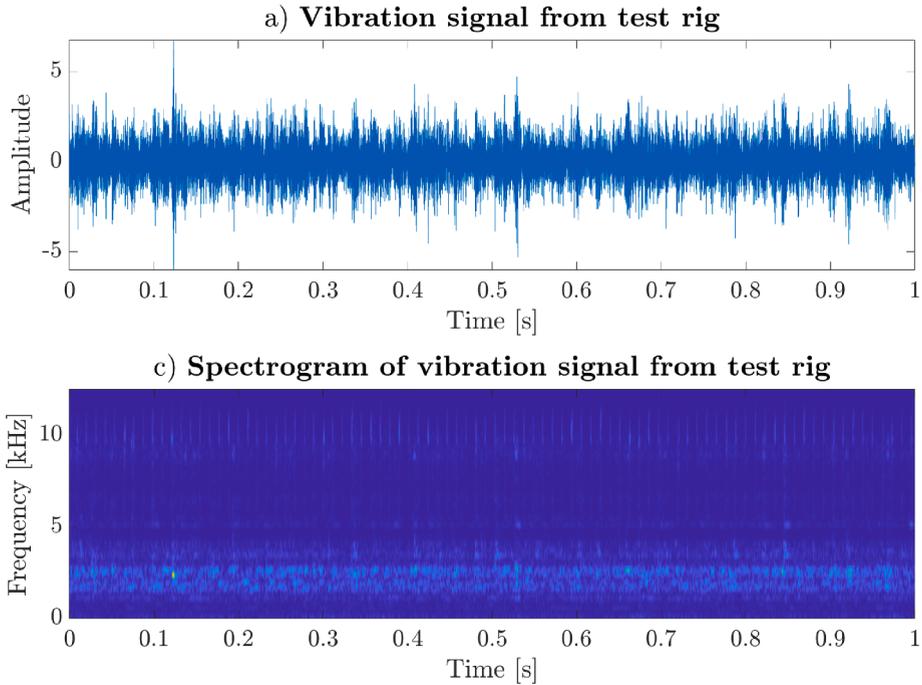


Fig. 2. Recorded vibration signal and its spectrogram.

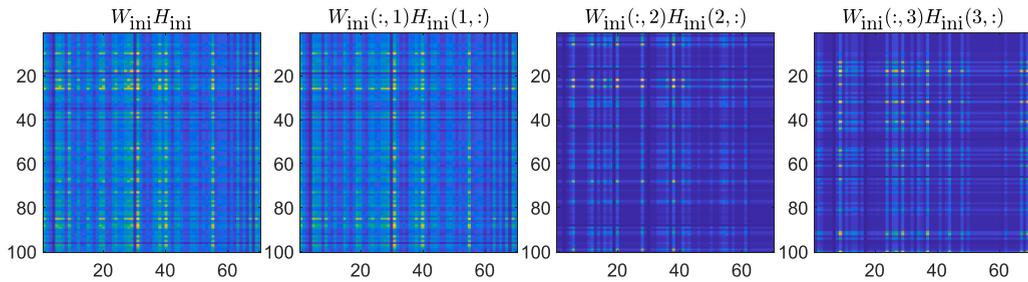


Fig. 3. Initial factors generated using $(1.5|\mathcal{N}| + 0.5)/2$. Left: $\mathbf{W}_0 \mathbf{H}_0$. Right: each rank-1 component in $\mathbf{W}_0 \mathbf{H}_0$.

decomposition (NDSVD) of the matrix \mathbf{X} [23]. For real-dataset, we initialize the initial factor matrices as follows. Let $|\mathcal{N}|$ denote the truncated Gaussian distribution (i.e., each entries of the matrix follows $\mathcal{N}(0, 1)$ and then perform a clipping to replace negative values to 0), we generate both $\mathbf{W}_0, \mathbf{H}_0$ using $(1.5|\mathcal{N}| + 0.5)/2$. See Fig. 3.

Simulation. We perform 100 Monte Carlo simulations, where each run uses a different random data matrix. At the start of each run, the initial value of λ^0 was selected randomly following a uniform distribution $\mathcal{U}[0, 1]$.

Normalization. We perform a normalization step on each column of \mathbf{W} , \mathbf{w}_k for all $k = 1, \dots, r$:

$$\mathbf{w}_k = \mathbf{w}_k / \max(\mathbf{w}_k); \quad \underline{\mathbf{h}}_k = \underline{\mathbf{h}}_k * \max(\mathbf{w}_k).$$

Termination. All the algorithms were allowed to run a maximum number of iterations of 500 for the synthetic dataset and 100 for the real one, with an early termination tolerance of 10^{-6} on the relative fitting error using the IS-divergence D_0 , defined as

$$\left| D_0(\mathbf{X}, \mathbf{W}^{k+1} \mathbf{H}^{k+1}) - D_0(\mathbf{X}, \mathbf{W}^k \mathbf{H}^k) \right| / \left| D_0(\mathbf{X}, \mathbf{W}^k \mathbf{H}^k) \right| \leq 10^{-6},$$

where k indicates the iteration of the outer loop of the algorithm. For the inner loop of the bi-level approach on iteration counter t , we stop at the maximum number of (inner) $T = 4$ iterations. We remark that Problem (3) is nonconvex, thus we do not have a theoretical guarantee that the algorithm’s convergence can avoid poor local minima. We only empirically show that the local minima reached are stable across different random datasets.

Table 2
 p -values of pairwise comparisons for SIR coefficients on (\mathbf{W}, \mathbf{H}) . In the table $\epsilon = 10^{-16}$.

	MU	MU $\lambda = 0.1$	MU $\lambda = 0.5$
MU $\lambda = 0.1$	$8.1 \cdot 10^{-5}, 6.3 \cdot 10^{-10}$	–	–
MU $\lambda = 0.5$	$< 2\epsilon, < 2\epsilon$	$< 2\epsilon, < 2\epsilon$	–
SHINBO	0.14, 0.0042	$5.4 \cdot 10^{-5}, 1.2 \cdot 10^{-9}$	$< 2\epsilon, < 2\epsilon$

Evaluation. We evaluate the algorithms according to different criteria on synthetic and real-world datasets. For both datasets, we plot the convergence of the algorithms by looking at the behavior of the Response function.

For the synthetic dataset, we evaluate the quality of the factorization with respect to the identification problem using the Signal-to-Interference Ratio (SIR) measure [24] on both matrices.² To highlight the effectiveness of the proposed method, we conduct a statistical comparison of SIR values across algorithms using the Kruskal-Wallis test, followed by a post-hoc multiple comparison based on the Mann-Whitney test [25] with Benjamini-Hochberg (BH) correction [26], maintaining a significance level of $\alpha = 0.05$. We also investigate the sparsity of the results obtained for the synthetic dataset using the following measure of sparsity for a generic matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$:

$$Sp(\mathbf{A}) := (1 - \|\mathbf{A} - 10^{-6} \mathbf{E}\|_0 / mn) 100\%,$$

where $\|\cdot\|_0$ denotes the number of non-zeros elements. To demonstrate the advantages of the proposed method, we analyze sparsity values across algorithms using the same statistical test as for the evaluation of the SIR (Kruskal-Wallis test, post-hoc Mann-Whitney test with BH and $\alpha = 0.05$).

On the real-world dataset, we check the goodness of the factorization by quantifying the impulsive and cyclic behavior of the signal under analysis, using a modified envelope spectrum-based indicator (ENVSI) [27,28] on time profiles. ENVSI can be expressed as a spectrum-based indicator (SBI):

$$SBI := \frac{\sum_{i=1}^{M_1} AIS_i^2}{\sum_{k=1}^{M_2} S_k^2},$$

where AIS_i is the magnitude of the i th harmonic of the estimated signal in the frequency domain, S_k is the magnitude in the k th frequency bin in the spectrum of the time profile, M_1 is the number of harmonics to be analyzed (assuming a periodic signal), and M_2 is the number of frequency bins to calculate the total energy. SBI is zero if there are no impulsive components in the time profile, whereas a larger SBI occurs when the impulses in the time profile are stronger (which corresponds to the amplitudes in the spectrum), and the noise is weaker. In the experiments, the number of harmonics M_1 is set to 6, and it was experimentally found to be sufficient for demonstrating the impulsive and periodic behavior of the SOI representing the time profile in the analyzed application. To demonstrate the superiority of the proposed method, we statistically compare ENVSI values across algorithms using a Kruskal-Wallis test and a multiple comparison based on the Mann-Whitney test with BH correction, with a statistical significance level $\alpha = 0.05$.

Computer. All the experiments were conducted in MATLAB 2021a and executed on a machine with an i7 octa-core processor and 16GB of RAM.

5.2. Results on synthetic dataset

Referring to Fig. 4, the proposed algorithm demonstrates superior performance in terms of the convergence rate of Response function compared to the other methods. We remark that, despite SHINBO exhibiting similar, rapid convergence to unpenalized MU for the first 450 iterations (on average), the resulting decompositions have different SIR values. Referring to Fig. 5, we observe that SHINBO, by operating on the vector λ across its components, achieves the highest SIR values on both matrices, confirming its effectiveness in the identification problem. At the same time, the results show higher sparsity on \mathbf{H} and lower sparsity on \mathbf{W} , indicating that the automatically selected hyperparameter behaves as expected by enforcing sparsity primarily on the rows of \mathbf{H} . On average, SHINBO outperforms the comparison algorithms with about 10% higher SIR and 5% higher sparsity on \mathbf{H} .

These results are also confirmed by the statistical comparisons. The Kruskal-Wallis tests present p -values lesser than 10^{-14} and the details of the pairwise comparison with a BH comparison are presented in the Tables 2 and 3. Table 2 reports the results of pairwise comparisons for the SIR coefficients on (\mathbf{W}, \mathbf{H}) . The significantly low p -values indicate that SHINBO consistently achieves higher SIR values compared to MU under different penalization settings, thereby confirming its superior accuracy in handling the identification problem in matrix reconstruction. Similarly, Table 3 presents the p -values obtained from pairwise comparisons of sparsity coefficients on (\mathbf{W}, \mathbf{H}) . The results show that SHINBO exhibits statistically significant differences with respect to MU variants, highlighting its ability to balance accuracy and sparsity in both synthetic and real datasets.

² This measure, computed between the estimated signals and the true signals, is a log-scale similarity measure (expressed in decibels), often used in signal processing applications, and its higher value indicates a higher similarity level.

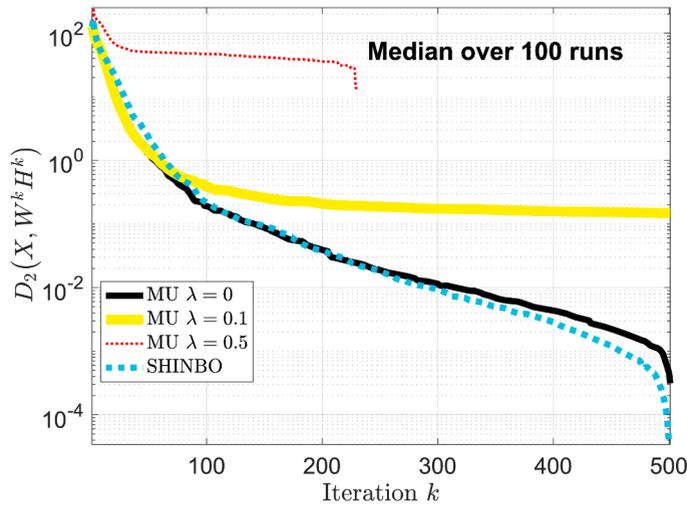


Fig. 4. Behavior of the response function (outer problem) for the synthetic dataset.

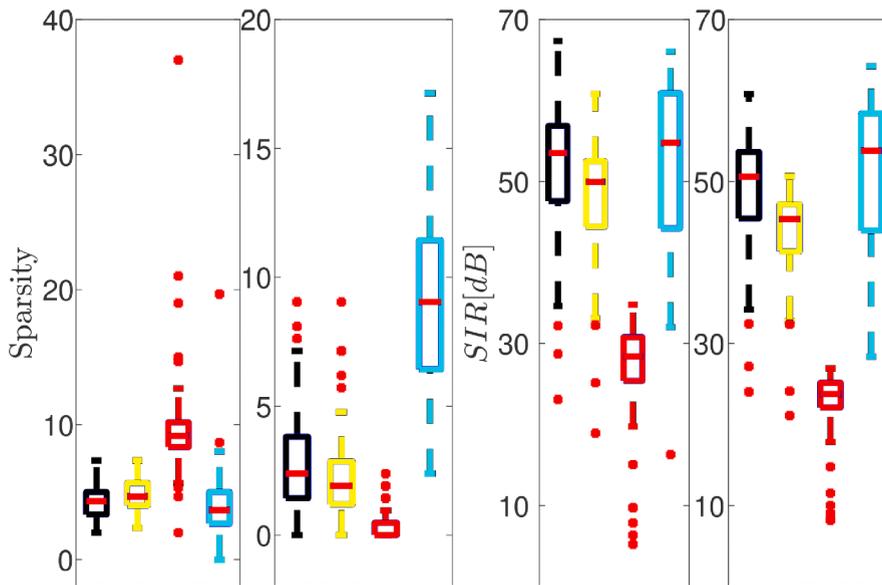


Fig. 5. Results of sparsity and SIR on the factor matrices on the synthetic dataset. The left two columns are the sparsity of W (left) and H (right) of the four algorithms: from left to right are MU with $\lambda = 0$, MU with $\lambda = 0.1$, MU with $\lambda = 0.5$, and SHINBO. The right two columns are SIR of W (left) and H (right) of the four algorithms: from the left to right are MU with $\lambda = 0$, MU with $\lambda = 0.1$, MU with $\lambda = 0.5$, and SHINBO.

Table 3
p-values of pairwise comparisons for sparsity coefficients on (W, H) .

	MU	MU $\lambda = 0.1$	MU $\lambda = 0.5$
MU $\lambda = 0.1$	0.00036, 0.0075	–	–
MU $\lambda = 0.5$	$< 2\epsilon, < 2\epsilon$	$< 2\epsilon, < 2\epsilon$	–
SHINBO	0.023, $< 2\epsilon$	$3.12 \cdot 10^{-6}, < 2\epsilon$	$< 2\epsilon, < 2\epsilon$

Sensitivity to the rank. We performed a sensitivity analysis with respect to the rank r for the synthetic dataset. Tables 4-5 show that across different values of r , the SHINBO algorithm performs better on SIR average than the other compared methods. We note that SHINBO also has the highest variance among all tested methods. Finally, we observed that the mean \pm std values of SIR decreases as rank r increases for all method. All these observations can be explained as follows.

- First, MU $\lambda = 0.5$ has the smallest variance across all methods, but also at the same time has the lowest SIR on average, meaning the method has the worst performance.

Table 4

Mean±std of SIR values of H for the synthetic dataset, w.r.t. values of r .

rank	MU	MU $\lambda = 0.1$	MU $\lambda = 0.5$	SHINBO
4	39.096 ± 14.023	35.899 ± 11.641	17.457 ± 7.164	44.734 ± 17.894
5	29.830 ± 12.216	28.183 ± 11.050	14.769 ± 6.545	35.801 ± 16.903
6	23.423 ± 9.159	22.576 ± 8.694	11.849 ± 5.028	26.930 ± 13.764

Table 5

Mean±std of SIR values of W for the synthetic dataset, w.r.t. values of r .

rank	MU	MU $\lambda = 0.1$	MU $\lambda = 0.5$	SHINBO
4	41.836 ± 17.689	39.183 ± 15.542	20.736 ± 10.644	46.108 ± 20.023
5	30.665 ± 15.257	29.489 ± 13.976	17.347 ± 9.014	36.755 ± 18.479
6	23.879 ± 10.719	23.317 ± 10.242	14.153 ± 6.501	27.983 ± 15.461

Table 6

Mean±std of SIR values for H , with respect to different noise.

ϵ	MU	MU $\lambda = 0.1$	MU $\lambda = 0.5$	SHINBO
0.01	27.3259 ± 4.7506	27.3571 ± 4.5884	20.9436 ± 4.4491	27.6639 ± 5.1397
0.05	19.5527 ± 3.7249	19.5402 ± 3.7146	16.3236 ± 4.3636	19.2795 ± 4.1752
0.1	15.6428 ± 3.3477	15.6939 ± 3.2599	12.6351 ± 4.5464	15.3858 ± 3.3559

Table 7

Mean±std of SIR values for W , with respect to different noise.

ϵ	MU	MU $\lambda = 0.1$	MU $\lambda = 0.5$	SHINBO
0.01	28.3267 ± 4.7552	28.4009 ± 4.7298	23.8437 ± 6.1478	29.0103 ± 4.8775
0.05	19.8840 ± 3.6120	19.8576 ± 3.6338	17.3408 ± 5.4629	19.3763 ± 4.1957
0.1	15.7853 ± 3.1266	15.8039 ± 3.1181	13.2717 ± 5.1220	15.6706 ± 3.1455

- The low variance of MU $\lambda = 0.5$ is contributed to the fact that the method has a relatively larger λ compared to the baseline MU method. Recall that the diversity measure (2), such as $J(A) = \sum_{i=1}^n \|A_{i\cdot}\|_1^2$, is applied on H , thus a higher λ in MU will made the method focusing more on minimizing the column norm of H . A solution with small column-wise norm will naturally has smaller variance, and thus this explains why MU $\lambda = 0.5$ has the smallest variance. In general, if λ is huge (e.g. $\lambda = 10, 100, \dots$), then we can expect that MU with such value will have even smaller variance on the SIR results. As Problem (3) contains two terms, thus a large λ will lead to the MU focusing on column norm minimization instead of the factorization, leading to a lower SIR in reconstruction.
- Among the MU methods, MU with $\lambda = 0$ has the highest SIR on average and also the highest variance. Compared with MU, SHINBO has a higher SIR on average in this case.
- Note that SHINBO needs to find more components of $\lambda = (\lambda_1, \lambda_2, \dots)$, instead of just one λ for MU. For example, for $r = 3$, SHINBO produces a value $\lambda = (0.0142, 0.0292, 0.0305)$. With an increasing r , the number of components of λ also increase, making Problem (3) harder to solve for SHINBO. Under such situation, SHINBO is still able to have on average higher SIR.

Robustness under measurement noise. To assess robustness to measurement noise, we corrupted the clean data X by adding i.i.d. zero-mean Gaussian noise and then projecting onto the nonnegative orthant. Formally, we generated

$$Y = \max(X + \epsilon \mathcal{N}(0, 1), 0),$$

where $\epsilon > 0$ denotes the noise standard deviation. The SIR metrics on both factors in Tables 6–7 show a gradual degradation of SIR as ϵ increases (i.e., as the noise level grows), without abrupt performance breakdowns. Across all tested noise conditions, SHINBO remains comparable to the MU baselines, often slightly superior at a low noise level (e.g., $\epsilon = 0.01$), and on par at higher noise, thus substantiating its robustness to measurement noise.

5.3. Results on real-world dataset

We evaluate the proposed IS-NMF model with row-wise penalty adaptation with SHINBO on a vibration signal from a faulty rolling bearing (Fig. 2). Following the experimental setup, NMF is applied to the signal spectrogram, and the temporal activations (rows of H) are shown in Fig. 6(a). One activation (here, the second component) exhibits a sparse, periodic, spike-like pattern characteristic of the SOI, i.e., fault-induced impulsivity in the bearing. Because of NMF’s permutation ambiguity, the SOI can appear in any row; the remaining rows capture latent components associated with broadband and transient disturbances.

Fig. 6(b) displays the envelope spectra of the activations from Fig. 6(a). The first harmonic at approximately 91Hz aligns with the shaft rotational speed, corroborating that the second component corresponds to the SOI. To quantify spectral leakage and

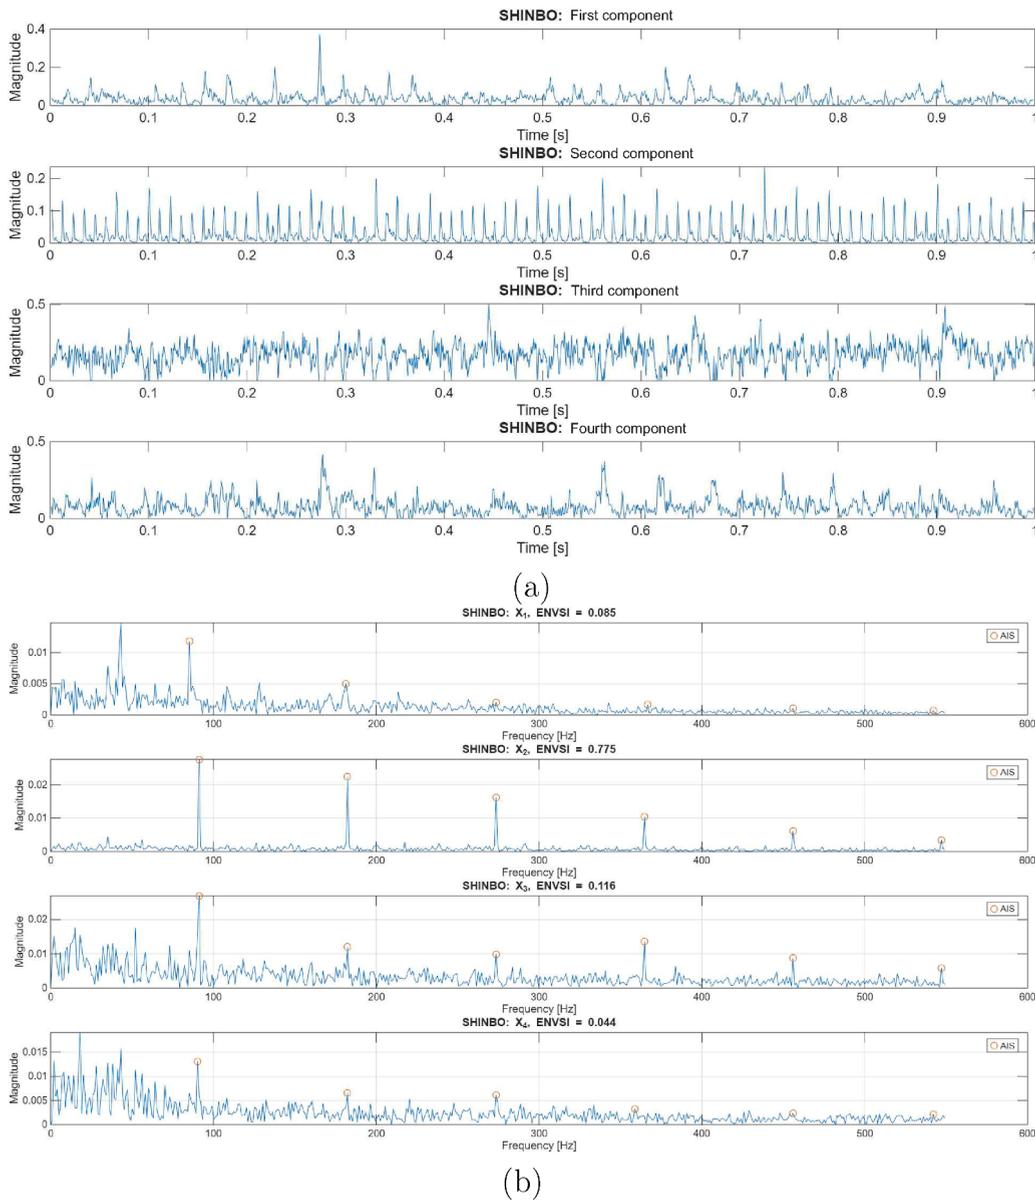


Fig. 6. (a) Time-domain components (the rows of H) estimated with SHINBO from the faulty signal. (b) Envelope spectrum of the estimated time-domain components. The respective ENVSI scores are given in the subtitles.

identifiability, the ENVSI metric was used. Fig. 7 summarizes the quantitative results. The right panel shows that SHINBO attains the highest ENVSI, with values above 0.77 and fewer outliers, indicating reliable SOI recovery consistent with our prior knowledge about the physical state of the tested rolling bearing. The left panel reports the convergence of the Response function, where the proposed method outperforms the baselines.

These results are also confirmed by the statistical comparisons. The Kruskal-Wallis tests present p -values lesser than the fixed significant level (0.04), and the details of the pairwise comparison with a BH comparison are presented in Table 8, which summarizes the results for the ENVSI coefficient. The p -values indicate no significant differences among the MU variants, whereas SHINBO shows statistically significant improvements with respect to all MU configurations, further confirming its effectiveness in capturing relevant structural information.

The concept is that SHINBO is designed to effectively identify which components of λ are associated with noise and which are linked to the true sparse signal. The goal is to apply greater penalization to certain components, thereby filtering out the noise while preserving the component representing the SOI. The results indicate that SHINBO successfully isolates the meaningful signal while suppressing irrelevant noise.

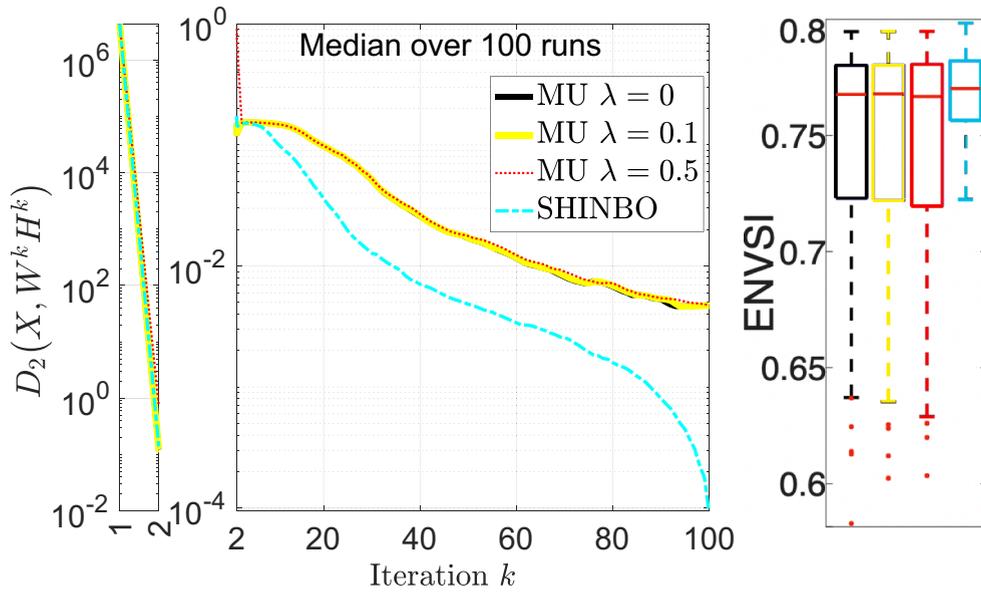


Fig. 7. Left: Behavior of the Response function (outer problem) for the real dataset. Right: The ENVSI score of H of the four algorithms: from left to right are MU with $\lambda = 0$, MU with $\lambda = 0.1$, MU with $\lambda = 0.5$, and SHINBO.

Table 8
p-values results of pairwise comparisons for ENVSI coefficients.

	MU	MU $\lambda = 0.1$	MU $\lambda = 0.5$
MU $\lambda = 0.1$	0.98	–	–
MU $\lambda = 0.5$	0.98	0.98	–
SHINBO	0.04	0.04	0.04

6. Conclusion

In this work, we addressed the critical challenge of selecting penalty hyperparameters in NMF by introducing SHINBO, a novel algorithm that employs a bi-level optimization framework to adaptively tune row-dependent penalties. The core idea is that SHINBO does not require setting the penalty coefficient value as a fixed parameter in advance; instead, it automatically optimizes the hyperparameter during the process. Compared to standard algorithms, where the penalty coefficient value is user-dependent, SHINBO consistently achieves high SIR, sparsity, and ENVSI on synthetic and real datasets in a fully automatic manner. This demonstrates that the proposed method can solve the problem in a one-shot way, without relying on any prior parameter tuning, while still providing optimal results in terms of: i) identification: SHINBO yields optimal SIR for both reconstructed matrices; ii) Sparsity: it respects the constraints imposed by the penalizations on the matrix H ; iii) Spectral analysis: it ensures an effective analysis of the signal spectrum. By focusing on the IS divergence, SHINBO proves highly effective for extracting low spectral density components in spectrograms, particularly in the presence of noise. The ability of the algorithm to enforce sparsity constraints and dynamically adjust penalties ensures a more precise separation of meaningful signals from noisy disturbances.

Through experiments on both synthetic and real-world datasets, SHINBO demonstrated its superior performance compared to traditional NMF methods. For real-world applications, such as noninvasive vibration-based fault detection in rolling bearings, SHINBO excelled at isolating sparse, periodic signal components in high-frequency subbands, even when heavily masked by broader noise.

Overall, the results highlight SHINBO’s potential to significantly improve signal recovery in complex, noise-dominated environments. By tackling the hyperparameter selection problem with an adaptive, data-driven approach, SHINBO not only advances the field of NMF but also provides a robust tool for applications requiring precise spectral decomposition and noise suppression. Future work will explore the scalability of SHINBO for larger datasets and its adaptability to other domains with similar challenges.

Data availability

The authors do not have permission to share data.

Funding

N.D.B., F.E., and L.S. are partially supported by “INdAM - GNCS Project”, CUP: E53C24001950001. N.D.B. and F.E. are supported by Piano Nazionale di Ripresa e Resilienza (PNRR), Missione 4 “Istruzione e Ricerca”-Componente C2 Investimento 1.1, “Fondo per il Programma Nazionale di Ricerca e Progetti di Rilevante Interesse Nazionale”, Progetto PRIN-2022 PNRR, P2022BLN38, Computational approaches for the integration of multi-omics data. CUP: H53D23008870001.

Declaration of generative AI and AI-assisted technologies

The authors used ChatGPT 5.1 (OpenAI) in order to revise the English grammar. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Declaration of competing interest

The authors have no relevant financial interest to disclose.

Acknowledgement

The authors thank Professor Radoslaw Zimroz from Faculty of Geoenvironment, Mining and Geology at Wroclaw University of Science and Technology, for providing the real data collected in his laboratory. N.D.B., F.E., L.S. are members of the Gruppo Nazionale Calcolo Scientifico - Istituto Nazionale di Alta Matematica (GNCS-INdAM).

References

- [1] N. Gillis, Nonnegative matrix factorization, *arXiv:2012.03889*, <https://doi.org/10.4236/ojs.2015.57077> 2020.
- [2] M. Corsetti, E. Fokoué, Nonnegative Matrix Factorization with Zellner Penalty, Technical Report, arXiv preprint, *arXiv:2012.03889*, <https://doi.org/10.4236/ojs.2015.57077>, 2020.
- [3] H. Lantéri, C. Theys, C. Richard, Nonnegative matrix factorization with regularization and sparsity-enforcing terms, in: *Multi-Sensor Adaptive Processing*, IEEE, 2011, pp. 97–100, <https://doi.org/10.1109/CAMSAP.2011.6136055>
- [4] R. Fabregat, N. Pustelnik, P. Gonçalves, P. Borgnat, Solving NMF with Smoothness and Sparsity Constraints Using Palm, Technical Report, arXiv preprint, <https://doi.org/10.48550/arXiv.1910.14576>, 2019.
- [5] N.D. Buono, F. Esposito, L. Selicato, R. Zdunek, Bi-level algorithm for optimizing hyperparameters in penalized nonnegative matrix factorization, *Appl. Math. Comput.* 457 (2023) 128184.
- [6] N.D. Buono, F. Esposito, L. Selicato, R. Zdunek, Penalty hyperparameter optimization with diversity measure for nonnegative low-rank approximation, *Appl. Numer. Math.* 208 (2025) 189–204.
- [7] C. Févotte, J. Idier, Algorithms for nonnegative matrix factorization with the β -divergence, *Neural Comput.* 23 (9) (2011) 2421–2456.
- [8] R. Kompass, A generalized divergence measure for nonnegative matrix factorization, *Neural Comput.* 19 (3) (2007) 780–791.
- [9] J. Wodecki, P. Kruczek, A. Bartkowiak, R. Zimroz, A.W. omańska, Novel method of informative frequency band selection for vibration signal using nonnegative matrix factorization of spectrogram matrix, *Mech. Syst. Signal Process.* 130 (2019) 585–596.
- [10] J. Wodecki, A. Michalak, R. Zimroz, T. Barszcz, A.W. omańska, Impulsive source separation using combination of nonnegative matrix factorization of bi-frequency map, spatial denoising and Monte Carlo simulation, *Mech. Syst. Signal Process.* 127 (2019) 89–101.
- [11] J. Wodecki, A. Michalak, R. Zimroz, A.W. omańska, Separation of multiple local-damage-related components from vibration data using nonnegative matrix factorization and multichannel data fusion, *Mech. Syst. Signal Process.* 145 (2020) 106954.
- [12] C. Yang, Interval riccati equation-based and non-probabilistic dynamic reliability-constrained multi-objective optimal vibration control with multi-source uncertainties, *J. Sound Vib.* 595 (2025) 118742. <https://www.sciencedirect.com/science/article/pii/S0022460X24005042>.
- [13] C. Yang, Q. Wang, W. Lu, Y. Li, Integrated uncertain optimal design strategy for truss configuration and attitude-vibration control in rigid-flexible coupling structure with interval uncertainties, *Nonlinear Dyn.* 113 (3) (2025) 2215–2238.
- [14] M. Gabor, R. Zdunek, R. Zimroz, J. Wodecki, A. Wylomanska, Non-negative tensor factorization for vibration-based local damage detection, *Mech. Syst. Signal Process.* 198 (2023) 110430.
- [15] V. Leplat, N. Gillis, A.M. Ang, Blind audio source separation with minimum-volume beta-divergence NMF, *IEEE Trans. Signal Process.* 68 (2020) 3400–3410.
- [16] X. Fu, K. Huang, N.D. Sidiropoulos, On identifiability of nonnegative matrix factorization, *IEEE Signal Process. Lett.* 25 (3) (2018) 328–332.
- [17] S.F. Cotter, B.D. Rao, K. Engan, K. Kreutz-Delgado, Sparse solutions to linear inverse problems with multiple measurement vectors, *IEEE Trans. Signal Process.* 53 (7) (2005) 2477–2488.
- [18] C. Févotte, N. Bertin, J.-L. Durrieu, Nonnegative matrix factorization with the itakura-saito divergence: with application to music analysis, *Neural Comput.* 21 (3) (2009) 793–830.
- [19] T. Evgeniou, M. Pontil, D. Spinellis, N. Nassuphis, J. Suykens, M. Signoretto, A. Argyriou, Regularized robust portfolio estimation, regularization, optimization, kernels, and support vector machines, CRC Press, Boca Raton, FL, 2015, pp. 237–256.
- [20] A. Beck, First-order methods in optimization, SIAM, 2017.
- [21] L. Franceschi, M. Donini, P. Frasconi, M. Pontil, Forward and Reverse Gradient-Based Hyperparameter Optimization, *ICML, PMLR*, 1165–1173, 2017.
- [22] D. Maclaurin, D. Duvenaud, R. Adams, Gradient-based hyperparameter optimization through reversible learning, in: *Proc. of ICML*, 2015, pp. 2113–2122.
- [23] C. Boutsidis, E. Gallopoulos, SVD based initialization: a head start for nonnegative matrix factorization, *Pattern Recognit.* 41 (4) (2008) 1350–1362.
- [24] A. Cichocki, R. Zdunek, A.H. Phan, Amari Si Nonnegative Matrix and Tensor Factorizations: Applications to Exploratory Multi-Way Data Analysis and Blind Source Separation, John Wiley & Sons, 2009.
- [25] E.L. Lehmann, J.P. Romano, G. Casella, *Testing Statistical Hypotheses*, 3, Springer, 1986.
- [26] Y. Benjamini, Y. Hochberg, Controlling the false discovery rate: a practical and powerful approach to multiple testing, *J. R. Stat. Soc. Series B (Methodol.)* 57 (1) (1995) 289–300.
- [27] M. Gabor, R. Zdunek, R. Zimroz, A. Wylomanska, Bearing damage detection with orthogonal and nonnegative low-rank feature extraction, *IEEE Trans. Ind. Inform.*, 2023, <https://doi.org/10.1109/TII.2023.3300455>
- [28] Y. Berrouche, G. Vashishtha, S. Chauhan, R. Zimroz, Local damage detection in rolling element bearings based on a single ensemble empirical mode decomposition, *Knowl. Based Syst.* 301 (2024) 112265.