



HAL
open science

Estimation with Low-Rank Time-Frequency Synthesis Models

Cédric Févotte, Matthieu Kowalski

► **To cite this version:**

Cédric Févotte, Matthieu Kowalski. Estimation with Low-Rank Time-Frequency Synthesis Models. IEEE Transactions on Signal Processing, 2018, 66 (15), pp.4121 - 4132. 10.1109/TSP.2018.2844159 . hal-01680655v2

HAL Id: hal-01680655

<https://hal.science/hal-01680655v2>

Submitted on 12 Jun 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Estimation with Low-Rank Time-Frequency Synthesis Models

Cédric Févotte, *Senior Member, IEEE*, and Matthieu Kowalski

Abstract—Many state-of-the-art signal decomposition techniques rely on a low-rank factorization of a time-frequency (t-f) transform. In particular, nonnegative matrix factorization (NMF) of the spectrogram has been considered in many audio applications. This is an *analysis* approach in the sense that the factorization is applied to the squared magnitude of the analysis coefficients returned by the t-f transform. In this paper we instead propose a *synthesis* approach, where low-rankness is imposed to the synthesis coefficients of the data signal over a given t-f dictionary (such as a Gabor frame). As such we offer a novel modeling paradigm that bridges t-f synthesis modeling and traditional analysis-based NMF approaches. The proposed generative model allows in turn to design more sophisticated multi-layer representations that can efficiently capture diverse forms of structure. Additionally, the generative modeling allows to exploit t-f low-rankness for compressive sensing. We present efficient iterative shrinkage algorithms to perform estimation in the proposed models and illustrate the capabilities of the new modeling paradigm over audio signal processing examples.

I. INTRODUCTION

MATRIX factorization methods currently enjoy a large popularity in machine learning and signal processing. In signal processing, the input data is usually a time-frequency (t-f) transform of some original time series $x(t)$. For example, in the audio setting, nonnegative matrix factorization (NMF) is commonly used to decompose magnitude or power spectrograms into elementary components [1]; the spectrogram \mathbf{P} is approximately factorized into \mathbf{WH} , where \mathbf{W} is the dictionary matrix collecting spectral patterns in its columns and \mathbf{H} is the activation matrix. The approximate \mathbf{WH} is generally of lower rank than \mathbf{P} , unless additional constraints are imposed on the factors. NMF is at the core of classical source separation systems such as [2], [3].

The spectrogram \mathbf{P} is usually obtained from the short-time Fourier transform \mathbf{Y} . The coefficients y_{fn} of \mathbf{Y} are the inner products of $x(t)$ with t-f atoms $\phi_{fn}(t)$, where f and n index frequencies and time frames, respectively, and a common choice is $\mathbf{P} = |\mathbf{Y}|^2$. The STFT coefficients are so-called *analysis* coefficients. As such, spectral decomposition by NMF can be viewed as a low-rank time-frequency *analysis* procedure. Leveraging on the potential of *synthesis* models as opposed to analytical ones (see, e.g., [4]–[7]), we propose to explore a dual view of the usual NMF approach and present a new paradigm that we name *low-rank time-frequency synthesis* (LRTFS). In this new paradigm, the signal is decomposed as

$$x(t) = \sum_{fn} \alpha_{fn} \phi_{fn}(t) + e(t) \quad (1)$$

where the synthesis coefficients $\{\alpha_{fn}\}$ are endowed with a low-rank structure such that $|\alpha_{fn}|^2 \approx [\mathbf{WH}]_{fn}$. Formulation (1) provides a generative representation of the raw data $x(t)$ and extends the modeling capacities of standard NMF-based signal decomposition towards more advanced multi-layer hybrid decompositions. Having a generative model of the raw data (instead of its transform) is also useful for some inverse problems such as compressive sampling, an application that will be considered in the paper.

The low-rankness of the synthesis coefficients $\{\alpha_{fn}\}$ is induced through a probabilistic model named *Gaussian Composite Model* (GCM) [8]. The GCM underlies Itakura-Saito NMF, a baseline method that will be recalled in Section II. Section III-A presents our new paradigm LRTFS in the general case of complex-valued signals. It also describes an alternate minimization algorithm for maximum joint likelihood estimation of the parameters. Section III-B shows how the methodology for complex signals can be adapted to real-valued signals. Section III-C discusses how LRTFS relates to other temporal models with low-rank spectrograms and/or structured variance [9]–[15]. Section IV describes how LRTFS can accommodate more advanced multi-layer decompositions in which every layer can have its own t-f resolution or structure (e.g., a sparse instead of low-rank time-frequency structure). Section V describes a new approach to compressive sampling, which exploits latent low-rank time-frequency structure instead of sparsity, with superior results for the considered type of data. The article is illustrated throughout with experiments on audio signals (the presented methodology is however not limited to audio signals). In particular, we use a running piano toy example to illustrate every stage of our contributions.

This article unifies and continues our work presented in the conference papers [16], [17]. In particular, it provides more detailed experiments and presents the following novel methodological contributions: the case of real-valued signals (which require to properly handle the Hermitian symmetry of their synthesis coefficients) is now rigorously treated in Section III-B, algorithm accelerations are presented in Section III-A2, and the concept of compressive LRTFS presented in Section V is entirely novel.

II. A BASELINE: ITAKURA-SAITO NMF AND THE GAUSSIAN COMPOSITE MODEL (GCM)

NMF was originally designed in a deterministic setting [18]: a measure of fit between \mathbf{P} and \mathbf{WH} is minimized with respect to (w.r.t) \mathbf{W} and \mathbf{H} . Choosing the “right” measure for a specific type of data and task is not straightforward. Furthermore, NMF-based spectral decompositions often arbitrarily discard

C. Févotte is with IRIT, Université de Toulouse, CNRS, France.
M. Kowalski is with L2S, Université Paris-Sud, Gif-sur-Yvette, France.

phase information: only the magnitude of the complex-valued short-time Fourier transform (STFT) is considered. To remedy these limitations, a generative probabilistic latent factor model of the STFT, the GCM, was proposed in [8]. It is defined by

$$y_{fn} \sim N_c(0, [\mathbf{WH}]_{fn}), \quad (2)$$

where N_c refers to the circular complex-valued normal distribution.¹ As shown by Eq. (2), in the GCM the STFT is assumed centered and its variance has a low-rank structure. Many temporal waveforms (such as audio signals) can be assumed centered and this remains true for their Fourier coefficients by linearity of the transformation. This explains the zero-mean assumption in the GCM on the one hand. The low-rank variance structure on the other hand underlies a composite signal structure that makes the model relevant for decomposition task. Indeed, introducing the latent complex-valued components y_{kfn} , Eq. (2) is equivalent to

$$y_{fn} = \sum_k y_{kfn}, \quad (3)$$

$$y_{kfn} \sim N_c(0, w_{fk}h_{kn}). \quad (4)$$

The latent component \mathbf{Y}_k with coefficients $\{y_{kfn}\}_{fn}$ reflects the contribution of the spectral pattern \mathbf{w}_k , the k^{th} column of \mathbf{W} , amplitude-modulated in time by the activation coefficients of the k^{th} row of \mathbf{H} .

Under these assumptions, the negative log-likelihood $-\log p(\mathbf{Y}|\mathbf{W}, \mathbf{H})$ is equal, up to a constant, to the Itakura-Saito (IS) divergence $D_{\text{IS}}(\mathbf{P}|\mathbf{WH})$ between the power spectrogram $\mathbf{P} = |\mathbf{Y}|^2$ and \mathbf{WH} . The IS divergence between nonnegative matrices \mathbf{A} and \mathbf{B} is defined by

$$D_{\text{IS}}(\mathbf{A}|\mathbf{B}) = \sum_{ij} \frac{a_{ij}}{b_{ij}} - \log \frac{a_{ij}}{b_{ij}} - 1. \quad (5)$$

The GCM is a step forward from traditional NMF approaches that fail to provide a valid generative model of the STFT itself – other approaches have only considered probabilistic models of the magnitude spectrogram under Poisson or multinomial assumptions, see [1] for a review. Still, the GCM is not yet a generative model of the raw signal $x(t)$ itself, but of its STFT. LRTFS fills in this ultimate gap.

III. LOW-RANK TIME-FREQUENCY SYNTHESIS (LRTFS)

In this section we first present LRTFS for complex-valued signals, closely following [16]. Then we rigorously address the case of real-valued signals represented as a complex-valued linear combination of complex-valued t-f atoms (such as Gabor atoms) with Hermitian symmetry. Finally, we discuss relevant connections with the state-of-the-art and illustrate the potential of LRTFS on an audio example.

¹A random variable x has distribution $N_c(x|\mu, \lambda) = (\pi\lambda)^{-1} \exp(-(|x - \mu|^2/\lambda))$ if and only if its real and imaginary parts are independent and with distribution $N(\Re[\mu], \lambda/2)$ and $N(\Im[\mu], \lambda/2)$, respectively.

A. Complex-valued signals

1) *Model*: Let $x(t)$ denote a complex-valued signal of length T and $\{\phi_{fn}(t)\}_{f=1..F, n=1..N}$ denote a dictionary of complex-valued t-f atoms of length T . LRTFS is defined as follows. For $t = 1, \dots, T$, $f = 1, \dots, F$, $n = 1, \dots, N$:

$$x(t) = \sum_{fn} \alpha_{fn} \phi_{fn}(t) + e(t), \quad (6)$$

$$\alpha_{fn} \sim N_c(0, [\mathbf{WH}]_{fn}), \quad (7)$$

$$e(t) \sim N_c(0, \lambda), \quad (8)$$

where $\{\alpha_{fn}\}$ are the complex-valued synthesis coefficients, \mathbf{W} and \mathbf{H} are nonnegative matrices of sizes $F \times K$ and $K \times N$, respectively, and $e(t)$ is an additive complex-valued residual term with Gaussian distribution $N_c(0, \lambda)$. The synthesis coefficients $\{\alpha_{fn}\}$ are furthermore assumed independent given \mathbf{W} and \mathbf{H} . The synthesis coefficients are dual of the analysis coefficients, defined by $y_{fn} = \sum_t x(t) \phi_{fn}^H(t)$, where \cdot^H denotes conjugate transpose. IS-NMF assumes that the *analysis* coefficients follow a GCM, see Eq. (2). In contrast, LRTFS assumes that the *synthesis* coefficients follow a GCM, as given by Eq. (7). As announced, LRTFS provides a generative model of the raw data $x(t)$, where IS-NMF only provides a generative model of the transformed data \mathbf{Y} .

Let us denote by \mathbf{x} and \mathbf{e} the column vectors of size T with coefficients $x(t)$ and $e(t)$, respectively. Let \mathcal{I} be an arbitrary one-to-one “vectorizing” mapping from $(f, n) \in \{1, \dots, F\} \times \{1, \dots, N\}$ to $m = \mathcal{I}(f, n) \in \{1, \dots, M\}$, where $M = FN$. We denote by $\boldsymbol{\alpha}$ the column vector of dimension M with coefficients $\alpha_m = \alpha_{\mathcal{I}(f, n)} = \alpha_{fn}$. We are abusing the notations by indexing the synthesis coefficients by either m (unstructured vectorized form) or (f, n) (matrix form where f indexes frequencies and n indexes time frames). It should be understood that m and (f, n) are in one-to-one correspondence and the meaning should be clear from the context. The notation α_m discards the inherent t-f structure of the coefficients while the notation α_{fn} makes it explicit. Despite abusing, this convention allows to significantly reduce cluttering in the following. Similarly, we denote by Φ the matrix of size $T \times M$ with columns $\phi_m = \phi_{fn}$, where ϕ_{fn} is the column vector of size T with coefficients $\phi_{fn}(t)$. We denote by \mathbf{v} the column vector of dimension M with coefficients $v_m = v_{fn} \stackrel{\text{def}}{=} [\mathbf{WH}]_{fn}$. We will sometimes write $\mathbf{v} = \text{vect}[\mathbf{WH}]$, where $\text{vect}[\cdot]$ refers to the vectorizing operator induced by \mathcal{I} . Equipped with these notations, we may write Eq. (6) and (7) as

$$\mathbf{x} = \Phi \boldsymbol{\alpha} + \mathbf{e}, \quad (9)$$

$$\boldsymbol{\alpha} \sim N_c(\mathbf{0}, \text{diag}(\mathbf{v})), \quad (10)$$

$$\mathbf{e} \sim N_c(\mathbf{0}, \lambda \mathbf{I}_T). \quad (11)$$

Ignoring the low-rank structure of \mathbf{v} , Eqs. (9)-(11) resemble sparse Bayesian learning (SBL), as introduced in [19], [20], where it is shown that marginal likelihood estimation of the variance induces sparse solutions of \mathbf{v} (and as a consequence, of $\boldsymbol{\alpha}$). The essential difference between our model and SBL is that the coefficients are no longer unstructured in LRTFS.

Indeed, in SBL, each coefficient α_m has a free variance parameter v_m . This property is fundamental to the sparsity-inducing effect of SBL [19]. In contrast, in LRTFS, the variances are now tied together and such that $v_m = v_{fn} = [\mathbf{WH}]_{fn}$.

2) *Maximum joint likelihood estimation* : We now address the estimation of \mathbf{W} , \mathbf{H} and α and possibly λ in LRTFS. We consider maximum joint likelihood estimation (MJLE), also referred to as type-I maximum likelihood estimation in [20]. MJLE relies on the minimization of the following objective function:

$$C_{\text{JL}}(\alpha, \mathbf{W}, \mathbf{H}, \lambda) \stackrel{\text{def}}{=} -\log p(\mathbf{x}, \alpha | \mathbf{W}, \mathbf{H}, \lambda) \quad (12)$$

$$= -\log p(\mathbf{x} | \alpha, \lambda) - \log p(\alpha | \mathbf{W}, \mathbf{H}) \quad (13)$$

$$= \frac{1}{\lambda} \|\mathbf{x} - \Phi \alpha\|_2^2 + \sum_{fn} \left[\frac{|\alpha_{fn}|^2}{[\mathbf{WH}]_{fn}} + \log [\mathbf{WH}]_{fn} \right] + cst \quad (14)$$

$$= \frac{1}{\lambda} \|\mathbf{x} - \Phi \alpha\|_2^2 + D_{\text{IS}}(|\alpha|^2 | \mathbf{v}) + \log(|\alpha|^2) + cst \quad (15)$$

where $cst = T \log \lambda + (T + M) \log \pi$ and we recall that $\mathbf{v} = \text{vect}[\mathbf{WH}]$.

Another possible estimation procedure for LRTFS is maximum marginal likelihood estimation (MMLE), also referred to as type-II maximum likelihood estimation in [20]. It relies on the minimization of $-\log p(\mathbf{x} | \mathbf{W}, \mathbf{H}, \lambda)$, i.e., involves the marginalization of α from the joint likelihood, following the principle of SBL. We considered MMLE for LRTFS in [16] and presented a valid EM algorithm. However our implementation does not scale with the dimensions involved in signal processing, as it requires the estimation of the diagonal elements of the inverse of a $M \times M$ matrix. Large-scale algorithms for MMLE are left as future work.

3) *Alternate minimization algorithm for MJLE* : We now describe an alternate minimization algorithm that returns stationary points of $C_{\text{JL}}(\theta)$, where $\theta = \{\alpha, \mathbf{W}, \mathbf{H}, \lambda\}$. The optimization of α given the other parameters reduces to

$$\min_{\alpha \in \mathbb{C}^M} \frac{1}{\lambda} \|\mathbf{x} - \Phi \alpha\|_2^2 + \sum_{fn} \frac{|\alpha_{fn}|^2}{[\mathbf{WH}]_{fn}} \quad (16)$$

which defines a convex ridge regression problem. The problem has the closed-form solution

$$\hat{\alpha} = \left[\Phi^H \Phi + \lambda \text{diag}(\mathbf{v})^{-1} \right]^{-1} \Phi^H \mathbf{x} \quad (17)$$

which involves the resolution of a linear system of size $M \times M$. The linear system can be reduced to dimension T thanks to the Woodbury identity, but this is still too large in typical signal processing applications. Computing Eq. (17) can be done efficiently with a numerical optimization procedure and several options are available, such as conjugate gradient descent (GCD), expectation-minimization (EM), forward-backward optimization or majorization-minimization. The latter three are closely related and lead in the present case to an iterative shrinkage algorithm (ISA) [21], [22]. We used in our implementation a complex-valued version of ISA, similar to the complex-valued cases treated in [23], [24], and using the acceleration described in [25]. This leads to a simple and parameter-free implementation with satisfactory

speed of convergence. This in particular results in a faster algorithm than the original EM algorithm presented in our initial contribution [16]. Moreover, ISA leads in our case to a simpler algorithm than GCD, with same theoretical and practical speed of convergence [26]. The resulting updates are given in Algorithm 1. The value of the inverse step-size L should be set to the maximum eigenvalue of $\Phi^H \Phi$, i.e., the squared spectral norm of Φ . If this value is not available in closed form or difficult to compute, a larger value $L \geq \|\Phi\|_2^2$ is also permissible but will result in smaller step sizes. In Algorithm 1, the operations $\mathbf{A} \circ \mathbf{B}$, $\mathbf{A}^{\circ p}$ and $\frac{\mathbf{A}}{\mathbf{B}}$ denote entry-wise multiplication, exponentiation and division, respectively.

The optimization of \mathbf{W} and \mathbf{H} given α reduces to

$$\min_{\mathbf{W}, \mathbf{H} \geq 0} \sum_{fn} D_{\text{IS}}(|\alpha_{fn}|^2 | [\mathbf{WH}]_{fn}) \quad (18)$$

which defines a IS-NMF problem with input matrix $\mathbf{S} = [|\alpha_{fn}|^2]_{fn}$. This a non-convex problem that is generally approached with alternating updates of \mathbf{W} and \mathbf{H} and majorization-minimization (MM) [27]. This results in the multiplicative updates given in Algorithm 1.

Finally, the optimization of λ given α is trivially given by $\hat{\lambda} = \|\mathbf{x} - \Phi \alpha\| / T$. However, the MJLE setting is known to be inefficient for the estimation of both the variance parameters of α and of \mathbf{e} , with either $\Phi \hat{\alpha}$ or $\hat{\mathbf{e}}$ capturing most of the signal variance. As such, though the estimation of λ is possible in principle, we will consider λ to be a fixed hyper-parameter in the following.

The objective function C_{JL} being non-convex and because we are using an alternate minimization algorithm, the output of Algorithm 1 depends on the initialization. In all simulations we initialized the synthesis coefficients α with the analysis coefficients $\Phi^H \mathbf{x}$. The matrices \mathbf{W} and \mathbf{H} are initialized using the absolute values of the complex SVD of the synthesis coefficients [28]. Finally, a tempering strategy with warm restart is used to speed up convergence for small target values of λ . The hyper-parameter λ is set to an arbitrarily large value in the first iterations and is then gradually decreased to its target value, as proposed in [29]. Convergence of the main and inner loops is monitored using the relative difference in norm between successive parameter iterates, see Section III-D.

4) *Reconstruction of the latent components*: Algorithm 1 outputs estimates of α , \mathbf{W} , \mathbf{H} . The approximate signal can directly be recovered from the estimated synthesis coefficients as $\hat{\mathbf{x}} = \Phi \hat{\alpha}$. LRTFS further assumes that the synthesis coefficients follow a GCM, see Eq. (7). As such, α_{fn} may be written as a sum of Gaussian latent components, such that $\alpha_{fn} = \sum_k \alpha_{kfn}$, with $\alpha_{kfn} \sim N_c(0, w_{fk} h_{kn})$. Denoting by α_k the column vector of dimension M with coefficients $\{\alpha_{kfn}\}_{fn}$, Eq. (9) may be written as

$$\mathbf{x} = \sum_k \Phi \alpha_k + \mathbf{e} = \sum_k \mathbf{c}_k + \mathbf{e}, \quad (19)$$

where $\mathbf{c}_k = \Phi \alpha_k$. The component \mathbf{c}_k is the ‘‘temporal expression’’ of spectral pattern \mathbf{w}_k , the k^{th} column of \mathbf{W} . Given estimates of α , \mathbf{W} and \mathbf{H} , the components may be

Algorithm 1: Alternate minimization for LRTFS

```

Set  $L = \|\Phi\|_2^2$  (or a larger value)
Compute the synthesis coefficients  $\mathbf{y} = \Phi^H \mathbf{x}$  (with
matrix form  $\mathbf{Y}$ )
Set  $\alpha^{(0)} = \mathbf{y}$ 
Initialize  $\mathbf{W}^{(0)}$  and  $\mathbf{H}^{(0)}$  with the absolute values of the
complex SVD of  $\mathbf{Y}$ 
Set  $i = 0$ 
repeat
  %% Update  $\mathbf{W}$  and  $\mathbf{H}$  with MM
  Compute spectrogram  $\mathbf{S}^{(i)} = \left[ |\alpha_{fn}^{(i)}|^2 \right]_{fn}$ 
  Initialize inner loop:  $\mathbf{W} = \mathbf{W}^{(i)}$ ,  $\mathbf{H} = \mathbf{H}^{(i)}$ 
  repeat
     $\mathbf{W} \leftarrow \mathbf{W} \circ \frac{[\mathbf{S}^{(i)} \circ (\mathbf{W}\mathbf{H})^{\circ-2}] \mathbf{H}^T}{[(\mathbf{W}\mathbf{H})^{\circ-1}] \mathbf{H}^T}$ 
     $\mathbf{H} \leftarrow \mathbf{H} \circ \frac{\mathbf{W}^T [\mathbf{S}^{(i)} \circ (\mathbf{W}\mathbf{H})^{\circ-2}]}{\mathbf{W}^T [(\mathbf{W}\mathbf{H})^{\circ-1}]}$ 
  until convergence;
  Leave inner loop:  $\mathbf{W}^{(i+1)} = \mathbf{W}$ ,  $\mathbf{H}^{(i+1)} = \mathbf{H}$ 
  Set  $\mathbf{v}^{(i+1)} = \text{vect}[\mathbf{W}^{(i+1)} \mathbf{H}^{(i+1)}]$ 

  %% Update  $\alpha$  with accelerated ISA
  Initialize inner loop:  $\mathbf{a}^{(0)} = \mathbf{z}^{(0)} = \alpha^{(i)}$ 
  Set  $j = 0$ 
  repeat
    % Descend
     $\mathbf{z}^{(j+1/2)} = \mathbf{a}^{(j)} + \frac{1}{L} \Phi^H (\mathbf{x} - \Phi \mathbf{a}^{(j)})$ 
    % Shrink
     $\mathbf{z}^{(j+1)} = \frac{\mathbf{v}^{(i+1)}}{\mathbf{v}^{(i+1)} + \lambda/L} \circ \mathbf{z}^{(j+1/2)}$ 
    % Accelerate
     $\mathbf{a}^{(j+1)} = \mathbf{z}^{(j+1)} + \frac{j+1}{j+5} (\mathbf{z}^{(j+1)} - \mathbf{z}^{(j)})$ 
     $j \leftarrow j + 1$ 
  until convergence;
  Leave inner loop:  $\alpha^{(i+1)} = \mathbf{z}^{(j+1)}$ 
until convergence;

```

reconstructed a posteriori in various ways. A natural choice is $\hat{\alpha}_k^{\text{MMSE}} = \Phi \hat{\alpha}_k^{\text{MMSE}}$ with

$$\hat{\alpha}_k^{\text{MMSE}} \stackrel{\text{def}}{=} \mathbb{E}[\alpha_k | \mathbf{x}, \hat{\theta}] = \mathbb{E}[\alpha_k | \hat{\alpha}, \hat{\mathbf{W}}, \hat{\mathbf{H}}]. \quad (20)$$

The coefficients of $\hat{\alpha}_k^{\text{MMSE}}$ are given by

$$\hat{\alpha}_{kfn}^{\text{MMSE}} = \frac{\hat{w}_{fk} \hat{h}_{kn}}{[\hat{\mathbf{W}}\hat{\mathbf{H}}]_{fn}} \hat{\alpha}_{fn}. \quad (21)$$

Using this estimate, the latent components are reconstructed by applying a t-f dependent ‘‘Wiener mask’’ to the synthesis coefficients. This procedure and the expression of $\hat{\alpha}_{kfn}^{\text{MMSE}}$ is analog to the standard Wiener estimate of the latent components in IS-NMF applied to $|\mathbf{Y}|^2$ [8] and given by

$$\hat{y}_{kfn}^{\text{MMSE}} = \frac{\hat{w}_{fk} \hat{h}_{kn}}{[\hat{\mathbf{W}}\hat{\mathbf{H}}]_{fn}} y_{fn}. \quad (22)$$

The estimate $\hat{\alpha}$ is used as an intermediate variable in the expression of $\hat{\alpha}_{kfn}^{\text{MMSE}}$ given by Eq. (21). Another possible estimate, which marginalizes α , is

$$\hat{\alpha}_k = \mathbb{E}[\alpha_k | \mathbf{x}, \hat{\mathbf{W}}, \hat{\mathbf{H}}, \hat{\lambda}], \quad (23)$$

where $\hat{\mathbf{v}}_k$ is the vector of dimension M with coefficients $\{\hat{w}_{fk} \hat{h}_{kn}\}_{fn}$. The input of the estimator is now the raw data \mathbf{x} which may be more sensible. The expression of this alternative estimate can be derived in closed-form but the resulting expression involves the large-scale inversion of $T \times T$ matrices, which is hardly feasible in practice.

B. Real-valued signals

1) *Model:* In many signal processing settings the data is a real-valued signal $x(t)$ expressed as a linear combination of complex-valued t-f atoms with Hermitian symmetry. More specifically, the dictionary and synthesis coefficients are such that $\phi_{fn} = \phi_{(F-f)n}^*$ and $\alpha_{fn} = \alpha_{(F-f)n}^*$ for $f = 1, \dots, F/2$ (assuming F to be even-valued for simplicity), where $*$ denotes conjugation. Under this particular structure, we have

$$\sum_{f=1}^F \sum_{n=1}^N \alpha_{fn} \phi_{fn}(t) = \sum_{f=1}^{F/2} \sum_{n=1}^N 2\Re[\alpha_{fn} \phi_{fn}(t)] \quad (24)$$

and we define real-valued LRTFS (rLRTFS) as follows. For $t = 1, \dots, T$, $f = 1, \dots, F/2$, $n = 1, \dots, N$:

$$x(t) = \sum_{f=1}^{F/2} \sum_{n=1}^N 2\Re[\alpha_{fn} \phi_{fn}(t)] + e(t), \quad (25)$$

$$\alpha_{fn} \sim N_c(0, [\mathbf{W}\mathbf{H}]_{fn}), \quad (26)$$

$$e(t) \sim N(0, \lambda). \quad (27)$$

Note how F now runs from 1 to $F/2$ instead of 1 to F . The synthesis coefficients α_{fn} remain complex-valued and the residual $e(t)$ becomes real-valued. \mathbf{W} and \mathbf{H} are nonnegative matrices of sizes $F/2 \times K$ and $K \times N$, respectively.

Let us now denote by $\underline{\alpha}$ and \mathbf{v} the vectors of dimension $M/2$ with coefficients α_{fn} and $v_{fn} = [\mathbf{W}\mathbf{H}]_{fn}$, respectively, and by $\underline{\Phi}$ the matrix of dimension $T \times M/2$ with columns ϕ_{fn} , for $f = 1, \dots, F/2$ and $n = 1, \dots, N$. With these notations we have $\underline{\alpha} = [\underline{\alpha}^T, \underline{\alpha}^H]^T$, $\underline{\Phi} = [\underline{\Phi}, \underline{\Phi}^*]$ and $\underline{\Phi} \underline{\alpha} = 2\Re[\underline{\Phi} \underline{\alpha}]$. Consequently, we may write Eq. (25)-(27) as

$$\mathbf{x} = 2\Re[\underline{\Phi} \underline{\alpha}] + \mathbf{e}, \quad (28)$$

$$\underline{\alpha} \sim N_c(\mathbf{0}, \text{diag}(\mathbf{v})), \quad (29)$$

$$\mathbf{e} \sim N(\mathbf{0}, \lambda \mathbf{I}_T). \quad (30)$$

2) *Estimation:* The MJLE objective function for rLRTFS writes

$$C_{\text{JL}}^{\Re}(\underline{\alpha}, \mathbf{W}, \mathbf{H}, \lambda) \stackrel{\text{def}}{=} -\log p(\mathbf{x}, \underline{\alpha} | \mathbf{W}, \mathbf{H}, \lambda) \quad (31)$$

$$= \frac{1}{2\lambda} \|\mathbf{x} - 2\Re[\underline{\Phi} \underline{\alpha}]\|_2^2 \quad (32)$$

$$+ \sum_{f=1}^{F/2} \sum_{n=1}^N \left[\frac{|\alpha_{fn}|^2}{[\mathbf{W}\mathbf{H}]_{fn}} + \log [\mathbf{W}\mathbf{H}]_{fn} \right] + cst \quad (33)$$

$$= \frac{1}{2\lambda} \|\mathbf{x} - 2\Re[\underline{\Phi} \underline{\alpha}]\|_2^2 + D_{\text{Is}}(|\underline{\alpha}|^2 | \mathbf{v}) + \log(|\underline{\alpha}|^2) + cst \quad (34)$$

where $cst = \frac{T}{2} \log(2\pi\lambda) + \frac{M}{2} \log \pi$. Using an alternate minimization setting like in Section III-A3, the updates of \mathbf{W} and \mathbf{H} are virtually unchanged. They amount to IS-NMF of the matrix form of the synthesis spectrogram $|\underline{\alpha}|^2$

(of size $F/2 \times N$). The update of λ is easily given by $\hat{\lambda} = \|\mathbf{x} - 2\Re[\underline{\Phi}\underline{\alpha}]\|_2^2/T$, but here again we prefer to treat λ as an hyper-parameter. The update of $\underline{\alpha}$ involves the following minimization problem:

$$\min_{\underline{\alpha} \in \mathbb{C}^{M/2}} F(\underline{\alpha}) \stackrel{\text{def}}{=} \frac{1}{2\lambda} \|\mathbf{x} - \Re[\underline{\Phi}\underline{\alpha}]\|_2^2 + \sum_{f=1}^{F/2} \sum_{n=1}^N \frac{|\alpha_{fn}|^2}{[\mathbf{WH}]_{fn}}. \quad (35)$$

The problem defined by Eq. (35) has a closed-form solution, with a less simpler expression than Eq. (17). The solution is still computationally demanding and the following numerical procedure is preferable.

Theorem 1 (Iterative shrinking algorithm for rLRTFS). *Let $L = \|\underline{\Phi}\|_2^2$ (with $\underline{\Phi} = [\underline{\Phi}, \underline{\Phi}^*]$) and $\underline{\alpha}^{(0)}$ be an initial estimate. The following sequence of updates converge to the global solution of problem (35):*

$$\underline{\alpha}^{(j+1/2)} = \underline{\alpha}^{(j)} + \frac{1}{L} \underline{\Phi}^H (\mathbf{x} - 2\Re[\underline{\Phi}\underline{\alpha}^{(j)}]), \quad (36)$$

$$\underline{\alpha}^{(j)} = \frac{\mathbf{v}}{\mathbf{v} + \lambda/L} \circ \underline{\alpha}^{(j+1/2)}. \quad (37)$$

Proof. The proof consists in reformulating Eq. (35) as a quadratic optimization problem over the real and imaginary parts of $\underline{\alpha}$ and applying ISA. Let $\mathbf{A} = 2[\Re[\underline{\Phi}], -\Im[\underline{\Phi}]]$, $\mathbf{b} = [\Re[\underline{\alpha}]^T \Im[\underline{\alpha}]^T]^T$ and $\mathbf{c} = \frac{1}{2}[\mathbf{v}^T \mathbf{v}^T]^T$. Then we may write

$$F(\underline{\alpha}) = F(\mathbf{b}) = \frac{1}{2\lambda} \|\mathbf{x} - \mathbf{A}\mathbf{b}\|_2^2 + \frac{1}{2} \sum_{m=1}^M \frac{b_m^2}{c_m}. \quad (38)$$

Denoting $L_{\mathbf{A}} = \|\mathbf{A}\|_2^2$, the ISA update for problem (38) writes [21], [22]

$$\mathbf{b}^{(j+1/2)} = \mathbf{b}^{(j)} + \frac{1}{L_{\mathbf{A}}} \mathbf{A}^T (\mathbf{x} - \mathbf{A}\mathbf{b}^{(j)}), \quad (39)$$

$$\mathbf{b}^{(j)} = \frac{\mathbf{c}}{\mathbf{c} + \lambda/L_{\mathbf{A}}} \circ \mathbf{b}^{(j+1/2)}. \quad (40)$$

Using the identities $\mathbf{A}\mathbf{b} = 2\Re[\underline{\Phi}\underline{\alpha}]$ and

$$\mathbf{A}^T \mathbf{e} = 2 \begin{bmatrix} \Re[\underline{\Phi}^H \mathbf{e}] \\ \Im[\underline{\Phi}^H \mathbf{e}] \end{bmatrix}, \quad (41)$$

Eqs. (39)-(40) can be rearranged in complex form as

$$\underline{\alpha}^{(j+1/2)} = \underline{\alpha}^{(j)} + \frac{2}{L_{\mathbf{A}}} \underline{\Phi}^H (\mathbf{x} - 2\Re[\underline{\Phi}\underline{\alpha}^{(j)}]) \quad (42)$$

$$\underline{\alpha}^{(j)} = \frac{\mathbf{v}}{\mathbf{v} + 2\lambda/L_{\mathbf{A}}} \circ \underline{\alpha}^{(j+1/2)}. \quad (43)$$

To complete the proof we only need to show that $L = 2L_{\mathbf{A}}$. Let $\underline{\alpha} = [\underline{\alpha}_1^T, \underline{\alpha}_2^T]^T$ be an eigenvector of $\underline{\Phi}^H \underline{\Phi}$ with maximum eigenvalue $L = \|\underline{\Phi}\|_2^2$. By definition we have

$$\begin{bmatrix} \underline{\Phi}^H \underline{\Phi} & \underline{\Phi}^H \underline{\Phi}^* \\ \underline{\Phi}^T \underline{\Phi} & \underline{\Phi}^T \underline{\Phi}^* \end{bmatrix} \begin{bmatrix} \underline{\alpha}_1 \\ \underline{\alpha}_2 \end{bmatrix} = L \begin{bmatrix} \underline{\alpha}_1 \\ \underline{\alpha}_2 \end{bmatrix}. \quad (44)$$

By taking the conjugate of Eq. (44), we easily show that $[\underline{\alpha}_2^H, \underline{\alpha}_1^H]^T$ is also an eigenvector with eigenvalue L . It follows that $[\underline{\alpha}_1 + \underline{\alpha}_2^H, \underline{\alpha}_2 + \underline{\alpha}_1^H]^T$ is also an eigenvector, which happens to have a Hermitian structure. We may thus impose $\underline{\alpha}_2 = \underline{\alpha}_1^*$

Algorithm 2: Alternate minimization for rLRTFS

Set $L = \|\underline{\Phi}\|_2^2$ (or a larger value)

Compute the synthesis coefficients $\mathbf{y} = \underline{\Phi}^H \mathbf{x}$ (with matrix form \mathbf{Y})

Set $\underline{\alpha}^{(0)} = \mathbf{y}$

Initialize $\mathbf{W}^{(0)}$ and $\mathbf{H}^{(0)}$ with the absolute values of the complex SVD of \mathbf{Y}

Set $i = 0$

repeat

 %% Update \mathbf{W} and \mathbf{H} with MM

 Compute spectrogram

$$\mathbf{S}^{(i)} = \left[|\underline{\alpha}_{fn}^{(i)}|^2 \right]_{f=1, \dots, F/2, n=1, \dots, N}$$

 Initialize inner loop: $\mathbf{W} = \mathbf{W}^{(i)}$, $\mathbf{H} = \mathbf{H}^{(i)}$

repeat

$$\mathbf{W} \leftarrow \mathbf{W} \circ \frac{[\mathbf{S}^{(i)} \circ (\mathbf{W}\mathbf{H})^{\circ-2}] \mathbf{H}^T}{[(\mathbf{W}\mathbf{H})^{\circ-1}] \mathbf{H}^T}$$

$$\mathbf{H} \leftarrow \mathbf{H} \circ \frac{\mathbf{W}^T [\mathbf{S}^{(i)} \circ (\mathbf{W}\mathbf{H})^{\circ-2}]}{\mathbf{W}^T [(\mathbf{W}\mathbf{H})^{\circ-1}]}$$

until convergence;

 Leave inner loop: $\mathbf{W}^{(i+1)} = \mathbf{W}$, $\mathbf{H}^{(i+1)} = \mathbf{H}$

 Set $\mathbf{v}^{(i+1)} = \text{vect}[\mathbf{W}^{(i+1)} \mathbf{H}^{(i+1)}]$

 %% Update $\underline{\alpha}$ with accelerated ISA

 Initialize inner loop: $\mathbf{a}^{(0)} = \mathbf{z}^{(0)} = \underline{\alpha}^{(i)}$

 Set $j = 0$

repeat

 % Descend

$$\mathbf{z}^{(j+1/2)} = \mathbf{a}^{(j)} + \frac{1}{L} \underline{\Phi}^H (\mathbf{x} - 2\Re[\underline{\Phi}\mathbf{a}^{(j)}])$$

 % Shrink

$$\mathbf{z}^{(j+1)} = \frac{\mathbf{v}^{(i+1)}}{\mathbf{v}^{(i+1)} + \lambda/L} \circ \mathbf{z}^{(j+1/2)}$$

 % Accelerate

$$\mathbf{a}^{(j+1)} = \mathbf{z}^{(j+1)} + \frac{j+1}{j+5} (\mathbf{z}^{(j+1)} - \mathbf{z}^{(j)})$$

$j \leftarrow j + 1$

until convergence;

 Leave inner loop: $\underline{\alpha}^{(i+1)} = \mathbf{z}^{(j+1)}$

until convergence;

and as such $\underline{\alpha} = [\underline{\alpha}^T, \underline{\alpha}^H]^T$. Then we have the following series of equivalences:

$$\underline{\Phi}^H \underline{\Phi} \underline{\alpha} = L \underline{\alpha} \iff \underline{\Phi}^H \underline{\Phi} \underline{\alpha} + \underline{\Phi}^H \underline{\Phi}^* \underline{\alpha}^* = L \underline{\alpha} \quad (45)$$

$$\iff 2\underline{\Phi}^H \Re[\underline{\Phi} \underline{\alpha}] = L \underline{\alpha} \quad (46)$$

$$\iff \frac{1}{2} \mathbf{A}^T \mathbf{A} \mathbf{b} = L \mathbf{b}. \quad (47)$$

As such, the spectra of $\underline{\Phi}^H \underline{\Phi}$ and $\mathbf{A}^T \mathbf{A}$ coincide up to a factor 2 and we have $L = 2L_{\mathbf{A}}$, which concludes the proof. ■

3) *Comments about implementation:* Eq. (36) and (37) can be accelerated like before and this results in the general procedure summarized in Algorithm 2. As compared to Algorithm 1, $\underline{\alpha}$ is essentially replaced by $\underline{\alpha}$, of size half, and the expression of $\mathbf{z}^{(i+1)}$ is changed with Eq. (36). Although the same notations are used for convenience, \mathbf{Y} , $\mathbf{S}^{(i)}$ and \mathbf{W} become matrices with $F/2$ rows.

Eq. (36) can be read as follows. The operation $2\Re[\underline{\Phi}\underline{\alpha}]$ consists of reconstructing an approximation $\hat{\mathbf{x}}$ of \mathbf{x} based on the current synthesis coefficients $\underline{\alpha}$. The operation $\underline{\Phi}^H \mathbf{e}$ then consists in computing the analysis coefficients (restricted to

“positive” frequencies, i.e., $f = 1, \dots, F/2$) of the current residual $\mathbf{e} = \mathbf{x} - \hat{\mathbf{x}}$. When Φ is a tight Gabor frame, these operations can be efficiently performed with dedicated time-frequency libraries, such as the MATLAB & Python Large Time-Frequency Analysis Toolbox (LTFAT) [30].² When the Gabor frame is tight, i.e., $\Phi\Phi^H\mathbf{x} = \mathbf{x}$, Φ has a unit spectral norm and we may set $L = 1$. A MATLAB implementation of Algorithm 2 is available online.³

Finally, given estimates of α , \mathbf{W} and \mathbf{H} , latent component coefficients $\hat{\alpha}_k$ may be reconstructed like in Eq. (21), and then $\hat{\mathbf{c}}_k = 2\Re[\Phi\hat{\alpha}_k]$.

C. Related work

The closest to our work are probably the recent papers by Kameoka [9], [10] which addresses temporal models of the form $\mathbf{x} = \sum_k \mathbf{c}_k$, like Eq. (19), where the spectrograms of the latent components are approximately rank-one. In essence (and slightly simplifying) these papers address optimization problems of the form

$$\min_{\mathbf{c}_k, \mathbf{W}, \mathbf{H}} \sum_{fkn} D([\Phi^H \mathbf{c}_k]^2]_{fn} | w_{fk} h_{kn}) \quad \text{s.t.} \quad \mathbf{x} = \sum_k \mathbf{c}_k \quad (48)$$

where $|\Phi^H \mathbf{c}_k|^2$ is the spectrogram of \mathbf{c}_k , indexed by f and n according to the convention of Section III-A1, and $D(\cdot)$ is a divergence between nonnegative matrices (either the generalized Kullback-Leibler divergence or the quadratic cost in [9], [10]). Though very elegant in our opinion, the approaches of [9], [10] are still analysis-based and do not yet provide a fully generative synthesis-based model like LRTFS.

Another related trend of work are the approaches of [11]–[13] which essentially model $x(t)$ as a sum of variance-structured Gaussian processes. Using our notations, the model in [13] sets $N = T$ and assumes $x(t) = \sum_f \sqrt{[\mathbf{W}\mathbf{H}]_{ft}} \Re[s_f(t)e^{-j2\pi ft/F}]$ where $s_f(t)$ is a complex-valued Gaussian autoregressive sequence. In [11], [12], $x(t)$ is modeled as a short-time stationary process. The signal is segmented into overlapping temporal frames \mathbf{x}_n of size P , like in the first stage of a STFT. Each temporal frame \mathbf{x}_n is then assumed to follow a multivariate Gaussian distribution with covariance $\mathbf{R}_n = \sum_k h_{kn} \Sigma_k$, where Σ_k is a full covariance matrix of size $P \times P$ (real-valued and symmetric). Estimation then consists in estimating the set of parameters $\{\Sigma_k\}_k$ and \mathbf{H} from the entire set of temporal frames (the whole approach is coined PSDTF in [12]). When it is further assumed that Σ_k is the covariance matrix of a real-valued latent stationary process, Σ_k becomes circulant and is diagonalized by the discrete Fourier transform (DFT). In that case, it can be shown that PSDTF specializes to IS-NMF [11], [12]. PSDTF remains close to the analysis view of IS-NMF: the raw data is segmented into overlapping frames and each frame is individually assigned a covariance model. In contrast, LRTFS provides a generative model of the entire signal $x(t)$,

²The specific commands being of the like $\alpha = \text{dgtreal}(\mathbf{x}, \text{'parameters'})$ and $\mathbf{x} = \text{idgtreal}(\alpha, \text{'parameters'})$, where dgt stands for discrete Gabor transform.

³<https://www.irit.fr/~Cedric.Fevotte/extras/tsp2018/>. Future references to online material refer to this same url.

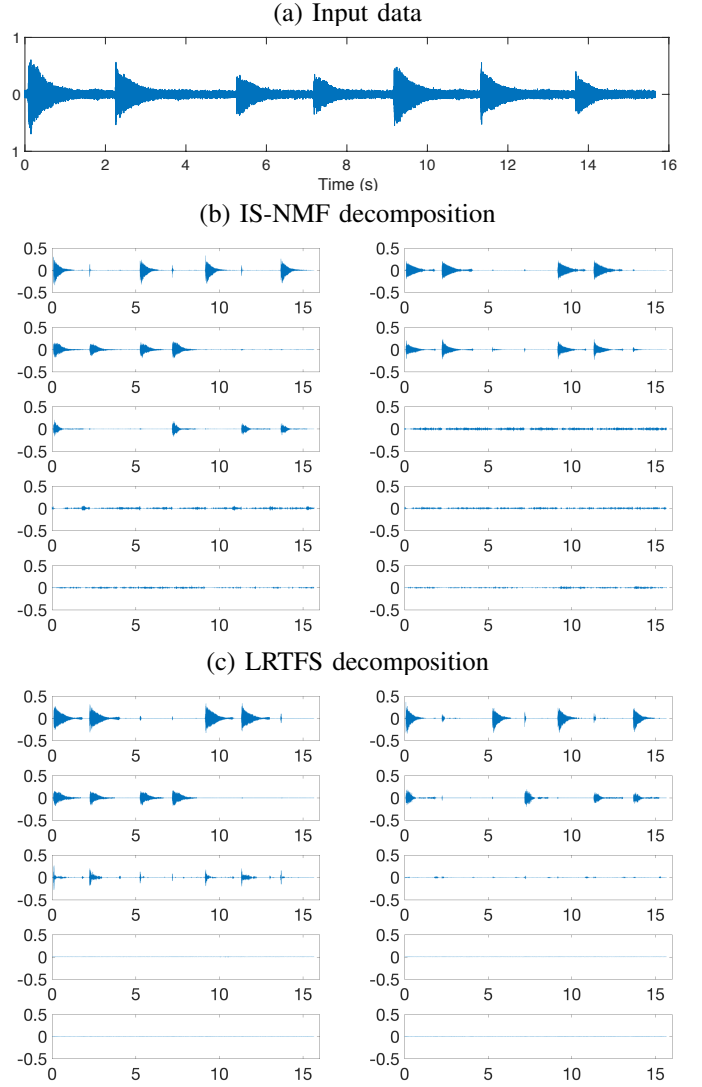


Fig. 1. Decomposition of a piano sequence consisting of four notes. The subplots in (b) display the latent components obtained by STFT inversion of Eq. (22). The subplots in (c) display the latent components $\hat{\mathbf{c}}_k^{\text{MSE}}$. The components are displayed by decreasing energy (from left to right and top to bottom).

assumed to be a linear combination of elementary t-f bricks endowed with a low-rank variance structure. In LRTFS, we assumed the synthesis coefficients to be conditionally mutually independent, i.e., $p(\alpha|\mathbf{W}, \mathbf{H}) = \prod_{fn} (\alpha_{fn} | [\mathbf{W}\mathbf{H}]_{fn})$. We could very well consider more sophisticated models similar to PSDTF or extensions [14], [15] which assume some correlation across frequencies or frames. Papers [11]–[13] describe NMF-related generative probabilistic models rooted in time series analysis while LRTFS offers a different perspective, rooted in the sparse approximation literature. In particular, LRTFS can be used with any time-frequency dictionary Φ , can easily accommodate multi-layer variants (see Section IV) or be considered for inverse problems (see Section V).

D. Example

We illustrate the performance of LRTFS compared to standard IS-NMF using the piano example used in [8]. The

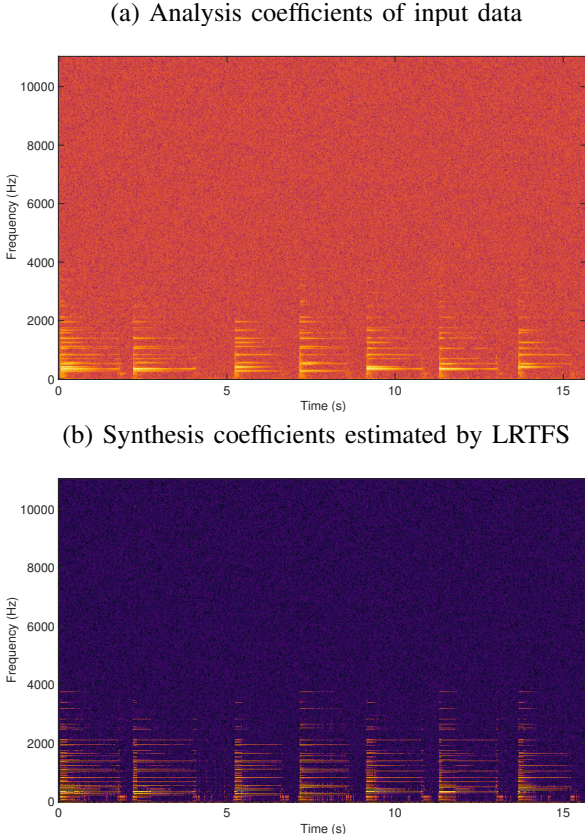


Fig. 2. Time-frequency analyses of the piano sequence displayed in Fig. 1 (a). Subplot (a) displays the squared magnitude of the analysis coefficients given by the STFT, i.e., the power spectrogram $|y_{fn}|^2$. Subplot (b) displays the squared magnitude of synthesis coefficients α estimated by LRTFS, i.e., $|\alpha_{fn}|^2$. Common dB scale is used on the two subplots.

sequence has a simple structure: four notes are played together at once in the first measure and are then played by pairs in all possible combinations in the subsequent measures. The duration is 15.6 s and the sampling rate 22050 Hz. In noise-free conditions and with appropriate initialization, standard IS-NMF is able to extract six identifiable latent components from this musical signal: the temporal expression of the four individual notes in a set of four components, the transient parts produced by the hammer hitting the strings in a fifth component and the sound produced by the sustain pedal when it is released in a sixth component [8]. We here consider a noisy example using additive white Gaussian noise with 20 dB input Signal to Noise Ratio (SNR). The resulting signal is displayed in Fig. 1 (a). A tight Gabor dictionary (with Hermitian symmetry) built on a Hann window of 1024 samples (46 ms) with 50% overlap is used for Φ . IS-NMF is applied to the analysis power spectrogram $|\underline{\Phi}\mathbf{x}|^2$ displayed in Fig. 2 (a). The number of latent components is arbitrarily set to $K = 10$ for both IS-NMF and rLTFS and the two methods are run from the same initialization (based on the SVD of \mathbf{Y} , see Algorithm 2). Iteration of the main and inner loops is stopped when the relative error between two successive parameter iterates falls under 10^{-5} . rLTFS is run with 30 different values of λ logarithmically equally spaced between 10^{-1} and 10^{-6} . The initialization described in Algorithm 2 was used for the first

value $\lambda = 10^{-1}$ and warm restart was used for the subsequent experiments. The CPU time for the total 30 experiments is ~ 5 min using a MATLAB implementation running on a Intel Core i5 processor. We show results corresponding to the value of λ that maximizes the output SNR given by

$$10 \log \frac{\|\hat{\mathbf{x}} - \mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2}. \quad (49)$$

Decomposition results are reported in Fig. 1. LRTFS is able to recover the four notes in the first four components, like standard IS-NMF in the noise free case, while the fifth component recovers the transient components produced by the hammer and the sustain pedal, corresponding to the last two components estimated by standard IS-NMF in the noise free case. As expected, the remaining five components are inaudible because of the denoising performed by LRTFS. In this noisy setting, IS-NMF fails to recover this transient part and splits the first note into two components. The input noise is spread over the five remaining components. Audio files are available online. Fig. 2 (b) displays the squared magnitude of the synthesis coefficients estimated by LRTFS. Denoising is clearly illustrated by the recovery of high frequencies.

IV. MULTI-LAYER LRTFS

Besides the advantage of modeling the raw signal itself, and not its STFT, another major strength of LRTFS is that it offers the possibility of multi-layer modeling. This means we may envisage models of the form

$$\mathbf{x} = \mathbf{x}_a + \mathbf{x}_b + \mathbf{e} = \Phi_a \alpha_a + \Phi_b \alpha_b + \mathbf{e} \quad (50)$$

where $\mathbf{x}_a = \Phi_a \alpha_a$ and $\mathbf{x}_b = \Phi_b \alpha_b$ are referred to as *layers*. This setting covers a variety of situations. Φ_a and Φ_b may be equal with α_a and α_b having a different structure. For example, α_a may follow a GCM like before and α_b may be given a sparsity-inducing prior. In such a case, multi-layer LRTFS offers a synthesis perspective to sparse + low-rank spectrogram decompositions, such as those presented in [31]–[33] which propose variants of robust principal component analysis (RPCA) [34] for spectral unmixing. Even more interestingly, the time-frequency dictionaries Φ_a and Φ_b may be chosen with different t-f resolutions. This yields so-called hybrid or morphological decompositions [35], [36], in which each layer may capture specific resolution-dependent structures. A typical audio example is transient + tonal decomposition: transient components are by nature adequately represented by a t-f dictionary with short time resolution while tonal components (such as the sustained parts of musical notes) are better represented by a t-f dictionary with larger time resolution (and as a consequence, finer frequency resolution). A variety of priors can be considered for α_a and α_b , such as frequency grouping for the transient synthesis coefficients and temporal grouping for the tonal synthesis coefficients [37].

A. Sparse and low-rank time-frequency synthesis

We consider a special case of multi-layer LRTFS that illustrates the potential of the synthesis approach. We present the methodology in the complex case for simplicity, but the

results can readily be adapted to the real case following the procedure described in Section III.

1) *Model*: Let Φ_a and Φ_b be time-frequency dictionaries consisting of atoms $\phi_{fn}^a(t)$ and $\phi_{fn}^b(t)$ with common dimension T and t-f pavings of size $F_a \times N_a$ and $F_b \times N_b$, respectively. We consider the following model, for $t = 1, \dots, T$:

$$x(t) = \sum_{f=1}^{F_a} \sum_{n=1}^{N_a} \alpha_{fn}^a \phi_{fn}^a(t) + \sum_{f=1}^{F_b} \sum_{n=1}^{N_b} \alpha_{fn}^b \phi_{fn}^b(t) + e(t) \quad (51)$$

$$\alpha_{fn}^a \sim N_c(0, [\mathbf{WH}]_{fn}), f = 1, \dots, F_a, n = 1, \dots, N_a \quad (52)$$

$$\alpha_{fn}^b \sim N_c(0, v_{fn}^b), f = 1, \dots, F_b, n = 1, \dots, N_b \quad (53)$$

$$e(t) \sim N_c(0, \lambda) \quad (54)$$

where $\{\alpha_{fn}^a\}$ and $\{\alpha_{fn}^b\}$ are the complex-valued synthesis coefficients, \mathbf{W} and \mathbf{H} are nonnegative matrices of sizes $F_a \times K$ and $K \times N_a$, respectively, $\{v_{fn}^b\}$ are nonnegative variance parameters and $e(t)$ is an additive complex-valued residual term. Eq. (51) is nothing but the scalar form of Eq. (50). Eq. (52) defines a GCM, while Eq. (53) defines the sparse-inducing prior that is used in SBL. Like before, we denote by \mathbf{v}^a and \mathbf{v}^b the column vectors with coefficients $[\mathbf{WH}]_{fn}$ and v_{fn}^b , respectively. Both \mathbf{v}^a and \mathbf{v}^b are parameters of a hierarchical variance model. Notice however how \mathbf{v}^b is a *free* parameter, while \mathbf{v}^a is structured through \mathbf{W} and \mathbf{H} . Overall, Eq. (51)-(54), define a multi-layer LRTFS model with latent low-rank t-f structure for layer \mathbf{x}_a and latent sparse t-f structure for layer \mathbf{x}_b .

Note that a wavelet dictionary could alternatively be used to represent the sparse layer: the estimation procedure presented next would apply in the exact same way. LRTFS however has to be supported by a regular t-f lattice and is not compatible with wavelets (in which time resolution decreases with frequency). LRTFS can however accommodate constant- Q t-f representations (constant time resolution, logarithmic frequency resolution) provided it can be inverted (near-accurate synthesis operator are proposed in [38], [39]).

B. Estimation

The negative log-likelihood of the data and parameters in model (51)-(54) is given by

$$\begin{aligned} -\log p(\mathbf{x}, \alpha_a, \alpha_b | \mathbf{W}, \mathbf{H}, \mathbf{v}^b) &= \frac{1}{\lambda} \|\mathbf{x} - \Phi_a \alpha_a - \Phi_b \alpha_b\|_2^2 \\ &+ D_{\text{IS}}(|\alpha_a|^2 | \mathbf{v}_a) + \log(|\alpha_a|^2) \\ &+ D_{\text{IS}}(|\alpha_b|^2 | \mathbf{v}_b) + \log(|\alpha_b|^2) + cst \end{aligned} \quad (55)$$

where $cst = T \log \lambda + (T + F_a N_a + F_b N_b) \log \pi$. Unfortunately, and similarly to the difficulty of estimating λ raised in Section III-A2, MJLE fails to evenly distribute the signal variance onto the two layers, and one of the two layers takes it all in practice. Such a problem can be mitigated using MMLE instead of MJLE, but again, MMLE is too costly in our setting. To solve this issue we introduce an extra hyper-parameter μ

that balances the contributions of each layer and propose to optimize the following objective

$$\begin{aligned} C_{\text{SLR}}(\boldsymbol{\theta}) &\stackrel{\text{def}}{=} \frac{1}{\lambda} \|\mathbf{x} - \Phi_a \alpha_a - \Phi_b \alpha_b\|_2^2 \\ &+ \mu [D_{\text{IS}}(|\alpha_a|^2 | \mathbf{v}_a) + \log(|\alpha_a|^2)] \\ &+ (1 - \mu) [D_{\text{IS}}(|\alpha_b|^2 | \mathbf{v}_b) + \log(|\alpha_b|^2)] + cst, \end{aligned} \quad (56)$$

where $0 \leq \mu \leq 1$, $\boldsymbol{\theta} = \{\alpha_a, \alpha_b, \mathbf{W}, \mathbf{H}, \mathbf{v}_b\}$ is the set of latent variables and parameters and SLR stands for ‘‘sparse + low-rank’’.

We may again find a stationary point of $C_{\text{SLR}}(\boldsymbol{\theta})$ by alternate minimization. The update of \mathbf{v}_b is trivially given by $\mathbf{v}_b = |\alpha_b|^2$. The update of \mathbf{W} and \mathbf{H} amounts to finding an IS-NMF of the synthesis spectrogram $|\alpha_{fn}^a|^2$ like in Algorithm 1. The synthesis coefficients α_a and α_b may be updated jointly via ridge regression over the joint dictionary $[\Phi_a, \Phi_b]$. This leads to the following updates

$$\hat{\mathbf{e}}^{(j)} = \mathbf{x} - \Phi_a \alpha_a^{(j)} - \Phi_b \alpha_b^{(j)} \quad (58)$$

$$\alpha_a^{(j+1/2)} = \alpha_a^{(j)} + \frac{1}{L} \Phi_a^H \hat{\mathbf{e}}^{(j)} \quad (59)$$

$$\alpha_b^{(j+1/2)} = \alpha_b^{(j)} + \frac{1}{L} \Phi_b^H \hat{\mathbf{e}}^{(j)} \quad (60)$$

$$\alpha_a^{(j+1)} = \frac{\mathbf{v}_a}{\mathbf{v}_a + \lambda/L} \circ \alpha_a^{(j+1/2)} \quad (61)$$

$$\alpha_b^{(j+1)} = \frac{\mathbf{v}_b}{\mathbf{v}_b + \lambda/L} \circ \alpha_b^{(j+1/2)} \quad (62)$$

where the inverse-step size should satisfy $L \geq \|[\Phi_a, \Phi_b]\|_2^2$. A convenient choice is $L = \|\Phi_a\|_2^2 + \|\Phi_b\|_2^2$. Eq. (58) computes the current residual, Eqs. (59) and (60) produce a step in the descent direction and Eqs. (61) and (62) shrink the resulting iterates.

C. Example

We use exactly the same data and setting as in Section III-D but we now add a sparse layer $\Phi_b \alpha_b$ to the LRTFS layer. Φ_b is set to be a tight Gabor dictionary built on a Hann window of 128 samples (6 ms) with 50% overlap. Φ_a is set as in Section III-D. The parameter μ was experimentally fixed to $\mu = 0.05$, and λ was again chosen among logarithmically spaced vales. Fig. 3 displays the 10 latent components characterizing the tonal layer and the transient layer. The components of the tonal layer are similar to those obtained from the single-layer LRTFS decomposition of Fig. 1. The fourth component captures part of the hammer attacks (especially from the first, most energetic note) with the shortest resolution components relegated to the transient layer $x^b(t)$ as expected. Audio files are available online.

V. COMPRESSIVE LRTFS

A striking advantage of LRTFS is that it may be used as a source model in inverse problems. For instance, LRTFS has been used in multichannel source separation in [40]. We here consider compressive sensing (CS) in which a source signal $x(t)$ must be recovered from $S \ll T$ random projections.




Fig. 3. Two-layer decomposition of the piano sequence displayed in Fig. 1.

Traditionally, CS exploits the sparsity of the synthesis coefficients of $x(t)$ onto a suitable dictionary. In this section we show that sparsity can be efficiently replaced with low-rankness, for the class of signals considered.

A. Model

Let us denote by $\mathbf{x} \in \mathbb{C}^T$ the vector source signal. The source is assumed to be sensed through the given linear operator $\mathbf{A} \in \mathbb{C}^{S \times T}$ ($S < T$), with output $\mathbf{b} \in \mathbb{C}^S$. We assume the following observation model:

$$\mathbf{b} = \mathbf{A}\mathbf{x} + \mathbf{e} \quad (63)$$

$$= \mathbf{A}\Phi\boldsymbol{\alpha} + \mathbf{e} \quad (64)$$

where $\Phi \in \mathbb{C}^{T \times M}$ is a given dictionary, $\boldsymbol{\alpha}$ are the synthesis coefficients of \mathbf{x} , and \mathbf{e} is a residual term that accounts for noise or model errors. Where traditional CS assumes some form of sparsity for $\boldsymbol{\alpha}$, we assume the synthesis coefficients to have the LRTFS low-rank structure described by Eq. (7). Like in traditional CS settings, we assume \mathbf{A} to be a random matrix. Finally, we assume \mathbf{e} to follow a complex Gaussian distribution like in Eq. (8).

B. Estimation

MJLE amounts to minimizing the following objective function:

$$\begin{aligned} C_{\text{CS}}(\boldsymbol{\alpha}, \mathbf{W}, \mathbf{H}) &= -\log p(\mathbf{b}, \boldsymbol{\alpha} | \mathbf{W}, \mathbf{H}, \lambda) \\ &= \frac{1}{\lambda} \|\mathbf{b} - \mathbf{A}\Phi\boldsymbol{\alpha}\|_2^2 + D_{\text{IS}}(|\boldsymbol{\alpha}|^2 | \mathbf{W}\mathbf{H}) + \log(|\boldsymbol{\alpha}|^2) + cst \end{aligned} \quad (65)$$

where $cst = T \log \lambda + (T + M) \log \pi$. The problem of optimizing $C_{\text{CS}}(\boldsymbol{\alpha}, \mathbf{W}, \mathbf{H})$ is equivalent to the one of optimizing $C_{\text{JL}}(\boldsymbol{\alpha}, \mathbf{W}, \mathbf{H})$ given by Eq. (15). In the complex case, the methodology developed in Section III-A3 can be readily

applied by replacing Φ with $\mathbf{M} = \mathbf{A}\Phi$. The spectral norm of \mathbf{M} may be difficult to derive or compute and we may set $L = \|\mathbf{A}\|_2 \|\Phi\|_2$ thanks to the inequality

$$\|\mathbf{A}\Phi\|_2 \leq \|\mathbf{A}\|_2 \|\Phi\|_2. \quad (66)$$

In the real case, i.e., when $\mathbf{x} \in \mathbb{R}^T$, the methodology developed in Section III-B may again be applied by assuming $\mathbf{A} \in \mathbb{R}^{S \times T}$ and replacing Φ with $\mathbf{M} = \mathbf{A}\Phi$. Posterior to estimation, an estimate of the original source is given by $\hat{\mathbf{x}} = \Phi\hat{\boldsymbol{\alpha}}$.

Note that we have addressed compressive sampling of real or complex-valued signals by exploiting a latent NMF-type t-f structure, which is different from compressive sampling of non-negative signals, a topic addressed for example in [41].

C. Example

We evaluate the recovery accuracy of the piano sequence used in Sections III-D and IV-C using a number of measurements S varying increasingly from $T/100$ to $T/10$. For this experiment, the length of the sequence remains 15.6 s but the sampling rate has been fixed at 11025 Hz because of the memory and computational complexities. The Gabor parameters have been adjusted accordingly with a Hann window of length 512 samples (46 ms) with 50% overlap.

We compare CS recovery methods based on LRTFS, SBL and ℓ_1 regularization, using a common alternating minimization setting (only the shrinkage or thresholding operators are changed). Note that we here consider type-I SBL (equivalent to MJLE) and not type-II (which again does not scale with the dimensions of our problem). The algorithms are initialized with $\boldsymbol{\alpha} = \mathbf{0}_{M \times 1}$. The first IS-NMF step of the LRTFS estimation was initialized with the absolute value of the complex-SVD as explained in Section III-A3. LRTFS was applied with $K = 10$ and the hyper-parameter λ was incrementally decreased from 10^3 to 10^{-2} . In addition, we provide the performance results of two oracles. In the first oracle, the vector of variances in Eq. (29) is set to the power spectrogram of the ground-truth uncompressed signal $x(t)$ (updates of \mathbf{W} , \mathbf{H} are thus removed). In the second oracle, we set the matrices \mathbf{W} , \mathbf{H} to their estimates returned by IS-NMF applied to the power spectrogram of $x(t)$. These two oracles allow one to evaluate the remaining gap between adaptive and optimal CS recovery.

Estimation accuracy was measured by means of output SNR. The results are displayed in Fig. 4 and show that LRTFS-based recovery improves accuracy by several dBs as compared to sparsity-based methods. This means that for this type of signals which are endowed with a strong low-rank t-f structure, there is a significant gain in exploiting low-rankness instead of mere unstructured sparsity for CS. Such a recovery approach is made possible thanks to the generative design of LRTFS. Fig. 5 displays the estimated components $\hat{\mathbf{c}}_k$ returned by LRFTS. It is interesting to note that only 4 components are meaningful. The first two notes are well recovered, like in the experiment of Section III-D, see Fig. 1 (b), while the two other notes are mixed in the third component. The fourth component still captures some transient information. We also run experiments for various values of the rank K in the case $S = 5T/100$. The recovery results appeared very robust to this parameter.




Fig. 4. Recovery of a compressively sensed piano sequence using LRTFS, SBL and ℓ_1 regularization, compared with two oracles.




Fig. 5. Latent components of the compressively sensed piano sequence recovered by LRTFS. The temporal components are displayed by decreasing energy (from left to right and top to bottom).

For $K \in \{5, 8, 10, 15, 20, 30\}$ the largest difference in the output SNRs was less than 0.5 dB. This robustness is partially explained by the fact that LRTFS tends to shrink irrelevant components, as explained by Eq. (14) and illustrated by Fig. 1. We believe that the deterministic initialization provided by SVD is another explanation.

Finally, we run the same CS experiment using the first 12 s of the song *Mamavatu* from Susheela Raman. The excerpt contains acoustic guitar and drums. Output SNRs are displayed on Fig. 6. Again, LRTFS recovery outperforms ℓ_1 regularization and SBL by several dBs which confirms the potential of the proposed model for audio inverse problems, where t-f low-rankness is a valid assumption.

VI. CONCLUSION

We have presented a new modeling paradigm that bridges t-f synthesis modeling and traditional analysis-based approaches. The proposed generative model allows in turn to design more sophisticated multi-layer representations that can efficiently capture diverse forms of structure. Additionally, the generative modeling allows to exploit NMF-like structure for compressive sensing which, to the best of our knowledge, is entirely new. Maximum joint likelihood estimation in the proposed models




Fig. 6. Recovery of the compressively sensed Mamavatu sequence using LRTFS, SBL and ℓ_1 regularization, compared with two oracles.

can be efficiently addressed using state-of-the-art iterative shrinkage and NMF algorithms. They can be efficiently implemented thanks to dedicated time-frequency analysis/synthesis packages. In this paper, we also addressed the modeling and decomposition of real signals in a rigorous way, which was missing from our preliminary contributions and appeared more tricky than initially expected.

The MLJE objective function (15) induced by the proposed generative modeling suggests more general problems of the form

$$C(\alpha, \mathbf{W}, \mathbf{H}, \lambda) = \frac{1}{\lambda} \|\mathbf{x} - \Phi\alpha\|_2^2 + D(|\alpha|^p | \mathbf{v}) \quad (67)$$

where $\mathbf{v} = \text{vec}[\mathbf{WH}]$, $D(\cdot)$ is an arbitrary divergence between nonnegative numbers and p is an arbitrary exponent. $D = D_{\text{IS}}$ and $p = 2$ follow naturally from the GCM assumptions but other choices could be more suitable for other families of signals or images. Such problems do not seem to have been addressed yet in the literature and offer stimulating optimization problems. The exact reconstruction case $\lambda = 0$ is also very interesting in itself. Another challenging line of research is the design of workable large-scale optimization algorithms for type-II maximum marginal likelihood estimation. As known from [20], such an estimator would be robust to the joint estimation of λ and \mathbf{v} , something in which MJLE fails in practice. The low-rank structure used in Eq. (7) to model the variance of the synthesis coefficients could be changed for more complex structures, such as neural architectures. This has been considered to model STFT synthesis coefficients in audio applications [42], using for example variational auto-encoders for training [43]. The framework presented in this paper can readily accommodate such variants, in particular in the multi-layer setting in which a layer can be assigned a pre-trained variance (for a specific class of signals such as speech) and another layer can be endowed with a free adaptive low-rank variance.

ACKNOWLEDGMENT

Cédric Févotte acknowledges support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No 681839 (project FACTORY).

REFERENCES

- [1] P. Smaragdis, C. Févotte, G. Mysore, N. Mohammadiha, and M. Hoffman, "Static and dynamic source separation using nonnegative factorizations: A unified view," *IEEE Signal Processing Magazine*, vol. 31, no. 3, pp. 66–75, May 2014.
- [2] T. Virtanen, "Monaural sound source separation by non-negative matrix factorization with temporal continuity and sparseness criteria," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 15, no. 3, pp. 1066–1074, Mar. 2007.
- [3] A. Ozerov and C. Févotte, "Multichannel nonnegative matrix factorization in convolutive mixtures for audio source separation," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 18, no. 3, pp. 550–563, Mar. 2010.
- [4] M. Elad, P. Milanfar, and R. Rubinstein, "Analysis versus synthesis in signal priors," *Inverse problems*, vol. 23, no. 3, p. 947, 2007.
- [5] P. Balazs, M. Doerfler, M. Kowalski, and B. Torrèsani, "Adapted and adaptive linear time-frequency representations: a synthesis point of view," *IEEE Signal Processing Magazine*, vol. 30, no. 6, pp. 20–31, 2013.
- [6] P. Sprechmann, R. Litman, T. B. Yakar, A. M. Bronstein, and G. Sapiro, "Supervised sparse analysis and synthesis operators," in *Advances in Neural Information Processing Systems*, 2013, pp. 908–916.
- [7] S. Nam, M. E. Davies, M. Elad, and R. Gribonval, "The cosparsity analysis model and algorithms," *Applied and Computational Harmonic Analysis*, vol. 34, no. 1, pp. 30–56, 2013.
- [8] C. Févotte, N. Bertin, and J.-L. Durrieu, "Nonnegative matrix factorization with the Itakura-Saito divergence. With application to music analysis," *Neural Computation*, vol. 21, no. 3, pp. 793–830, Mar. 2009.
- [9] H. Kameoka, "Multi-resolution signal decomposition with time-domain spectrogram factorization," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015.
- [10] —, "Complex NMF with the generalized Kullback-Leibler divergence," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017.
- [11] A. Liutkus, R. Badeau, and G. Richard, "Gaussian processes for under-determined source separation," *IEEE Transactions on Signal Processing*, vol. 59, no. 7, pp. 3155–3167, July 2011.
- [12] K. Yoshii, R. Tomioka, D. Mochihashi, and M. Goto, "Infinite positive semidefinite tensor factorization for source separation of mixture signals," in *Proc. International Conference on Machine Learning (ICML)*, 2013.
- [13] R. E. Turner and M. Sahani, "Time-frequency analysis as probabilistic inference," *IEEE Transactions on Signal Processing*, vol. 62, no. 23, pp. 6171–6183, Dec 2014.
- [14] A. Liutkus and K. Yoshii, "A diagonal plus low-rank covariance model for computationally efficient source separation," in *Proc. IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, 2017.
- [15] K. Yoshii, "Correlated tensor factorization for audio source separation," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2018.
- [16] C. Févotte and M. Kowalski, "Low-rank time-frequency synthesis," in *Advances in Neural Information Processing Systems (NIPS)*, Dec. 2014.
- [17] —, "Hybrid sparse and low-rank time-frequency signal decomposition," in *Proc. European Signal Processing Conference (EUSIPCO)*, Nice, France, Sep. 2015.
- [18] D. D. Lee and H. S. Seung, "Learning the parts of objects with nonnegative matrix factorization," *Nature*, vol. 401, pp. 788–791, 1999.
- [19] M. E. Tipping, "Sparse Bayesian learning and the relevance vector machine," *Journal of Machine Learning Research*, vol. 1, pp. 211–244, 2001.
- [20] D. P. Wipf and B. D. Rao, "Sparse bayesian learning for basis selection," *IEEE Transactions on Signal Processing*, vol. 52, no. 8, pp. 2153–2164, Aug. 2004.
- [21] M. Figueiredo and R. Nowak, "An EM algorithm for wavelet-based image restoration," *IEEE Transactions on Image Processing*, vol. 12, no. 8, pp. 906–916, 2003.
- [22] 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.
- [23] L. Chaàri, J.-C. Pesquet, A. Benazza-Benyahia, and P. Ciuciu, "A wavelet-based regularized reconstruction algorithm for SENSE parallel MRI with applications to neuroimaging," *Medical Image Analysis*, vol. 15, no. 2, pp. 185–201, 2011.
- [24] A. Florescu, E. Chouzenoux, J.-C. Pesquet, P. Ciuciu, and S. Ciochina, "A majorize-minimize memory gradient method for complex-valued inverse problems," *Signal Processing*, vol. 103, pp. 285–295, 2014.
- [25] A. Chambolle and C. Dossal, "On the convergence of the iterates of the fast iterative shrinkage/thresholding algorithm," *Journal of Optimization Theory and Applications*, vol. 166, no. 3, pp. 968–982, 2015.
- [26] A. N. A. Ben-Tal and A. Nemirovski, "Optimization III: Convex analysis, Nonlinear programming theory, Standard nonlinear programming algorithms," Lecture Notes, 2004.
- [27] C. Févotte and J. Idier, "Algorithms for nonnegative matrix factorization with the beta-divergence," *Neural Computation*, vol. 23, no. 9, pp. 2421–2456, Sep. 2011.
- [28] J. Becker, M. Menzel, and C. Rohlfing, "Complex SVD initialization for NMF source separation on audio spectrograms," *Proc. Deutsche Jahrestagung für Akustik (DAGA)*, 2015.
- [29] E. T. Hale, W. Yin, and Y. Zhang, "Fixed-point continuation for ℓ_1 -minimization: Methodology and convergence," *SIAM Journal on Optimization*, vol. 19, no. 3, pp. 1107–1130, 2008.
- [30] Z. Pruyva, P. L. Sondergaard, N. Holighaus, C. Wiesmeyer, and P. Balazs, "The large time-frequency analysis toolbox 2.0," in *Sound, Music, and Motion, Lecture Notes in Computer Science*. Springer, 2014, pp. 419–442.
- [31] P.-S. Huang, S. D. Chen, P. Smaragdis, and M. Hasegawa-Johnson, "Singing-voice separation from monaural recordings using robust principal component analysis," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2012.
- [32] Z. Chen and D. P. W. Ellis, "Speech enhancement by sparse, low-rank, and dictionary spectrogram decomposition," in *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2013.
- [33] C. Sun, Q. Zhu, and M. Wan, "A novel speech enhancement method based on constrained low-rank and sparse matrix decomposition," *Speech Communication*, vol. 60, pp. 44–55, 2014.
- [34] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *Journal of ACM*, vol. 58, no. 1, pp. 1–37, 2009.
- [35] L. Daudet and B. Torrèsani, "Hybrid representations for audiophonic signal encoding," *Signal Processing*, vol. 82, no. 11, pp. 1595 – 1617, 2002.
- [36] J.-L. Starck, Y. Moudden, J. Bobin, M. Elad, and D. Donoho, "Morphological component analysis," in *Optics & Photonics*, 2005.
- [37] M. Kowalski, "Sparse regression using mixed norms," *Applied and Computational Harmonic Analysis*, vol. 27, no. 3, pp. 303–324, 2009.
- [38] J. C. Brown, "Calculation of a constant Q spectral transform," *The Journal of the Acoustical Society of America*, vol. 89, no. 1, pp. 425–434, 1991.
- [39] D. Fitzgerald, M. Cranitch, and M. T. Cychowski, "Towards an inverse constant Q transform," in *Proc. Audio Engineering Society Convention*, 2006.
- [40] S. Leglaive, R. Badeau, and G. Richard, "Separating time-frequency sources from time-domain convolutive mixtures using non-negative matrix factorization," in *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2017.
- [41] P. D. O'Grady and S. T. Rickard, "Compressive sampling of non-negative signals," in *Proc. IEEE Workshop on Machine Learning for Signal Processing (MLSP)*, 2008.
- [42] Y. Bando, M. Mimura, K. Itoyama, K. Yoshii, and T. Kawahara, "Statistical speech enhancement based on probabilistic integration of variational autoencoder and non-negative matrix factorization," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2018.
- [43] D. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proc. International Conference on Learning Representations*, 2004.

PLACE
PHOTO
HERE

Cédric Févotte is a CNRS senior researcher at Institut de Recherche en Informatique de Toulouse (IRIT). Previously, he has been a CNRS researcher at Laboratoire Lagrange (Nice, 2013-2016) & Télécom ParisTech (2007-2013), a research engineer at Mist-Technologies (the startup that became Audionamix, 2006-2007) and a postdoc at University of Cambridge (2003-2006). He holds MEng and PhD degrees in EECS from École Centrale de Nantes. His research interests concern statistical signal processing and machine learning, for inverse problems and

source separation. He was a member of the IEEE Machine Learning for Signal Processing technical committee (2012-2018) and is a member of SPARS steering committee since 2018. He has been a member of the editorial board of the IEEE Transactions on Signal Processing since 2014, first as an associate editor and then as a senior area editor (from 2018). In 2014, he was the co-recipient of an IEEE Signal Processing Society Best Paper Award for his work on audio source separation using multichannel nonnegative matrix factorization. He is the principal investigator of the European Research Council project FACTORY (New paradigms for latent factor estimation, 2016-2021).

PLACE
PHOTO
HERE

Matthieu Kowalski received the engineering degree in computer science from the Université de Technologie de Compiègne in 2005, and the master degree in Mathematics Vision and Learning from the Ecole Normale Supérieure, Cachan, the same year. He received the PhD degree in applied mathematics from the University of Provence in 2008. His thesis was axed on sparse time-frequency decompositions. He is now an associate professor at the University of Paris-Sud, in the L2S Lab, and his research focuses on Inverse Problems and structured sparse

approximations. He is an elected member of the SPARS Steering committee since 2013.