

RESEARCH

Open Access



Sparse pursuit and dictionary learning for blind source separation in polyphonic music recordings

Sören Schulze^{1*}  and Emily J. King²

Abstract

We propose an algorithm for the blind separation of single-channel audio signals. It is based on a parametric model that describes the spectral properties of the sounds of musical instruments independently of pitch. We develop a novel sparse pursuit algorithm that can match the discrete frequency spectra from the recorded signal with the continuous spectra delivered by the model. We first use this algorithm to convert an STFT spectrogram from the recording into a novel form of log-frequency spectrogram whose resolution exceeds that of the mel spectrogram. We then make use of the pitch-invariant properties of that representation in order to identify the sounds of the instruments via the same sparse pursuit method. As the model parameters which characterize the musical instruments are not known beforehand, we train a dictionary that contains them, using a modified version of Adam. Applying the algorithm on various audio samples, we find that it is capable of producing high-quality separation results when the model assumptions are satisfied and the instruments are clearly distinguishable, but combinations of instruments with similar spectral characteristics pose a conceptual difficulty. While a key feature of the model is that it explicitly models inharmonicity, its presence can also still impede performance of the sparse pursuit algorithm. In general, due to its pitch-invariance, our method is especially suitable for dealing with spectra from acoustic instruments, requiring only a minimal number of hyperparameters to be preset. Additionally, we demonstrate that the dictionary that is constructed for one recording can be applied to a different recording with similar instruments without additional training.

Keywords: Blind source separation, Unsupervised learning, Dictionary learning, Pitch-invariance, Pattern matching, Sparsity, Stochastic optimization, Adam, Orthogonal matching pursuit

1 Introduction

1.1 Problem definition and approach

Source separation concerns the recovery of signals X_1, \dots, X_c from a mixture $X = X_1 + \dots + X_c$. We speak of *blind separation* when no specific prior information to characterize the sources of the signals is provided, especially not in the form of labeled training data. However, we do make *structural* assumptions about the signals; in our case, we assume that they follow the typical characteristics of tones from wind and string instruments.

In order to exploit this structure, it is helpful to regard a time-frequency representation (*spectrogram*), which subdivides the problem into smaller time frames and highlights the frequency characteristics of the signal. One simple spectrogram is obtained via the modulus of the short-time Fourier transform (STFT) (cf. [1]). However, in the STFT spectrogram, different pitch of the instrument tones manifests in linear scaling of the distances between the peaks on the frequency axis, which makes it computationally hard to identify the tones in the spectrum.

Thus, we apply a novel *sparse pursuit* algorithm that represents the time frames of the STFT spectrogram via a limited number of peaks, under the assumption that they originate from sinusoidal signals in the recording.

*Correspondence: sschulze@uni-bremen.de

¹AG Computational Data Analysis, Faculty 3, University of Bremen, Bibliothekstr. 5, 28359 Bremen, Germany
Full list of author information is available at the end of the article

We then place these identified peaks in a new spectrogram that has a logarithmic frequency axis and is therefore *pitch-invariant* (cf. Section 3). On this, we apply a *dictionary learning* algorithm, where the dictionary contains the learned relative amplitudes of the harmonics for each instrument. In an alternating loop, we identify the sounds of the instruments by now applying the sparse pursuit algorithm on time frames of the log-frequency spectrogram using the current value of the dictionary and then update the dictionary based on that identification. Both the problem of finding the peaks in the STFT spectrogram and the problem of finding the patterns representing the instrument sounds are generally underdetermined (cf. Section 1.3), so sparsity plays a crucial role in their regularization.

After training has finished, we apply the sparse pursuit algorithm on the entire log-frequency spectrogram in order to obtain the separated spectrograms, and after masking with the original mixture spectrogram, we employ the algorithm by Griffin and Lim [2] in order to convert them back into time-domain signals, using the phase of the original spectrogram as the initial value. The overall procedure is displayed in Fig. 1.

The novelty of our approach lies in the combination of pitch-invariant representations with a sparse pursuit algorithm during training: Provided that the characteristics of the sounds of the instruments are sufficiently stable, the relative amplitudes of the harmonics saved in the dictionary represent the sounds of the instruments at any arbitrary pitch, without making assumptions about their tuning or range. At the same time, the use of a log-frequency axis enables us to match the spectrogram with the modeled patterns of these sounds in an efficient manner, and due to a non-linear optimization step, the parameters are locally optimal on a continuous scale. As the outcome of the training is sometimes sensitive with respect to the initial dictionary, we typically use the method in an ensemble setting. Apart from the sparsity condition, there is no

need to hand-tune any hyperparameters for a specific recording.

The sparse pursuit algorithm that we propose is designed to match a generic sampled spectrum with shifted non-negative continuous patterns. While it was developed with audio frequency spectra in mind, it may be used in other signal processing applications as well.

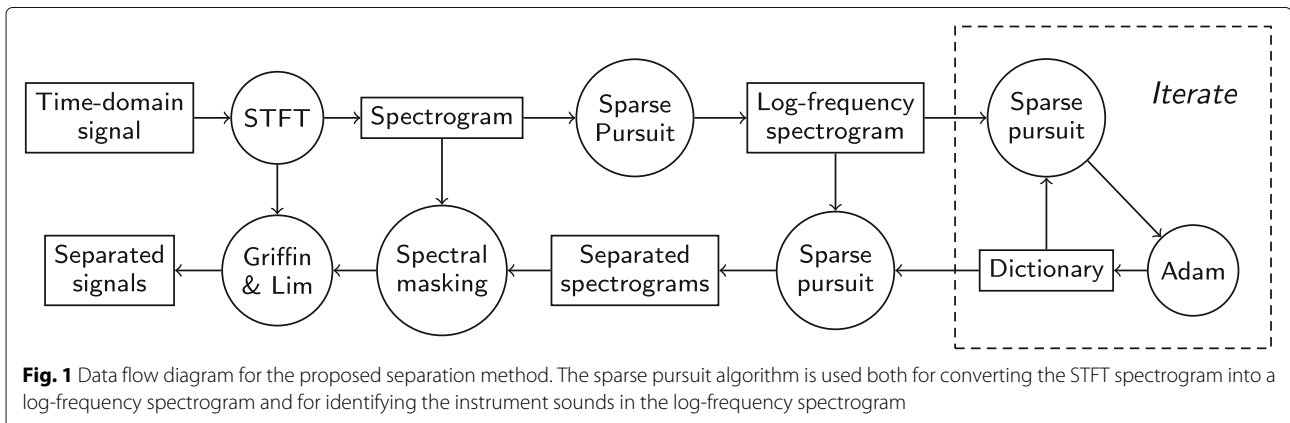
1.2 Related work

During the past years, audio source separation has become a very wide field, now incorporating a number of fundamentally different applications and approaches. A thorough overview can be found in books on the subject that have recently appeared [3–5].

The first instance of learning the harmonic structure of musical instruments via non-negative matrix factorization (NMF) [6] on spectrograms was by Smaragdis and Brown [7] for the purpose of polyphonic music transcription. This approach was then applied to audio source separation by Wang and Plumbley [8]. The algorithm learns a dictionary where each atom represents one instrument at a specific pitch. By estimating the tones of the instruments at specific points in time, it is thereby possible to reconstruct the contributions of the individual instruments. An overview of single-channel NMF-based methods can be found in [9].

In many cases, a single musical instrument can generate different sounds which are perceptually similar and only vary in the pitch of the tones. Using the *constant-Q transform* (CQT) [10] as a log-frequency spectrogram, Fitzgerald et al. [11] use non-negative tensor factorization to generate a dictionary containing the frequency spectra of different instruments, which can be shifted on a fixed grid of semitones in order to apply them to different notes. This approach was later refined by Jaiswal et al. [12–14].

The advantage of this representation is that it can be applied to a large variety of musical instruments, as long as pitch-invariance is fulfilled. The drawback is that it



requires the instruments to be tuned precisely to a known equal-temperament scale, which makes it impractical for real-world recordings with acoustic instruments.

Alternatively, the source separation problem on spectrograms can be formulated in probabilistic terms, which is done in the method of probabilistic latent component analysis (PLCA) [15, 16]. Here, the entire spectrogram is regarded as a probability distribution, which is then decomposed via expectation maximization (EM) into marginal distributions that depend on latent variables. In its original form, both the model and the numerics are identical to NMF, but it can be argued that the probabilistic notation is more powerful and especially beneficial when incorporating priors.

The latent variables can be chosen so that separation via PLCA is also pitch-invariant [17, 18], and it is also possible to model the harmonics explicitly [19–21]. Those algorithms operate in the discrete domain, so they effectively perform non-negative tensor factorization. In this formulation, the approach was pioneered by [22] for application in multiple pitch estimation (MPE).

Duan et al. [23] also follow a probabilistic approach, but with a more explicit model of the spectral structure of the harmonics of the instruments. They first use a peak detection algorithm in order to find the potential frequencies for the harmonics. Using a greedy maximum-likelihood model, the fundamental frequencies are estimated, and the harmonic patterns are clustered in order to assign them to certain instruments. This approach is interesting because it allows the representation of tones without a predetermined tuning.

In our algorithm, we apply a more advanced tone model that during optimization incorporates inharmonicity (cf. [24]) and also deviations in the width of the peaks, which may occur in case of volume changes. While we also preselect peaks, we only do so in order to generate a pitch-invariant log-frequency spectrogram that is suitable for wideband signals.

For narrowband signals, the CQT could be used instead. Alternatively, one could employ the mel spectrogram (cf. [3]) or the method proposed in [25], which combines favorable properties from both time-frequency representations. However, the resolution of any spectrogram that was computed via classical means is ultimately limited by the Heisenberg uncertainty principle (cf. [1, 26]).

The pitch-invariance property of the representation is important since it allows us to locate the sounds of the instruments via cross-correlation, making the determination of the fundamental frequencies much easier. However, rather than *explaining* the peaks in the spectrogram via a parametric model of the harmonic structure of the instruments via clustering, we use stochastic optimization to train a dictionary containing the relative amplitudes of the harmonics in order to *reproduce* their sounds.

In our model, we aim to be parsimonious in the number of parameters and, following the spirit of blind separation, also in the assumptions on the data. Therefore, we regard each time frame of the spectrogram independently. However, models that take the time axis into account do exist. Smaragdis [27] introduced *NMFD* (non-negative matrix factor deconvolution), which is NMF with convolution in time (again, a form of tensor factorization), and Schmidt and Mørup [28] combined time- and pitch-invariant approaches to *NMF2D* (non-negative matrix factor two-dimensional deconvolution). Virtanen [29] added a temporal sparsity criterion, and later, in [30], a temporal continuity objective. Blumensath and Davies [31] operate completely in the time domain, without any time-frequency representation.

The musical score that matches the piece in the recording is also a valuable piece of information, as it resolves ambiguities about the fundamental frequencies. Hennequin et al. [32] first proposed a pitch-invariant model that can accommodate local variation from predetermined tuning via gradient descent, but the authors faced the problem that this approach did not work on a global scale. Therefore, in [33], they use the score to give the algorithm hints about the approximate frequencies and thereby reduce the optimization problem to a local one. One of the main challenges in score-informed separation is the alignment of the score with the audio recording. For this, a combined approach has recently been proposed by Munoz-Montoro et al. [34].

Due to the growing interest in deep learning among the machine learning community, it is also applied to audio source separation in a supervised manner. However, this approach requires labeled training data. Huang et al. [35] proposed a deep recurrent neural network architecture and achieved respectable results. In the SiSEC (Signal Separation Evaluation Campaign) 2018 [36], different state-of-the-art algorithms were compared, and the reference implementation *Open-Unmix* [37] was determined as the overall winner. The network operates on magnitude spectrograms and combines different kinds of layers, including long short-term-memory (LSTM) units. Its performance was recently surpassed by Défossez et al. [38], whose network is based on LSTMs and convolutions, but operates directly in the time (i.e., waveform) domain.

Due to their good performance, supervised deep learning methods currently dominate the focus of many researchers. They make only very mild explicit prior structural assumptions on the data and instead rely on training to sort out the separation process. Thus, whenever appropriate training data is available, they make a very powerful and versatile tool.

Naturally, using more prior information in a machine learning problem typically improves the quality of the results. Conversely, purely blind approaches can only

work under very controlled conditions, and they have therefore received relatively little attention in recent years. We aim to show that progress on this problem is nevertheless still possible, and that even blind separation can profit from the modern machine learning techniques that have been developed.

Our sparse pursuit algorithm is a greedy approximation to ℓ_0 sparsity, based on concepts from orthogonal matching pursuit (OMP) [39] and subspace pursuit [40] while making use of the pitch-invariance of the time-frequency representation. However, a similar problem has been formulated in an ℓ_1 setting as *convolutional sparse coding* for image processing [41]. While it is relatively fast, the drawback of this method is that it is still limited to discrete convolutions. In *continuous basis pursuit* [42], this problem is approached by either Taylor or polar interpolation. *Beurling LASSO* [43, 44] first solves the sparse representation problem in the *dual* space, but finding the corresponding primal solution generally remains a challenge. Whereas the general advantage of ℓ_1 -based formulations lies in their convexity, greedy methods allow for a more flexible optimization step while keeping the dimensionality low.

1.3 The musical role of sparsity

The representation of the time frames of a spectrogram of a music recording with a pitch-invariant dictionary is in general not unique. If we consider wind and string instruments, their sound is dominated by a linear combination of sinusoids, which show up as horizontal lines in the spectrogram. Thus, there exists a trivial solution that assumes a single sinusoidal instrument which plays a large number of simultaneous tones. While this solution is valid, it is undesirable, as no separation is performed at all.

A similarly trivial solution is to construct different instruments for each time frame of the spectrogram. This, however, leaves us with the problem of matching the constructed instruments with the actual instruments. This process would need to be done either manually or via an appropriate clustering algorithm, such as the one used in [23]. Also, instruments which play harmonically related notes may be mistaken for a single instrument, and this case would need special consideration.

In order to attain meaningful solutions, we thus decide to limit both the total number of instruments and the number of tones that are assumed to be played at the same time. The former is controlled by the layout of the dictionary, while the latter is a sparsity condition that requires the use of appropriate algorithms.

The constraints imposed by these numbers are supposed to encourage solutions that will appear meaningful to a human listener. Good results can be achieved if both numbers are known and sufficiently low, but blind separation meets its conceptual limits in case of very polyphonic

works such as orchestral symphonies. One particularly difficult instrument would be the pipe organ, where the combination of organs stops blurs the borders of what should be considered a single instrument (cf. [24, 45]).

1.4 Structure of this paper

In Section 2, we propose a novel general-purpose sparse pursuit algorithm that matches a sampled spectrum with non-negative continuous patterns. The algorithm is a modified version of orthogonal matching pursuit (OMP) [39] with a non-linear optimization step for refinement.

In Section 3, we use this algorithm in order to convert an STFT magnitude spectrogram into a wideband pitch-invariant log-frequency spectrogram. In Section 4, we explain how we use the same algorithm (with slightly different parameter choices) and a dictionary representation of the harmonics in order to identify patterns of peaks related to the sounds of musical instruments in time frames of the spectrogram. Due to the non-linear optimization, we can identify the fundamental frequency, the width of the Gaussian, and the inharmonicity individually for each tone on a continuous scale.

In Section 5, we expound the learning algorithm: For the dictionary update, we employ a modified version of Adam [46], which is a popular stochastic gradient descent algorithm that was initially developed for the training of deep neural networks. Our modifications adapt this algorithm to dictionary learning, preserving the relative scaling of certain components of the gradient and periodically resetting parts of the dictionary as needed. In Section 6, we explain how we use the trained dictionary in order to perform the separation and obtain the audio signals for the separated instruments.

In Section 7, we apply our algorithm on mixtures that we recorded using acoustic instruments as well as on samples from the literature. We evaluate the performance of the overall algorithm via standard measures and discuss the results. We also provide spectrograms of the separation result.

A pseudo-code implementation of the algorithm as well as an additional elaboration about the choice of the time-frequency representation can be found in the appendix.

2 Sparse pursuit algorithm for shifted continuous patterns

For both the transformation of the spectrogram and the identification of instruments inside the spectrogram, we need an algorithm to approximate a non-negative discrete mixture spectrum $Y[s] \geq 0, s \in \mathbb{Z}$, via shifted versions of continuous patterns $y_{\eta,\theta}(s) \geq 0, s \in \mathbb{R}$. The exact meaning of the variables depends on the specific application, but in general, $\eta \in \{0, \dots, N_{\text{pat}} - 1\}$ is a discrete index, and $\theta \in \mathbb{R}^{N_{\text{par}}}$ is a set of continuous, real-valued parameters. The

fixed values $N_{\text{pat}}, N_{\text{par}} \in \mathbb{N}$ specify the number of patterns and the number of parameters in θ .

Mathematically speaking, we aim to identify amplitudes $a_j > 0$, shifts μ_j , indices η_j , and parameter sets θ_j such that:

$$Y[s] \approx \sum_j a_j y_{\eta_j, \theta_j}(s - \mu_j), \quad (1)$$

for $s \in \mathbb{Z}$.

For a preliminary intuition, y_{η_j, θ_j} can be understood as the spectrum of the instrument with the number η_j , and θ_j can contain additional parameters that influence the exact shape of the pattern, like the width of the peaks and the inharmonicity.

In order to formalize the approximation, we define a *loss function* to be minimized. The first natural choice for such a loss function is the ℓ_2 distance, but it is not ideal for use in magnitude frequency spectra, as it focuses very much on the high-volume parts of the spectrum, and the same applies to other ℓ_p (quasi-)distances for $p > 0$.

This problem is often approached by use of the β -divergence (cf. [9, 47]), which puts a high penalty on “unexplained” peaks in the spectrum. However, it is asymmetric, and while it is natural in NMF-based methods, it is difficult to integrate in the algorithm that we propose.

Instead, we remain with ℓ_2 , but we *lift* low-volume parts of the spectrum via a concave power function:

$$\begin{aligned} L(Y, (a_j), (\mu_j), (\eta_j), (\theta_j)) \\ = \sum_s \left((Y[s] + \delta)^q - \left(\delta + \sum_j a_j y_{\eta_j, \theta_j}(s - \mu_j) \right)^q \right)^2, \end{aligned} \quad (2)$$

with $q \in (0, 1]$, where $\delta > 0$ is a small number merely used to ensure differentiability.

Furthermore, we impose the *sparsity condition* that every value of η_j may only occur at most $N_{\text{spr}} \in \mathbb{N}$ times in the linear combination.

Minimizing L is a highly non-convex and partly combinatorial problem, so we cannot hope to reach the perfect solution. Instead, we follow a *greedy* approach, using ideas from orthogonal matching pursuit (OMP) [39] and sub-space pursuit [40].

We start with an empty index set \mathcal{J} and then run the following steps in a loop:

1. Compute the (discrete) cross-correlation between the residual

$$r[s] = Y[s]^q - \left(\sum_j a_j y_{\eta_j, \theta_j}(s - \mu_j) \right)^q \quad (3)$$

(i.e., the lifted difference between the raw spectrum and the current reconstruction) and the sampled

patterns. Assume a default parameter set θ_{nil} , and with

$$\rho[\mu, \eta] = \sum_i \frac{r[i] (y_{\eta, \theta_{\text{nil}}}[i - \mu])^q}{\|y_{\eta, \theta_{\text{nil}}}[\cdot]\|_2^q}, \quad (4)$$

preselect the $N_{\text{pre}} \in \mathbb{N}$ combinations

$(\mu, \eta) \in \mathbb{Z} \times \{0, \dots, N_{\text{pat}} - 1\}$ with the greatest $\rho[\mu, \eta]$, equip them with indices, and add those to the index set \mathcal{J} . For each preselected pair (μ_j, η_j) , initialize $a_j = (\rho[\mu_j, \eta_j] / \|y_{\eta_j, \theta_{\text{nil}}}[\cdot]\|_2^q)^{1/q}$. Skip the combinations for which a_j is non-positive. If none are left, terminate.

2. Do non-linear optimization on a_j, μ_j , and $\theta_j, j \in \mathcal{J}$, in order to minimize L , where $a_j \geq 0$ and $\theta_j \in \Omega_\theta$ with $\Omega_\theta \subseteq \mathbb{R}^{N_{\text{par}}}$.
3. For each $\eta = 0, \dots, N_{\text{pat}} - 1$, find the indices $j \in \mathcal{J}$ where $\eta_j = \eta$, and remove all but those with the N_{spr} highest amplitudes a_j such that, in the end, each pattern η is represented at most N_{spr} times in the index set \mathcal{J} .
Re-run the non-linear optimization procedure on the now smaller index set \mathcal{J} .
4. If the loss L has decreased by less than the factor of $1 - \lambda$ compared to the previous iteration, with $\lambda \in (0, 1]$, restore all the values from the previous iteration and return them as the result.
Otherwise, if the count of iterations has reached $N_{\text{itr}} \in \mathbb{N}$, return the current parameters. If this is not the case, do another iteration.

The hyperparameters N_{pat} and N_{spr} determine the number of given patterns and the maximum number of times that any pattern can be selected for the representation of a spectrum. Both are assumed to be known from the application. For the q exponent, we usually pick $q = 1/2$, as this is the lowest one to keep L convex in a_j , which is beneficial to the optimization procedure. In some cases, better results can be achieved by choosing the value of q even lower, but this also increases the chance of divergence.

Further, the hyperparameters λ and N_{itr} are safeguards to limit the runtime of the algorithm, such that the loop is not run indefinitely with marginal improvement in the non-linear optimization step. They also mitigate the problem of overfitting. The value of λ should be chosen slightly below 1; in practice, we find that $\lambda = 0.9$ yields good results. We limit the number of iterations to $N_{\text{itr}} = 2N_{\text{spr}}N_{\text{pat}}$, which is twice the overall sparsity level. The loop typically terminates due to insufficient decrease in L , not by exceeding N_{itr} .

The value for θ_{nil} should be determined so that the point-wise difference $y_{\eta_j, \theta_j} - y_{\eta_j, \theta_{\text{nil}}}$ is as close to 0 as possible over a reasonable range of θ_j . This is because the cross-correlation in (4) is always computed using $y_{\eta_j, \theta_{\text{nil}}}$ while the value of the loss function (2) depends on y_{η_j, θ_j} .

Thus, if the difference is too large, a suboptimal η_j may be selected. This especially becomes a problem when inharmonicity is considered.

As continuous functions are highly correlated with slightly shifted versions of themselves, we typically choose $N_{\text{pre}} = 1$ in order to avoid the preselection of the same pattern multiple times for one feature in the spectrum.

The choice of the non-linear optimization algorithm is not critical, as long as it supports box bounds. We decided to employ the L-BFGS-B algorithm [48–50], which is fast even for high-dimensional problems.

Figure 2 provides an illustrative example of the sparse pursuit algorithm. The input is displayed in Fig. 2a: It consists of a discrete spectrum Y and two continuous patterns $y_{0,\theta}, y_{1,\theta}$. For simplicity, we assume that these patterns are perfectly constant, so they do not depend on any additional parameters (therefore, $\theta \in \mathbb{R}^0$), and we set the exponent to $q = 1$ (cf. (2),(3),(4)). The algorithm selects $\eta_0 = 0$ and $\eta_1 = 1$ one after another and finds appropriate amplitudes $a_0, a_1 > 0$ and shifts $\mu_0, \mu_1 \in \mathbb{R}$ such that the superposition of these patterns matches the discrete spectrum Y within numerical precision ($L(Y, a_0, a_1, \mu_0, \mu_1, \eta_0, \eta_1) = 0$), as is displayed in Fig. 2b.

The patterns used for this example are purely synthetic, but similar patterns will appear both in the computation of the pitch-invariant spectrogram and in the separation of the instrument sounds, and they could also originate from other physical phenomena.

3 Computation of the pitch-invariant spectrogram

A spectrogram is a function defined on the time-frequency plane that is supposed to indicate to what extent a certain frequency is present in the recording at a given point in time.

The “canonical” time-frequency representation is the spectrogram obtained from the modulus of the STFT (cf. [1]), which is defined via:

$$\mathcal{V}_w X(t, f) = \int_{-\infty}^{\infty} X(\tau) w(\tau - t) e^{-i2\pi f \tau} d\tau. \quad (5)$$

One particularly popular window with very favorable properties is the Gaussian:

$$w(t) = \frac{1}{\sqrt{2\pi\zeta^2}} \exp(-t^2/(2\zeta^2)), \quad \zeta > 0. \quad (6)$$

For a sinusoidal signal $X(t) = a \exp(i2\pi \nu t)$ with amplitude $a \geq 0$, this results in a horizontal line in the spectrogram:

$$\mathcal{V}_w X(t, f) = a \mathcal{F} w(f - \nu) e^{-i2\pi(f - \nu)t}, \quad (7)$$

and

$$\mathcal{F} w(f - \nu) = \exp(-(f - \nu)^2/(2\sigma^2)) \quad (8)$$

with standard deviation $\sigma = 1/(2\pi\zeta)$, where \mathcal{F} is the unitary Fourier transform. In practice, we use an FFT-computed sampled version $Z[f, t] = |\mathcal{V}_w X(t/T, f/F)|$, where $T, F > 0$ are time and frequency units. While X has a sampling frequency of $f_s = 48$ kHz, we want the time resolution of Z to be lower by a factor of 256; thus, $1/T = 256/f_s = 5.3$ ms. Further, we set $\zeta = 1024/f_s$ and cut w at $\pm 6\zeta$, yielding $1/F = f_s/(12 \cdot 1024) = 3.90625$ Hz.

Note that contrary to the definition of the spectrogram in [1], we do *not* square the magnitude of the STFT, as we require positive homogeneity: If the signal X is multiplied by a positive factor, then we need $Z[f, t]$ to be multiplied by the same factor.

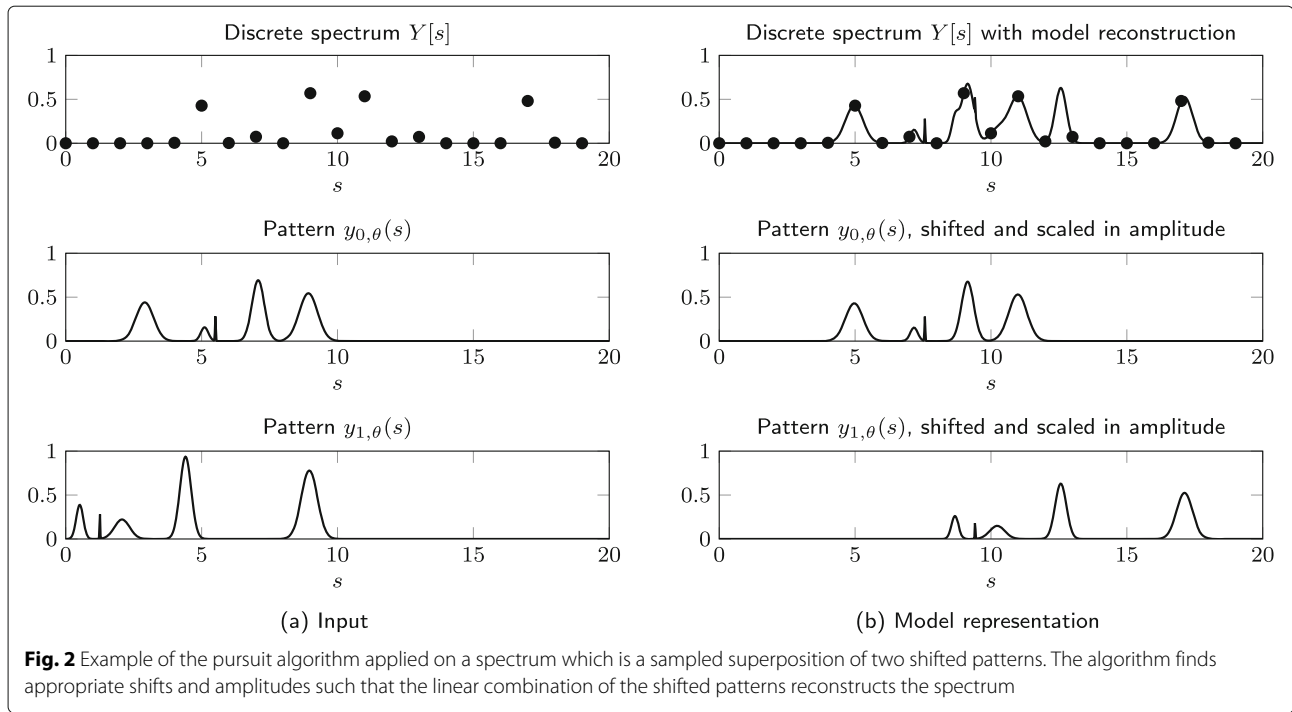
The problem is that the STFT spectrogram is not *pitch-invariant*: We would like a representation where varying the pitch of the tone of an instrument *shifts* the pattern, but, for instance, changing the pitch of a tone by an octave *scales* it by a factor of 2 on a linear frequency axis, which is a different distance depending on the original pitch of the tone.¹

In order to achieve pitch-invariance, one needs a representation with a logarithmic frequency axis. However, a naive transform of the modulus of the STFT would not only influence the position of the horizontal lines, but also their width. In order to overcome this problem, there exist two classical approaches:

- The mel spectrogram (cf. [3]) performs a logarithmic transform on the frequency axis of the STFT spectrogram and then applies smoothing along that axis in order to keep the widths consistent. The frequency range that can be represented by this approach is limited by the Heisenberg uncertainty principle, which states that one cannot have arbitrarily good time and frequency resolution at the same time.
- The constant-Q transform [10] is a discrete wavelet transform and can thus be understood as an STFT with differently dilated windows for each frequency. While it keeps the width of the horizontal lines constant on a logarithmic frequency axis, the time resolution will vary for different frequencies. This is problematic, as it results in simultaneously starting sinusoids first appearing in different time frames of the spectrogram.

As was shown by [51], the constant-Q transform can be turned into a mel spectrogram by applying additional smoothing along the time axis, but it is not possible to overcome the limitations of the Heisenberg uncertainty principle by classical means.

¹When we speak of *pitch*, we refer to the ability of one musical instrument to generate tones whose harmonics have the same relative amplitudes but different locations on the frequency axis, whether that be linear or logarithmic.



For narrowband signals, this is not a problem; the above methods can and have been used in order to provide a time-frequency representation for audio source separation. However, as we experimentally show in the appendix, the time-frequency resolution of the mel spectrogram is too low for the data that we consider, leading to significantly inferior quality of the separation. Instead, as we already have the algorithm from Section 2 at hand, we can use it as another way to transform the linear-frequency STFT spectrogram into a pitch-invariant log-frequency spectrogram. Since this method gives us sharp frequency values, we are no longer constrained by the Heisenberg uncertainty principle. On wideband signals, this “super-resolution” gives us an advantage in the subsequent separation procedure.

We set $Y = Z[\cdot, t]$ and assume a single Gaussian pattern

$$y_{0,\theta}(s) = \exp(-s^2/(2F^2\sigma^2)), \quad \theta = (\sigma), \quad (9)$$

with $N_{\text{pat}} = 1$ and $\theta_{\text{nil}} = (1/(2\pi\zeta))$.

Since the number of Gaussian peaks in a spectrum can be high, we set $N_{\text{spr}} = 1000$ to make sure they can all be represented. This makes the algorithm rather slow, so we choose $q = 1$ in order to bring L closer to a quadratic objective; as we aim to represent the spectrum with very low overall error, there is no need to lift certain features of the spectrum.

To reduce the number of iterations, we also set $N_{\text{pre}} = 1000$. However, this comes with the aforementioned problem that the algorithm would select a lot of neighboring shifts. Thus, instead of computing the cross-correlation,

we simply select the 1000 largest local maxima of the residual that satisfy $r[i] \geq r[i+k]$ for $|k| \leq 3$ and assume their heights as initial values for the amplitudes.

To allow for high-detail representation, we set $\lambda = 1$. The maximum number of iterations is $N_{\text{itr}} = 20$, but the algorithm often terminates before that.

After having identified the Gaussian peaks in the sampled STFT magnitude spectrogram $Z[f, t]$, we resynthesize them in another magnitude spectrogram $U[\alpha, t]$, applying a logarithmic frequency transform $\alpha(f) = \alpha_0 \log_2(f/f_0)$ to the mean frequencies μ_j , $j \in \mathcal{J}$. With $f_0 = 20 \text{ Hz}/f_s \cdot 12 \cdot 1024 = 5.12$ and $\alpha_0 = 1024/10 = 102.4$, we can, assuming a sampling frequency of $f_s = 48 \text{ kHz}$ and $\alpha = \{0, \dots, 1023\}$, represent 10 octaves from 20 Hz to 20.48 kHz.

The algorithm can also be used without modification for compact disc (CD) recordings with a sampling frequency of $f_s = 44.1 \text{ kHz}$. In this case, the represented audio frequency range consists of the 10 octaves from 18.375 Hz to 18.816 kHz.

For Fig. 3, we performed different transforms on an excerpt of a commercial recording of a piece for violin and piano. The mel spectrogram in Fig. 3a had to be cut off at 530 Hz in order to maintain a constant time-log-frequency resolution. The constant-Q transform in Fig. 3b can represent lower frequencies, but its time-log-frequency resolution varies with frequency: Clearly, the tones with lower frequencies have a wider time spread in the representation than those with higher frequencies, giving an inconsistent image in the individual time frames.

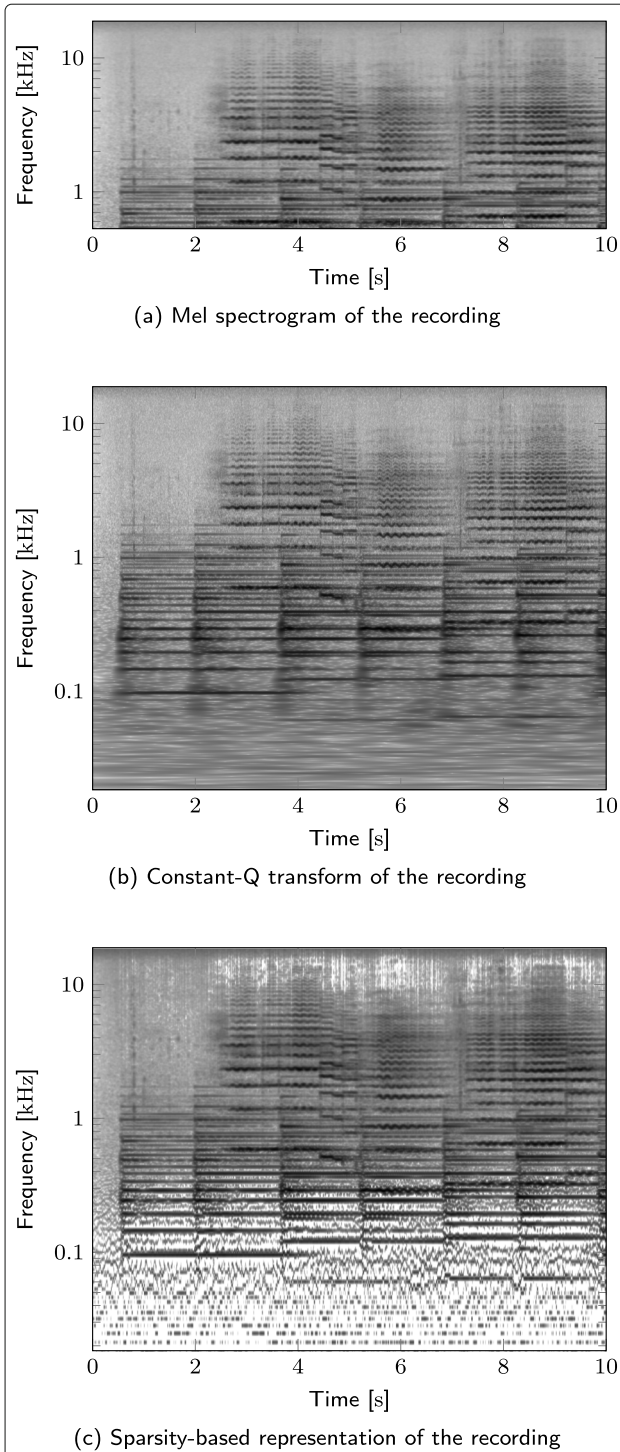


Fig. 3 Log-frequency spectrograms of the beginning of the 1st mvmt. of the sonata no. 1 for violin and piano by Johannes Brahms (op. 100). The grayscale axis is logarithmic and normalized to a dynamic range of 100 dB for each plot. Performance by Itzhak Perlman and Vladimir Ashkenazy. Remastered CD recording by EMI Classics, 1999

Our proposed sparsity-based transform in Fig. 3c does not have this problem: It aligns the tones properly along the time axis like the mel spectrogram, but it can represent much lower frequencies.

As our proposed representation is specifically designed for sinusoids, it largely fails to represent other sounds; in this case, however, this is even beneficial, as it removes portions of the spectrogram that do not correspond to the tones that we aim to represent (creating the white regions in Fig. 3c). From this perspective, we can say that it *denoises* the spectrogram.²

However, it should be kept in mind that the uncertainty principle cannot be “tricked” arbitrarily; if two sinusoids have very low and very similar frequencies, their representations in the STFT spectrogram will overlap greatly, and our algorithm may fail to tell them apart. On the other hand, if a peak is slightly perturbed, it may also occur that the algorithm will identify one single sinusoid as two.

Some parts of the noise do get mistaken for sinusoids and are thus carried over to the log-frequency spectrogram. In the low frequencies, this creates the illusion of sparsity in the log-frequency spectrogram, causing horizontal lines that do not belong to the music to appear in Fig. 3c. Their vertical positions correspond to the transformed frequencies of the pixels in the linear-frequency spectrogram. However, we do not consider these artifacts as a problem from the algorithm, as the noise was already present in the STFT spectrogram. Our algorithm merely creates the white space between the lines.

4 Model representation of the spectrogram

In the previous section, we have described how to obtain a discrete log-frequency spectrogram $U[\alpha, t]$, $\alpha, t \in \mathbb{Z}$ from an audio signal that contains the superposed sound of the musical instruments. Now, the goal is to represent $U[\alpha, t]$ via a parametric model of the sounds of the individual instruments, while the parameter values that characterize the instruments are not known beforehand.

A simple model for the tone production of many musical instruments (particularly string and wind instruments) is the wave equation, which has sinusoidal solutions (the *harmonics*) at frequencies $f_h = hf_1^\circ$, $h = 1, \dots, N_{\text{har}}$, where $f_1^\circ > 0$ is the *fundamental frequency* and $N_{\text{har}} \in \mathbb{N}$ is the number of harmonics to be considered. However, many string instruments (especially the piano in its high notes) have non-negligible stiffness in their strings, leading to a fourth-order equation which has solutions $f_h = (1 + bh^2)^{1/2}hf_1^\circ$, $h = 1, \dots, N_{\text{har}}$, with the inharmonicity parameter $b \geq 0$ (cf. [24]).

²To our separation algorithm, anything non-sinusoidal is noise. This does not imply, however, that these parts of the signal are undesirable for a human listener.

Neglecting any negative frequencies, we model our time-domain signal for the j th tone as a linear combination of complex exponentials:

$$x_j(t) = \sum_{h=1}^{N_{\text{har}}} a_{j,h} \cdot e^{i2\pi(f_{j,h}t + \varphi_{j,h})}, \quad (10)$$

with amplitudes $a_{j,h} \geq 0$ and phase values $\varphi_{j,h} \in [0, 2\pi)$. This could locally be interpreted as an extension of the McAulay-Quatieri model [52].

We assume that the images of these sinusoids superpose linearly in the spectrograms. In reality, this is not the case in the presence of non-constructive interference (*beats*), but if we accept the error introduced by this common simplification, we can set $\varphi_{j,h} = 0$, apply (7) and (8), and approximate $Z[f, t]$ via:

$$z[f, t] := \sum_{j,h} a_{j,h,t} \cdot \exp\left(-\frac{(f - f_{j,h,t})^2}{2F^2\sigma_{j,t}^2}\right), \quad (11)$$

where $a_{j,h,t}$ is the amplitude of the h th harmonic of the j th tone in the t th time frame, and $f_{j,h,t}$ is the respective frequency. For the log-frequency spectrogram $U[\alpha, t]$, this transforms to the following approximation:

$$u[\alpha, t] := \sum_{j,h} a_{j,h,t} \cdot \exp\left(-\frac{(\alpha - \alpha_{j,h,t})^2}{2F^2\sigma_{j,t}^2}\right), \quad (12)$$

with $\alpha_{j,h,t} = \alpha(f_{j,h,t}) = \alpha((1 + b_{j,t}h^2)^{1/2}h) + \alpha(f_{j,1,t}^\circ)$.

We further make the simplifying assumption that the sound of a musical instrument is constant over the duration of a tone and that the relation of the amplitudes of the harmonics is constant with respect to pitch and volume. We thus save the relative amplitudes of the instruments in a *dictionary*, which is a matrix $D \in [0, 1]^{N_{\text{har}} \times N_{\text{pat}}}$. Introducing an overall amplitude $a_{j,t}$ for each tone, we can express $a_{j,h,t} = D[h, \eta_{j,t}] a_{j,t}$, where $\eta_{j,t}$ is the instrument by which the tone is played. For practical acoustic instruments, this assumption is never fully satisfied, so the deviation between the modeled amplitudes and the true amplitudes introduces a certain error. However, we will later apply a *spectral masking* step (Section 6.1) that restores the amplitudes of each harmonic directly from the recording in order to mitigate this error in the final output.

Our pursuit algorithm can now be applied to (12) by setting the patterns as:

$$y_{\eta_{j,t}, \theta_{j,t}}(\alpha) = \sum_h D[h, \eta_{j,t}] \cdot \exp\left(-\frac{(\alpha - \alpha((1 + b_{j,t}h^2)^{1/2}h))^2}{2F^2\sigma_{j,t}^2}\right), \quad (13)$$

with $\theta_{j,t} = (\sigma_{j,t}, b_{j,t})$ and $\mu_{j,t} = \alpha(f_{j,1,t}^\circ)$, according to the notation from (1) with time dependency added. The initial value is $\theta_{\text{nil}} = (1/(2\pi\zeta), 0)$.

As the patterns now depend on the dictionary, this dependency is carried over to the loss function (2), which we thus denote as L_D .

5 Dictionary learning

5.1 Scheme

In order to train the dictionary, we pursue a stochastic alternating-optimization approach. First the dictionary is initialized; for each $\eta = 0, \dots, N_{\text{pat}} - 1$, we generate a uniformly distributed random vector $d \in [0, 1]^{N_{\text{har}}}$ and an exponent e that is Pareto-distributed with a scale parameter of 1/2 (to make sure that $e \geq 1$, guaranteeing a minimum decay rate), and we set $D[h, \eta] = d[h]/h^e$.

Given an initial dictionary, a random time frame $U[\cdot, t]$ of the log-frequency spectrogram of the recording is chosen, and the sparse pursuit algorithm is applied on it. Afterwards, the gradient $\nabla_D L_D$ of the dictionary-dependent loss function is computed with the parameters from the sparse pursuit algorithm, and this is used to update the dictionary in order to reduce the loss. The process is repeated $N_{\text{trn}} \in \mathbb{N}$ times, which is the number of training iterations as specified by the user.

We set the number of patterns to be generated from the dictionary to twice the expected number of instruments in the recording ($N_{\text{pat}} = 2N_{\text{ins}}$, $N_{\text{ins}} \in \mathbb{N}$), allowing for some redundancy during the training.

5.2 Dictionary update

Classically, dictionary learning is performed via techniques like NMF [6, 53], K-SVD [54], or tensor factorization (cf. [5]). However, the first two methods do not account for the pitch-invariant structure of our data. Tensor factorization does, but only for a fixed number of frequency shifts. Moreover, all of these methods become slow when the amount of data is large.

While the use of stochastic gradient descent for dictionary learning has been common for many years (cf., e.g., [55]), new methods have been arising very recently due to their applications in deep learning. One of the most popular methods for this purpose is Adam [46]. Its underlying idea is to treat the gradient as a random variable and, for each component, compute unbiased estimates \hat{v}_1, \hat{v}_2 for the first and second moments, and choose the step size proportional to $\hat{v}_1/\sqrt{\hat{v}_2}$. If the derivative of the i th component is constant, then $\hat{v}_1[i]/\sqrt{\hat{v}_2[i]} = \pm 1$, in which case a large step size can be used. If the derivative oscillates a lot, however, then $\hat{v}_1[i]/\sqrt{\hat{v}_2[i]}$ will also be small and thereby dampen the oscillation in that direction.

The standard formulation of Adam is completely independent of the scale of the derivatives. This makes it easy to control the absolute step size of the components, but it destroys the Landweber regularization property of gradient descent, which automatically decreases the step size

for components whose partial derivative is small, taking into account the scaling of different harmonics.

Our first modification to Adam is that while we still estimate the first moments for each dictionary entry (i.e., for each instrument and for each harmonic), we only compute one second moment estimate for each instrument, which is the arithmetic mean over the all the estimates for the harmonics. With this, we restore the regularization property and prevent excessive change of the components that have small values.

Furthermore, we require all entries in the dictionary to be non-negative, since negative harmonic amplitudes would violate our model assumptions. For consistency, we also require that no entries be larger than 1, so we end up with the box constraint that $D[h, \eta] \in [0, 1]$ for $h = 1, \dots, N_{\text{har}}, \eta = 0, \dots, N_{\text{pat}} - 1$. To enforce this, we project each component to $[0, 1]$ after the end of a step.

Finally, we have to tackle the problem that due to the stochastic nature of the optimization procedure, dictionary entries for a particular supposed instrument may diverge to a point where they will not be used by the identification algorithm anymore and thus not contribute to the separation. For this purpose, we track the sum of the amplitudes associated with a specific instrument in the past. In regular intervals, we sort the instruments in the dictionary by the ratio of the amplitude sum versus the number of iterations since its initialization (minus a small head start that benefits new instrument entries); then, we prune the dictionary by reinitializing the entries for those supposed instruments where the ratio is lowest, leaving the N_{ins} instruments with the highest ratios intact.

Concerning the parameters for moment estimation and parameter update in Adam, the default values (cf. the description of the pseudo-code in the appendix) have turned out to be a good choice for the majority of applications. In our case, a step-size of $\kappa = 10^{-3}$ means that if the gradient is constant, the dominant component will go from 0 to 1 in the dictionary within less than 1000 iterations, which is fast enough if $N_{\text{trn}} \geq 10000$. While lowering κ is a common way to improve training accuracy, this did not appear to have any effect in our applications.

6 Separation and resynthesis

After the dictionary has been trained by alternating between identification and dictionary update, we represent the entire recording by running the identification/pursuit algorithm on each time frame $U[\cdot, t]$ for $t = 0, \dots, n - 1$ (where n is the number of time frames in the spectrogram) with those N_{ins} instruments in the dictionary that were left intact after the latest pruning. This time, however, we need a linear-frequency spectrogram, since this is much easier to convert back into a time-domain signal, so we apply the reverse

transformation $f(\alpha) = f_0 2^{\alpha/\alpha_0}$ on the means of the Gaussians and reconstruct the spectrogram for the η th instrument via:

$$z_\eta[f, t] := \sum_{\substack{j, h \\ \eta_{j, t} = \eta}} a_{j, h, t} \cdot \exp\left(-\frac{(f - f_{j, h, t})^2}{2F^2\sigma_{j, t}^2}\right), \quad (14)$$

which is the model from (11) limited to one instrument.

For the generation of the time-domain signal, we use the classical algorithm by Griffin and Lim [2], which iteratively approximates the signal whose corresponding STFT magnitude spectrogram is (in the ℓ_2 sense) closest to the given one. As initial value, we give the phase of the STFT of the original signal.

While more sophisticated phase retrieval methods have been developed recently (e.g., [56]), the algorithm by Griffin and Lim is well-established, robust, and simple.

6.1 Spectral masking

As an optional post-processing step, we can mask the spectrograms from the dictionary representation with the spectrogram from the original recording. This method was proposed in [12, 13]:

$$\tilde{z}_\eta[f, t] := \frac{z_\eta[f, t]}{z[f, t]} \cdot Z[f, t]. \quad (15)$$

In practice, a tiny value is added to the denominator in order to avoid division by zero.

With this procedure, we make sure that the output spectrograms do not have any artifacts at frequencies that are not present in the original recording. Another benefit is mentioned in [12]: In cases where the sound of an instrument is not perfectly invariant with respect to pitch and volume, the masking can correct this.

A potential drawback with masking is that destructive interference in the original spectrogram may alter the spectrograms of the isolated instruments.

From a statistical perspective, spectral masking can also be regarded as a (trivial) *Wiener filter* (cf. [3]). In this case, one would regard the *squared* magnitude spectrograms in the fraction in (15) and treat them as power spectra that give priors for the frequency distribution of the signals. However, we consider this perspective problematic, as the masks are in fact generated from the data itself, which is already subject to interference, and squaring the spectra could exacerbate the error.

7 Experimental results and discussion

We generate the log-frequency spectrogram as specified in Section 3. For the dictionary, we use $N_{\text{har}} = 25$ harmonics.

Table 1 Performance measures for the best-case run of the separation of recorder and violin

Method	Mask	Instrument	SDR	SIR	SAR
Ours	No	Recorder	12.9	32.5*	12.9
		Violin	7.1	24.1*	7.2
	Yes	Recorder	15.1*	32.4	15.2*
		Violin	11.9*	23.8	12.2*
[23]	—	Recorder	10.6	21.4	11.0
		Violin	5.8	18.4	6.1

Best numbers are marked

7.1 Performance measures

Vincent et al. [57] define the signal-to-distortion ratio (SDR), the signal-to-interference ratio (SIR), and the signal-to-artifacts ratio (SAR). These ℓ_2 -based measures have become the de facto standard for the performance evaluation of blind audio source separation³.

The SDR is an “overall” performance measure that incorporates all kinds of errors in the reconstructed signal; it yields a value of $-\infty$ if the original signal and the reconstructed signal are uncorrelated. The SIR is similar, but it ignores any artifacts that are uncorrelated with the original signals. The SAR only measures the artifacts and ignores interference; it is constant with respect to permutations of the original signals. Those measures are independent of the scale of the reconstruction, but they are very sensitive to phase mismatch, as the projection of a sinusoid on its 90° -shifted copy will be zero, even though the signals are otherwise identical. In order to find the right mapping between the synthesized and the original signals, the synthesized signals are permuted such that the mean SIR over all instruments is maximized.

Another method for the performance evaluation of audio source separation is given by the PEASS [59, 60], which define the overall perceptual score (OPS), the target-related perceptual score (TPS), the interference-related perceptual score (IPS), and the artifacts-related perceptual score (APS), which are computed using psychoacoustically motivated measures and were trained via empirical listening experiments. The OPS and IPS correspond conceptually to the SDR and SIR, but the artifacts as measured via the SAR are subdivided into the TPS, which accounts for the misrepresentation of the original signal itself, and the APS, which only comprises the remaining error that does not originate from misrepresentation or interference. The values of the scores range from 0 (worst) to 100 (best).

7.2 Separation of recorder and violin sounds

In order to generate a realistic separation scenario, we chose the 8th piece from the 12 Basset Horn Duos by Wolfgang A. Mozart (K. 487) in an arrangement by Alberto Gomez Gomez for two recorders⁴. The upper part was played on a soprano recorder, and the lower part was played on a violin. These instruments are easily distinguishable, as the recorder has an almost sinusoidal sound, while the sound of the violin is sawtooth-like, with strong harmonics [24].

The instrument tracks were recorded separately in an apartment room ($RT_{60} \approx 0.4$ s) with an audio recorder at a distance of approximately 1 m to the instrument, while a metronome/“play-along” track was provided via headphones. Evenness of the tone was favored over musical expression. We combined the tracks by adding the two digital signals with no post-processing other than adjustment of volume and overall timing and let the algorithm run with $N_{\text{trn}} = 100000$ training iterations⁵, with $N_{\text{ins}} = 2$ and $N_{\text{spr}} = 1$.

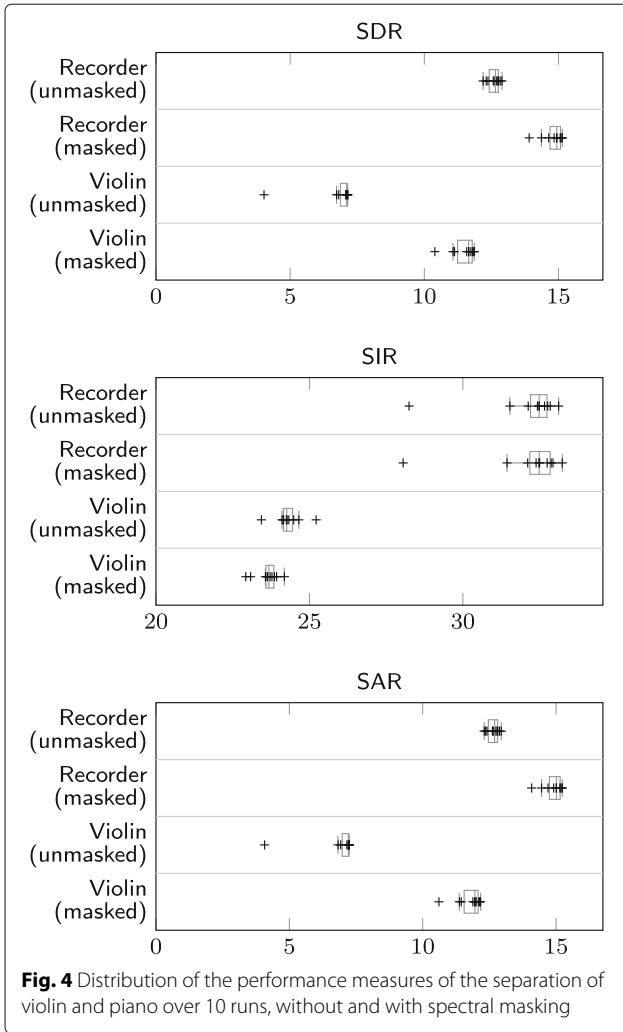
This procedure was performed with random seeds $0, \dots, 9$. For comparison, we further applied the algorithm developed in [23] on our data. We found that their method is sensitive with respect to hyperparameters, and we searched for those values that optimize separation performance for this piece, but we could only achieve marginal improvement over the defaults provided in the code. For application of this algorithm, we downsampled the audio data to 22050 Hz, as this is the sampling frequency that the algorithm was designed to operate on. The best-case results for both algorithms are presented in Table 1, and the distribution over all 10 runs of our algorithm is displayed in Fig. 4.

Our criterion for the best run in our algorithm was the mean SDR over both instruments. This was achieved by a random seed of 7 for this sample. When the algorithm is used in a real-world scenario in which the original tracks are not available, the performance measures are

³In the meantime, version 3.0 of the *BSS Eval* software package has become available, which employs a slightly different definition that includes time shifts. However, for comparability with [12–14, 23], we are using the original measures as implemented in version 2.1 [58].

⁴[https://imslp.org/wiki/12_Horn_Duos_K.487/496a_\(Mozart,_Wolfgang_Amadeus\)](https://imslp.org/wiki/12_Horn_Duos_K.487/496a_(Mozart,_Wolfgang_Amadeus))

⁵We already achieve similarly good performance with $N_{\text{trn}} = 10000$ iterations, but more iterations make the result more consistent with respect to initial values.



unknown to the user. In this case, the user can select the “best-sounding” result from all 10 candidates, perhaps guided by the value of the loss function as a proxy measure. The notion of ensemble learning does not apply to the algorithm in [23], as it is a clustering method and does not have an initial dictionary. Instead, we there consider the result that we achieve with the hand-optimized parameters as best-case.

With our algorithm, the recorder is universally better represented than the violin, and spectral masking leads to considerable improvements in SDR and SAR especially for the violin. This complies with the explanation in [12] that spectral masking helps represent instruments with more diverse spectra, such as the violin, which has 4 different strings and a sound that is very sensitive to technique. When we compare the outcomes in pairs without and with spectral masking over the random seeds 0, . . . , 9 respectively, the improvement in SDR achieved by spectral masking is statistically significant at $p_{\text{Recorder}} = p_{\text{Violin}} = 9.8 \times 10^{-4}$ in a one-sided Wilcoxon signed-rank

Table 2 PEASS scores for the best-case run of the separation of recorder and violin

Method	Mask	Inst.	OPS	TPS	IPS	APS
Ours	No	Rec.	34*	31	70*	38*
		Vln.	34*	27	71*	37*
	Yes	Rec.	25	64	39	38
		Vln.	13	100*	33	52
[23]	—	Rec.	28	84*	26	35
		Vln.	32	19	71	30

Best numbers are marked

test [61]⁶, as for each dictionary, spectral masking leads to a consistent improvement of the separation result.

The algorithm from [23] reacts in a similar way, yielding better performance for the recorder than for the violin. However, the working principle is different: Rather than trying to represent both instruments, it clusters the peaks from the spectrum in order to make out a “dominant” instrument, while the second “instrument” is just the collection of residual peaks. In our example, the violin was identified as the dominant instrument, but nonetheless the representation of the recorder is better. However, our algorithm provides superior performance for both instruments, even without spectral masking.

For phase reconstruction, we used merely one iteration (i.e., only one magnitude adjustment and one projection) of the Griffin-Lim algorithm in order to preserve the phase of the original spectrogram as much as possible.

The aural impression of the results with different random seeds is largely very similar. While some artifacts and interference are audible, the generated audio data provides a good aural representation of the actually played tracks. The only tone⁷ that is misidentified over a long period of time is a recorder tone that interferes with the even-numbered harmonics of the violin tone that is played at the same time and is one octave lower. In this case, the third harmonic of the violin tone is erroneously identified as the recorder tone.

The PEASS scores for the same runs and parameters are given in Table 2. Surprisingly, the results without spectral masking are now mostly preferred. Our only explanation is that as discussed in Section 6.1, spectral masking can cause interference in overlapping tones, as can be seen in the drop in both SIR and IPS. While the SDR still increases overall with spectral masking, this interference might have a large negative impact on the OPS. However, we did not

⁶ Briefly speaking, the Wilcoxon signed-rank test has the null hypothesis that the differences in the pairs are symmetrically distributed around 0. For this, the sum of the signed ranks of the differences is computed. In the one-sided test, the acceptance region for this sum is asymmetric.

⁷ which occurs 4 times in total, due to repetitions of the passage

find this discrepancy in most of the other samples, so it does not appear to be a general pattern.

Spectrograms of the original recording and the synthesized representations (with the random seed of 7 that maximizes the SDR) are displayed in Fig. 5. The original spectrogram contains broad-spectrum components (“noise”) that do not fit the dictionary model and thus cannot be represented, so they are not found in the output spectrograms. The choice of $N_{\text{har}} = 25$ must be regarded as a compromise: Although the sound of the violin could be represented more accurately with an even higher number of harmonics, this would increase both the computation time of the algorithm and also the number of variables to be trained. The incorrectly identified recorder tone corresponds to the rightmost set of horizontal lines in Fig. 5b. It is not audible when the synthesized audio files are mixed back together.

Since spectral masking is only applied on the linear-frequency spectrograms, its effects cannot be seen in Fig. 5.

7.3 Separation of clarinet and piano sounds

We recorded the same piece on clarinet and piano using the same set-up as for recorder and violin, except that the instruments were played in a rehearsal hall ($RT_{60} \approx 1.4$ s). The algorithm was also run under the same conditions. The distribution of the results over random seeds $0, \dots, 9$ is displayed in Fig. 6. The best-case results of our algorithm with a random seed value of 6 as well as those for the algorithm from [23] (with again, the data downsampled to 22050 Hz) are presented in Table 3.

The separation quality with our algorithm is much worse than the for recorder and violin, and representation of the piano is especially problematic. We have several explanations for this:

1. Piano tones exhibit non-negligible inharmonicity, which makes it harder to identify them in the spectrum. Even though our model incorporates this inharmonicity, cross-correlation does not.
2. Compared to the rather steady tone of recorder, violin, and clarinet, the piano tone has a very characteristic onset (*attack*), which exhibits different spectral characteristics than the rest of the tone.

This raises the question whether our algorithm can represent piano tones at all. In order to test this, we ran it on the original piano track. The result was very stable, with a maximum SDR of 8.7 dB for a random seed of 9 – without spectral masking, as this would not make sense for a single instrument. In Fig. 7, we show a time frame from the spectrogram within the first tone of the piano ($t = 100$). The fundamental frequency was identified by the algorithm as $f_1^\circ = 441.8$ Hz and the inharmonicity as $b = 5.3 \times 10^{-4}$.

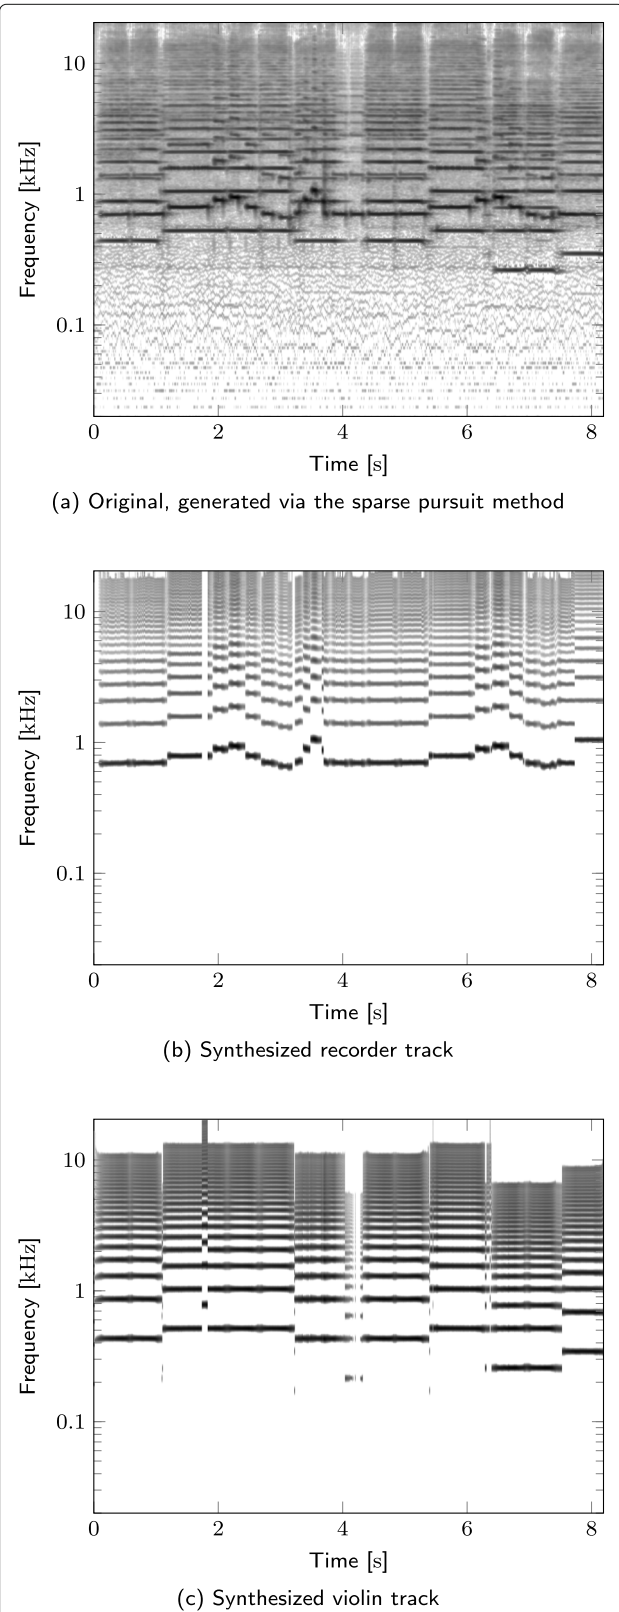
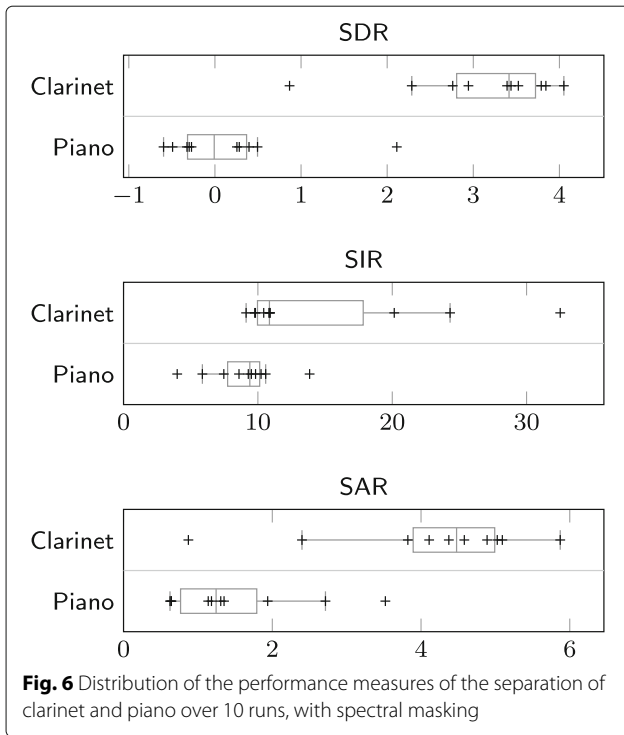


Fig. 5 Log-frequency spectrograms for beginning of the recorded piece and the synthesized tracks. The grayscale axis is logarithmic and normalized to a dynamic range of 100 dB for each plot



In Fig. 7a, the original spectrum is displayed with the predicted frequencies of the harmonics when inharmonicity is neglected, and the deviation upward from the 5th harmonic becomes clearly recognizable. In Fig. 7b, the computed inharmonicity is incorporated, and so the predicted frequencies of the harmonics match those from the original spectrum almost perfectly. Figure 7c represents the reconstructed spectrogram time frame as returned by the separation algorithm with all the other parameters considered, but without spectral masking.

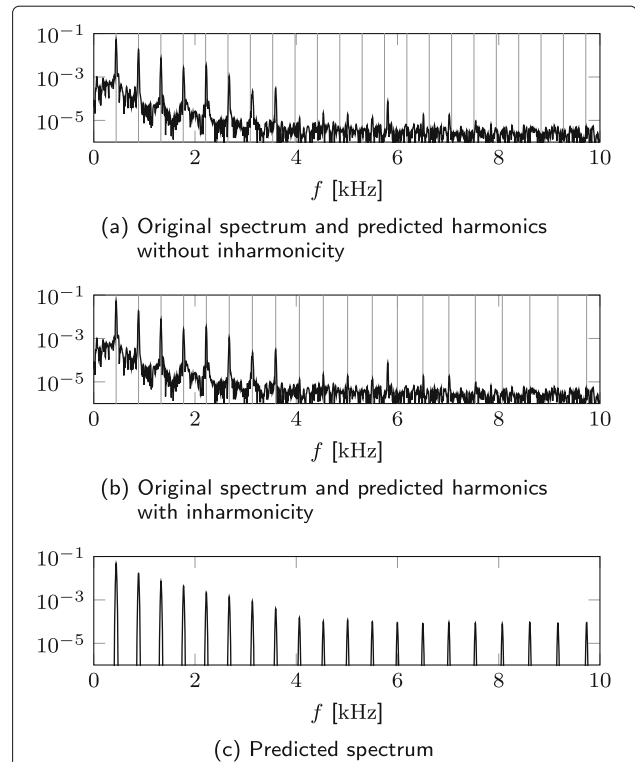
Thus, our algorithm does not have any issue *representing* the piano tones; the difficulty in this case is to identify them in the presence of the clarinet tones.

The algorithm from [23] performs comparatively well. This is, again, due to the different approach: Rather than trying to represent both instruments, this algorithm only finds the clarinet tones as the dominant cluster and assigns the remaining parts of the spectrum to the piano. Thus, even though their model cannot represent

Table 3 Performance measures for the best-case run of the separation of clarinet and piano, with spectral masking

Method	Instrument	SDR	SIR	SAR
Ours	Clarinet	4.1	24.3*	4.1
	Piano	2.1	9.3	3.5
[23]	Clarinet	6.7*	21.3	6.9*
	Piano	5.5*	16.4*	5.9*

Best numbers are marked



the piano, as it does not include inharmonicity at all, it can still separate it under the assumption that the clarinet is modeled correctly. However, for this recording, it is essential to hand-tune the hyperparameters: Those that were used for the separation of recorder and violin still work reasonably well for clarinet and piano, but with the default values, the algorithm fails.

In terms of the PEASS (Table 4), the results of both algorithms achieve very similar overall scores. While our reconstruction of the piano sound is inferior in terms of TPS, the interference and artifacts are evaluated as perceptually less severe.

Table 4 PEASS scores for the best-case run of the separation of clarinet and piano, with spectral masking. Best numbers are marked

Method	Instrument	OPS	TPS	IPS	APS
Ours	Clarinet	39*	56*	68*	46*
	Piano	25	37	59*	31*
[23]	Clarinet	39	46	62	41
	Piano	26*	87*	30	22

Best numbers are marked

7.4 Generalization experiment

Usually, we train our dictionary on the same audio recording that we aim to separate. In this experiment, however, our goal is to ascertain whether a dictionary that was trained on one recording can be used for the separation of another recording without additional training.

Under the recording conditions specified in Section 7.2, we recorded the traditional tune “Frère Jacques” with B♭ tin whistle and viola in the key of E♭ major as well as with C tin whistle and violin in the key of F major. The violin and viola were offset by two bars compared to the tin whistles in order to create a musically realistic *canon*. The lowest frequency of the B♭ tin whistle was measured as 463 Hz, and the lowest frequency of the C tin whistle was measured as 534 Hz. Thus, they do *not* fit in the same equal-temperament tuning, and the intervals on these instruments are not very consistent, either. Their tuning was mimicked by ear when recording the viola/violin tracks.

First, the separation was performed with random seeds 0, . . . , 9 on the recording with B♭ tin whistle and viola. Then the dictionaries obtained from this separation were used on the recording with C tin whistle and violin without any further training. The experiment was repeated vice versa with the recordings permuted.

For the viola and B♭ whistle combination, the dictionary from the run with seed 8 was optimal, but that from a seed of 0 was best when applying on the violin and C whistle recording. Vice versa, when training on the violin and C whistle recording, the seed of 0 was also ideal for separation of that recording, but the dictionary from a random seed of 2 was better when applying on the B♭ whistle and viola recording. All the best-case numbers are presented in Table 5.

Overall, the performance figures are similar to those from recorder and violin, as could be expected because

those are similar instruments. To our surprise, the performance in the generalization even sometimes exceeds that from direct training and separation.

For a better analysis, we gathered the data from seeds 0, . . . , 29 and displayed the distribution in Fig. 8. This reveals a paradox: Maximum SDR performance for each instrument is achieved on a dictionary that was trained on the recording with C tin whistle and violin. At the same time, when comparing the performance of each instrument over all random seeds pairwise between the recordings that the dictionaries were trained on, the Wilcoxon signed rank sums for each instrument indicate a better performance when training on the recording with B♭ tin whistle and viola. Thus, while the former recording yields a better-performing best-case dictionary with a sufficient number of runs, the training is also more likely to fail than with the latter recording.

We conclude that as intended, the model does not overfit to the specific recording, but it instead provides a dictionary that can be applied to a different recording even if slightly different instruments are used and the key is changed (confirming pitch-invariance). For a practical scenario, this means that if a dictionary for a specific combination of instruments is already available, it can be applied to other similar recordings, which saves computation time.⁸ In fact, re-using a well-trained dictionary can lead to superior separation results than training on the recording itself.

7.5 Comparison on other data

To our knowledge, there exists no standard benchmark database with the kind of samples that our algorithm is designed for. While the *BASS-dB* set [62] was created with blind source separation in mind, it contains instruments which violate the structural assumptions that we make about the sounds, and the polyphony levels are not sufficiently controlled. A similar issue occurs with the databases that are used for supervised learning, such as in the SiSEC 2018 [36].

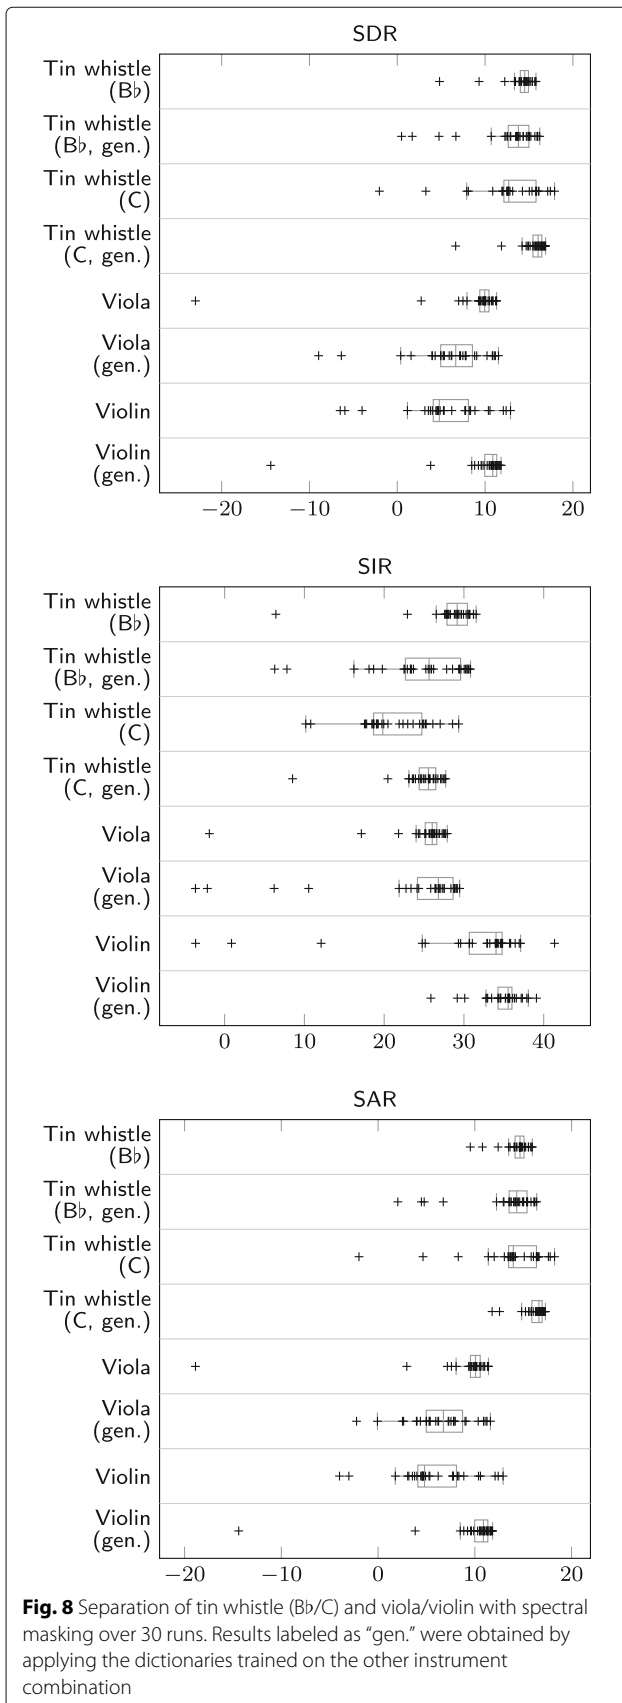
For score-informed separation, the *Bach10* [63] and *URMP* [64] databases are popular, which contain recordings of melodic acoustic instruments. In terms of polyphony and similarity of the instruments in these samples, one cannot expect to obtain reasonable performance from blind separation on most of the samples. However, a subset of the two-instrument recordings in URMP appeared to be usable, so we are incorporating it in our evaluation.

Table 5 Performance measures for the best-case run of the separation of B♭/C tin whistle and viola/violin, with spectral masking. Results indicated as “Orig.” were generated from the dictionary that was trained on that recording, while “Gen.” means that the dictionary was trained on the other recording

Mode	Instrument	SDR	SIR	SAR
Orig.	Tin whistle B♭	15.0	29.3	15.1
	Viola	10.5	26.9	10.6
	Tin whistle C	17.1*	27.0	17.6*
	Violin	12.1*	36.4*	12.1*
Gen.	Tin whistle B♭	15.9*	30.0*	16.1*
	Viola	11.2*	28.4*	11.3*
	Tin whistle C	16.7	27.3*	17.1
	Violin	11.6	34.5	11.6

Best numbers are marked

⁸ For the sample with B♭ tin whistle and viola which has a duration of 24 s, the computation of the log-frequency spectrogram lasted 137 min. Training took 212 min for each of the 10 dictionaries (with $N_{\text{trn}} = 100000$ iterations), while separation and resynthesis with a given dictionary were performed within 7 min. All computations were conducted on an Intel i5-4460 microprocessor using 2 cores for multiprocessing. Note that there is still significant potential for saving computation time by reducing redundancy in the sampling of the STFT and decreasing the number of training iterations.



Also, we were able to obtain the data used by Jaiswal et al. [12–14]. As it does not contain any samples with acoustic instruments, it is not ideal for evaluation of our method, but being able to perform the separation provides a proof of concept.

Further, we used the publicly available data from Duan et al. [23], which does contain a sample with acoustic instruments.

7.5.1 URMP

The URMP dataset [64] contains a total number of 44 audio samples arranged from classical music that were recorded using acoustic musical instruments. In many of these samples, the instruments are very similar, so we selected suitable samples based on the following criteria:

- No instrument should be duplicated.
- No two bowed string instruments should appear in one recording.
- No two brass instruments should appear in one recording.
- If two woodwinds appear together, one should be a reed instrument and the other one should not.

The samples with three or more instruments quickly turned out to be too difficult for our blind separation algorithm. From the total number of 11 duets, this therefore left us with 4 samples:

1. *Dance of the Sugar Plum Fairy* by P. Tchaikovsky with flute and clarinet,
2. *Jesus bleibet meine Freude* by J. S. Bach with trumpet and violin,
3. *March from Occasional Oratorio* by G. F. Handel with trumpet and saxophone,
4. *Ave Maria* by F. Schubert with oboe and cello.

Considering the combination of trumpet and saxophone, we were doubtful whether a separation would be possible. Even though the sound production principle is very different, their sound appears somewhat similar, which is supported by the roles of these instruments in jazz ensembles. We decided to include the sample anyway in order to see how the algorithm reacts.

Again, we are taking the best-case number from 10 runs with $N_{\text{trn}} = 100000$ training iterations, and for comparison, we are using the algorithm from [23] with hand-optimized hyperparameters on the data (as downsampled to 22050 Hz). The results with the classical measures are shown in Table 6.

The piece for flute and clarinet was challenging for both algorithms (perhaps because both instruments are woodwinds). The algorithm from [23] isolated the clarinet as the dominant instrument but only achieved inferior performance on it, whereas the residual has good resemblance with the flute track. On the piece with

Table 6 Performance measures for the best-case runs over a selection of samples from the URMP [64] dataset

Method	Instrument	SDR	SIR	SAR
Ours	Flute	2.4	9.5	3.9*
	Clarinet	6.2*	25.3*	6.3*
	Trumpet	5.3*	16.6*	5.7*
	Violin	7.7*	25.1*	7.8*
	Trumpet	-2.4	1.1	2.7*
	Saxophone	0.1	22.5*	0.2
	Oboe	6.3*	17.0*	6.8*
	Cello	4.2*	17.1*	4.5
[23]	Flute	3.4*	19.6*	3.6
	Clarinet	2.1	5.9	5.4
	Trumpet	—	—	—
	Violin	—	—	—
	Trumpet	1.2*	9.4*	2.3
	Saxophone	6.9*	17.2	7.4*
	Oboe	-0.8	13.1	-0.4
	Cello	3.4	6.4	7.3*

Best numbers are marked

trumpet and violin, our algorithm performed quite well, but the algorithm from [23] got stuck in an apparently endless loop, so we could not get a comparison result. With the piece for trumpet and saxophone, which we had already considered problematic beforehand, our algorithm failed to give an acceptable result in terms of SDR and SIR (in contrast to the PEASS evaluation, as we will discuss later). The compared algorithm gives better figures when separating the trumpet as the dominant instrument, but the result cannot be considered good, either; however, the residual signal gives a decent separation of the saxophone track. By contrast, in the piece with oboe and cello, the algorithm from [23] separated the cello as the dominant instrument comparatively well, whereas it failed on the oboe. For both instruments, the results from our algorithm are better.

As before, it turned out that adjustment of the hyperparameters for every sample was crucial in application of the algorithm from [23], as the clustering depends on the amount of variation in the sound of the dominant instrument as well as on the similarity of the sounds of both instruments.

The corresponding PEASS scores are given in Table 7. The main difference is that our separation of the trumpet in the third piece that received very bad SDR/SIR/SAR values was given very good perceptual scores, mostly exceeding those of the compared method. Listening to the separated trumpet tracks ourselves, we find that while

Table 7 PEASS scores for the best-case runs over a selection of samples from the URMP [64] dataset

Method	Instrument	OPS	TPS	IPS	APS
Ours	Flute	28	46	66*	29
	Clarinet	36*	58*	71	39*
	Trumpet	30*	67*	47*	36*
	Violin	31*	33*	69*	36*
	Trumpet	47*	69*	63	54*
	Saxophone	24	23	70*	15
	Oboe	18*	7	60*	7
	Cello	30*	42*	58	42*
[23]	Flute	35*	75*	38	46*
	Clarinet	27	28	76*	25
	Trumpet	—	—	—	—
	Violin	—	—	—	—
	Trumpet	42	52	64*	46
	Saxophone	26*	72*	22	59*
	Oboe	15	54*	28	19*
	Cello	20	16	67*	22

Best numbers are marked

ours certainly has issues, large parts are much more usable than the SDR suggests, and we can understand why one would perceive the errors as less disruptive than in the track that was isolated by the algorithm from [23].

We believe that one key challenge with this dataset is that the instruments were played with the mindset of a musical performance, and thus there is more variation in playing technique than with our own samples.

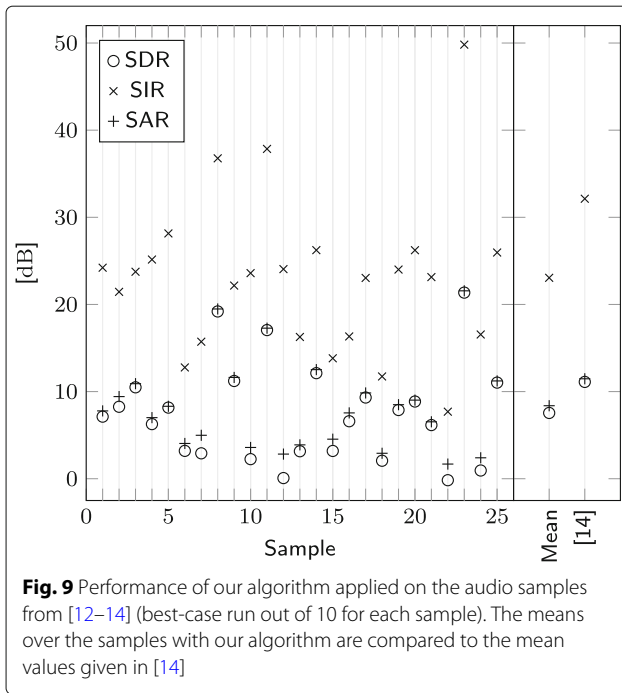
7.5.2 Jaiswal et al.

We ran our algorithm on the data that was used in [12–14], which consists of computer-synthesized samples with two instruments, each playing one tone at a time. Due to the large number of samples, and since we are only interested in best-case numbers, we set $N_{\text{trn}} = 10000$ and selected the best result (in terms of mean SDR) out of 10 runs (with random seeds 0, ..., 9) for each sample. No further adjustments to our algorithm were conducted. The performance measures are displayed in Fig. 9 and Table 8.

It can be seen that for certain samples, our algorithm performs very well, while for others, it fails to produce acceptable results. When comparing the means, our algorithm is inferior to [12–14].⁹

Our explanation for this is that our algorithm assumes much looser constraints on the data that it gets, as it

⁹We could not compare the performance on the individual samples, as those numbers are not available to us.



accepts arbitrary tones in the audible range. By contrast, in [12–14], the expected fundamental frequencies for the instruments are hardcoded in the algorithm due to prior knowledge. In [12], 7 values are allowed per sample, while in [13], this number was individually adjusted to 4–9 values for each sample in order to achieve maximum performance figures; in [14], those were 5–12 values. Further, the algorithms can exploit the fact that the tone ranges for the respective instruments in the samples were chosen to have little or no overlap. In the case of no overlap, such distinctive information would even make it possible to separate instruments with identical frequency spectra, but this would violate our notion of blind separation.

As can be seen in Table 8, the individual adjustments that were conducted in [13] had a much greater effect on the performance than the algorithmic improvements in [14].

Table 8 Comparison of our algorithm to [12–14] on the data used therein (means over all instruments and all samples in the best cases)

Method	SDR	SIR	SAR
[12]	8.9	23.7	9.7
[13]	10.9	25.4	11.5
[14]	11.1*	32.1*	11.5*
Ours	7.5	23.1	8.4

Best numbers are marked

Applying the algorithms in [12–14] to our data would not be meaningful, as those algorithms require, due to their data representation, perfectly consistent equal temperament tuning, which wind instruments and string instruments without frets do not satisfy.

We conclude that the out-of-the-box performance of our algorithm is on average inferior to the figures in [12–14] on the samples used therein, but this is compensated by its vastly greater flexibility, which enables it to operate on real-world acoustic signals and eliminates the need for prior specification of the tuning or range of the instruments.

7.5.3 Duan et al.

From the data used in [23], we selected the samples that we deemed suitable for our algorithm, skipping the ones that contain human voice components, as those cannot be represented by our model.

The three samples that we therefore consider are composed as follows:

1. Acoustic oboe and acoustic euphonium,
2. Synthesized piccolo and synthesized organ,
3. Synthesized piccolo, synthesized organ, and synthesized oboe.

The original samples are sampled at $f_s = 22050$ Hz. We upsampled them to $f_s = 44100$ Hz in order to apply them to our algorithm. We again ran the algorithm with $N_{\text{trn}} = 10000$ iterations and picked the best-case runs from random seeds 0, ..., 9, respectively. The results are displayed in Table 9.

Table 9 Performance measures for the best-case runs of different instrument combinations, with spectral masking. Instruments labeled as “s.” are synthetic, those labeled as “a.” are acoustic

Method	Instrument	SDR	SIR	SAR
[23]	Oboe (a.)	8.7	25.8	8.8
	Euphonium (a.)	4.6	14.5	5.3
	Piccolo (s.)	14.2*	27.9*	14.4*
	Organ (s.)	11.8*	25.1*	12.1*
	Piccolo (s.)	6.5*	20.0	6.7*
	Organ (s.)	6.6*	17.3	7.1*
	Oboe (s.)	9.0*	21.9*	9.2*
Ours	Oboe (a.)	18.6*	33.6*	18.8*
	Euphonium (a.)	14.7*	31.5*	14.7*
	Piccolo (s.)	11.2	25.9	11.3
	Organ (s.)	10.1	20.7	10.5
	Piccolo (s.)	4.2	24.8*	4.3
	Organ (s.)	6.0	20.0*	6.3
	Oboe (s.)	5.3	12.4	6.4

Best numbers are marked

The main goal of our algorithm was to provide good performance for acoustic instruments, and in fact, on the combination of two acoustic instruments, it exceeds the original performance of the compared method by roughly 10 dB in SDR. For the synthetic instruments, the performance achieved by the algorithm in [23] is mostly superior, while our algorithm still attains acceptable performance for piccolo and organ, and we demonstrate that it can at least in principle also be applied to combinations of more than two instruments.

The corresponding PEASS scores for the separated tracks are given in Table 10. Here, in the example with two acoustic instruments, the separation of the oboe track by the algorithm in [23] receives a higher OPS and IPS, suggesting that the overall quality of our separation is perceptually worse and this is at least partly caused by interference. However, according to our own listening opinion, the result from our algorithm matches the original signal very well and contains no audible interference while the result from the compared algorithm contains very obvious interference and also other representation errors, so we cannot explain the outcome of this evaluation. On the other hand, with the synthetic instruments, it is now often our algorithm that is preferred.

8 Conclusion and future work

We developed a novel algorithm to represent discrete mixture spectra as sparse shifted linear combinations

of analytically given non-negative continuous patterns. We applied this algorithm to spectrograms of audio recordings, first to convert an STFT magnitude spectrogram into a log-frequency spectrogram, then to identify patterns of peaks related to the sounds of musical instruments in the context of a dictionary learning algorithm based on Adam, a method that originates from the field of deep learning.

This led us to an algorithm to perform blind source separation on polyphonic music recordings of wind and string instruments, making only minimal structural assumptions about the data. In its model, the spectral properties of the musical instruments are fixed and pitch-invariant. Thus, instruments that satisfy this assumption can be represented irrespectively of their tuning. The only parameters that have to be known a-priori are the number of instruments and an upper bound for the sparsity level.

When applied to recordings of appropriate acoustic instruments, the performance of our algorithm surpasses that of comparable literature. Further, we show that once a dictionary has been trained on a certain combination of instruments, it can be applied to combinations of “related” instruments, even if those have a different tuning.

We note, however, that blind source separation always needs favorable data: Representing other kinds of instruments would require a different model, and instruments with a pronounced attack sound are also problematic. The sound of the instruments must be sufficiently pitch- and volume-invariant with only little overall variation in the harmonic structure, and the sparsity level must be rather strict.

While the pitch-invariant spectrogram substantially facilitates the identification of the instrument sounds, it has a lower resolution in the high frequencies, and therefore some information from the STFT spectrogram is lost. Also, any phase information is lost completely. Despite an inharmonicity parameter being included in our model, instruments with strong inharmonicity are problematic to identify.

Overall, while our algorithm appears to work well in certain settings, the framework that was created in order to bring the computational complexity under control is not very flexible. Instead of the hand-crafted pursuit algorithm, one may also consider the application of a neural network for identification, while still doing blind separation via a parametric model.

In our application, the frequency shift of the spectrum is caused by a change in pitch. Another common source for frequency shifts in various areas (such as communication technology or astronomy) is the Doppler effect. We believe that this could open new applications for our pursuit algorithm and potentially also the dictionary learning algorithm. Specifically, the pursuit algorithm could be

Table 10 PEASS scores for the best-case runs of different instrument combinations, with spectral masking. Instruments labeled as “s.” are synthetic, those labeled as “a.” are acoustic

Method	Instrument	OPS	TPS	IPS	APS
[23]	Oboe (a.)	24*	33	82*	9
	Euphonium (a.)	24	66	43*	5
	Piccolo (s.)	48*	74	59	54
	Organ (s.)	41	86	73	87
	Piccolo (s.)	22	67*	35	35*
	Organ (s.)	28	63*	60*	58
	Oboe (s.)	44*	70*	58	57
Ours	Oboe (a.)	19	99*	44	66*
	Euphonium (a.)	34*	70*	38	60*
	Piccolo (s.)	24	83*	69*	77*
	Organ (s.)	79*	93*	86*	87
	Piccolo (s.)	27*	29	56*	32
	Organ (s.)	38*	53	50	52
	Oboe (s.)	20	61	68*	80*

Best numbers are marked. The APS in the fourth row for each method was a perfect tie

used as an alternative to continuous basis pursuit [42], which is advertised as a method for radar and geological seismic data and has been used for the analysis of neural spike trains [65].

Appendix

We provide pseudo-code with a description of the implementation details as well as some additional figures with commentary.

Pseudo-code

We will now present the mentioned algorithms in more detail via pseudo-code. First, Algorithm 1 describes the sparse pursuit/identification algorithm. It takes as arguments the dictionary, the sample vector, a selector function, the sparsity levels, and the sum of the previous amplitudes for each pattern (which will become important for dictionary pruning). In the non-linear optimization step, it calls the L-BFGS-B minimizer to minimize the loss L_D with respect to the given parameters.

Algorithm 1 Sparse identification algorithm

```

function PURSUIT( $D, Y, \text{select}, N_{\text{pre}}, N_{\text{spr}}, A$ )
   $\mathcal{J} \leftarrow \emptyset$ 
   $r \leftarrow Y^q$ 
  loop  $N_{\text{itr}}$  times
     $a_j \leftarrow 0, \theta_j \leftarrow \theta_{\text{nil}}$ 
    for  $j \in \{1, \dots, N_{\text{spr}} + N_{\text{pre}}\} \setminus \mathcal{J}$ 
       $\mathcal{J}_{\text{new}} \leftarrow \text{sort}(\{1, \dots, N_{\text{spr}} + N_{\text{pre}}\} \setminus \mathcal{J})[1, \dots, N_{\text{pre}}]$ 
       $a_{\mathcal{J}_{\text{new}}}, \mu_{\mathcal{J}_{\text{new}}}, \eta_{\mathcal{J}_{\text{new}}} \leftarrow \text{select}(r, y_0, \dots, y_{N_{\text{pat}}-1}, \mathcal{J}_{\text{new}}, N_{\text{pre}})$ 
       $\mathcal{J} \leftarrow \mathcal{J} \cup \mathcal{J}_{\text{new}}$ 
       $a_{\mathcal{J}}, \mu_{\mathcal{J}}, \theta_{\mathcal{J}} \leftarrow \text{bfgs}(L_D, Y, a_{\mathcal{J}}, \mu_{\mathcal{J}}, \eta_{\mathcal{J}}, \theta_{\mathcal{J}})$ 
       $a_{\mathcal{J}} \geq 0, \theta_{\mathcal{J}} \in \Omega_{\theta}$ 
    for  $\eta = 0, \dots, N_{\text{pat}} - 1$  do
       $\mathcal{J}_{\eta} \leftarrow \left( \arg \text{sort}_{j \in \{1, \dots, N_{\text{spr}} + N_{\text{pre}}\}, \eta_j = \eta} a_j \right)$ 
       $[1, \dots, N_{\text{spr}}]$ 
       $\mathcal{J} \leftarrow \bigcup_{\eta \in \{0, \dots, N_{\text{pat}}-1\}} \mathcal{J}_{\eta}$ 
       $a_{\mathcal{J}}, \mu_{\mathcal{J}}, \theta_{\mathcal{J}} \leftarrow \text{bfgs}(L_D, Y, a_{\mathcal{J}}, \mu_{\mathcal{J}}, \eta_{\mathcal{J}}, \theta_{\mathcal{J}})$ 
       $a_{\mathcal{J}} \geq 0, \theta_{\mathcal{J}} \in \Omega_{\theta}$ 
       $r \leftarrow Y^q - \left( \sum_{j \in \mathcal{J}} a_j y_{\eta_j, \theta_j} (\cdot - \mu_j) \right)^q$ 
       $\theta \leftarrow \|r\|_2$ 
      if  $\|r\|_2 \geq \lambda \theta$  then
        restore values from previous iteration
      break
    for  $\eta = 0, \dots, N_{\text{pat}} - 1$  do
       $A[\eta] \leftarrow A[\eta] + \sum_{\eta_j = \eta} a_j$ 
  return  $\mathcal{J}, a_{\mathcal{J}}, \mu_{\mathcal{J}}, \eta_{\mathcal{J}}, \theta_{\mathcal{J}}, A$ 

```

The `sel_xcorr` function in Algorithm 2 is used in the separation. It selects up to N_{pre} patterns based on cross-correlation, and it computes their discrete amplitudes and shifts. In the implementation, this is accelerated via

the FFT convolution theorem. The `sel_peaks` function in Algorithm 3 ignores the patterns, and it simply returns the N_{pre} largest local maxima with dominance N_{dom} (typically, $N_{\text{dom}} = 3$).

The dictionary learning algorithm (Algorithm 4) is largely identical to the original formulation of Adam (with values $\beta_1 = 0.9, \beta_2 = 0.999, \varepsilon = 10^{-8}$, and a step-size of $\kappa = 10^{-3}$), except that v_2 is averaged over all the harmonics for one instrument. It counts the number of training iterations $\tau[\eta]$ for each instrument η individually. The dictionary is initialized by the function in Algorithm 5, which creates a new dictionary column with random values. The function in Algorithm 6 removes seldom-used instruments in the dictionary by comparing their average amplitude but with a head start which is half the length of the pruning interval: $\tau_0 = N_{\text{prn}}/2$.

The `logspect` function in Algorithm 7 takes an STFT magnitude spectrogram and applies the sparse pursuit algorithm in order to convert it into a log-frequency spectrogram with a height of $m = 1024$. Finally, the separate function in Algorithm 8 performs the overall separation procedure. It prunes the dictionary every $N_{\text{prn}} = 500$ steps.

Algorithm 2 Selector function based on cross-correlation

```

function SEL_XCORR( $r, y_0, \dots, y_{N_{\text{pat}}-1}, \mathcal{J}, N_{\text{pre}}$ )
   $\rho[\mu, \eta] \leftarrow \sum_{i=0}^{m-1} r[i] (y_{\eta, \theta_{\text{nil}}}[i - \mu])^q / \|y_{\eta, \theta_{\text{nil}}}[\cdot]\|_2^q$ 
  for  $\mu = 0, \dots, m - 1$  and  $\eta = 0, \dots, N_{\text{pat}} - 1$ 
     $(\mu_{\mathcal{J}}, \eta_{\mathcal{J}}) \leftarrow \arg \text{sort}_{(\mu, \eta)} (-\rho[\mu, \eta])[: N_{\text{pre}}]$ 
   $a_{\mathcal{J}} \leftarrow (\rho[\mu_{\mathcal{J}}, \eta_{\mathcal{J}}] / \|y_{\eta_{\mathcal{J}}, \theta_{\text{nil}}}[\cdot]\|_2^q)^{1/q}$ 
   $\mathcal{J} \leftarrow \{j \in \mathcal{J} : a_j > 0\}$ 
  return  $a_{\mathcal{J}}, \mu_{\mathcal{J}}, \eta_{\mathcal{J}}$ 

```

Algorithm 3 Selector function based on peaks

```

function SEL_PEAKEs( $r, \_, \mathcal{J}, N_{\text{pre}}$ )
   $\mu_{\mathcal{J}} \leftarrow \arg \text{sort}_{\mu}$ 
   $\{a_{\mu} : r[\mu] \geq r[\mu + k], |k| \leq N_{\text{dom}}\}[: N_{\text{pre}}]$ 
   $\mathcal{J} \leftarrow \{j \in \mathcal{J} : r[j] > 0\}$ 
  return  $r_{\mathcal{J}}, \mu_{\mathcal{J}}, 0$ 

```

Benefits over the Mel spectrogram

In Fig. 3, we compared our log-spectrogram that was computed via the sparse pursuit method to the mel spectrogram and the constant-Q transform. We concluded in Section 3 that as the CQT uses windows of different length for different frequencies, it is not a good choice for our dictionary representation.

The mel spectrogram does not have this particular problem, but the Heisenberg uncertainty principle constrains the time-log-frequency resolution according to the

Algorithm 4 Dictionary learning function

```

function ADAM( $D, \tau, v_1, v_2, g$ )
  for  $\eta = 0, \dots, N_{\text{pat}} - 1$  do
     $\tau[\eta] \leftarrow \tau[\eta] + 1$ 
     $v_1[\cdot, \eta] \leftarrow \beta_1 \cdot v_1[\cdot, \eta] + (1 - \beta_1) \cdot g[\cdot, \eta]$ 
     $v_2[\eta] \leftarrow \beta_2 \cdot v_2[\eta] + (1 - \beta_2) \cdot \text{mean}(g[\cdot, \eta]^2)$ 
     $\hat{v}_1[\cdot, \eta] \leftarrow v_1[\cdot, \eta] / (1 - \beta_1^{\tau[\eta]})$ 
     $\hat{v}_2[\eta] \leftarrow v_2[\eta] / (1 - \beta_2^{\tau[\eta]})$ 
     $D[\cdot, \eta] \leftarrow D[\cdot, \eta] - \kappa \cdot \hat{v}_1[\cdot, \eta] / (\sqrt{\hat{v}_2[\eta]} + \varepsilon)$ 
     $D[\cdot, \eta] \leftarrow \max(0, \min(1, D[\cdot, \eta]))$ 
  return  $D, \tau, v_1, v_2$ 

```

Algorithm 5 Dictionary initialization function

```

function INIT()
   $e \leftarrow \text{Par}(1, 0.5)$ 
  for  $h = 1, \dots, N_{\text{har}}$  do
     $d[h] \leftarrow \mathcal{U}[0, 1]$ 
     $d[h] \leftarrow d[h] / h^e$ 
  return  $d[\cdot]$ 

```

Algorithm 6 Dictionary pruning function

```

function PRUNE( $A, D, \tau$ )
   $\mathcal{I} \leftarrow \left( \arg \text{sort}_{\eta \in \{0, \dots, N_{\text{pat}} - 1\}} A[\eta] / (\tau[\eta] - \tau_0) \right)$ 
   $[0, \dots, N_{\text{ins}}]$ 
   $\tau[\mathcal{I}^c] = 0, v_1[\cdot, \mathcal{I}^c] = 0, v_2[\mathcal{I}^c] = 0, A[\mathcal{I}^c] = 0$ 
  for  $\eta \in \mathcal{I}^c$  do
     $D[\cdot, \eta] \leftarrow \text{INIT}()$ 
  return  $\mathcal{I}, \tau, v_1, v_2, A, D$ 

```

Algorithm 7 Log-spectrogram generation function

```

function LOGSPECT( $Z, N_{\text{pre}}, N_{\text{spr}}$ )
  for  $t = 0, \dots, n - 1$  do
     $\mathcal{J}_t, a_{\mathcal{J}_t, t}, \mu_{\mathcal{J}_t, t}, \eta_{\mathcal{J}_t, t}, \theta_{\mathcal{J}_t, t} \leftarrow$ 
       $\leftarrow \text{PURSUIT}([1], Z[\cdot, t], \text{sel\_peaks}, 1, N_{\text{spr}}, \_)$ 
     $U[\alpha, t] \leftarrow \sum_{j, h} a_{j, h, t} \exp(-(\alpha - \alpha(\mu_{\mathcal{J}_t, t}))^2 / (2F^2 \sigma_{j, t}^2))$ 
    for  $\alpha = 0, \dots, m - 1, t = 0, \dots, n - 1$ 
  return  $U$ 

```

lowest frequency to be represented. In Fig. 3a, we cut the spectrogram at 530 Hz (which corresponds to 577 Hz when compensating for the different sampling frequency), but for our sample with recorder and violin, this is not sufficient, as it contains notes as low as c'. Thus, we chose the lowest frequency as 200 Hz, sacrificing some resolution.

We computed the mel spectrogram on this sample and ran the separation algorithm 10 times with $N_{\text{trn}} = 100000$ training iterations in order to obtain a fair comparison.

Algorithm 8 Dictionary learning and separation function

```

function SEPARATE( $U, N_{\text{pre}}, N_{\text{spr}}$ )
  for  $\eta = 0, \dots, N_{\text{pat}} - 1$  do
     $D[\cdot, \eta] \leftarrow \text{INIT}()$ 
     $\tau[\eta] \leftarrow 0, v_1[\cdot, \eta] \leftarrow 0, v_2[\eta] \leftarrow 0, A[\eta] \leftarrow 0$ 
  loop a multiple of  $N_{\text{prn}}$  times
     $t \leftarrow \text{random}(\{0, \dots, n - 1\})$ 
     $\mathcal{J}, a_{\mathcal{J}}, \mu_{\mathcal{J}}, \eta_{\mathcal{J}}, \theta_{\mathcal{J}}, A \leftarrow$ 
       $\leftarrow \text{PURSUIT}(D, U[\cdot, t], \text{sel\_xcorr}, 1, N_{\text{spr}}, A)$ 
     $g \leftarrow \nabla_D L_D(Y, a_{\mathcal{J}}, \mu_{\mathcal{J}}, \eta_{\mathcal{J}}, \theta_{\mathcal{J}})$ 
     $D, \tau, v_1, v_2 \leftarrow \text{ADAM}(D, \tau, v_1, v_2, g)$ 
    if  $\min(\tau) \bmod N_{\text{prn}} = 0$  then
       $\mathcal{I}, \tau, v_1, v_2, A, D \leftarrow \text{PRUNE}(A, D, \tau)$ 
  for  $t = 0, \dots, n - 1$  do
     $\mathcal{J}_t, a_{\mathcal{J}_t, t}, \mu_{\mathcal{J}_t, t}, \eta_{\mathcal{J}_t, t}, \theta_{\mathcal{J}_t, t} \leftarrow$ 
       $\leftarrow \text{PURSUIT}(D[\cdot, \mathcal{I}], U[\cdot, t], \text{sel\_xcorr}, 1, N_{\text{spr}}, \_)$ 
  return  $\{\mathcal{J}_t, a_{\mathcal{J}_t, t}, \mu_{\mathcal{J}_t, t}, \eta_{\mathcal{J}_t, t}, \theta_{\mathcal{J}_t, t} : t = 0, \dots, n - 1\}$ 

```

The performance figures are given in Table 11 and Fig. 10, and the results from the best-case run with a random seed of 7 are displayed in Fig. 11. It can be seen that the performance does not reach what we achieved with a spectrogram generated via the sparse pursuit method (cf. Fig. 5 and Table 1).

Using again a one-sided Wilcoxon signed-rank test, we find that without spectral masking, the SDR when using the mel spectrogram is worse at $p_{\text{Recorder}} = 9.8 \times 10^{-4}$ and $p_{\text{Violin}} = 2.0 \times 10^{-3}$. With spectral masking applied, we achieve $p_{\text{Recorder}} = p_{\text{Violin}} = 9.8 \times 10^{-4}$, as for each random seed $0, \dots, 9$, the results from our representation are consistently better.

We thus conclude that our use of the sparse pursuit algorithm for generating a log-frequency spectrogram provides a notable benefit for the subsequent processing.

Log-frequency spectrograms of the instrument tracks

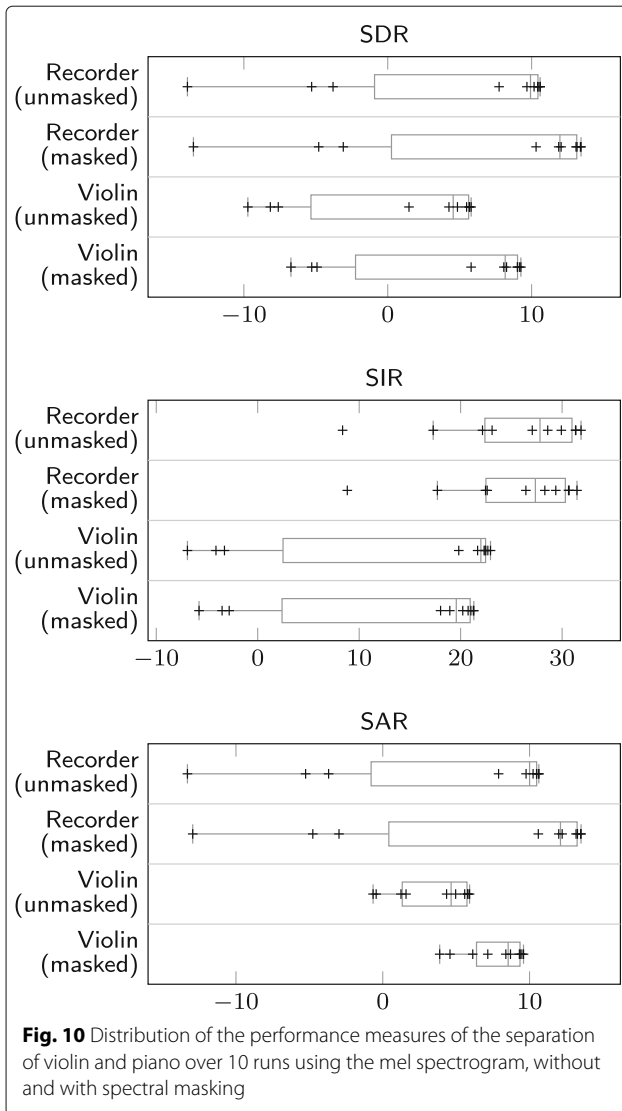
For additional comparison, we computed the log-frequency spectrograms of the original (Fig. 12) and the computed (Fig. 13) instrument tracks.

It should be noted that these spectrograms are not used anywhere in the computation or evaluation process, and

Table 11 Performance measures for the best-case run of the separation of recorder and violin using the mel spectrogram

Mask	Instrument	SDR	SIR	SAR
No	Recorder	10.6	31.9*	10.6
	Violin	5.8	22.5*	5.9
Yes	Recorder	13.4*	31.5	13.5*
	Violin	9.3*	21.0	9.6*

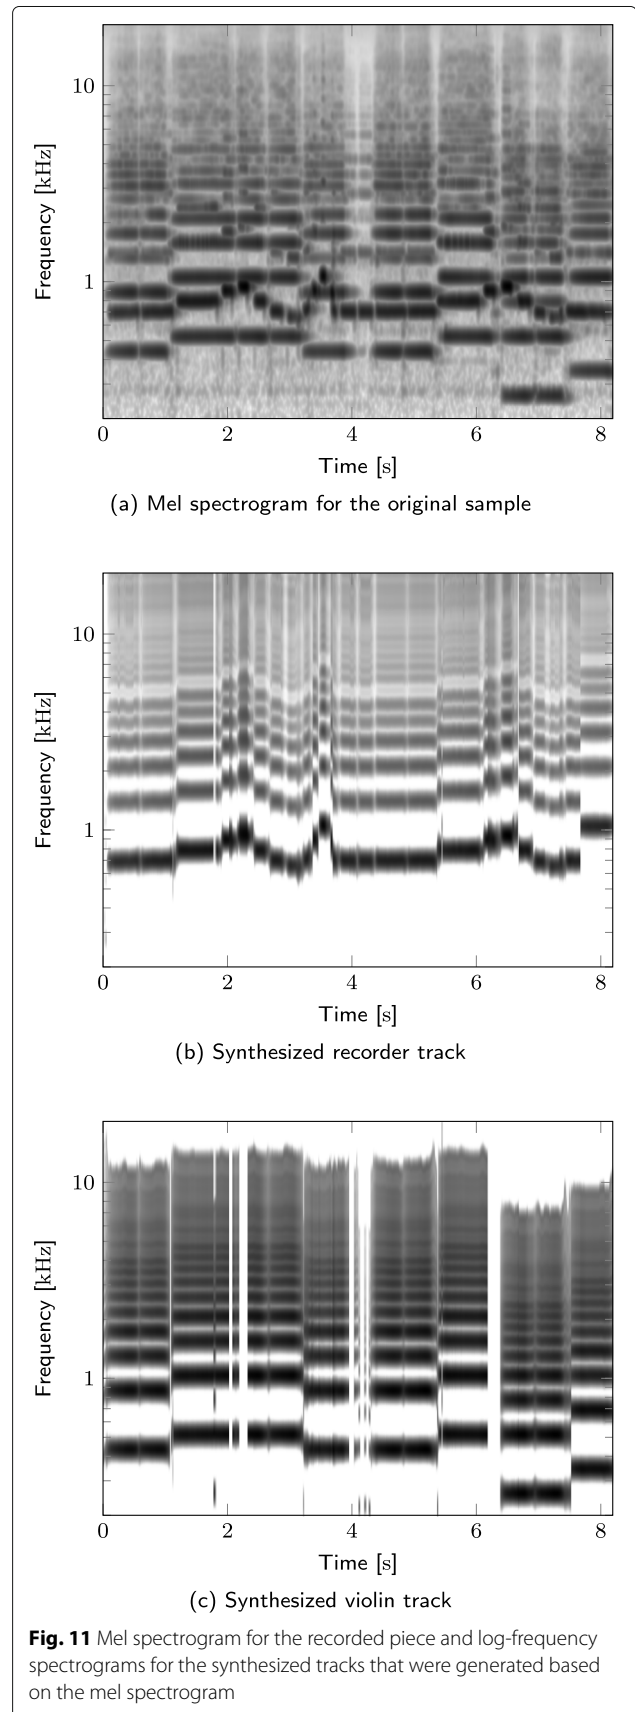
Best numbers are marked




due to artifacts from the sparse pursuit algorithm, they are not an accurate representation of the time-domain signal.


Nevertheless, two effects can be seen when comparing Fig. 13 to Fig. 5:

1. Due to spectral masking, the harmonics now have different intensities.
2. The Griffin-Lim phase reconstruction algorithm *smoothes* some of the artifacts that were introduced by the sparse pursuit algorithm. This is because not every two-dimensional image is actually a valid spectrogram that corresponds to an audio signal; instead, the Griffin-Lim algorithm aims to find an audio signal whose spectrogram is as close as possible to the given image, and it uses the phase of the original mixed sample as the initial value.







(a) Recorder




(b) Violin




(c) Mixture (copy of Figure 5a)

Fig. 12 Sparsity-derived log-frequency spectrograms of the original instrument samples

(a) Recorder



(b) Violin



(c) Mixture

Fig. 13 Sparsity-derived log-frequency spectrograms of the separated audio tracks and its mixture

Abbreviations

Adam: Adaptive moment estimation; APS: Artifacts-related perceptual score; BASS-dB: Blind Audio Source Separation Evaluation Database; CD: Compact disc; EM: Expectation maximization; FFT: Fast Fourier transform; CQT: Constant-Q transform; ICA: Independent component analysis; IPS: Interference-related perceptual score; K-SVD: K-singular value decomposition; LASSO: Least absolute shrinkage and selection operator; LSTM: Long short-term memory; L-BFGS-B: Limited-memory Broyden-Fletcher-Goldfarb-Shanno with bounds; MPE: Multiple pitch estimation; NMF: Non-negative matrix factorization; NMFD: Non-negative matrix factor deconvolution; NMF2D: Non-negative matrix factor two-dimensional deconvolution; OMP: Orthogonal matching pursuit; OPS: Overall perceptual score; PEASS: Perceptual Evaluation methods for Audio Source Separation; PLCA: Probabilistic latent component analysis; SISEC: Signal Separation Evaluation Campaign; SAR: Signal-to-artifacts ratio; SDR: Signal-to-distortion ratio; SIR: Signal-to-interference ratio; STFT: Short-time Fourier transform; TPS: Target-related perceptual score; URMP: University of Rochester Multi-Modal Music Performance Dataset

Acknowledgements

The authors would like to thank Bernhard G. Bodmann, Gitta Kutyniok, and Monika Dörfler for engaging discussions on the subject and Kara Tober for playing the clarinet samples. Further, we thank the anonymous reviewers, whose valuable and constructive comments helped us improve the manuscript.

Authors' contributions

SS devised the algorithm, wrote the source code, and drafted the manuscript. EJK supervised the research and revised the manuscript. All authors read and approved the final manuscript.

Funding

The first author acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Projektnummer 281474342/GRK2224/1. Open Access funding enabled and organized by Projekt DEAL.

Availability of data and materials

The software is presented along with audio samples and the original input data on the institute website¹⁰. The source code is available on GitHub¹¹ under the GNU General Public License (version 3).

Competing interests

The authors declare that they have no competing interests.

Author details

¹AG Computational Data Analysis, Faculty 3, University of Bremen, Bibliothekstr. 5, 28359 Bremen, Germany. ²Mathematics Department, Colorado State University, 1874 Campus Delivery, 111 Weber Bldg, 80523 Fort Collins, CO, USA.

Received: 21 May 2020 Accepted: 1 December 2020

Published online: 28 January 2021

References

1. K. Gröchenig, *Foundations of Time-Frequency Analysis*. (Birkhäuser, Basel, 2001). <https://doi.org/10.1007/978-1-4612-0003-1>
2. D. Griffin, J. Lim, Signal estimation from modified short-time Fourier transform. *IEEE Trans. Acoust. Speech Signal Process.* **32**(2), 236–243 (1984). <https://doi.org/10.1109/TASSP.1984.1164317>
3. E. Vincent, T. Virtanen, S. Gannot (eds.), *Audio Source Separation and Speech Enhancement* (Wiley, Chichester, 2018)
4. S. Makino (ed.), *Audio Source Separation* (Springer, Cham, 2018)
5. J.-T. Chien, *Source Separation and Machine Learning*. (Academic Press, London, 2018)
6. D. D. Lee, H. S. Seung, Learning the parts of objects by non-negative matrix factorization. *Nature*. **401**(6755), 788–791 (1999). <https://doi.org/10.1038/44565>
7. P. Smaragdis, J. C. Brown, in *Applications of Signal Processing to Audio and Acoustics, 2003 IEEE Workshop On*. Non-negative matrix factorization for polyphonic music transcription (IEEE, 2003), pp. 177–180. <https://doi.org/10.1109/ASPA.2003.1285860>
8. B. Wang, M. D. Plumbley, in *Proc. UK Digital Music Research Network (DMRN) Summer Conf.* Musical audio stream separation by non-negative matrix factorization, (Glasgow, 2005)
9. C. Févotte, E. Vincent, A. Ozerov, in *Audio Source Separation*. Single-channel audio source separation with NMF: divergences, constraints and algorithms (Springer, Cham, 2018), pp. 1–24. https://doi.org/10.1007/978-3-319-73031-8_1
10. J. C. Brown, Calculation of a constant Q spectral transform. *J. Acoust. Soc. Am.* **89**(1), 425–434 (1991). <https://doi.org/10.1121/1.400476>
11. D. Fitzgerald, M. Cranitch, E. Coyle, in *Statistical Signal Processing, 2005 IEEE/SP 13th Workshop On*. Shifted non-negative matrix factorisation for sound source separation (IEEE, Bordeaux, 2005), pp. 1132–1137. <https://doi.org/10.1109/SSP.2005.1628765>
12. R. Jaiswal, D. Fitzgerald, D. Barry, E. Coyle, S. Rickard, in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference On*. Clustering NMF basis functions using shifted NMF for monaural sound source separation (IEEE, Prague, 2011), pp. 245–248. <https://doi.org/10.1109/ICASSP.2011.5946386>
13. R. Jaiswal, D. Fitzgerald, E. Coyle, S. Rickard, in *22nd IET Irish Sig. and Sys. Conf.* Shifted NMF using an efficient constant-Q transform for monaural sound source separation (IET, Dublin, 2011)
14. R. Jaiswal, D. Fitzgerald, E. Coyle, S. Rickard, in *24th IET Irish Signals and Systems Conference*. Towards shifted NMF for improved monaural separation (IET, Dublin, 2013)
15. P. Smaragdis, B. Raj, M. Shashanka, in *Advances in Models for Acoustic Processing Workshop (NIPS)*. A probabilistic latent variable model for acoustic modeling, (Vancouver, 2006)
16. P. Smaragdis, B. Raj, M. Shashanka, in *International Conference on Independent Component Analysis and Signal Separation*. Supervised and semi-supervised separation of sounds from single-channel mixtures (Springer, London, 2007), pp. 414–421. https://doi.org/10.1007/978-3-540-74494-8_52
17. P. Smaragdis, B. Raj, M. Shashanka, in *Acoustics, Speech and Signal Processing (ICASSP), 2008 IEEE International Conference On*. Sparse and shift-invariant feature extraction from non-negative data (IEEE, Las Vegas), pp. 2069–2072. <https://doi.org/10.1109/ICASSP.2008.4518048>
18. B. Fuentes, R. Badeau, G. Richard, in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference On*. Adaptive harmonic time-frequency decomposition of audio using shift-invariant PLCA (IEEE, Prague, 2011), pp. 401–404. <https://doi.org/10.1109/ICASSP.2011.5946425>
19. B. Fuentes, A. Liutkus, R. Badeau, G. Richard, in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference On*. Probabilistic model for main melody extraction using constant-Q transform (IEEE, Kyoto, 2012), pp. 5357–5360. <https://doi.org/10.1109/ICASSP.2012.6289131>
20. B. Fuentes, R. Badeau, G. Richard, in *2012 Proceedings of the 20th European Signal Processing Conference (EUSIPCO)*. Blind harmonic adaptive decomposition applied to supervised source separation (IEEE, Bucharest, 2012), pp. 2654–2658
21. B. Fuentes, R. Badeau, G. Richard, Harmonic adaptive latent component analysis of audio and application to music transcription. *IEEE Trans. Audio Speech Lang. Process.* **21**(9), 1854–1866 (2013). <https://doi.org/10.1109/TASL.2013.2260741>
22. E. Vincent, N. Bertin, R. Badeau, Adaptive harmonic spectral decomposition for multiple pitch estimation. *IEEE Trans. Audio Speech Lang. Process.* **18**(3), 528–537 (2009). <https://doi.org/10.1109/TASL.2009.2034186>
23. Z. Duan, Y. Zhang, C. Zhang, Z. Shi, Unsupervised single-channel music source separation by average harmonic structure modeling. *IEEE Trans. Audio Speech Lang. Process.* **16**(4), 766–778 (2008). <https://doi.org/10.1109/TASL.2008.919073>
24. N. H. Fletcher, T. D. Rossing, *The Physics of Musical Instruments*. (Springer, New York, 2012)
25. S. Schulze, E. J. King, A frequency-uniform and pitch-invariant time-frequency representation. *PAMM*. **19** (2019). <https://doi.org/10.1002/PAMM.201900374>

¹⁰<https://www.math.colostate.edu/~king/software.html#Musisep>

¹¹<https://github.com/rgcda/Musisep>

26. G. B. Folland, A. Sitaram, The uncertainty principle: A mathematical survey. *J. Fourier Anal. Appl.* **3**(3), 207–238 (1997). <https://doi.org/10.1007/BF02649110>
27. P. Smaragdis, in *International Conference on Independent Component Analysis and Signal Separation*. Non-negative matrix factor deconvolution; extraction of multiple sound sources from monophonic inputs (Springer, Granada, 2004), pp. 494–499. https://doi.org/10.1007/978-3-540-30110-3_63
28. M. N. Schmidt, M. Mørup, in *International Conference on Independent Component Analysis and Signal Separation*. Nonnegative matrix factor 2-D deconvolution for blind single channel source separation (Springer, 2006), pp. 700–707. https://doi.org/10.1007/11679363_87
29. T. Virtanen, in *ISCA Tutorial and Research Workshop (ITRW) on Statistical and Perceptual Audio Processing*. Separation of sound sources by convolutive sparse coding, (Jeju, 2004)
30. T. Virtanen, Monaural sound source separation by nonnegative matrix factorization with temporal continuity and sparseness criteria. *IEEE Trans. Audio Speech Lang. Process.* **15**(3), 1066–1074 (2007). <https://doi.org/10.1109/TASL.2006.885253>
31. T. Blumensath, M. Davies, Sparse and shift-invariant representations of music. *IEEE Trans. Audio Speech Lang. Process.* **14**(1), 50–57 (2005). <https://doi.org/10.1109/TSA.2005.860346>
32. R. Hennequin, R. Badeau, B. David, NMF with time-frequency activations to model nonstationary audio events. *IEEE Trans. Audio Speech Lang. Process.* **19**(4), 744–753 (2010). <https://doi.org/10.1109/TASL.2010.2062506>
33. R. Hennequin, B. David, R. Badeau, in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference On*. Score informed audio source separation using a parametric model of non-negative spectrogram (IEEE, Prague, 2011), pp. 45–48. <https://doi.org/10.1109/ICASSP.2011.5946324>
34. A. J. Munoz-Montoro, J. J. Carabias-Orti, P. Vera-Candeas, F. J. Canadas-Quesada, N. Ruiz-Reyes, Online/offline score informed music signal decomposition: application to minus one. *J. Audio Speech Music Proc.* **2019**, 1–30 (2019). <https://doi.org/10.1186/s13636-019-0168-6>
35. P.-S. Huang, M. Kim, M. Hasegawa-Johnson, P. Smaragdis, Joint optimization of masks and deep recurrent neural networks for monaural source separation. *IEEE/ACM Trans. Audio Speech Lang. Process.* **23**(12), 2136–2147 (2015). <https://doi.org/10.1109/TASLP.2015.2468583>
36. F.-R. Stöter, A. Liutkus, N. Ito, in *International Conference on Latent Variable Analysis and Signal Separation*. The 2018 signal separation evaluation campaign (Springer, Guildford, 2018), pp. 293–305. https://doi.org/10.1007/978-3-319-93764-9_28
37. F.-R. Stöter, S. Uhlich, A. Liutkus, Y. Mitsufuji, Open-Unmix – a reference implementation for music source separation. *J. Open Source Softw.* **4**(41) (2019). <https://doi.org/10.21105/joss.01667>
38. A. Défossez, N. Usunier, L. Bottou, F. Bach, Music source separation in the waveform domain. *arXiv preprint arXiv:1911.13254* (2019)
39. J. A. Tropp, A. C. Gilbert, Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Trans. Inf. Theory.* **53**(12), 4655–4666 (2007). <https://doi.org/10.1109/TIT.2007.909108>
40. W. Dai, O. Milenkovic, Subspace pursuit for compressive sensing signal reconstruction. *IEEE Trans. Inf. Theory.* **55**(5), 2230–2249 (2009). <https://doi.org/10.1109/TIT.2009.2016006>
41. H. Bristow, A. Eriksson, S. Lucey, in *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference On*. Fast convolutional sparse coding (IEEE, 2013), pp. 391–398
42. C. Ekanadham, D. Tranchina, E. P. Simoncelli, Recovery of sparse translation-invariant signals with continuous basis pursuit. *IEEE Trans. Signal Process.* **59**(10), 4735–4744 (2011). <https://doi.org/10.1109/TSP.2011.2160058>
43. Y. De Castro, F. Gamboa, Exact reconstruction using Beurling minimal extrapolation. *J. Math. Anal. Appl.* **395**(1), 336–354 (2012). <https://doi.org/10.1016/j.jmaa.2012.05.011>
44. P. Catala, V. Duval, G. Peyré, in *J. Phys.: Conf. Ser.*, vol. 904. A low-rank approach to off-the-grid sparse deconvolution (IOP Publishing, Paris, 2017). <https://doi.org/10.1088/1742-6596/904/1/012015>
45. W. H. Barnes, *The Contemporary American Organ*, 8th edn. (J. Fischer and Bro., Glen Rock, 1964)
46. D. P. Kingma, J. Ba, Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
47. C. Févotte, J. Idier, Algorithms for nonnegative matrix factorization with the β -divergence. *Neural Comput.* **23**(9), 2421–2456 (2011). https://doi.org/10.1162/NECO_a_00168
48. R. H. Byrd, P. Lu, J. Nocedal, C. Zhu, A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput.* **16**(5), 1190–1208 (1995). <https://doi.org/10.1137/0916069>
49. C. Zhu, R. H. Byrd, P. Lu, J. Nocedal, Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Trans. Math. Softw.* **23**(4), 550–560 (1997). <https://doi.org/10.1145/279232.279236>
50. J. L. Morales, J. Nocedal, Remark on “Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound constrained optimization”. *ACM Trans. Math. Softw.* **38**(1), 7 (2011). <https://doi.org/10.1145/2049662.2049669>
51. M. Dörfler, T. Grill, R. Bammer, A. Flexer, Basic filters for convolutional neural networks: Training or design? *Neural Comput. Appl.*, 1–14 (2018). <https://doi.org/10.1007/S00521-018-3704-X>
52. R. J. McAulay, T. F. Quatieri, Speech analysis/synthesis based on a sinusoidal representation. *IEEE Trans. Acoust. Speech Signal Process.* **34**(4), 744–754 (1986). <https://doi.org/10.1109/TASSP.1986.1164910>
53. D. D. Lee, H. S. Seung, in *Adv. Neural Inf. Process. Syst.* Algorithms for non-negative matrix factorization, (Denver, 2001), pp. 556–562
54. M. Aharon, M. Elad, A. Bruckstein, K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans. Signal Process.* **54**(11), 4311–4322 (2006). <https://doi.org/10.1109/TSP.2006.881199>
55. M. Aharon, M. Elad, Sparse and redundant modeling of image content using an image-signature-dictionary. *SIAM J. Imaging Sci.* **1**(3), 228–247 (2008). <https://doi.org/10.1137/07070156X>
56. G. E. Pfander, P. Salanevich, Robust phase retrieval algorithm for time-frequency structured measurements. *SIAM J. Imaging Sci.* **12**(2), 736–761 (2019). <https://doi.org/10.1137/18M1205522>
57. E. Vincent, R. Gribonval, C. Févotte, Performance measurement in blind audio source separation. *IEEE Trans. Audio Speech Lang. Process.* **14**(4), 1462–1469 (2006). <https://doi.org/10.1109/TSA.2005.858005>
58. C. Févotte, R. Gribonval, E. Vincent, BSS_EVAL toolbox user guide – revision 2.0. Technical Report 1706, IRISA (2005). http://bass-db.gforge.inria.fr/bss_eval/
59. V. Emiya, E. Vincent, N. Harlander, V. Hohmann, Subjective and objective quality assessment of audio source separation. *IEEE Trans. Audio Speech Lang. Process.* **19**(7), 2046–2057 (2011). <https://doi.org/10.1109/TASL.2011.2109381>
60. E. Vincent, in *International Conference on Latent Variable Analysis and Signal Separation*. Improved perceptual metrics for the evaluation of audio source separation (Springer, 2012), pp. 430–437. https://doi.org/10.1007/978-3-642-28551-6_53
61. R Core Team, *R: A Language and Environment for Statistical Computing*. (R Foundation for Statistical Computing, Vienna, 2017)
62. E. Vincent, R. Gribonval, C. Févotte, BASS-dB: the blind audio source separation evaluation database. <http://www.irisa.fr/metis/BASS-dB/>. Accessed 23 Apr 2020
63. Z. Duan, B. Pardo, Soundprism: An online system for score-informed source separation of music audio. *IEEE J. Sel. Top. Signal Process.* **5**(6), 1205–1215 (2011). <https://doi.org/10.1109/JSTSP.2011.2159701>
64. B. Li, X. Liu, K. Dinesh, Z. Duan, G. Sharma, Creating a multitrack classical music performance dataset for multimodal music analysis: Challenges, insights, and applications. *IEEE Trans. Multimed.* **21**(2), 522–535 (2018). <https://doi.org/10.1109/TMM.2018.2856090>
65. C. Ekanadham, D. Tranchina, E. P. Simoncelli, A unified framework and method for automatic neural spike identification. *J. Neurosci. Methods.* **222**, 47–55 (2014). <https://doi.org/10.1016/j.jneumeth.2013.10.001>

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.