Audio Inpainting via ℓ_1 -Minimization and Dictionary Learning

Shristi Rajbamshi Georg Tauböck Nicki Holighaus Peter Balazs Acoustic Research Institute Acoustic Research Institute Acoustic Research Institute Acoustic Research Institute Austrian Academy of Sciences Austrian Academy of Sciences Austrian Academy of Sciences Austrian Academy of Sciences Vienna, Austria Vienna, Austria Vienna, Austria Vienna, Austria shristi.rajbamshi@oeaw.ac.at georg.tauboeck@oeaw.ac.at nicki.holighaus@oeaw.ac.at peter.balazs@oeaw.ac.at

Abstract—Audio inpainting refers to signal processing techniques that aim at restoring missing or corrupted consecutive samples in audio signals. Prior works have shown that ℓ_1 minimization with appropriate weighting is capable of solving audio inpainting problems, both for the analysis and the synthesis models. These models assume that audio signals are sparse with respect to some redundant dictionary and exploit that sparsity for inpainting purposes. Remaining within the sparsity framework, we utilize dictionary learning to further increase the sparsity and combine it with weighted ℓ_1 -minimization adapted for audio inpainting to compensate for the loss of energy within the gap after restoration. Our experiments demonstrate that our approach is superior in terms of signal-to-distortion ratio (SDR) and objective difference grade (ODG) compared with its original counterpart.

Index Terms—Audio inpainting, error concealment, sparse representation, dictionary learning, proximal algorithms, Chambolle-Pock algorithm.

I. INTRODUCTION

Audio signals are often prone to distortions resulting in modification or loss of information at certain sections. These localized distortions occur mostly during recording, transmission or storage and may span from several milliseconds to few seconds. They may be caused by, e.g., impulsive noise, packet loss, or scratches on the storage device. In the field of audio signal processing, the task of recovering such a lost section aka *gap* while ensuring that the audio artifact is almost imperceptible is known as *audio inpainting* [1]. We note that this recovery has also been referred to in the literature as audio interpolation/extrapolation, waveform substitution, and, more generally, for arbitrary (i.e., not necessarily audio) signals as *error concealment*.

Over the years, several audio inpainting methods have been proposed, based on various approaches like auto regression by Janssen [2], Bayesian estimation [3], convex optimization [4]–[6], sinusoidal modeling [7], similarity graphs [8], and deep neural networks [9]–[11]. Whereas auto regression, Bayesian estimation and convex optimization are applicable to short gaps (less than 100 ms), sinusoidal modeling, similarity graphs and deep neural networks work well for longer gaps (more than 100 ms).

This work was supported by the Vienna Science and Technology Fund (WWTF) project MA16-53 (INSIGHT) and the Austrian Science Fund (FWF) project I 3067-N30 (MERLIN) and Y 551-N13 (FLAME).

In this paper, we focus on audio inpainting approaches which are motivated by the observation that real-world audio signals may be sparsely represented in some overcomplete dictionary [1], [4], [5]. The cited methods exploit sparsity to fill the gap by employing optimization techniques while utilizing the information present around the gap. There are many different possibilities to formulate an audio inpainting problem as a sparsity-based optimization task. For example, by greedy heuristics such as OMP [12] and convex optimization via ℓ_1 -minimization [5].

A. Notation

Roman letters A, B, ..., a, b, ..., and a, b, ... designate matrices, vectors, and scalars, respectively. Sets will be denoted by calligraphic letters such as S. We use the notation A_S to indicate the column submatrix of A consisting of the columns indexed by S. A_{kl} denotes the element in the *k*th row and *l*th column of A. The superscripts ^T and ^H represent transposition and Hermitian transposition, respectively. I_N denotes the $N \times N$ identity matrix. We write for the indicator function $\chi_s(\mathbf{x})$, where

$$\chi_{\mathcal{S}}(\mathbf{x}) = \begin{cases} 0 & \mathbf{x} \in \mathcal{S} \\ +\infty & \text{otherwise.} \end{cases}$$

For a vector $\mathbf{u} = (u_0, u_1, \cdots, u_{N-1})^\mathsf{T}$ its support (i.e., indices of non-zero entries) is denoted by $\mathsf{supp}(\mathbf{u})$. Its ℓ_0 -"norm" and the ℓ_1 -norm are given as $\|\mathbf{u}\|_0 = \mathsf{card}(\mathsf{supp}(\mathbf{u}))$ (i.e., the cardinality of the support) and $\|\mathbf{u}\|_1 = \sum_n |u_n|$, respectively. Finally, we denote by $\|\mathbf{X}\|_{1,1} = \sum_{m,n} |X_{m,n}|$ the entry-wise ℓ_1 -norm of a matrix X.

II. MOTIVATION AND CONTRIBUTION

This work is inspired by two model-based optimization approaches for audio inpainting: *SParse Audio INpainter (SPAIN)* with dictionary learning [4], see also [13], and (weighted) ℓ_1 -minimization based audio inpainting [5]. Note that both methods use the Gabor transform (see Section III) to obtain a sparse representation of the original audio signal.

The optimization based inpainting methods are quite efficient in filling short gaps, but as the gaps tend to get longer, they suffer from energy loss within the gap after restoration. As a means to tackle this problem, remaining in the framework of sparse audio inpainting, in [5], various weighting methods were combined with ℓ_1 -minimization, whereas in [4], a dictionary learning method was proposed and combined with a modified version of the SPAIN algorithm, i.e., SPAIN-MOD [4]. This combination of dictionary learning and SPAIN-MOD will be referred to as SPAIN-LEARNED henceforth. SPAIN-LEARNED deforms the underlying Gabor dictionary such that the sparsity of the analysis coefficients of the resulting dictionary is further enhanced. A short description on how the Gabor dictionary is deformed, i.e, how the dictionary is learned from Gabor coefficients is given in VII-A. SPAIN-LEARNED exhibits excellent reconstruction performance compared to other techniques like (plain) SPAIN [6] and weighted ℓ_1 methods as presented in [5].

Motivated by these observations, we propose to slightly modify a selected weighted ℓ_1 -method from [5] as well as to incorporate the dictionary learning technique presented in [4]. As we will demonstrate, this modified weighted ℓ_1 -method exhibits significantly improved reconstruction performance compared with the original setting and achieves comparable results as the (currently) state-of-the-art SPAIN-LEARNED algorithm.

III. GABOR SYSTEMS

In this work, we will exploit time-frequency (TF) sparsity, specifically with respect to Gabor dictionaries. In a finite discrete setting, the TF coefficients are obtained by applying the discrete Gabor transform (DGT) [14] to the input audio signal. Note that the DGT is also commonly known as short-time Fourier transform (STFT).

For a window function $g \in \mathbb{R}^L$ and time and frequency hop sizes a and b, the collection of TF modulated versions of g, is called a *Gabor system*, i.e.,

$$\mathcal{G} = \{ \mathbf{g}[\cdot - na] e^{i2\pi mb \cdot /L} \}_{m,n} = \{ \mathbf{g}_p \}_{p=0,1,\dots,NM-1}.$$

Here, m = 0, 1, ..., M-1 is the frequency modulation index, n = 0, 1, ..., N-1 is the time index, and

$$p := p_{N,M}(m,n) = n + mN \tag{1}$$

is the combined time-frequency index. The matrix $G_S \in \mathbb{C}^{L \times NM}$, defined by setting its *p*-th column equal to g_p , is denoted as *synthesis operator* of the Gabor system \mathcal{G} . With the *analysis operator* $G_A = G_S^H$, the Gabor coefficients of a signal $x \in \mathbb{R}^L$ with respect to \mathcal{G} are given by

$$c_p = (\mathbf{G}_{\mathbf{A}}\mathbf{x})_p = \mathbf{g}_p^{\mathsf{H}} \cdot \mathbf{x},$$

where $c = (c_0, c_1, \dots, c_{NM-1})^T$ is called the coefficient vector of x and contains P = MN elements. An appropriate choice of g as well as parameters a and M will ensure that G_A has a bounded left-inverse, or equivalently, that \mathcal{G} is a *frame* for \mathbb{C}^L [15]. In other words, it allows perfect reconstruction from the redundant coefficients c. In that case, the composition $F = G_S G_A$, called the *frame operator*, is invertible and perfect reconstruction is achieved by the Gabor system generated from the *dual window* $F^{-1}g$. Most DGTs used in practice form a *tight* frame, i.e., $F = AI_L$, for some A > 0, such that the dual system is generated by (1/A)g. In this contribution, we restrict to such tight frames in order to remain consistent with the reference paper [5] for comparison later on.

Since we assume that $\mathbf{x}, \mathbf{g} \in \mathbb{R}^L$, it is easy to see that we have

$$c_{n+mN} = \overline{c_{n+(M-m)N}}, \text{ for all } m \in [1, \lceil M/2 \rceil - 1].$$
(2)

Hence, it is sufficient to compute only the first P' = NM' values of c, where $M' = \lfloor M/2 \rfloor + 1$, see also [14], [16]. Similarly, the analysis and synthesis operators can, for most purposes, be truncated at the NM'-th row and column, respectively. We will denote these truncated matrices by \tilde{G}_A and \tilde{G}_S and will refer to the corresponding transform as real DGT.

IV. PROBLEM FORMULATION

For an audio signal $x \in \mathbb{R}^L$, let the indices of the *reliable* and *unreliable* samples (= signal vs. gap resp.) be known. The goal of audio inpainting is to recover a signal, $y \in \mathbb{R}^L$ from $x \in \mathbb{R}^L$ such that the gap is filled (in a 'meaningful' way) and is also consistent with the reliable part of x. The (convex) set S_x consisting of all such possible signals can expressed as,

$$\mathcal{S}_{\mathbf{x}} \triangleq \left\{ \mathbf{y} \in \mathbb{R}^{L} : \mathbf{M}_{\mathbf{R}}\mathbf{y} = \mathbf{M}_{\mathbf{R}}\mathbf{x} \right\},$$

where $M_R : \mathbb{R}^L \to \mathbb{R}^L$ is a binary masking operator that sets unreliable samples to zero and keeps only the reliable samples.

A formulation of the sparse audio inpainting problem for an analysis model is,

$$\underset{z}{\arg\min} \|\tilde{G}_{A}z\|_{0} \quad \text{s.t.} \quad z \in \mathcal{S}_{x}. \tag{3}$$

The ℓ_0 -minimization problem described in (3) is NP hard. Hence, instead, one normally opts for its convex relaxation, i.e., ℓ_1 -minimization.

V. Audio inpainting with weighted $\ell_1\text{-minimization}$

Introducing a *weighting vector*

$$\tilde{\mathbf{w}} = (\tilde{w}_0, \dots, \tilde{w}_{NM'-1})^\mathsf{T} \in \mathbb{R}^{NM'},\tag{4}$$

see also [17], a sparsity-based formulation of audio inpainting with ℓ_1 -minimization is given as,

$$\underset{\mathbf{z}}{\arg\min} \|\tilde{\mathbf{w}} \odot \tilde{\mathbf{G}}_{\mathbf{A}} \mathbf{z}\|_{1} \quad \text{s.t.} \quad \mathbf{z} \in \mathcal{S}_{\mathbf{x}}. \tag{5}$$

where the symbol \odot denotes the element-wise product.

We will discuss the choice of weighting vector \tilde{w} in the next section. Instead of (5), we can consider the unconstrained problem

$$\arg\min_{\mathbf{z}} \|\tilde{\mathbf{w}} \odot \tilde{\mathbf{G}}_{\mathbf{A}} \mathbf{z}\|_1 + \chi_{\mathcal{S}_{\mathbf{x}}}(\mathbf{z}), \tag{6}$$

which can be solved using a proximal algorithm [18], [19]. Proximal methods are efficient in iteratively solving complex large-scale convex minimization problems by solving a series of more, but smaller and simpler convex optimization problems.

Note that in (6), if the real DGT, \tilde{G}_A , and corresponding (short) weighting vector $\tilde{w} \in \mathbb{R}^{NM'}$ were replaced by the

full DGT, G_A, and another (long) weighting vector, $w = (w_0, \ldots, w_{MN-1})^{\mathsf{T}} \in \mathbb{R}^{NM}$, i.e.,

$$\arg\min_{\mathbf{w}} \|\mathbf{w} \odot \mathbf{G}_{\mathbf{A}}\mathbf{z}\|_{1} + \chi_{\mathcal{S}_{\mathbf{x}}}(\mathbf{z}), \tag{7}$$

then (7) would be identical to the weighted ℓ_1 -minimization problem solved in [5] by means of the Chambolle-Pock (CP) [20]. However, in order to be compatible with our dictionary learning framework according to [4], see also Subsection VII-A, which is explicitly tailored to the real-valuedness of audio signals, we will consider (6) instead.

VI. CHOICE OF WEIGHTING VECTOR

Larger coefficients are often penalized more heavily than smaller coefficients in ℓ_1 -based minimization problems. In audio inpainting, this behaviour causes an energy loss within the gap in the restored signal [21], [22].

Besides this universal effect of ℓ_1 -minimization, the choice of the overcomplete dictionary, i.e., in our case, the Gabor transform also contributes to the energy drop phenomenon. This is because the coefficients corresponding to the window that overlap with the gap carry significantly less information about the reliable signal in comparison to the coefficients obtained from the reliable parts. In [5], in order to compensate for this loss, the Gabor atoms are weighted. Smaller weights are assigned to atoms carrying more unreliable information so that they contribute less to the objective function, thereby getting less penalized during minimization.

There are several approaches to obtain a weighting vector w, such as support-based, ℓ_1 -norm-based, ℓ_2 -norm-based, and energy-based approaches, see [5]. In this work, we will only consider *energy-based* weighting, as in pilot tests we found this choice to be the most promising. If g_p is a Gabor atom and M_Rg_p is its part corresponding to the reliable part of the signal, then energy-based weighting according to [5] assigns the weights $w = (w_0, \ldots, w_{P-1})^T$ for (7) with

$$w_p = \frac{\|\mathbf{M}_{\mathbf{R}}\mathbf{g}_p\|_2^2}{\|\mathbf{g}_p\|_2^2}.$$
 (8)

Similarly, for (6), the weights introduced in (4), $\tilde{w} = (\tilde{w}_0, \ldots, \tilde{w}_{P'-1})^{\mathsf{T}}$ are simply given by (8) with $\tilde{w}_p = w_p$.

VII. LEARNED WEIGHTED ℓ_1 -MINIMIZATION

A. Dictionary Learning

In [4], we increased TF sparsity by means of a sparsifying dictionary, obtained by a suitable deformation of the Gabor dictionary. The applied deformation matrix was learned from the TF coefficients in the neighborhood \mathcal{N} of the gap, via a basis optimization method developed in [23], [24].

Succinctly, the dictionary learning technique in [4] solves

$$\hat{\mathbf{D}} = \underset{\mathbf{D}\in\mathcal{D}}{\arg\min} \left\| \mathbf{D}\mathbf{X}_{\mathcal{N}} \right\|_{1,1},\tag{9}$$

where $\mathcal{N} \subset \{0, \ldots, N-1\}$ is a temporal neighborhood of the gap and $X_{\mathcal{N}} \in \mathbb{C}^{M' \times |\mathcal{N}|}$ denotes the Gabor analysis coefficients in that neighborhood, arranged in matrix form, i.e., $(X_{\mathcal{N}})_{m+1,n+1} = (\tilde{G}_A x)_{n+mN}$, for $n \in \mathcal{N}$ and $m \in$ $\{0, \ldots, M' - 1\}$. Further, $\mathcal{D} \in \mathbb{C}^{M' \times M'}$ represents a set of unitary deformation matrices with a special structure taking into account the conjugate-symmetry (2) underlying the real DGT (see [4] for a detailed description).

B. Weighted ℓ_1 -Minimization with Learned Dictionary

Utilizing conjugate-symmetry and including the learned dictionary, we can modify the optimization problems defined in (5) and (6). To this end, we define the following¹ block matrix

$$\hat{\mathbf{D}}_{\text{block}} := \begin{bmatrix} \hat{D}_{1,1}\mathbf{I}_N & \cdots & \hat{D}_{1,M'}\mathbf{I}_N \\ \vdots & \ddots & \vdots \\ \hat{D}_{M',1}\mathbf{I}_N & \cdots & \hat{D}_{M',M'}\mathbf{I}_N \end{bmatrix}$$

and replace (5) and (6) by the following minimization problems, respectively,

$$\arg\min_{z} \|w_{L} \odot \hat{D}_{block} \tilde{G}_{A} z\|_{1} \quad \text{s.t} \quad z \in \mathcal{S}_{x}$$

$$\arg\min_{z} \|w_{L} \odot \hat{D}_{block} \tilde{G}_{A} z\|_{1} + \chi_{\mathcal{S}_{x}}(z), \quad (10)$$

where $w_{L} \in \mathbb{R}^{P'}_{+}$ denotes a weighting vector that depends on the learned dictionary.

The elements of the deformed Gabor dictionary can be expressed as $\hat{\mathbf{g}}_p = (\tilde{\mathbf{G}}_{\mathbf{S}} \hat{\mathbf{D}}_{\text{block}}^{\mathsf{H}})_{p+1}, p \in \{0, \dots, P'-1\}$. Then, the weighting vector corresponding to the learned dictionary is given by $\hat{\mathbf{w}} = (\hat{w}_0, \dots, \hat{w}_{P'-1})^{\mathsf{T}}$ with

$$\hat{w}_p = \frac{\|\mathbf{M}_{\mathbf{R}}\hat{\mathbf{g}}_p\|_2^2}{\|\hat{\mathbf{g}}_p\|_2^2}$$

As in the original setting [5], we can also solve (10) using the CP proximal algorithm. The learned weighted ℓ_1 -minimization algorithm for inpainting (CP-LEARNED) is summarized in Alg. 1.

Note that we define $\operatorname{clip}_{w_L}(z)$, $\operatorname{proj}_{\mathcal{S}_x}(z)$, and stopping criterion essentially in the same way as in [5],

$$\begin{aligned} \text{clip}_{w_L}\left(z\right) &= z - \arg(z) \odot \max(|z| - w_L, 0),\\ \text{proj}_{\mathcal{S}_r}\left(z\right) &= \Re\left((I - M_R)z + M_Rx\right), \end{aligned}$$

by, additionally, taking into account that S_x is real-valued. The stopping criterion is

$$\|\tilde{G}_{A}z^{(i)} - \tilde{G}_{A}z^{(i-1)}\|_{2}^{2} \leq \epsilon \|\tilde{G}_{A}z^{(i-1)}\|_{2}^{2}$$

where $\epsilon > 0$ is the tolerance.

VIII. SIMULATION RESULTS AND DISCUSSION

In this section, we compare the audio inpainting performance of our learned ℓ_1 -minimization technique, CP-LEARNED, with the original method, CP, presented in [5]. We also compare our approach with other audio inpainting techniques such as A-SPAIN-MOD [4], A-SPAIN-LEARNED [4], A-SPAIN [6] and JANSSEN [2]. 'A-' stands for analysis variant. Note that A-SPAIN uses a frame-wise DFT dictionary

¹Note that this special structure is required to cope with the matrix arrangement/stacking according to (1).

Algorithm 1: Learned weighted ℓ_1 -minimization (CP-LEARNED)

Input: x, M_R, \hat{D}_{block} , \tilde{G}_A , \tilde{G}_S , w_L, σ , τ **Output:** \hat{v}

- 1 choose $\tau, \sigma > 0$ satisfying $\tau \sigma \|\tilde{G}_{S}\hat{D}_{block}^{H}\| \leq 1$ 2 choose primal variable $p^{(0)} \in \mathbb{C}^{L}$ and dual variable $q^{(0)} \in \mathbb{C}^{NM'}$ arbitrarily 3 set output variable $z^{(0)} = p^{(0)}$ 4 set iteration counter i = 0
- 5 repeat

10 until stopping criterion met 11 return $\hat{\mathbf{y}} = \text{proj}_{\mathcal{S}_{-}} \left(\mathbf{z}^{(i+1)} \right)$

(redundancy 4) whereas A-SPAIN-LEARNED and A-SPAIN-MOD use a Gabor dictionary (redundancy 4). JANSSEN uses auto-regressive model of order $p = \min(3H+2, w_g/3)$, where H denotes the number of missing samples within the current frame (window), w_g is the window length, and the number of iterations was set to 50. In order to allow for a valid and comparable assessment, we consider essentially the same setup as in [4] and [5].

As a means to compare the performance, we use the signalto-distortion ratio (SDR) [1] defined as,

$$\text{SDR}(\mathbf{x}_{\text{orig}}, \mathbf{x}_{\text{inp}}) = 10 \, \log_{10} \frac{\|\mathbf{x}_{\text{orig}}\|_2^2}{\|\mathbf{x}_{\text{orig}} - \mathbf{x}_{\text{inp}}\|_2^2} \quad \text{[dB]},$$

where x_{orig} and x_{inp} denote original and inpainted signal within the gaps, respectively. Note that higher SDR values imply better signal restoration.

For our experiments, we used a collection of ten music recordings chosen from the EBU SOAM dataset [25]. All chosen signals were sampled at 44.1 kHz and possessed different levels of sparsity with respect to the original Gabor dictionary. In each test instance, the input was a signal with 5 gaps at random positions. The lengths of these gaps ranged from 5 ms up to 85 ms. For fixed lengths, the results over all ten signals containing the 5 gaps were averaged.

Furthermore, we also calculated the PEMO-Q values [26], which utilize a human auditory system model to assess the quality of the restored signal. Thus, the PEMO-O criterion is closer to the subjective evaluation than the SDR. PEMO-Q consists of two quantities: *objective difference grade* (ODG) and perceptual similarity measure (PSM). For simplicity, we restrict to ODG in this paper. ODG can be interpreted as the degree of perceptual similarity of x_{orig} and x_{inp}. Basically, ODG quantifies the perceptual impact of audio artifacts in the restored signal and its value ranges from -4 (very annoying) up to 0 (imperceptible).



GABOR PARAMETERS

Throughout our experiments, we used the fast implementation of Gabor transforms available in the LTFAT toolbox [14], [16], and adopted its time-frequency conventions. The Gabor parameters used in our experiment are summarized in Table I. The dictionary learning procedure proposed, see [4] for algorithmic details, is run for at most $iter_{max} = 20$ iterations with off-diagonal parameter d = 1, and the remaining parameters were set to $\rho_{\text{start}} = 1$ and $\varepsilon = 2^{-20}$. Finally, all SPAIN variants used the iteration parameters s = t = 1 [4].

Fig. 1(a) shows the inpainting performance in terms of SDR of all tested algorithms. We can observe that CP-LEARNED performs significantly better than CP, i.e., the original weighted ℓ_1 -minimization method. Even more, it beats all other methods except A-SPAIN-LEARNED, which is slightly superior for gap lengths below 45 ms. Remarkably, CP-LEARNED even outperforms the otherwise best performing A-SPAIN-LEARNED algorithm for longer gaps (\geq 45 ms) although the observed improvements seem to be not very pronounced.

Fig. 1(b) shows the overall inpainting performance in terms of ODG. For smaller gap lengths (< 35 ms), the performance of CP-LEARNED is superior to CP and a bit worse compared to other competing methods, whereas for longer gap lengths (> 45 ms), CP-LEARNED's performance relative to the other methods improves, exceeding all the other methods except A-SPAIN-LEARNED for gap lengths above 65 ms.

IX. CONCLUSION

We presented an approach for audio inpainting that combines weighted ℓ_1 -minimization with a dictionary learning framework in order to mitigate the energy drop phenomenon within the gap in the restored signal. This approach learns a sparsifying dictionary from the Gabor coefficients obtained from the reliable parts near the gap ultimately resulting in increased sparsity of the analysis coefficients. Additionally, the learned dictionary is also used for computing the weights. Moreover, by exploiting the real-valuedness of audio signals, we only used M' = |M/2| + 1 coefficients instead of the full M coefficients for the inpainting task unlike in the original setting. This is noteworthy because utilizing half of the coefficients implies reduced computational effort. Our experimental results demonstrated that the proposed method, i.e, weighted ℓ_1 -minimization combined with dictionary learning, yields large performance gains and significantly outperforms its original counterpart for all gap lengths. Finally, comparison with A-SPAIN-LEARNED demonstrated that CP-LEARNED has comparable (in some instances even better, in other slightly



Fig. 1. Performance comparison of different audio inpainting algorithms in terms of (a) SDR and (b) ODG for varying gap lengths.

worse) performance as (than) the current state-of-the-art inpainting algorithm.

REFERENCES

- A. Adler, V. Emiya, M. Jafari, M. Elad, R. Gribonval, and M. Plumbley, "Audio inpainting," *IEEE Trans. Audio, Speech, and Language Process*ing, vol. 20, no. 3, pp. 922–932, March 2012.
- [2] A. Janssen, R. Veldhuis, and L. Vries, "Adaptive interpolation of discrete-time signals that can be modeled as autoregressive processes," *IEEE Trans. Ac., Sp., Sign. Proc.*, vol. 34, no. 2, pp. 317–330, 1986.
- [3] S. J. Godsill and P. J. Rayner, "A bayesian approach to the restoration of degraded audio signals," *IEEE Transactions on Speech and Audio Processing*, vol. 3, no. 4, pp. 267–278, 1995.
- [4] G. Tauböck, S. Rajbamshi, and P. Balazs, "Dictionary learning for sparse audio inpainting," *IEEE Journal of Selected Topics in Signal Processing*, vol. 15, no. 1, pp. 104–119, 2021.
- [5] O. Mokrý and P. Rajmic, "Audio inpainting: Revisited and reweighted," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2906–2918, 2020.
- [6] O. Mokrỳ, P. Záviška, P. Rajmic, and V. Veselỳ, "Introducing SPAIN (SParse Audio INpainter)," in 2019 27th European Signal Processing Conference (EUSIPCO). IEEE, 2019, pp. 1–5.
- [7] M. Lagrange, S. Marchand, and J.-B. Rault, "Long interpolation of audio signals using linear prediction in sinusoidal modeling," *Journal of the Audio Engineering Society*, vol. 53, no. 10, pp. 891–905, 2005.
- [8] N. Perraudin, N. Holighaus, P. Majdak, and P. Balazs, "Inpainting of long audio segments with similarity graphs," *IEEE/ACM Transactions* on Audio, Speech, and Language Processing, vol. 26, no. 6, pp. 1083– 1094, 2018.
- [9] A. Marafioti, N. Holighaus, P. Majdak, N. Perraudin *et al.*, "Audio inpainting of music by means of neural networks," in *Audio Engineering Society Convention 146*. Audio Engineering Society, 2019.
- [10] A. Marafioti, N. Perraudin, N. Holighaus, and P. Majdak, "A context encoder for audio inpainting," *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 27, no. 12, pp. 2362–2372, Dec. 2019.
- [11] A. Marafioti, P. Majdak, N. Holighaus, and N. Perraudin, "GACELA: A generative adversarial context encoder for long audio inpainting of music," *IEEE Journal of Selected Topics in Signal Processing*, vol. 15, no. 1, pp. 120–131, 2021.
- [12] A. Adler, V. Emiya, M. G. Jafari, M. Elad, R. Gribonval, and M. D. Plumbley, "Audio inpainting," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 3, pp. 922–932, 2011.

- [13] G. Tauböck, S. Rajbamshi, and P. Balazs, "Sparse audio inpainting: A dictionary learning technique to improve its performance," in *Audio Engineering Society Convention 149*, Oct 2020. [Online]. Available: http://www.aes.org/e-lib/browse.cfm?elib=20939
- [14] P. L. Søndergaard, B. Torrésani, and P. Balazs, "The linear time frequency analysis toolbox," *Int. J. Wavel., Multires. a. Inf. Process.*, vol. 10, no. 04, p. 1250032, 2012.
- [15] O. Christensen, Frames and bases: An introductory course. Springer Science & Business Media, 2008.
- [16] Z. Průša, P. L. Søndergaard, N. Holighaus, C. Wiesmeyr, and P. Balazs, "The Large Time-Frequency Analysis Toolbox 2.0," in *Sound, Music, and Motion*, ser. LNCS. Springer International Publishing, 2014, pp. 419–442.
- [17] P. Balazs, J.-P. Antoine, and A. Grybos, "Weighted and controlled frames: Mutual relationship and first numerical properties," *International Journal of Wavelets, Multiresolution and Information Processing*, vol. 8, no. 1, pp. 109–132, January 2010. [Online]. Available: http://arxiv.org/abs/math.FA/0611729
- [18] P. L. Combettes and J.-C. Pesquet, "Proximal splitting methods in signal processing," in *Fixed-point algorithms for inverse problems in science* and engineering. Springer, 2011, pp. 185–212.
- [19] L. Condat, "A generic proximal algorithm for convex optimization—application to total variation minimization," *IEEE Signal Processing Letters*, vol. 21, no. 8, pp. 985–989, 2014.
 [20] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for
- [20] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *Journal of mathematical imaging and vision*, vol. 40, no. 1, pp. 120–145, 2011.
- [21] T. Hastie, R. Tibshirani, and M. Wainwright, *Statistical learning with sparsity: the lasso and generalizations*. CRC press, 2015.
- [22] E. J. Candes, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted ℓ¹ minimization," *Journal of Fourier analysis and applications*, vol. 14, no. 5-6, pp. 877–905, 2008.
- [23] G. Tauböck and F. Hlawatsch, "Compressed sensing based estimation of doubly selective channels using a sparsity-optimized basis expansion," in 2008 16th European Signal Processing Conference, 2008, pp. 1–5.
- [24] G. Tauböck, F. Hlawatsch, D. Eiwen, and H. Rauhut, "Compressive estimation of doubly selective channels in multicarrier systems: Leakage effects and sparsity-enhancing processing," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 255–271, Apr. 2010.
- [25] "EBU SQAM CD: Sound quality assessment material recordings for subjective tests," 2008, https://tech.ebu.ch/publications/sqamcd.
- [26] R. Huber and B. Kollmeier, "PEMO-Q–A new method for objective audio quality assessment using a model of auditory perception," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 6, pp. 1902–1911, 2006.