

Audio Inpainting

Amir Adler, Valentin Emiya, Maria G. Jafari, Michael Elad, Rémi Gribonval and Mark D. Plumbley

Abstract—

We propose the **Audio Inpainting** framework that recovers portions of audio data distorted due to impairments such as impulsive noise, clipping, and packet loss. In this framework, the distorted data are treated as missing and their location is assumed to be known. The signal is decomposed into overlapping time-domain frames and the restoration problem is then formulated as an inverse problem per audio frame. Sparse representation modeling is employed per frame, and each inverse problem is solved using the **Orthogonal Matching Pursuit** algorithm together with a discrete cosine or a Gabor dictionary. The **Signal-to-Noise Ratio** performance of this algorithm is shown to be comparable or better than state-of-the-art methods when blocks of samples of variable durations are missing. We also demonstrate that the size of the block of missing samples, rather than the overall number of missing samples, is a crucial parameter for high quality signal restoration. We further introduce a constrained **Matching Pursuit** approach for the special case of audio declipping that exploits the sign pattern of clipped audio samples and their maximal absolute value, as well as allowing the user to specify the maximum amplitude of the signal. This approach is shown to outperform state-of-the-art and commercially available methods for audio declipping in terms of **Signal-to-Noise Ratio**.

Index Terms—Inpainting, clipping, sparse representation, matching pursuit.

I. INTRODUCTION

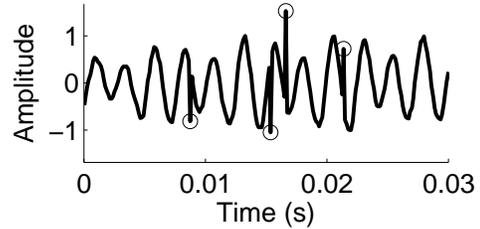
Speech and music signals are often subject to localized distortions, where the intervals of distorted samples are surrounded by undistorted samples. Examples include impulsive noise or clicks (see Fig. 1a), clipping (see Fig. 1b), CD scratches, packet loss in cordless phones or Voice over IP (VoIP) and more. In such situations, the distorted samples can be treated as missing. A restoration algorithm is employed to reconstruct the missing samples, in a similar way as for image inpainting (see Fig. 1c). However, in the audio field, such problems have been treated separately and depending on the context, they have been referred to as audio interpolation [2]–[6], extrapolation [3], [7], [8], imputation [9], [10], induction [11], (bandwidth) extension [12]–[15] or concealment [16], [17].

Substantial effort has been focused on the restoration of audio signals corrupted by clicks due to old recordings or scratched CDs (see Fig. 1a). In this problem, intervals of corrupted samples – from 20 μ s to 4 ms [4] – occur at random locations. Typical approaches employ autoregressive (AR)

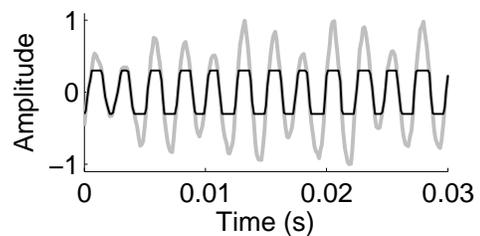
A. Adler and M. Elad are with the Computer Science Department, The Technion, Haifa 32000, Israel. V. Emiya and R. Gribonval are with INRIA, Centre Inria Rennes - Bretagne Atlantique, 35042 Rennes Cedex, France. M. G. Jafari and M. D. Plumbley are with Queen Mary University of London, Centre for Digital Music, Department of Electronic Engineering, London E1 4NS, U.K.

Part of this work has been presented at the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) in 2011 [1].

This work was supported by the EU Framework 7 FET-Open project FP7-ICT-225913- SMALL: Sparse Models, Algorithms and Learning for Large-Scale data.



(a) Speech signal corrupted by clicks (circles).



(b) Clipped version (black) of a speech signal (gray).



(c) The image inpainting problem: recovery of locally-hidden pixels.

Fig. 1. Examples of restoration problems related to inpainting.

modeling [2], [3], or Bayesian estimation to recover the corrupted samples [4]. Other methods utilize neural networks [18] or sinusoidal modeling [5], [8]. A related problem is automatic speech recognition in the presence of isolated noisy samples. This problem is treated in [10] with a compressive sensing approach in the spectrogram image domain, and by solving an l_1 regularized least squares problem.

Another important – though less often addressed – problem is audio clipping [6], [7], [19], which refers to the truncation of the waveform beyond a threshold when the maximum range in an acquisition system is exceeded (see Fig. 1b). The clipped samples are arranged in groups and their locations are determined by the amplitude of the signal (rather than being randomly spread). The declipping problem is particularly challenging for this reason and as the information carried by the highest-amplitude samples is completely absent.

Long intervals of samples may be lost during transmission

over cordless phones or in VoIP systems, where the problem is addressed using packet loss concealment algorithms [16], [17]. Missing intervals lengths are in the range of 5 ms to 60 ms, which are close to the typical duration for the pseudo-stationarity of speech signals. The low latency requirement in the VoIP case results in relatively simple algorithms; however, estimating missing packets in peer-to-peer repositories where several portions of a sound are not available from the network is a new application where higher quality reconstruction can be expected (as the latency requirement is less stringent).

Finally, instead of missing *samples*, the unreliable or missing audio data can be localized time-frequency regions [5], [9], [11], [14], [20], in classification applications like automatic speech recognition [9], [20] or source separation with time-frequency localized interference – the phrase “audio inpainting” has been used once in this specific case [11]. Bandwidth extension [12]–[15] is another important time-frequency-domain application, where high frequency content is estimated from the low frequency content in order to provide high quality audio.

In this paper, we present a unified framework for the restoration of distorted audio data, leveraging the concept of *Image Inpainting* [21]–[23]. In the proposed framework, termed *Audio Inpainting*, the distorted data is assumed missing and its location is assumed to be known a-priori. We further employ Sparse Representations (SR), which have been demonstrated to faithfully model audio signals [24], [25] and to address the image inpainting framework [22], [26], [27]. The proposed approach is directly based upon those prior works.

The contributions of this paper are four-fold:

- a) Audio inpainting is defined as an inverse problem, based upon the concept of image inpainting.
- b) A framework for audio inpainting in the time domain is proposed, based on sparse representations. It exploits two possible dictionaries (discrete cosine and Gabor) known to provide accurate sparse models for audio signals.
- c) The Orthogonal Matching Pursuit (OMP) algorithm for audio inpainting is adapted, in particular to deal with the properties of the Gabor dictionary.
- d) A constrained matching pursuit approach is applied to significantly enhance the performance for audio declipping problems.

This paper is organized as follows. In Section II, audio inpainting is formalized as an inverse problem. The proposed framework is introduced in Section III including the sparse models used for time-domain audio inpainting. The adaptation of the OMP algorithm for audio inpainting in the time domain and for audio declipping is presented in Section IV. Several experiments are proposed in Section V, while we discuss our findings and draw conclusions in Section VII.

II. AUDIO INPAINTING PROBLEM STATEMENT

We define audio inpainting as a general problem encountered in many applications: one observes a partial set of reliable audio data while the remaining unreliable data is either totally missing or highly degraded; the unreliable data is considered missing and it is estimated from the reliable data

portion. Optionally, some information about the missing data may be available, like in the clipping case explained below.

The general formulation of audio inpainting is given in Section II-A while several particular time-domain cases are detailed in Sections II-B and II-C.

A. Formulation of audio inpainting

We consider a vector $\mathbf{s} \in \mathbb{R}^L$ of audio data and an a-priori known partition $\{I^m, I^r\}$ of the support $I \triangleq \{1, 2, \dots, L\}$ of \mathbf{s} : $I^m \subset I$ and $I^r \triangleq I \setminus I^m$. We assume that the coefficients $\mathbf{s}(I^m)$ are either missing or masked by a severe distortion. Thus, the observed data $\mathbf{y} \in \mathbb{R}^L$ coincides with \mathbf{s} on I^r only. The audio inpainting problem is defined as the recovery of the coefficients $\mathbf{s}(I^m)$ based on the knowledge of:

- 1) the reliable data $\mathbf{y}^r \triangleq \mathbf{y}(I^r) = \mathbf{s}(I^r)$,
- 2) the partition $\{I^m, I^r\}$,
- 3) additional information about the observed signal,
- 4) and, optionally, information about the missing data (see *e.g.* in the case of clipping below).

In matrix form, the reliable data \mathbf{y}^r result from the linear model

$$\mathbf{y}^r = \mathbf{M}^r \mathbf{s}, \quad (1)$$

where $\mathbf{M}^r \in \{0, 1\}^{|I^r| \times L}$ is the so-called measurement matrix obtained from the $L \times L$ identity matrix \mathbf{I}_L by selecting the rows I^r associated with the reliable coefficients in \mathbf{s} . The notation $|I^r|$ above denotes the number of elements in I^r . In a similar way, the missing data to be recovered are $\mathbf{s}(I^m) = \mathbf{M}^m \mathbf{s}$, where $\mathbf{M}^m \in \{0, 1\}^{|I^m| \times L}$ consists of the rows I^m in \mathbf{I}_L .

In the general audio inpainting framework, audio data can be either samples in waveforms or coefficients in transforms like time-frequency representations. The problem formulation above can be used for multi-dimensional signals like multichannel waveforms or time-frequency coefficients, by simply considering the equivalent vectors (as images that are commonly vectorized [23]).

In the rest of this paper, we only consider the inpainting of missing samples in a single-channel waveform. The multi-dimensional case is discussed in the conclusion (see Sec. VII).

B. Inpainting samples distorted by impulsive noise

In the particular case of a signal corrupted by impulsive noise such as clicks (see Fig. 1a), I^m is a set of integers between 1 and L and must be estimated in a preliminary stage. One often considers that the distorted samples are corrupted by a Gaussian noise \mathbf{n} with high variance. Hence, the complete observed signal includes both the reliable samples \mathbf{y}^r and distorted ones \mathbf{y}^m :

$$\begin{cases} \mathbf{y}^r &= \mathbf{M}^r \mathbf{s} \\ \mathbf{y}^m &= \mathbf{M}^m \mathbf{s} + \mathbf{n}, \end{cases} \quad (2)$$

where the samples $\mathbf{M}^m \mathbf{s}$ in \mathbf{y}^m are masked by \mathbf{n} so that they are considered as unknown.

C. Inpainting intervals of missing samples

In the case where intervals of samples are missing, due to packet loss during transmission or to masking by audible interferences, I^m is composed of groups of consecutive integers: the samples $\mathbf{s}(I^m)$ are totally missing and one only observes $\mathbf{y}^r = \mathbf{M}^r \mathbf{s}$.

In the case of clipped signals, the samples to be estimated are also arranged in intervals of consecutive samples, as depicted in Fig. 1b. Their locations depend on the amplitude of the signal, such that

$$I^m \triangleq \{n \mid 1 \leq n \leq L, |\mathbf{s}(n)| \geq \theta_{\text{clip}}\}, \quad (3)$$

where θ_{clip} is the clipping level. One observes both the unclipped, reliable samples \mathbf{y}^r and the clipped, masked samples \mathbf{y}^m

$$\begin{cases} \mathbf{y}^r &= \mathbf{M}^r \mathbf{y} = \mathbf{M}^r \mathbf{s} \\ \mathbf{y}^m &= \mathbf{M}^m \mathbf{y} = \mathbf{M}^m \text{sign}(\mathbf{s}) \theta_{\text{clip}}, \end{cases} \quad (4)$$

where $\text{sign}(\cdot)$ is the element-wise sign function. As presented in the next sections, the information provided by \mathbf{y}^m , even though very crude – a sign (per sample) and the clipping level –, still substantially enhances the estimation performance.

III. TIME-DOMAIN FRAMEWORK AND MODELS

The proposed framework focuses on time-domain audio inpainting. It relies on a frame-based processing, as detailed in Section III-A and on the sparse representations modeling of audio signals, as presented in Section III-B. Two dictionaries used in this modeling are introduced in Section III-C.

A. Frame-based processing and reconstruction

As in many audio processing tasks, the signal is locally processed:

- by segmenting it into frames,
- by independently inpainting each frame,
- and by synthesizing the full restored signal using the overlap-add (OLA) method [28].

We decompose the signal into overlapping frames indexed by i , starting at time t_i and weighted by an analysis window \mathbf{w}_a with length N . By straightforwardly adapting to the local frames the problem statement defined for the full signal in Section II, the reliable samples in frame i can be written as

$$\mathbf{y}_i^r = \mathbf{M}_i^r \mathbf{s}_i \quad (5)$$

where \mathbf{M}_i^r is the measurement matrix of the i -th frame obtained from \mathbf{M}^r and $\mathbf{s}_i(t) \triangleq \mathbf{s}(t + t_i) \mathbf{w}_a(t)$ is the windowed frame defined for $0 \leq t \leq N - 1$. We also define the supports I_i^r and I_i^m of the reliable samples and of the missing or masked samples, respectively. Once the estimation $\hat{\mathbf{s}}_i$ of \mathbf{s}_i by some inpainting algorithm is achieved, the reconstruction of the full signal is obtained as

$$\hat{\mathbf{s}}(t) \triangleq \sum_i \mathbf{w}_s(t - t_i) \hat{\mathbf{s}}_i(t - t_i) \quad (6)$$

where \mathbf{w}_s is the synthesis window such that $\sum_i \mathbf{w}_s(t - t_i) \mathbf{w}_a(t - t_i) = 1, \forall t$. In the proposed approaches, we utilized 64ms-frames with 75% overlap, a rectangular window for \mathbf{w}_a and a sine window for \mathbf{w}_s .

B. Sparse Representations modeling of audio frames

In the Sparse Representations (SR) modeling framework [23], it is assumed that each frame is well approximated by a sparse linear combination of the columns of a (possibly overcomplete) dictionary:

$$\mathbf{s}_i \approx \mathbf{D} \mathbf{x}_i, \quad (7)$$

where $\mathbf{D} \in \mathbb{R}^{N \times K_D}$ is the dictionary, $N \leq K_D$ and $\mathbf{x}_i \in \mathbb{R}^{K_D \times 1}$ is the representation vector of the i -th frame. \mathbf{x}_i is assumed to be sparse, *i.e.* to have few non-zero coefficients compared to N . As a consequence, we can also utilize the SR model for the observed reliable samples in each frame

$$\mathbf{y}_i^r \triangleq \mathbf{M}_i^r \mathbf{s}_i \approx \mathbf{M}_i^r \mathbf{D} \mathbf{x}_i. \quad (8)$$

We propose to recover the unknown samples $\mathbf{s}_i(I_i^m)$ by estimating as $\hat{\mathbf{x}}_i$ the (sparse) representation vector of each frame, given only the clean observed samples (8) and limited side information (for the clipping case)

$$\hat{\mathbf{s}}_i(I_i^m) = \mathbf{M}_i^m \mathbf{D} \hat{\mathbf{x}}_i. \quad (9)$$

This formulation including the notion of sparsity was first introduced for image inpainting [22] with a global treatment with global transforms. Then, efforts were dedicated to work on local patches – similar to audio frames – and to introduce a learned dictionary to improve the inpainting results [26]; they have been improved [27] by better handling the reconstruction problem by learning the dictionary directly from the corrupted image, thereby getting a dictionary that is tightly related to the specific image content.

C. Dictionaries

We propose two options to choose a dictionary \mathbf{D} in which audio signals are sparse: the Discrete Cosine Transform dictionary, and a Gabor dictionary. Both are widely used for sparse models of audio signals [24], [25], [29]. Other fixed dictionaries such as multiscale DCT [30], or a learned dictionary [26] specific to particular inpainting tasks may also be interesting options.

1) *Discrete Cosine Transform (DCT) dictionary*: The first option consists in a windowed Discrete Cosine Transform (DCT) overcomplete dictionary $\mathbf{D}^c = [\mathbf{d}_0^c, \dots, \mathbf{d}_{K_c-1}^c]$, atom j being defined for $0 \leq j \leq K_c - 1$ and $0 \leq t \leq N - 1$ as

$$\mathbf{d}_j^c(t) \triangleq \mathbf{w}_d(t) \cos\left(\frac{\pi}{K_c} \left(t + \frac{1}{2}\right) \left(j + \frac{1}{2}\right)\right) \quad (10)$$

where K_c is the size of the DCT dictionary – *i.e.* the number of discrete frequencies – and \mathbf{w}_d is a weighting window set by the user. This choice is motivated by the wide use of windowed DCT atoms for sparse representation of audio signals [25]. However, DCT atoms can be generalized as Gabor atoms by introducing an initial phase in order to increase the potential fit to the signal structure, resulting in a sparser approximation.

2) *Gabor dictionary*: The second option aims at sparsely modeling arbitrary-phase sinusoidal components by using a Gabor dictionary $\mathbf{D}^g = \{\mathbf{d}_{(j,\varphi)}^g\}_{(j,\varphi) \in \Gamma}$ in which the atoms are indexed by a continuous set $\Gamma \triangleq \llbracket 0, K_g - 1 \rrbracket \times [0, 2\pi[$ and are defined as

$$\mathbf{d}_{j,\varphi}^g(t) \triangleq \mathbf{w}_d(t) \cos\left(\frac{\pi}{K_g}\left(t + \frac{1}{2}\right)\left(j + \frac{1}{2}\right) + \varphi\right), \quad (11)$$

where K_g is the size of the Gabor dictionary.

Note that in the current case of a continuously-indexed dictionary, eq. (7), (8) and (9) are still valid if we define

$$\mathbf{D}^g \mathbf{x}_i = \sum_{\substack{(j,\varphi) \in \Gamma \\ \mathbf{x}_i(j,\varphi) \neq 0}} \mathbf{d}_{j,\varphi}^g \mathbf{x}_i(j,\varphi) \quad (12)$$

where $\mathbf{x}_i = \{\mathbf{x}_i(j,\varphi)\}_{(j,\varphi) \in \Gamma}$. Indeed, eq. (12) is a finite sum since only a few coefficients in the sparse representation vector \mathbf{x}_i are non-zero. The algorithmic aspects of this decomposition will be addressed in Sections IV-B and IV-C.

IV. AUDIO INPAINTING ALGORITHMS BASED ON ORTHOGONAL MATCHING PURSUIT

For a given dictionary \mathbf{D} , we use the Orthogonal Matching Pursuit algorithm to perform the inpainting of an audio frame, as presented in Section IV-A. Some dictionary-dependent algorithmic stages are then detailed in Section IV-B and IV-C. An extension of the algorithm specific to declipping is finally detailed in Section IV-D.

A. Orthogonal Matching Pursuit (OMP) algorithm for inpainting

The approach emerges from the following optimization problem

$$\hat{\mathbf{x}}_i = \arg \min_{\mathbf{x}} \|\mathbf{x}\|_0 \quad \text{s.t.} \quad \|\mathbf{y}_i^r - \mathbf{M}_i^r \mathbf{D} \mathbf{x}\|_2^2 \leq \epsilon_i. \quad (13)$$

for a given approximation error threshold ϵ_i .

The l_0 pseudo-norm $\|\mathbf{x}\|_0$ counts the non-zeros components of the vector \mathbf{x} , leading to an NP-hard problem [31], [32]. Therefore, a direct solution of (13) is infeasible. An approximate solution is given by applying the Orthogonal Matching Pursuit (OMP) algorithm [24], [33], which successively approximates the sparsest solution. The inpainting OMP algorithm [23], detailed in Table I, is a slightly modified version of the classical OMP algorithm in the sense that all dictionary columns $\tilde{\mathbf{d}}_j$ are internally normalized to unit norm, using the diagonal matrix \mathbf{W} , due to the availability of only the clean samples. The algorithm stops iterating as soon as either the residual energy drops below the threshold ϵ_i^{OMP} or the maximum sparsity level K^{OMP} is exceeded.

B. Atom selection

When using the DCT dictionary, the algorithmic stage for the atom selection (14) at each iteration is well known: it consists of explicitly computing the correlation mentioned in eq. (14) of Table I, or of using a fast transform.

TABLE I
OMP INPAINTING ALGORITHM

$\mathbf{y}_i^r, \mathbf{M}_i^r, \mathbf{D} = \{\mathbf{d}_j\}_{j \in \Gamma}, K^{\text{OMP}}, \epsilon_i^{\text{OMP}}$	
Initialization:	
<ul style="list-style-type: none"> Dictionary $\tilde{\mathbf{D}} = \{\tilde{\mathbf{d}}_j\}_{j \in \Gamma} = \mathbf{M}_i^r \times \mathbf{D} \times \mathbf{W}$, where $\mathbf{W}_{jj'} = 0$ for $j \neq j'$ and $\mathbf{W}_{jj} = \ \mathbf{M}_i^r \mathbf{d}_j\ _2^{-1}$. Iteration counter $k = 0$ Support set $\Omega_0 = \emptyset$ Residual $\mathbf{r}_0 = \mathbf{y}_i^r$ 	
Sparse support selection and coefficients estimation:	
Repeat until $k = K^{\text{OMP}}$ or $\ \mathbf{r}_k\ _2 < \epsilon_i^{\text{OMP}}$	
<ul style="list-style-type: none"> Increment iteration counter $k = k + 1$ Select atom: find $j = \arg \max_{j \in \Gamma} \langle \mathbf{r}_{k-1}, \tilde{\mathbf{d}}_j \rangle \quad (14)$ Update Support $\Omega_k = \Omega_{k-1} \cup j$ Update current solution $\mathbf{x}_k = \arg \min_{\mathbf{u}} \ \mathbf{y}_i^r - \tilde{\mathbf{D}}_{\Omega_k} \mathbf{u}\ _2 \quad (15)$ Update Residual $\mathbf{r}_k = \mathbf{y}_i^r - \tilde{\mathbf{D}}_{\Omega_k} \mathbf{x}_k$ 	
Output: $\hat{\mathbf{x}}_i = \mathbf{W} \mathbf{x}_k$	

However, the atom selection (14) needs explaining in the case of the Gabor dictionary in order to deal with the continuous indexing. Without any approximation, the decomposition with continuously-indexed atoms \mathbf{d}^g can be expressed using pairs of atoms in a discrete dictionary with K_g frequency bins. Pairs of atoms can be either conjugate complex exponentials [24], [34], or pairs of cosine and sine at the same frequency and with a zero phase [29]. In order to use this latter option, we introduce sine atoms \mathbf{d}_j^s as

$$\mathbf{d}_j^s(t) \triangleq \mathbf{w}_d(t) \sin\left(\frac{\pi}{K_g}\left(t + \frac{1}{2}\right)\left(j + \frac{1}{2}\right)\right) \quad (16)$$

and we define the unit-norm version $\tilde{\mathbf{d}}_j^c$ and $\tilde{\mathbf{d}}_j^s$ of the atoms \mathbf{d}_j^c and \mathbf{d}_j^s , respectively, as described in Table I.

At each iteration k , selecting the best correlated Gabor atom \mathbf{d}_j^g (eq. (14)) is then equivalent to picking the pair $(\tilde{\mathbf{d}}_j^c, \tilde{\mathbf{d}}_j^s)$ such that

$$j = \arg \min_{j \in \llbracket 1, K_D/2 \rrbracket} \left\| \mathbf{r}_{k-1} - \tilde{\mathbf{d}}_j^c \hat{x}_j^c - \tilde{\mathbf{d}}_j^s \hat{x}_j^s \right\|_2^2 \quad (17)$$

where

$$\begin{cases} \hat{x}_j^c &= \frac{\langle \tilde{\mathbf{d}}_j^c, \mathbf{r}_{k-1} \rangle - \langle \tilde{\mathbf{d}}_j^c, \tilde{\mathbf{d}}_j^s \rangle \langle \tilde{\mathbf{d}}_j^s, \mathbf{r}_{k-1} \rangle}{1 - \langle \tilde{\mathbf{d}}_j^c, \tilde{\mathbf{d}}_j^s \rangle^2} \\ \hat{x}_j^s &= \frac{\langle \tilde{\mathbf{d}}_j^s, \mathbf{r}_{k-1} \rangle - \langle \tilde{\mathbf{d}}_j^s, \tilde{\mathbf{d}}_j^c \rangle \langle \tilde{\mathbf{d}}_j^c, \mathbf{r}_{k-1} \rangle}{1 - \langle \tilde{\mathbf{d}}_j^s, \tilde{\mathbf{d}}_j^c \rangle^2} \end{cases} \quad (18)$$

This particular selection stage has been proposed in [34, Appendix II] for conjugate Gabor chirp atoms and the use of blocks of coherent atoms in MP and OMP has been further studied in [35]. In the restricted case where atoms in a candidate pair are uncorrelated (*i.e.* $\langle \tilde{\mathbf{d}}_j^c, \tilde{\mathbf{d}}_j^s \rangle = 0$), eq. (17) can be simplified as $j = \arg \max_{j \in \llbracket 1, K_D/2 \rrbracket} \left\langle \tilde{\mathbf{d}}_j^c, \mathbf{r}_{k-1} \right\rangle^2 + \left\langle \tilde{\mathbf{d}}_j^s, \mathbf{r}_{k-1} \right\rangle^2$ and the resulting algorithm is equivalent to the

existing Modified Matching Pursuit [36] and Block OMP algorithms [37].

C. Solution update

When using the DCT dictionary, the solution update (15) performed at each iteration usually consists of a least-square projection.

In the case of the Gabor dictionary, once the best atom has been added to the set of atoms selected in previous iterations, the update of the current solution (15) can be performed by a least-square projection using the selected Gabor atoms $\{\mathbf{d}_{j,\varphi_j}^g\}_{j \in \Omega_k}$, their phases φ_j being fixed in the atom-selection stage.

However, this update can be improved by using the equivalent cosine and sine atoms $\{\mathbf{d}_j^c, \mathbf{d}_j^s\}_{j \in \Omega_k}$ in the least-square projection: not only amplitudes but also phases are thus updated at each iteration, leading to a better approximation of the signal. As far as we know, such an implementation of OMP with a Gabor dictionary has never been proposed before.

D. Algorithmic enhancements for inpainting clipped signals

1) *The ‘min’ declipping constraint:* Inpainting clipped signals can be performed with the algorithm presented in Section IV-A, by treating the clipped samples as completely unknown. However, extra information inherent to this problem can be integrated as additional constraints into equations (13). Constrained optimization approaches were also utilized in the case of l_1 -minimization for image desaturation [38] and of audio declipping based on a band-limited assumption [7].

Let θ_{clip} be the clipping level (which can be easily estimated as the maximum absolute value among the observed samples) and $\mathbf{M}_i^{\text{m}+}$ (resp. $\mathbf{M}_i^{\text{m}-}$) be the matrix such that $\mathbf{M}_i^{\text{m}+}\mathbf{s}_i$ (resp. $\mathbf{M}_i^{\text{m}-}\mathbf{s}_i$) is the vector of positive (resp. negative) clipped samples. The matrices $\mathbf{M}_i^{\text{m}+}$ and $\mathbf{M}_i^{\text{m}-}$ are known from the location and the sign of the clipped samples. The missing samples should satisfy the ‘min’ constraints

$$\mathbf{M}_i^{\text{m}+}\mathbf{s}_i \geq \theta_{\text{clip}} \text{ and } \mathbf{M}_i^{\text{m}-}\mathbf{s}_i \leq -\theta_{\text{clip}}. \quad (19)$$

2) *The ‘max’ declipping constraint:* The set of ‘min’ constraints can be further augmented by ‘max’ constraints, introducing an upper limit on the absolute value of the recovered samples θ_{max} , as follows

$$\mathbf{M}_i^{\text{m}+}\mathbf{s}_i \leq \theta_{\text{max}} \text{ and } \mathbf{M}_i^{\text{m}-}\mathbf{s}_i \geq -\theta_{\text{max}}. \quad (20)$$

The upper limit θ_{max} is an optional parameter that cannot be estimated automatically in a straightforward way but may be adjusted manually by the user.

3) *The ‘minmax’ constrained SR problem:* Using both sets of constraints, the ‘minmax’ declipping version of the l_0 -norm minimization problem (13) is given by

$$\hat{\mathbf{x}}_i = \arg \min_{\mathbf{x}} \|\mathbf{x}\|_0 \text{ s.t. } \begin{cases} \|\mathbf{y}_i^r - \mathbf{M}_i^r \mathbf{D} \mathbf{x}\|_2^2 & \leq \epsilon_i \\ \theta_{\text{max}} \geq \mathbf{M}_i^{\text{m}+} \mathbf{D} \mathbf{x} & \geq \theta_{\text{clip}} \\ -\theta_{\text{max}} \leq \mathbf{M}_i^{\text{m}-} \mathbf{D} \mathbf{x} & \leq -\theta_{\text{clip}} \end{cases} \quad (21)$$

where θ_{max} can be set to $+\infty$ if one does not want to use the ‘max’ constraint.

TABLE II

SUMMARY OF THE PROPOSED ALGORITHMS: EACH ROW INDICATES THE ALGORITHM USAGE (GENERAL INPAINTING OR DECLIPPING), DEPENDING ON POSSIBLE ADDITIONAL CONSTRAINTS, WHILE DICTIONARIES VARY ACROSS COLUMNS. ALGORITHM NOMENCLATURE APPEARS WITHIN QUOTES IN EACH CELL.

Additional specification	DCT	Gabor
Inpainting	‘OMP-C’ [Table I]	‘OMP-G’ [Table I]
Min-constraint declipping	‘OMP-C-min’ [Table I + eq. (22)]	‘OMP-G-min’ [Table I + eq. (22)]
Minmax-constraint declipping	‘OMP-C-minmax’ [Table I + eq. (23)]	‘OMP-G-minmax’ [Table I + eq. (23)]

4) *OMP declipping algorithm:* We propose approximate solutions by incorporating the constraints (19) and (20) into the final solution update stage of the OMP Inpainting algorithm. In other words, the OMP Inpainting algorithm presented in Table I is applied, in order to *select* the sparse support. Once the support Ω_k is selected, the sparse representation coefficients are *re-estimated* by solving the following constrained optimization problem:

$$\mathbf{x}_k = \arg \min_{\mathbf{u}} \|\mathbf{y}_i^r - \tilde{\mathbf{D}}_{\Omega_k} \mathbf{u}\|_2 \text{ s.t. } \begin{cases} \mathbf{M}_i^{\text{m}+} \mathbf{D} \mathbf{W} \mathbf{u} \geq \hat{\theta}_{\text{clip}} \\ \mathbf{M}_i^{\text{m}-} \mathbf{D} \mathbf{W} \mathbf{u} \leq -\hat{\theta}_{\text{clip}} \end{cases} \quad (22)$$

in the case of the ‘min’ constraint, or

$$\mathbf{x}_k = \arg \min_{\mathbf{u}} \|\mathbf{y}_i^r - \tilde{\mathbf{D}}_{\Omega_k} \mathbf{u}\|_2 \text{ s.t. } \begin{cases} \hat{\theta}_{\text{max}} \geq \mathbf{M}_i^{\text{m}+} \mathbf{D} \mathbf{W} \mathbf{u} \geq \hat{\theta}_{\text{clip}} \\ -\hat{\theta}_{\text{max}} \leq \mathbf{M}_i^{\text{m}-} \mathbf{D} \mathbf{W} \mathbf{u} \leq -\hat{\theta}_{\text{clip}} \end{cases} \quad (23)$$

for the case of the ‘minmax’ constraint. The constraints are linear, thus standard convex optimization solvers can be employed.

In theory, the solution of the constrained problem may not exist. We observed that this occurs very seldom in practice. Whenever no solution exists, the frame is restored using the unconstrained minimization $\mathbf{x}_k = \arg \min_{\mathbf{u}} \|\mathbf{y}_i^r - \tilde{\mathbf{D}}_{\Omega_k} \mathbf{u}\|_2$.

V. EXPERIMENTAL RESULTS

A summary of all versions of the algorithm presented in this paper is given in Table II. This section reports the major trends through different experiments. The performance measures are introduced in Section V-A. The test material and parameter settings are presented in Sections V-B and V-C. The global performance of all the proposed inpainting algorithms and a more detailed inpainting experiment are presented in Section V-D. Section V-E finally focuses on the case of clipping.

A. Performance measures

The performance can be assessed by the signal-to-noise ratio (SNR) computed on the full signals, defined by

$$\text{SNR}_{\text{full}}(\mathbf{s}, \hat{\mathbf{s}}) \triangleq 10 \log \frac{\|\mathbf{s}\|_2^2}{\|\mathbf{s} - \hat{\mathbf{s}}\|_2^2}. \quad (24)$$

In the reported experiments, the estimate is built by copying the observed signal and by only substituting the unreliable samples.

While SNR_{full} gives an overview of the global quality of the restored signal, it can be decomposed as

$$\text{SNR}_{\text{full}}(\mathbf{s}, \hat{\mathbf{s}}) = \text{SNR}_m(\mathbf{s}, \hat{\mathbf{s}}) + 10 \log \frac{\|\mathbf{s}\|_2^2}{\|\mathbf{s}(I^m)\|_2^2} \quad (25)$$

where

$$\text{SNR}_m(\mathbf{s}, \hat{\mathbf{s}}) \triangleq 10 \log \frac{\|\mathbf{s}(I^m)\|_2^2}{\|\mathbf{s}(I^m) - \hat{\mathbf{s}}(I^m)\|_2^2}. \quad (26)$$

SNR_m reflects the reconstruction performance per estimated sample and differs from SNR_{full} by an offset that does not depend on the inpainting algorithm. Indeed, the second term in Eq. (25) is a bias that reflects the degradation rate only. Thus, SNR_m will be preferred to show some detailed performance, without the influence of this bias, while SNR_{full} will be used to assess the global restoration quality.

Note that in the context of a perceptually-motivated evaluation of the results, SNR measures may be replaced by scores from listening tests or by objective measures. No subjective test specific to audio inpainting has been performed so far. Besides, existing objective measures for audio quality assessment are dedicated to applications such as coding [39], [40], source separation [41] or watermarking. Their use in the context of audio inpainting should be validated by listening tests. Thus, working on the perceptual evaluation of audio inpainting is an important future direction to consider.

B. The collection of tested signals

The experiments are conducted using three datasets:

- **Music@16kHz**: a set of music signals sampled at 16kHz, this sampling rate being a good trade-off between audio quality and computational requirements.
- **Speech@16kHz**: a set of speech signals sampled at 16kHz, *i.e.* high quality speech for which results can be compared to the previous case of music signals.
- **Speech@8kHz**: a set of speech signals sampled at 8kHz, representing phone-quality speech; this dataset was obtained by downsampling the previous 16kHz speech dataset.

Each dataset consists of ten 5-seconds signals from the 2008's Signal Separation Evaluation Campaign [42] and is freely available online¹. They include a large diversity of audio mixtures and isolated sources: male and female speech from different speakers, singing voice, pitched and percussive musical instruments.

In order to have comparable degradations among all signals in the clipping experiments (Section V-E), each original signal is normalized so that the maximum amplitude is 1.

¹<http://www.irisa.fr/metiss/vemiya/inpainting/> (this url may be changed to a more stable one by the submission of the final version of this paper)

C. Parameter settings

A specific training dataset was used to tune the parameters of the inpainting algorithms manually and without fine adjustment. The values of the tuned parameters are shown in Table III.

TABLE III
PARAMETER SETTINGS

Parameter	Value
Frame length	64 ms (<i>i.e.</i> $N \triangleq 512$ at 8kHz, $N \triangleq 1024$ at 16kHz)
Frame Overlap	75%
Analysis window \mathbf{w}_a	rectangular
Synthesis window \mathbf{w}_s	sine
Dictionary size	$K_c = 2N$ (DCT), $K_g = N$ (Gabor)
Atom weighting window \mathbf{w}_d	rectangular ($\mathbf{w}_d = \mathbf{w}_a$)
ϵ_i^{OMP}	$\epsilon \times I_i^r $ where $\epsilon \triangleq 10^{-6}$ is a fixed parameter and $ I_i^r $ is the number of reliable samples in the i th frame
K^{OMP}	$\frac{N}{4}$
$\hat{\theta}_{\text{clip}}$	$\ \mathbf{y}\ _{\infty}$
$\hat{\theta}_{\text{max}}$	$4\hat{\theta}_{\text{clip}}$

D. Inpainting experiments

1) Global effect of the duration of missing intervals:

The inpainting performance of the proposed algorithms was evaluated for variable durations of missing intervals of samples. Each experiment tested the performance with the entire collection of signals, for a fixed missing interval duration that repeated periodically every 100 ms. The missing intervals durations were in the range of a fraction of 1ms (corresponding to impulsive noise or clicks distortions) and up to 10 ms (corresponding to packet loss scenarios).

For comparison, we used the method by Janssen [2] based on linear prediction and a reconstruction method based on spline interpolation – the Matlab ‘interp1’ function. These methods are representatives of the two main families of state-of-the-art methods for interpolation of audio data and are able to handle multiple blocks of consecutive missing samples. In Janssen’s method, the autoregressive model order is set to $3N_{\text{miss}} + 2$, where N_{miss} is the number of missing samples in the current frame, as recommended by the authors. Note that the methods used for comparison are applied frame-wise, by embedding them in the frame-based processing scheme described in Section III-A.

The results are presented in Fig. 2. On average, the OMP algorithm with the Gabor dictionary provides an advantage of 1-2dB compared to the OMP with DCT dictionary. For Music@16kHz and Speech@8kHz these algorithms also outperform Janssen’s approach for short missing intervals of durations up to 1 ms. For durations above 1ms Janssen’s approach provides an advantage of 1-3dB. For Speech@16kHz, Janssen’s method performs better than the proposed ones, linear prediction being particularly well-adapted for speech. However, using more information with the proposed method

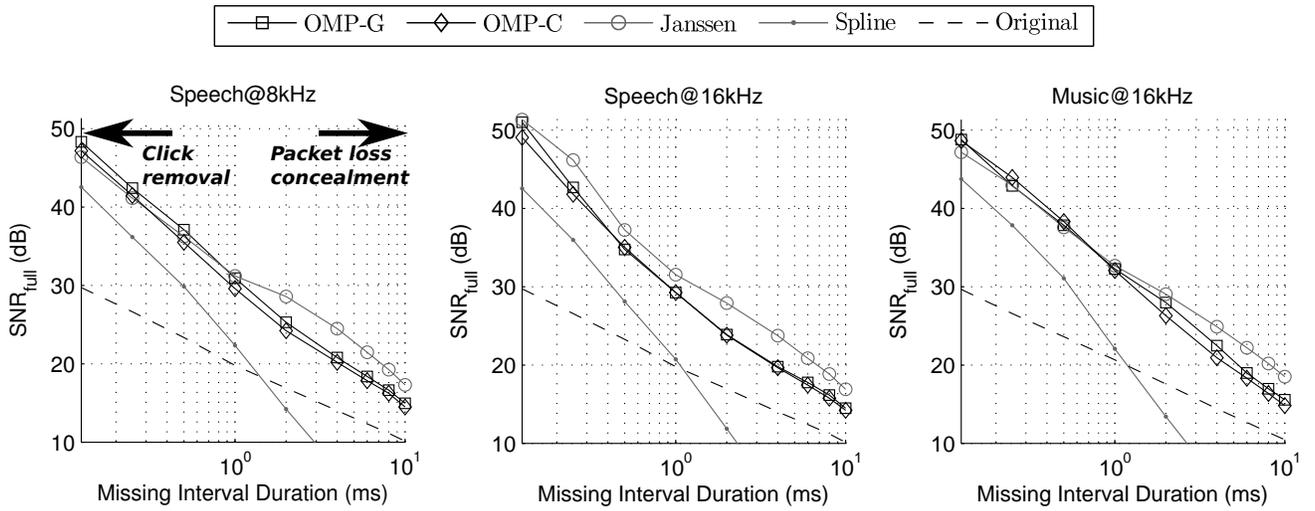


Fig. 2. Performance of inpainting algorithms as a function of the duration of missing intervals for each dataset (subfigures). The missing intervals were generated periodically every 100ms (a total of 50 equal duration missing intervals per signal).

can enhance the performance, as will be shown in the de-clipping experiment in Section V-E. The spline interpolation approach provides substantially worse results for all cases.

In terms of computational cost, the proposed approach is much slower than the methods used for comparison. The computational cost may be reduced by using fast transform when handling a DCT of a Gabor dictionary, or by using algorithms faster than OMP.

2) *Fine effect of the ‘topology’ of the missing samples:* The recovery or approximation performance of sparse approaches is often assessed as a function of the sparsity degree and the number of observations in the case of a random measurement matrix [43]. However, the latter assumption, recently highlighted in the compressed sensing framework [43], [44], does not hold in many audio inpainting applications: as introduced in the previous experiment, one must deal with blocks of consecutive missing samples. In this experiment, we investigate this assumption and assess empirical performance as a function of the randomness² and the consecutiveness of the location of the missing samples. The maximum randomness is achieved when the missing samples are isolated and distributed according to e.g. a uniform law. Conversely, when they are grouped, the missing samples in a given block are not randomly located, even if the blocks themselves may be randomly located. Hence the question: for a fixed number of missing samples, to which extent is the inpainting of few large blocks a more difficult problem than the inpainting of many small blocks (or isolated samples)?

The experimental protocol consists in the following steps:

- Choose a set of frames³
- Fix the number of missing samples N_{miss} ;
- For each $(a, b) \in \mathbb{N}^2$ such that $a \times b = N_{\text{miss}}$;
 - For each frame in the set,

²Here, randomness should be understood in a very common sense.

³The frames are randomly chosen in the datasets, only ensuring that the energy in the selected frames is high enough – i.e. down to -10 dB below the frame with maximum energy – to avoid silences.

- * Randomly generate a holes with length b ;
- * Recover the samples inside the holes from the samples outside the holes;
- * Compute SNR_m ;

– Average the SNR_m values w.r.t. all frames.

We use the OMP-G algorithm to recover the samples. The set of values for the number of missing samples N_{miss} is $\{12, 36, 60, 120, 180, 240\}$, which allow a large number of factorizations of the form $a \times b = N_{\text{miss}}$, $(a, b) \in \mathbb{N}^2$. For each test point (N_{miss}, a, b) , one thousand 64 ms frames from the 8kHz speech dataset are processed.

Results are presented in the left plot of Fig. 3. When the hole size is 1 – i.e. samples are randomly and uniformly distributed –, the recovery performance is very high with SNR_m values above 35 dB, including the case where the number of missing samples is high (e.g. $N_{\text{miss}} = 120$). When holes get larger, the performance significantly decreases: thus, inpainting a single 12-length hole happens to be a much more difficult problem than inpainting a frame with 120 isolated missing samples. Yet, SNR enhancement is still obtained for the largest holes (e.g. $\text{SNR}_m \approx 5\text{dB}$ at $N_{\text{miss}} = 100$).

The sensitivity of OMP-G to the stopping criteria was measured thanks to an oracle algorithm. It consists in applying the OMP-G algorithm with different values of $(K^{\text{OMP}}, \epsilon)$, and in selecting the set of parameters that gives the best performance for each frame independently. The tested parameters were $(K^{\text{OMP}}, \epsilon) \in \left\{ \frac{N}{2^{1.5}}, \frac{N}{2^2}, \dots, \frac{N}{2^{4.5}} \right\} \times \{10^{-10}, 10^{-9}, \dots, 10^{-1}\}$. Results are presented in the center plot of Fig. 3 and the difference between the oracle and blind systems is shown in the right plot. One can see that fixing parameters is a convenient, simple approximation that leads to suboptimal but satisfying performance compared to the oracle case. However, adapting the parameters to the frame to process may be worth studying: the difference between oracle and blind performance ranges from 4 to 10 dB in most of cases, showing a significant potential for improvement.

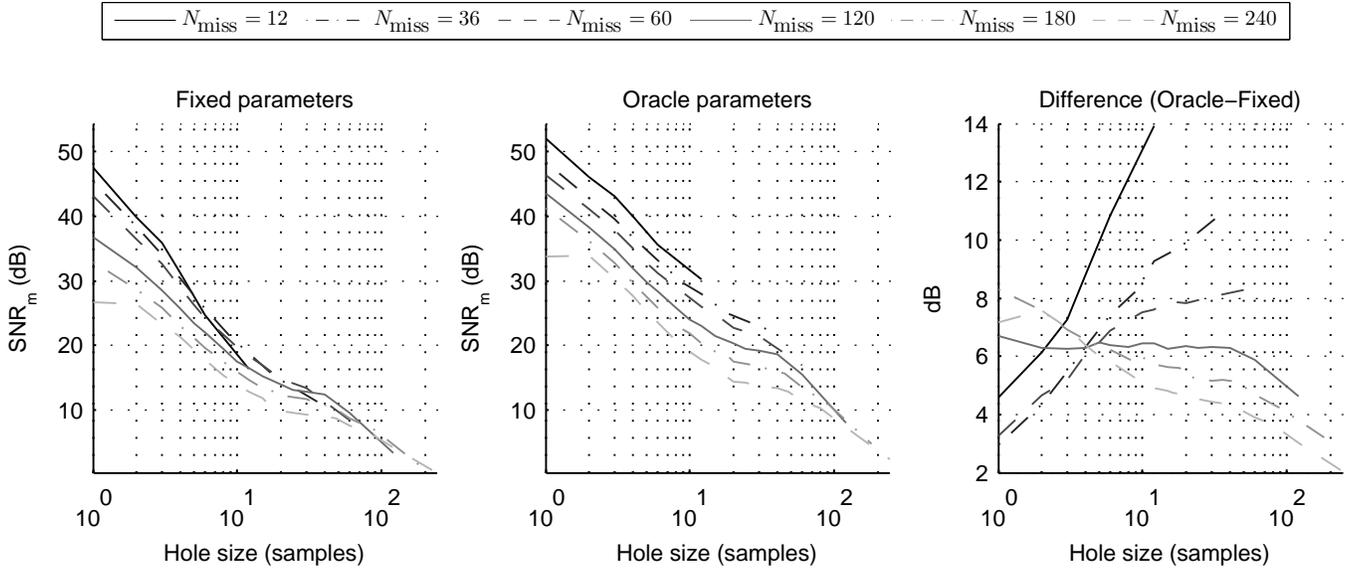


Fig. 3. Performance of the OMP-G algorithm, on the Speech@8kHz dataset, as a function of the hole size, for different values of the number N_{miss} of missing samples in a frame: estimation by the proposed algorithm with fixed parameters (left), ideal estimation with the best stopping parameters K^{OMP} and ϵ selected for each observed frame (center) and performance difference (right). The frame length is 512 samples and the hole size ranges from 1 sample (*i.e.* 0.12 ms, 0.2% of the frame) to 240 samples (*i.e.* 30 ms, 46.8% of the frame).

E. Declipping experiment

Clipping restoration is illustrated in Fig. 4 when the clipping level is 0.2. Here, the OMP-C-minmax algorithm is applied to an example of music signal, where one can observe that the reconstructed samples are close to the original signal.

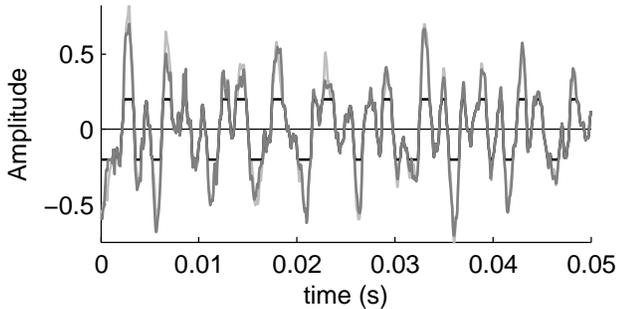


Fig. 4. Restoration of a music signal: original (light gray), clipped (black), estimate by the OMP-C-minmax algorithm (dark gray).

In a larger experiment, some of the proposed methods for restoring clipped signals are tested on the 3 datasets Speech@8kHz, Speech@16kHz and Music@16kHz. Each sound is artificially clipped with successive clipping levels, from 0.2 up to 0.9 with a 0.1-step. For this experiment, we selected the OMP-C-minmax, OMP-G, OMP-G-min and OMP-G-minmax algorithms after testing all the algorithms, since the results provide the most interesting conclusions (see below).

The performance of those algorithms are reported in Fig. 5, and show that:

- The use of the ‘min’ or ‘minmax’ declipping constraint

results in a large improvement of the SNR, on the average by 3 dB for OMP-G. A similar improvement has been obtained in the case of OMP-C. As in previous experiments, we see that methods based on SR, if efficient under random-measurement conditions [23], cannot straightforwardly recover partially-sampled signals when groups of missing samples are involved. But they are flexible enough to integrate additional constraints that leads to high performance.

- The minmax-constraint OMP-G-minmax algorithm reaches better results than the min-constraint OMP-G-min algorithm when the clipping level is 0.2. This corresponds to the range where the approximate value $\hat{\theta}_{\text{max}}$ is close to the actual maximum value as well as to the most degraded signals. A close analysis of the individual restored sounds reveals that large spikes are avoided thanks to the maximum value constraint. In a practical application, the maximum value $\hat{\theta}_{\text{max}}$ should be adjusted by the user until the best audio quality is achieved.
- The comparison between OMP-C-minmax and OMP-G-minmax shows that the initial-phase modeling by the Gabor dictionary improves the performance, as already observed in Section V-D1.

Performance comparison is obtained using two concurrent methods: the *ClipFix* Audacity plug-in based on cubic interpolation, the Cute Studio Declip commercial software⁴ and Janssen’s method [2] based on linear prediction. The OMP-G-minmax algorithm is compared against these methods and results are shown in Fig. 6. On the average, OMP-G-minmax

⁴<http://www.cutestudio.net/data/products/audio/seedeclip/>

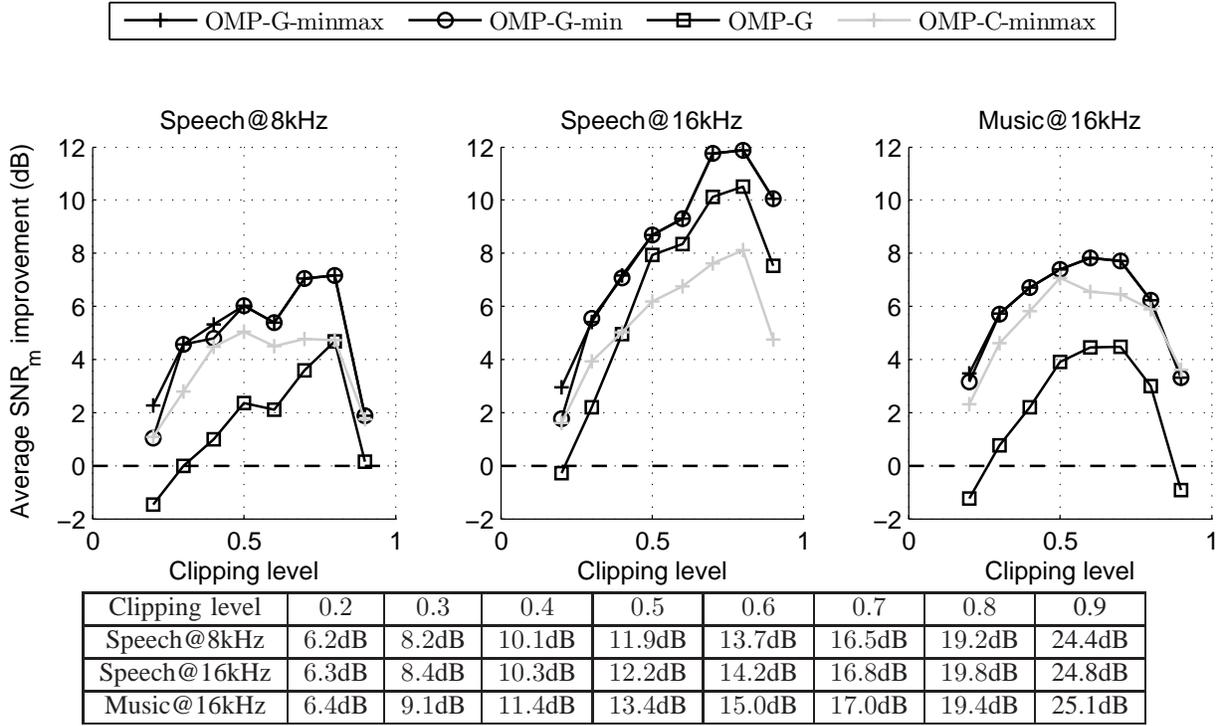


Fig. 5. Average SNR_m improvement, from the initial SNR of clipped signals to the SNR of restored signals, as a function of the clipping level: for each dataset – *i.e.* each subfigure –, the performance is presented as a function of the clipping level, for OMP-G, OMP-G-min, OMP-G-minmax and OMP-C-minmax.

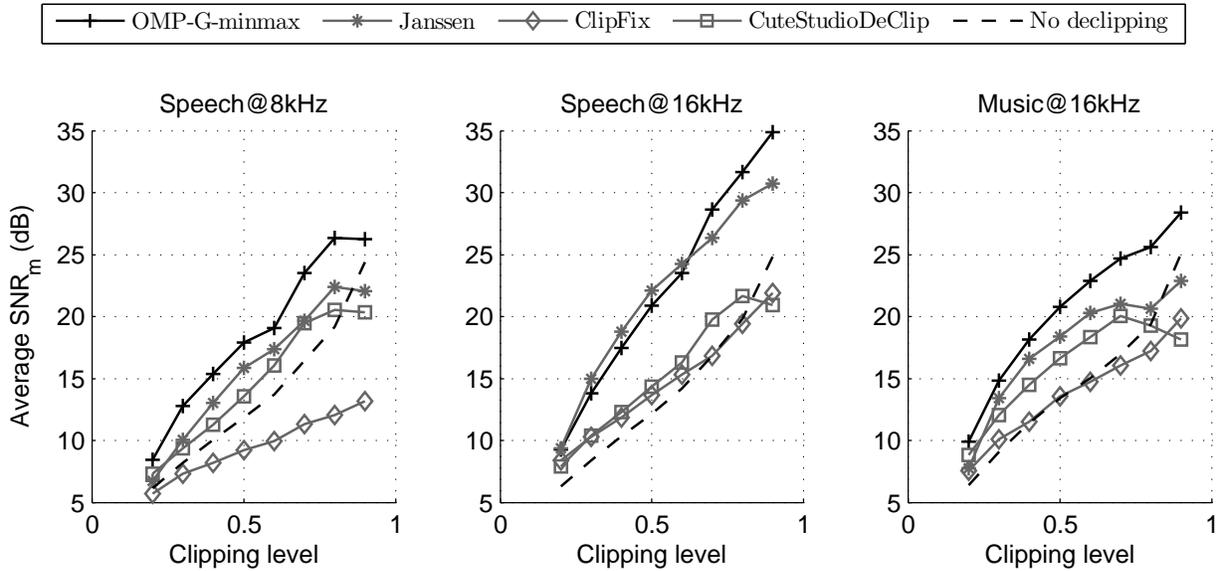


Fig. 6. Average SNR_m as a function of the clipping level: for each dataset, the performance is presented as a function of the clipping level, for BOMP-minmax, for Janssen’s approach [2] and for the spline interpolation (“Spline”). The initial SNR of the clipped signal is also plotted (“Clipped”).

outperform Janssen's method by 2.8 dB for the Speech@8kHz dataset; by 0.5 dB for the Speech@16kHz dataset; and by 3 dB for the Music@16kHz dataset. Lower performance is obtained from the Cute Studio Declip software, for which the underlying restoration technique is unknown. The *ClipFix* plug-in gives poor results, below all the reported ones.

VI. AUDIO INPAINTING TOOLBOX

We provide the Audio Inpainting Toolbox, which is dedicated to the development and usage of audio inpainting as well as to reproducible research. The current contents of the toolbox consist of the materials resulting from this paper. However, the modular design of the toolbox has been thought in order to facilitate new approaches and contributions. Modules include problem generation (*e.g.* general inpainting or clipping), solvers (*e.g.* OMP, OMP-minmax), audio data (*e.g.* speech, music), dictionaries (*e.g.* DCT, Gabor), and, finally, experiments, which are parameterized to call each of the other modules in order to integrate them into a single application-oriented sequence.

This toolbox is freely available⁵ as GPL Matlab code and Creative Commons license data.

VII. CONCLUSIONS

In this paper, we have presented the Audio Inpainting framework as the general problem of restoring distorted or missing audio data based on the available reliable data and on the prior knowledge of their location. We have defined Audio Inpainting as an inverse problem, and similarly to image inpainting approaches, we have proposed to use sparse representation methods to restore in the time domain the audio samples that are distorted or missing.

Using a frame-based processing of the audio signal, we have adapted the Orthogonal Matching Pursuit algorithm to address the Audio Inpainting problem, with either a discrete cosine or Gabor dictionary. The SNR performance of this algorithm has been shown to be comparable to or better than state-of-the-art methods when blocks of samples of variable durations were missing, and OMP with the Gabor dictionary has been found to give better results than OMP with DCT dictionary. Moreover, it has been shown that the size of the block of missing samples is more crucial for good signal restoration than the overall number of missing samples to estimate. For the special case of audio declipping, a constrained matching pursuit approach has been applied, that takes into account a priori and user-specified knowledge about the amplitude of the restored signal. This approach has been shown to significantly enhance the performance of the algorithm, which also outperforms state-of-the-art and commercially available methods for audio declipping in terms of Signal-to-Noise Ratio.

Based on the audio inpainting framework and on the baseline results presented in this paper, a number of future directions may be investigated. Technically, one may compare

the OMP-based methods to l_1 -minimization techniques, known to be another family of approaches to deal with sparse models. They are theoretically efficient but so far, we can only report preliminary performance that is lower than with greedy algorithms for audio inpainting⁶. Another perspective is the use of new sparse models for audio signals. In particular, structured sparse models and learned dictionary are promising directions. From an application point of view, time-frequency audio inpainting is a new investigation field for sparse approaches. Using the formulation of audio inpainting (see Section II-A) in the time-frequency domain, one must then introduce new dictionaries, targeting applications like source separation and bandwidth extension. Finally, the perceptually-based evaluation of audio inpainting should also be investigated in the context of each of the audio inpainting applications.

REFERENCES

- [1] A. Adler, V. Emiya, M. Jafari, M. Elad, R. Gribonval, and M. D. Plumbley, "A Constrained Matching Pursuit Approach to Audio Declipping," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, Prague, Czech Republic, May 2011.
- [2] A. Janssen, R. Veldhuis, and L. Vries, "Adaptive interpolation of discrete-time signals that can be modeled as autoregressive processes," *IEEE Trans. Acoustics, Speech and Sig. Proc.*, vol. 34, no. 2, pp. 317 – 330, apr 1986.
- [3] W. Etter, "Restoration of a discrete-time signal segment by interpolation based on the left-sided and right-sided autoregressive parameters," *IEEE Transactions on Signal Processing*, vol. 44, no. 5, pp. 1124 –1135, may 1996.
- [4] S. J. Godsill and P. J. W. Rayner, *Digital Audio Restoration - A Statistical Model-based Approach*. Springer-Verlag, 1998.
- [5] M. Lagrange and S. Marchand, "Long interpolation of audio signals using linear prediction in sinusoidal modeling," *Journal of the Audio Eng. Soc.*, vol. 53, pp. 891–905, 2005.
- [6] A. Dahimene, M. Noureddine, and A. Azrar, "A simple algorithm for the restoration of clipped speech signal," *Informatica*, vol. 32, pp. 183–188, 2008.
- [7] J. S. Abel and J. O. Smith, "Restoring a clipped signal," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, Toronto, Canada, May 1991.
- [8] R. C. Maher, "A method for extrapolation of missing digital audio data," in *95th AES Convention*, 1993.
- [9] M. Cooke, P. Green, L. Josifovski, and A. Vizinho, "Robust automatic speech recognition with missing and unreliable acoustic data," *Speech Communication*, vol. 34, no. 3, pp. 267 – 285, 2001.
- [10] J. Gemmeke, H. Van Hamme, B. Cranen, and L. Boves, "Compressive sensing for missing data imputation in noise robust speech recognition," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 4, no. 2, pp. 272 –287, 2010.
- [11] J. Le Roux, H. Kameoka, N. Ono, A. de Cheveigné, and S. Sagayama, "Computational auditory induction as a missing-data model-fitting problem with Bregman divergence," *Speech Communication*, vol. In Press, 2010.
- [12] M. Dietz, L. Liljeryd, K. Kjørting, and O. Kunz, "Spectral Band Replication, a novel approach in audio coding," in *Proc. of the 112th AES Convention*. Munich, Germany: Audio Engineering Society; 1999, May 2002.
- [13] E. Larsen and R. Aarts, *Audio bandwidth extension: application of psychoacoustics, signal processing and loudspeaker design*. Wiley, 2004.
- [14] P. Smaragdīs, B. Raj, and M. Shashanka, "Missing data imputation for spectral audio signals," in *Proc. of MLSP*, Grenoble, France, Sep. 2009.
- [15] M. Moussallam, P. Leveau, and S. M. Aziz Sbair, "Sound enhancement using sparse approximation with speclets," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, Mar. 2010, pp. 221 –224.

⁶While many people believe l_1 -minimization performs better than greedy algorithms, no theoretical result shows such a superiority in a general context [45], and the current works is a particular illustration where greedy algorithms perform better.

⁵<http://small-project.eu/software-data/>

- [16] C. Perkins, O. Hodson, and V. Hardman, "A survey of packet loss recovery techniques for streaming audio," *Network, IEEE*, vol. 12, no. 5, pp. 40–48, sep. 1998.
- [17] H. Ofir, D. Malah, and I. Cohen, "Audio Packet Loss Concealment in a Combined MDCT-MDST Domain," *Signal Processing Letters, IEEE*, vol. 14, no. 12, pp. 1032–1035, dec. 2007.
- [18] G. Cocchi and A. Uncini, "Subband neural networks prediction for on-line audio signal recovery," *IEEE Transactions on Neural Networks*, vol. 13, pp. 867–876, 2002.
- [19] S. J. Godsill, P. J. Wolfe, and W. N. W. Fong, "Statistical model-based approaches to audio restoration and analysis," *Journal of New Music Research*, vol. 30, no. 4, 2001.
- [20] J. Barker, *Computational auditory scene analysis: principles, algorithms, and applications*. IEEE Press/Wiley-Interscience, 2006, ch. Robust Automatic Speech Recognition, pp. 297–350.
- [21] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting," in *Proc. of 27th Conf. on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 2000, pp. 417–424.
- [22] M. Elad, J. L. Starck, P. Querre, and D. L. Donoho, "Simultaneous cartoon and texture image inpainting using morphological component analysis (mca)," *Applied and Computational Harmonic Analysis*, vol. 19, pp. 340–358, 2005.
- [23] M. Elad, *Sparse and Redundant Representations - From Theory to Applications in Signal and Image Processing*. Springer New-York, 2010.
- [24] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. On Signal Processing*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.
- [25] M. D. Plumbley, T. Blumensath, L. Daudet, R. Gribonval, and M. Davies, "Sparse representations in audio and music: from coding to source separation," *Proc. of the IEEE*, vol. 98, no. 6, 2010.
- [26] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An Algorithm for Designing Overcomplete Dictionaries for Sparse Representation," *IEEE Trans. On Signal Processing*, vol. 54, no. 11, pp. 4311–4322, Nov. 2006.
- [27] J. Mairal, M. Elad, and G. Sapiro, "Sparse representation for color image restoration," *IEEE Trans. on Image Processing*, vol. 17, no. 1, pp. 53–69, 2008.
- [28] D. Griffin and J. Lim, "Signal estimation from modified short-time fourier transform," *IEEE Trans. Acoustics, Speech and Sig. Proc.*, vol. 32, no. 2, pp. 236–243, Apr. 1984.
- [29] M. Davies and L. Daudet, "Sparse audio representations using the MCLT," *Signal Processing*, vol. 86, no. 3, pp. 457–470, 2006.
- [30] E. Ravelli, G. Richard, and L. Daudet, "Union of mdct bases for audio coding," *IEEE Trans. on Audio, Speech, and Language Processing*, vol. 16, no. 8, pp. 1361–1372, nov. 2008.
- [31] B. Natarajan, "Sparse approximate solutions to linear systems," *SIAM J. Computing*, vol. 25, no. 2, pp. 227–234, 1995.
- [32] G. Davis, S. Mallat, and M. Avellaneda, "Adaptive greedy approximations," *Constr. Approx.*, vol. 13, no. 1, pp. 57–98, 1997.
- [33] Y. Pati, R. Rezaifar, and P. Krishnaprasad, "Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition," in *Proc. Asilomar Conf. Signals, Systems, and Computers*, Nov. 1993, pp. 40–44 vol.1.
- [34] R. Gribonval, "Fast matching pursuit with a multiscale dictionary of gaussian chirps," *IEEE Trans. on Signal Processing*, vol. 49, no. 5, pp. 994–1001, May 2001.
- [35] L. Peotta and P. Vanderghenst, "Matching pursuit with block incoherent dictionaries," *Signal Processing, IEEE Transactions on*, vol. 55, no. 9, pp. 4549–4557, 2007.
- [36] R. Gribonval and E. Bacry, "Harmonic decomposition of audio signals with matching pursuit," *IEEE Trans. Signal Processing*, vol. 51, no. 1, pp. 101–111, 2003.
- [37] Y. Eldar and H. Bolcskei, "Block-sparsity: Coherence and efficient recovery," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, Taipei, Taiwan, 2009, pp. 2885–2888.
- [38] H. Mansour, R. Saab, P. Nasiopoulos, and R. Ward, "Color image desaturation using sparse reconstruction," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, Dallas, TX, USA, Mar. 2010.
- [39] *Method for objective measurements of perceived audio quality*, ITU-R Std. BS.1387, Dec. 1998.
- [40] R. Huber and B. Kollmeier, "Pemo-q - a new method for objective audio quality assessment using a model of auditory perception," *IEEE Trans. on Audio, Speech, and Language Processing*, vol. 14, no. 6, pp. 1902–1911, 2006.
- [41] V. Emiya, E. Vincent, N. Harlander, and V. Hohmann, "Subjective and objective quality assessment of audio source separation," *IEEE Trans. on Audio, Speech, and Language Processing*, vol. in press, 2011.
- [42] E. Vincent, S. Araki, and P. Bofill, "The 2008 signal separation evaluation campaign: A community-based approach to large-scale evaluation." Paraty, Brazil: Springer, Mar. 2009.
- [43] D. Donoho, "Compressed sensing," *Information Theory, IEEE Transactions on*, vol. 52, no. 4, pp. 1289–1306, April 2006.
- [44] E. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *Information Theory, IEEE Transactions on*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [45] J. Tropp, "Greed is good: algorithmic results for sparse approximation," *IEEE Transactions on Information Theory*, vol. 50, no. 10, pp. 2231–2242, Oct. 2004.



applications for audio, video and image processing. Amir Adler is the recipient of the 2011 Google Europe Doctoral Fellowship in Multimedia.



His research interests focus on audio signal processing and include sound modeling and indexing, sparse representations, source separation, quality assessment and applications to music and speech.



Her research interest are in the areas of blind signal separation, sparse representations and dictionary learning. She especially interested in the application of these and other techniques to speech and music processing and the analysis of other audio sounds, such as heart and bowel sounds.

Amir Adler received the B.Sc. (Cum Laude) degree in 1994 and the M.Eng. (Cum Laude) degree in 2001, both in Electrical Engineering, from the Technion, Israel. He is currently studying towards a PhD degree at the Computer Science Department, Technion, Israel.

From 2003 to 2006 he was Director of Signal Processing at GO Networks and from 2006 to 2008 he was Chief Technology Officer of the Wi-Fi Division at NextWave Wireless. His research interests include sparse and redundant representations with

Valentin Emiya graduated from Telecom Bretagne, Brest, France, in 2003 and received the M.Sc. degree in Acoustics, Signal Processing and Computer Science Applied to Music (ATIAM) at Ircam, France, in 2004. He received his Ph.D. degree in Signal Processing in 2008 at Telecom ParisTech, Paris, France. From 2008 to 2011, he is a post-doctoral researcher with the METISS group at INRIA, Centre Inria Rennes - Bretagne Atlantique, Rennes, France. He is now assistant professor (Maître de Conférences) in computer science at University Provence/Aix-Marseille I, Marseille, France.

Maria G. Jafari (S01M02) received the M.Eng. (Hons.) degree in electrical and electronic engineering from Imperial College London, U.K., in 1999 and the Ph.D. degree in signal processing from Kings College London, U.K. in 2003. From 2002 to 2004 she worked as a Research Associate at Kings College London, where her research focused on the application of signal processing to biomedical problems. In June 2004, she joined the Centre for Digital Music, Queen Mary University of London, U.K., where she is currently a Research Assistant.



Michael Elad received his B.Sc. (1986), M.Sc. (1988) and D.Sc. (1997) from the department of Electrical engineering at the Technion, Israel.

From 1988 to 1993 he served in the Israeli Air Force. From 1997 to 2000 he worked at Hewlett-Packard laboratories Israel as an R&D engineer. During 2000 to 2001 he headed the research division at Jigami corporation, Israel. During the years 2001 to 2003 Michael held a research-associate position at Stanford university. Since late 2003 Michael is a faculty member at the Computer-Science department at the Technion. On May 2007 Michael was tenured to an associate-professorship, and on July 2010 he was promoted to full-professorship.

Michael Elad works in the field of signal and image processing, specializing in particular on inverse problems, sparse representations and super-resolution. Michael received the Technion's best lecturer award six times, he is the recipient of the Solomon Simon Mani award for excellence in teaching in 2007, and he is also the recipient of the Henri Taub Prize for academic excellence (2008) and the Hershel-Rich prize (2010) for innovation. Michael is a senior IEEE member, and he is currently serving as an associate editor for SIAM journal on imaging sciences (SIIMS), and for the IEEE Transactions on Information Theory.



Rémi Gribonval graduated from École Normale Supérieure, Paris, France in 1997. He received the Ph. D. degree in applied mathematics from the University of Paris-IX Dauphine, Paris, France, in 1999, and his Habilitation à Diriger des Recherches in applied mathematics from the University of Rennes I, Rennes, France, in 2007. He is a Senior Member of the IEEE.

From 1999 until 2001 he was a visiting scholar at the Industrial Mathematics Institute (IMI) in the Department of Mathematics, University of South Carolina, SC. He is now a Research Director (Directeur de Recherche) with INRIA in Rennes, France, in the METISS group. His research interests include mathematical signal processing, machine learning, approximation theory and statistics. His research focuses on sparse approximation and applications to multichannel audio signal processing, with a particular emphasis in blind audio source separation and compressed sensing. Since 2002 he has been the coordinator of several national, bilateral and european research projects, and in 2008 he was elected a member of the steering committee for the international conference ICA on independent component analysis and source separation. He is the founder of the series of international workshops SPARS on Signal Processing with Adaptive/Sparse Representations. In 2011, he was awarded the Blaise Pascal Award in Applied Mathematics and Scientific Engineering from the French National Academy of Sciences. He has been awarded a starting investigator grant from the European Research Council in 2011.



Mark D. Plumbley (S'88–M'90) received the B.A. (Hons.) degree in electrical sciences in 1984 from the University of Cambridge, Cambridge, U.K., and the Ph.D. degree in neural networks in 1991, also from the University of Cambridge. From 1991 to 2001 he was a Lecturer at King's College London. He moved to Queen Mary University of London in 2002, and where he is now an EPSRC Leadership Fellow and Director of the Centre for Digital Music.

His research focuses on the automatic analysis of music and other audio sounds, including automatic music transcription, beat tracking, and audio source separation, and with interest in the use of techniques such as independent component analysis (ICA) and sparse representations.

Prof. Plumbley chairs the ICA Steering Committee, and is a member of the IEEE SPS TC on Audio and Acoustic Signal Processing.