



**HAL**  
open science

## Sparse off-the-grid computation of the zeros of STFT

Jean-Baptiste Courbot, Ali Moukadem, Bruno Colicchio, Alain Dieterlen

► **To cite this version:**

Jean-Baptiste Courbot, Ali Moukadem, Bruno Colicchio, Alain Dieterlen. Sparse off-the-grid computation of the zeros of STFT. IEEE Signal Processing Letters, 2023, pp.1-5. 10.1109/LSP.2023.3289769 . hal-04145192

**HAL Id: hal-04145192**

**<https://hal.science/hal-04145192>**

Submitted on 29 Jun 2023

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

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



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

# Sparse off-the-grid computation of the zeros of STFT

Jean-Baptiste Courbot, Ali Moukadem, Bruno Colicchio, Alain Dieterlen

**Abstract**—The Short Time Fourier Transform (STFT) of white noise present well-known properties regarding the distribution of its zeroes. In this letter, we propose a method to locate them beyond the usual time-frequency grid, in order to gain a better understanding on the underlying point process. To do so, we use a recent method based on off-the-grid sparse analysis, that aim at finding the best set of atomic shape forming the STFT modulus. We also improve it by ensuring the resulting combination of atoms minimizes a criterion based on the theoretical behavior of the set of zeros. Numerical experiments support our approach and show, with respect to discrete methods, an improved statistical power for the detection of signals in white noise.

**Index Terms**—Zeros of the STFT, off-the-grid sparse analysis, point process

## I. INTRODUCTION

The zeros of the Short-Time Fourier Transform (STFT) of white noise present a well-known structure which has been the topic of several studies, see *e.g.* [13], [14], [3], [4], [11], offering a promising statistical tool for filtering, detection or reconstruction of signal in presence of white noise. In this letter, we aim at providing an off-the-grid framework for the retrieval of their position beyond the sampling grid of the STFT.

### A. Zeros of the STFT

The STFT of a signal  $x \in L^2(\mathbb{R})$  is:

$$S_\gamma x(\tau, f) = \int_{\mathbb{R}} x(t) \overline{\gamma(t-\tau)} e^{-2i\pi ft} dt; \quad (1)$$

where  $\tau \in \mathbb{R}$  and  $f \in \mathbb{R}^*$  locate the STFT in time and frequency, and  $\gamma$  is an observation window. Assuming here and in the following that  $\gamma(t)$  is a Gaussian circular window, the STFT of (real or complex) white noise is equivalent to a planar Gaussian Analytic Function (GAF) [3]. So the zeros of the STFT coincide with the zeros of a GAF, whose theoretical properties are known [15], [12].

Numerically, the STFT is computed on a grid, so the zeros are only available from this discretization. In the literature, two methods retrieve the set of zeros from the STFT:

- In [3], [13], the authors search for local minima among 8-pixels neighborhoods. We refer to it as the *Local Neighborhood* (LN) methods afterwards.
- In [11], the authors introduced the Adaptive Minimal grid Neighbors (AMN) method which is proven to reach the accurate zero set, with a precision depending of the grid spacing.

Up to our knowledge, there is no method that allows to provide a gridless zero set recovery.

Submitted on March 24, 2022.

JBC, AM, BC and AD are within IRIMAS, UR 7499, Université de Haute-Alsace, Mulhouse, France. Contact: jean-baptiste.courbot@uha.fr.

### B. Off-the-grid sparse analysis

Sparse analysis aims at decomposing an observation  $\mathbf{y} \in \mathbb{R}^P$ ,  $P \in \mathbb{N}$  into elementary atoms, typically assuming:

$$\mathbf{y} = \sum_{n=1}^N \mathcal{G}(\mathbf{z}_n, \sigma_n, w_n) + \epsilon. \quad (2)$$

where  $\mathcal{G}$  is the atomic shape to consider,  $\mathbf{z}_n$  locates the  $n$ -th spike,  $\sigma_n$  its width, and  $w_n$  its weight. Sparse analysis aims at finding, from a given  $\mathbf{y}$ , the best values for the  $(\mathbf{z}_n, \sigma_n, w_n)$  together with their number.

Generally speaking, limiting the search space to a grid has known limitations, notably regarding the stability of the methods when refining the search grid [10]. More recent approaches [8], [5], [6] have recast the problem in a continuous fashion. In this framework, each atom is described by a weighted Dirac mass locating its parameters in the space of measures:

$$\mathbf{y} = \Phi \boldsymbol{\mu} + \epsilon \quad (3)$$

where  $\boldsymbol{\mu} = \sum_{n=1}^N w_n \delta_{\mathbf{z}_n, \sigma_n}$  and  $\Phi : \mathcal{M}(\mathcal{D}) \mapsto \mathbb{R}^P$  embeds  $\mathcal{G}$ , and  $\epsilon$  is a perturbation term. We assume that  $(\mathbf{z}_n, \sigma_n, w_n) \in \mathcal{D}$ , and that  $\mathcal{M}(\mathcal{D})$  is the Radon space of measure over  $\mathcal{D}$ , so that  $\boldsymbol{\mu} \in \mathcal{M}(\mathcal{D})$ .

This kind of approach, besides the more precise formulation of the problem, also offers recovery guarantees under mild conditions [9], yielding for instance the *Sliding Frank-Wolfe* (SFW) algorithm. For practical considerations, the choice of the regularization parameter is also of importance. Based on analogous discrete-valued problems [17], we proposed in [7] an *homotopy* method that allow to replace this choice by a noise-related criterion.

### C. Methodology

In this letter, our aim is to take the sparse off-the-grid approach to locate, in a continuous setting, the zeros of the STFT of white noise. To simplify the discussion regarding the symmetry w.r.t. the real axis in the real-valued case, we assume that the white noise is complex. In the following, we note  $S$  the STFT of a white Gaussian noise.

We address at first the problem modeling in Section II-A, before moving to algorithmic considerations in Section II-B. Then, we will present numerical results validating the approach in Section III.

## II. FINDING OFF-THE-GRID ZEROS

### A. Observation model and problem

*STFT topology near the zeros.* We aim at finding the localization of the zeros of the STFT over a continuous domain, modeling zeros as atomic shapes. To do so, we need to account

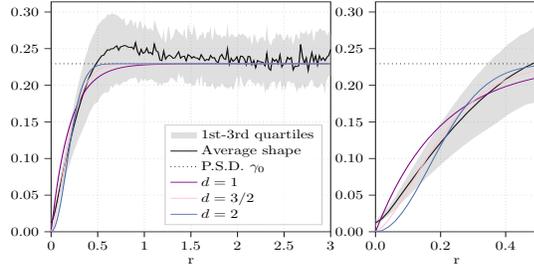


Fig. 1: Averaged radial profile in the vicinity of zeros, as determined by the discrete LN method [3]. In color are superimposed possible atomic shape fit  $\gamma_0 - e(\mathbf{z})$  (Eq. (5)).  $r = |\mathbf{z} - \mathbf{z}_0|$  depicts the distance in the complex plane from the zeros  $\mathbf{z}_0$ , and the average was obtained on  $10^3$  zeros.

for the shape of the STFT modulus  $|S|$  in the vicinity of zeros. While the *location* of the zeros has been the topic of many studies [13], [3], [11], to our knowledge only a few studies address the question of the vicinity of the zeros in white noise STFT. From [14, Chap.15] and originally from [20], we have that the analytic STFT admits a Weierstrass-Hadamard factorization from its corresponding Bargmann transform. At any point  $\mathbf{z} = \tau + if \in \mathbb{C}$ :

$$S(\mathbf{z}) = e^{-\frac{|\mathbf{z}|^2}{4}} \mathbf{z}^m e^{P_2(\mathbf{z})} \prod_n \left(1 - \frac{\mathbf{z}}{\mathbf{z}_n}\right) \exp\left(\frac{\mathbf{z}}{\mathbf{z}_n} + \frac{\mathbf{z}^2}{2\mathbf{z}_n^2}\right); \quad (4)$$

where  $m \geq 0$ ,  $P_2(\mathbf{z}) = C_0 + C_1\mathbf{z} + C_2\mathbf{z}^2$  and  $\mathbf{z}_n$  locates the  $n$ -th zero of the STFT in the complex plane. This factorisation is sometimes called the *Husimi* representation of the STFT of white noise. However, such decomposition is not helpful to recover the STFT from the positions of the zero only. Indeed,  $m$ ,  $C_0$ ,  $C_1$  and  $C_2$  are unknown, and the product involves possibly infinitely-many terms. The only available approximation so far, from [14, Chap.15], is that  $|S(\mathbf{z})| \propto |\mathbf{z} - \mathbf{z}_n|$  when  $\mathbf{z} \rightarrow \mathbf{z}_n$ .

Then, we adopt an empirical approach to choose an atomic shape that will represent the STFT in the vicinity of zeros. We consider for now elementary shape of the form:

$$e(\mathbf{z}) = \exp\left(-\frac{1}{2} \frac{|\mathbf{z} - \mathbf{z}_n|^d}{\sigma^d}\right); \quad (5)$$

with  $d = 1$  (Laplace kernel),  $d = 2$  (Gaussian) or  $d = 3/2$  (Generalized Gaussian). In Fig. 1 we depicts the average topology of the modulus of the STFT in the neighborhood of zeros, superimposed by various kernel fit. From this kind of observation we can not determine what would be the “best shape”, so numerical experiments considering these three options will allow, in Section III, to determine the best choice in terms of spike location.

*Observation model.* The average value of the STFT of white noise is its power spectrum density  $\gamma_0$  (see [14, Section 13.2] and Fig. 1). In order to handle positive spikes, we make the following convenience transformation:  $\mathbf{y} = 1 - |S|/\gamma_0$ . This does not modify the reasoning developed afterwards. So henceforth, we consider values of  $\mathbf{y}$  above zero as indicating a potential spike, and ignore  $\mathbf{y}$  below zero in order to avoid fitting regions without interest for spike location. Then, we model  $\mathbf{y}$

as in (2)(3), where  $\boldsymbol{\mu}$  locates each zero in the time-frequency plane.

In usual image-processing problems  $\epsilon$  encodes an additive noise originating from the data formation. In our context there is no *observation noise*, but an unknown discrepancy term between the spike mixture and the actual STFT.

*Inference.* Finding  $\boldsymbol{\mu}$  from  $\mathbf{y}$  can be made by forming the following *Beurling* LASSO (or BLASSO, see [1]) problem:

$$\arg \min_{\boldsymbol{\mu} \in \mathcal{M}(\mathcal{D})} C_\lambda(\boldsymbol{\mu}) = \arg \min_{\boldsymbol{\mu} \in \mathcal{M}(\mathcal{D})} \|\mathbf{y} - \Phi \boldsymbol{\mu}\|_2^2 + \lambda \|\boldsymbol{\mu}\|_{\text{TV}}. \quad (\mathcal{P}_\lambda)$$

$\lambda$  is a regularization parameter to set, and the operator  $\|\cdot\|_{\text{TV}}$  encodes the total variation of measures, so  $\|\boldsymbol{\mu}\|_{\text{TV}} = \sum_{n=1}^N |w_n|$ . This operator is analog, in our problem, to a  $\ell_1$  norm.

## B. Algorithms

*Finding zeros with SFW.* For a given  $\lambda$ , solving  $(\mathcal{P}_\lambda)$  is feasible using the Sliding Frank-Wolfe (SFW) algorithm [9]. Starting from an initial guess, SFW repeats at each iteration  $k$  the following:

- Compute the approximate *certificate*  $\eta_{[k]} : \mathcal{D} \rightarrow \mathbb{R}$  as:

$$\eta_{[k]} = \frac{1}{\lambda} \Phi^\top (\mathbf{y} - \Phi \boldsymbol{\mu}_{[k]}). \quad (6)$$

- Add a new atom to the support, maximizing  $\eta_{[k]}$ .
- Adjust only the weights  $\{w_n\}$ , minimizing  $C_\lambda$  from  $(\mathcal{P}_\lambda)$ .
- Adjust each parameter  $\{\boldsymbol{\theta}_n, w_n\}$  minimizing  $C_\lambda$ .

Then, SFW stops when  $\|\eta_{[k]}\|_\infty < 1$ . Theoretical foundations of SFW as well as implementation details can be found in [9].

*Homotopy algorithm.* Finding the good  $\lambda$  for a given  $(\mathcal{P}_\lambda)$  is not obvious. To circumvent this problem, we consider developing an homotopy algorithm as in [7]. However, the latter relied on a simple noise-related criterion that allowed practitioners to set some “target” residual variance. As mentioned, such criterion is not helpful in the context of the STFT.

Instead, we propose to make use of the theoretical knowledge of the behavior of the zeros within STFT of white noise. In particular, the theoretical pair correlation function  $g_0(r)$  is known [3], [14] (see Fig. 3). This functions depicts the likeliness that a pair of points occurs, depending on the distance  $r$  between them.  $g_0(r) < 1$  indicates a *repulsive* range,  $g_0(r) > 1$  distances are more likely to occurs, while  $g_0(r) \simeq 1$  indicates independence, as in a Poisson process (constant  $g_0(r) = 1$ ).

Based on  $g_0$ , we propose a  $\ell_g$  criterion measuring the cost of a point set  $\mathbf{z}$  seen through the theoretical  $g_0$ . Let us denote  $\mathcal{P}(\boldsymbol{\mu})$  the set of pair of atoms within  $\boldsymbol{\mu}$ . Then

$$\ell_g(\boldsymbol{\mu}) = \frac{1}{\text{Card}(\mathcal{P}(\boldsymbol{\mu}))} \sum_{(i,j) \in \mathcal{P}(\boldsymbol{\mu})} -\log(g_0(|\mathbf{z}_i - \mathbf{z}_j|)). \quad (7)$$

The intuition behind this choice is that as  $\lambda$  decreases,  $\ell_g(\boldsymbol{\mu})$  attains a unique minimum. Indeed, for a decreasing sequence  $\lambda_1 > \lambda_2 > \dots > \lambda_k$ , the corresponding solutions  $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \dots, \boldsymbol{\mu}_k$  contains more and more atoms. Because the estimation happens within a bounded space, increasing the number of atoms also increase the density, implying that on average  $g_0$  will be evaluated at a shorter range.

This postulate seems difficult to prove:

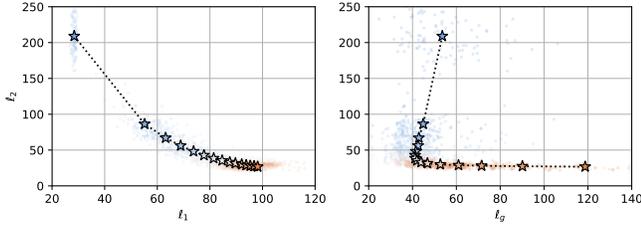


Fig. 2: Numerical values taken by the  $\ell_2$  and  $\ell_1$  norms ( $\|\mathbf{y} - \Phi\boldsymbol{\mu}\|_2^2$  and  $\|\boldsymbol{\mu}\|_{\text{TV}}$  respectively in  $(\mathcal{P}_\lambda)$ ), and the  $\ell_g$  criterion (Eq. (7)) over 15 steps of the homotopy algorithm (Alg. 1), without stopping criterion for the purpose of illustration. Each color represent a different step and the stars depicts the average values per step. Note that the left graph depicts the typical piecewise-linear Pareto frontier describing the  $\ell_2 / \ell_1$  tradeoff when  $\lambda$  decreases.

- Analytically, there is no explicit link between the pairwise point distance and the regularization parameter, so deriving  $\ell_g(\boldsymbol{\mu})$  w.r.t.  $\lambda$  is not obvious.
- The averaging operation is not interchangeable with  $g_0$ .

However, this postulate appears to be numerically valid, as shown in Fig. 2. These results allow to state that on average, the  $\ell_g$  criterion attains a unique minimum over the course of the homotopy algorithm, when applied on the modulus of STFT of white noise.

Starting from the zero set of a pixel-based method (as [3], [11]), this allows to design a homotopy algorithm that repetitively decreases  $\lambda$ , solves  $(\mathcal{P}_\lambda)$ , compute  $\ell_g$  and stops when it has attained a minimum. Formally, this implies that the homotopy algorithm solves:

$$\arg \min_{\boldsymbol{\mu}_\lambda \in \mathcal{M}(\mathcal{D})} \ell_g(\boldsymbol{\mu}_\lambda) \text{ subject to: } \boldsymbol{\mu}_\lambda \text{ is a solution to } (\mathcal{P}_\lambda) \quad (8)$$

The homotopy method we propose is formalized in Alg 1, and the Python code is available at [github.com/courbot/sparse-OTG-zeros-STFT](https://github.com/courbot/sparse-OTG-zeros-STFT). Note that the update rule for  $\lambda$  ensures that in the next homotopy step, a new atom will be added because the SFW stopping criterion will not be met anymore (see [7]). Note also that the *boosted* SFW introduced in [7] to accelerate the appending of atoms is not relevant for the purpose of this paper, as the starting value of the homotopy algorithm already contains most of the atoms of the solution.

### III. NUMERICAL EXPERIMENTS

#### A. Functional and summary statistics

Point process realizations can be summarized, and compared, by *summary statistics* that describe their behavior with a real number. To do so, it is common (as in [3]) to rely on some function of the inter-point distance  $r$ , as the pair correlation function  $g(r)$ , or the Ripley  $K$ -function [16] and its variance-stabilized version  $L$ , expressed as:

$$K(r) = 2\pi \int_0^r t g(t) dt \text{ and } L(r) = \sqrt{K(r)/\pi}; \quad (9)$$

or the empty space function  $F$ :

$$F(r) = p(\text{Card}(z \cap \mathcal{B}(r)) \geq 1); \quad (10)$$

#### Algorithm 1 Homotopy algorithm for zero location

**Require:**  $\mathbf{y}, \Phi$

**Ensure:** Estimation of  $\boldsymbol{\mu}^*$ , solution to (8).

Initialization:  $\boldsymbol{\mu}_{[0]}$  as found by a pixel-based method [3], [11], and  $\lambda_0 = \|\Phi^T \mathbf{y}\|_\infty$ .

**repeat** (iteration  $t$ )

1. Starting from  $\boldsymbol{\mu}_{[t-1]}$ , solve  $(\mathcal{P}_{\lambda_t})$  using the SFW algorithm to obtain  $\boldsymbol{\mu}_{[t]}$ .
2. Compute  $\ell_g(\boldsymbol{\mu}_{[t]})$  from (3).
3. Compute  $\max_{\mathcal{D}} \eta_{[t]}$  by local ascent.

4. Update  $\lambda_{t+1} = \frac{\lambda_t \max \eta_{[t]}}{1+c}$

**until**  $\arg \min_{1 \leq i \leq t} \{\ell_g(\boldsymbol{\mu}_{[i]})\} \neq \boldsymbol{\mu}_{[t]}$

Set  $(\boldsymbol{\mu}_{[t-1]})$  as the solution. **End** of the algorithm.

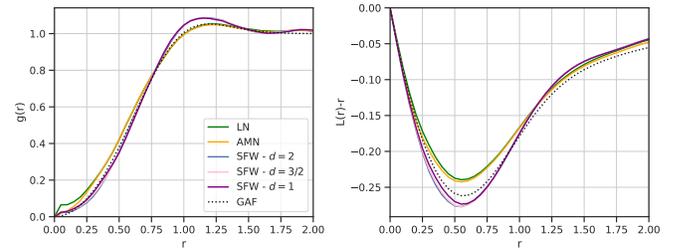


Fig. 3: Estimated  $g(r)$  and  $L(r) - r$  statistics averaged over 100 realizations of STFT of complex white noise (plain lines) and their theoretical counterpart  $g_0(r)$  and  $L_0(r) - r$  (dotted).

with  $\mathcal{B}(r)$  defining any ball of radius  $r$  within the domain of interest.

Following again [3], we use two summary statistics over  $L$ , using  $L_0$  as a known reference:

$$T_2^L = \sqrt{\int_{r_{\min}}^{r_{\max}} |L(r) - L_0(r)|^2 dr} \quad (11)$$

$$T_\infty^L = \sup_{r_{\min} < r < r_{\max}} |L(r) - L_0(r)| \quad (12)$$

Similar definitions holds for  $T_2^F$  and  $T_\infty^F$ , with the peculiarity that  $F_0$  is not known analytically (see *e.g.* [18]); so it is replaced by its numerical average from experiments on white noise.

Usually one sets  $r_{\min}$  to 0 and study the behavior of the summary statistics with respect to the value of  $r_{\max}$ .

#### B. Experimental variations

Average estimations of  $g(r)$  and  $L(r) - r$  are given in Fig. 3, and the corresponding summary statistics are reported in Table I. This allows to determine that a Laplace kernel (*i.e.*  $d = 1$  in Eq. (5)) is better suited to the problem than its alternatives. Besides, for the latter we observe on average lower  $T_2^L$  and  $T_\infty^L$  values than with the discrete methods [3], [11]. This suggest a better adedation of the H-SFW solutions with the GAF theoretical model.

Then, we evaluate the methods in their success to detect correctly that a signal contains more than noise, that is we test the null hypothesis  $\mathcal{H}_0$ : “ $\mathbf{y}$  contains only white Gaussian

	LN [3]	AMN [11]	H-SFW: $d = 2$	$d = 1.5$	$d = 1$
$T_2^L$	0.1768	0.1698	0.1625	0.1655	0.1616
$T_\infty^L$	0.0588	0.0563	0.0563	0.0575	0.0550

TABLE I: Averaged summary statistics of  $L$  over 100 realization of white noise at  $r_{\max} = 2$ , considering two discrete methods and 3 kernels for H-SFW.

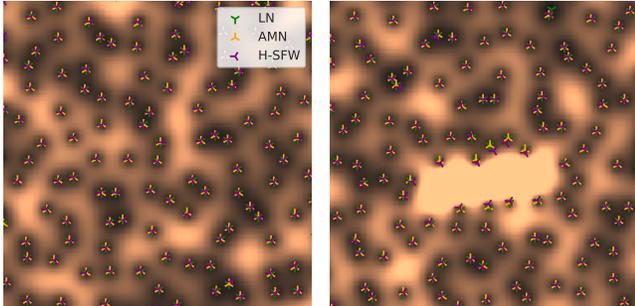


Fig. 4: STFT modulus for a white noise (left) and a white noise plus a chirp signal (right) with  $\text{SNR} = 12.5$ . The estimated zeros locations for LN [3], AMN [11] and H-SFW (this paper) are superimposed.

noise” against its alternative. For simplicity, we focus on signal containing white Gaussian noise plus a chirp signal. To do so we rely on a Monte Carlo envelope test [2], following:

- set a target radius  $r_{\max} > 0$  and  $r_{\min} = 0$ .
- compute  $K$  realizations of white noise as well as one realization of a noisy chirp signal.
- for each, compute the associated point set either by LN, AMN, or the H-SFW method proposed in this paper. An example of result is given in Fig. 4.
- For each point set compute the summary statistics  $T_2^L$ ,  $T_2^F$  and  $T_\infty^F$ ,  $T_\infty^L$ .
- for each summary statistic of a tested signal, if it is above the  $k$ -th greatest statistic of the white noise, we reject  $\mathcal{H}_0$ .

Typically, setting  $K = 99$  and  $k = 5$ , the significance level is  $\alpha = k/(K + 1) = 0.05$ . We can repeat the experiment several times and estimate the frequency at which  $\mathcal{H}_0$  is rejected when it is false: we can estimate empirically the power  $\beta$  of the test. We also repeat the experiment by:

- changing the target  $r_{\max}$
- changing  $r_{\max}$  together with the SNR. The latter is defined as  $A^2/2\sigma^2$ ,  $A$  being the chirp maximum amplitude, and  $\sigma^2$  the noise variance.

Fig. 5 and 6 respectively report the results for these two points.

### C. Results and discussion

From Figures 4, 5 and 6 we can make the following observations:

- The  $F$  functional statistic yields in most cases a higher detection power, generalizing the result shown in [3].
- The  $T_2$  summary statistic yields slightly better results than its  $T_\infty$  counterpart, improving the maximum power by 5%.

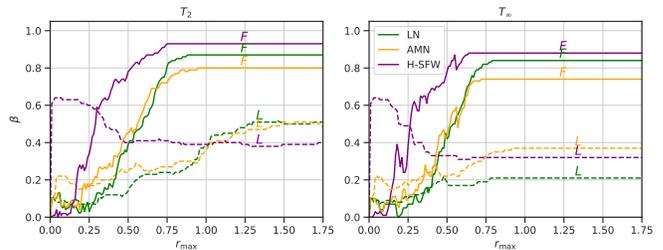


Fig. 5: Average empirical power of Monte Carlo envelope test, measured at  $\text{SNR} = 12.5$ , for the three considered methods using  $T_2$  (left) or  $T_\infty$  over the  $L$  function (dashed) or the  $F$  function (plain).

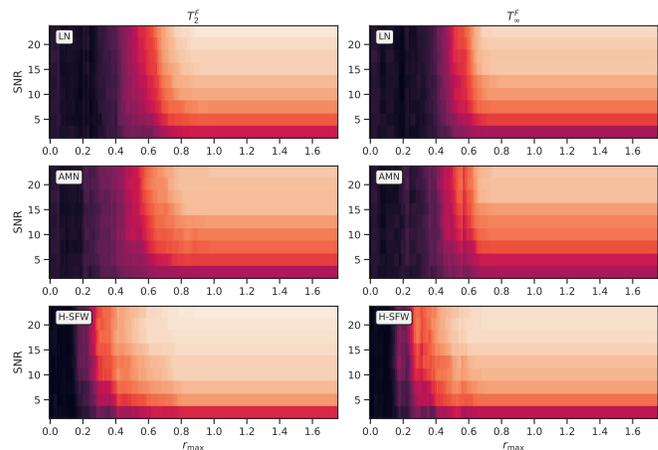


Fig. 6: Powers of the  $T_2^F$  and  $T_\infty^F$  Monte Carlo envelope test depending on the SNR and the choice of  $r_{\max}$ . The middle line of each plot ( $\text{SNR}=12.5$ ) can be found in plain in Fig. 5. The colormap ranges from 0 (dark) to 1 (light).

- Overall, at any SNR, a higher power is attainable using the H-SFW method introduced in this letter, with a 6% gain in maximum power.

Besides, the inspection of  $T_2^L$  and  $T_\infty^L$ , and examples as in Fig. 4, suggest that the H-SFW approach may uncover large deviations from the reference at close range ( $r_{\max} < 0.5$ ). We believe this is due to two phenomena: 1) a better (subpixel) location of the zeros; and 2) larger model discrepancies in the presence of signal.

## IV. CONCLUSION

This letter allowed us to introduce how off-the-grid sparse approaches can be used, and adapted, for the problem of zero localization in the STFT of white noise. Numerical experiments shown that the resulting point process better fit the theoretical GAF properties, and that we obtained an improved power for the detection test based on the zero pattern. Future works will consider the generalization to other transform (e.g. corresponding to non-planar GAF), adapted to the use of non-isotropic functional statistics. We will also consider how this method can help to use zeros as signal’s features, as in [19].

## REFERENCES

- [1] Jean-Marc Azais, Yohann De Castro, and Fabrice Gamboa. Spike detection from inaccurate samplings. *Applied and Computational Harmonic Analysis*, 38(2):177–195, 2015.
- [2] Adrian Baddeley, Peter J Diggle, Andrew Hardegen, Thomas Lawrence, Robin K Milne, and Gopalan Nair. On tests of spatial pattern based on simulation envelopes. *Ecological Monographs*, 84(3):477–489, 2014.
- [3] Rémi Bardenet, Julien Flamant, and Pierre Chainais. On the zeros of the spectrogram of white noise. *Applied and Computational Harmonic Analysis*, 48(2):682–705, 2020.
- [4] Rémi Bardenet and Adrien Hardy. Time-frequency transforms of white noises and gaussian analytic functions. *Applied and computational harmonic analysis*, 50:73–104, 2021.
- [5] Kristian Bredies and Hanna Katriina Pikkarainen. Inverse problems in spaces of measures. *ESAIM: Control, Optimisation and Calculus of Variations*, 19(1):190–218, 2013.
- [6] Emmanuel J Candès and Carlos Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on pure and applied Mathematics*, 67(6):906–956, 2014.
- [7] Jean-Baptiste Courbot and Bruno Colicchio. A fast homotopy algorithm for gridless sparse recovery. *Inverse Problems*, 37(2):025002, 2021.
- [8] Yohann De Castro and Fabrice Gamboa. Exact reconstruction using beurling minimal extrapolation. *Journal of Mathematical Analysis and applications*, 395(1):336–354, 2012.
- [9] Quentin Denoyelle, Vincent Duval, Gabriel Peyré, and Emmanuel Soubies. The sliding frank–wolfe algorithm and its application to super-resolution microscopy. *Inverse Problems*, 36(1):014001, 2019.
- [10] Vincent Duval and Gabriel Peyré. Exact support recovery for sparse spikes deconvolution. *Foundations of Computational Mathematics*, 15(5):1315–1355, 2015.
- [11] Luis Alberto Escudero, Naomi Feldheim, Günther Koliander, and José Luis Romero. Efficient computation of the zeros of the bargmann transform under additive white noise. *Foundations of Computational Mathematics*, pages 1–34, 2022.
- [12] Naomi D Feldheim. Zeroes of gaussian analytic functions with translation-invariant distribution. *Israel Journal of Mathematics*, 195(1):317–345, 2013.
- [13] Patrick Flandrin. Time–frequency filtering based on spectrogram zeros. *IEEE Signal Processing Letters*, 22(11):2137–2141, 2015.
- [14] Patrick Flandrin. *Explorations in time-frequency analysis*. Cambridge University Press, 2018.
- [15] John Ben Hough, Manjunath Krishnapur, Yuval Peres, et al. *Zeros of Gaussian analytic functions and determinantal point processes*, volume 51. American Mathematical Soc., 2009.
- [16] N Bert Loosmore and E David Ford. Statistical inference using the g or k point pattern spatial statistics. *Ecology*, 87(8):1925–1931, 2006.
- [17] Michael R Osborne, Brett Presnell, and Berwin A Turlach. A new approach to variable selection in least squares problems. *IMA journal of numerical analysis*, 20(3):389–403, 2000.
- [18] Barbara Pascal and Rémi Bardenet. A covariant, discrete time-frequency representation tailored for zero-based signal detection. *IEEE Transactions on Signal Processing*, 2022.
- [19] Pierre Rougé, Ali Moukadem, Alain Dieterlen, Antoine Boutet, and Carole Frindel. Generalizable features for anonymizing motion signals based on the zeros of the short-time fourier transform. *Journal of Signal Processing Systems*, pages 1–11, 2022.
- [20] Mikito Toda. Phase retrieval problem in quantum chaos and its relation to the origin of irreversibility i. *Physica D: Nonlinear Phenomena*, 59(1-3):121–141, 1992.