

UCLA

UCLA Previously Published Works

Title

Two-Dimensional Compact Variational Mode Decomposition

Permalink

<https://escholarship.org/uc/item/1v77d6q9>

Journal

JOURNAL OF MATHEMATICAL IMAGING AND VISION, 58(2)

ISSN

0924-9907

Authors

Zosso, Dominique
Dragomiretskiy, Konstantin
Bertozzi, Andrea L
[et al.](#)

Publication Date

2017

DOI

10.1007/s10851-017-0710-z

Peer reviewed

Two-Dimensional Compact Variational Mode Decomposition

Spatially Compact and Spectrally Sparse Image Decomposition and Segmentation

Dominique Zosso^{1,2}  · Konstantin Dragomiretskiy¹ · Andrea L. Bertozzi¹ · Paul S. Weiss³

Received: 9 November 2015 / Accepted: 8 February 2017 / Published online: 28 February 2017
© Springer Science+Business Media New York 2017

Abstract Decomposing multidimensional signals, such as images, into spatially compact, potentially overlapping modes of essentially wavelike nature makes these components accessible for further downstream analysis. This decomposition enables space–frequency analysis, demodulation, estimation of local orientation, edge and corner detection, texture analysis, denoising, inpainting, or curvature estimation. Our model decomposes the input signal into modes with narrow Fourier bandwidth; to cope with sharp region boundaries, incompatible with narrow bandwidth, we introduce binary support functions that act as masks on the narrow-band mode for image recomposition. L^1 and TV terms promote sparsity and spatial compactness. Constraining the support functions to partitions of the signal domain, we effectively get an image segmentation model based on spectral homogeneity. By coupling several sub-modes together with a single support function, we are able to decompose an image into several crystal grains. Our efficient

algorithm is based on variable splitting and alternate direction optimization; we employ Merriman–Bence–Osher-like threshold dynamics to handle efficiently the motion by mean curvature of the support function boundaries under the sparsity promoting terms. The versatility and effectiveness of our proposed model is demonstrated on a broad variety of example images from different modalities. These demonstrations include the decomposition of images into overlapping modes with smooth or sharp boundaries, segmentation of images of crystal grains, and inpainting of damaged image regions through artifact detection.

Keywords Image decomposition · Image segmentation · Spatio-spectral decomposition · Microscopy · Crystal grains · Artifact detection · Threshold dynamics · Variational methods · Sparse time–frequency analysis

Mathematics Subject Classification 68U10

This work was supported by the Swiss National Science Foundation (SNF) under Grants PBELP2-137727 and P300P2-147778, the UC Lab Fees Research Grant 12-LR-236660, and the W. M. Keck Foundation.

✉ Dominique Zosso
dominique.zosso@montana.edu

¹ Department of Mathematics, University of California, Los Angeles (UCLA), 520 Portola Plaza, Box 951555, Los Angeles, CA 90095-1555, USA

² Present Address: Department of Mathematical Sciences, Montana State University, 2-214 Wilson Hall, Box 172400, Bozeman, MT 59717-2400, USA

³ California NanoSystems Institute (CNSI), Department of Chemistry and Biochemistry, Department of Materials Science and Engineering, University of California, Los Angeles (UCLA), 570 Westwood Plaza, Building 114, Los Angeles, CA 90095, USA

1 Introduction

In this paper, we are interested in decomposing images $f: \mathbb{R}^n \rightarrow \mathbb{R}$ into ensembles of constituent modes (components) that have specific directional and oscillatory characteristics.¹ Put simply, the goal is to retrieve a small number K of modes $u_k: \mathbb{R}^n \rightarrow \mathbb{R}$ that each have a very limited bandwidth around their characteristic center frequency ω_k . As with most partitioning models, K is assumed to be known a priori. However, in analogy to methods used in other contexts, one can think of ways to include a model complexity

¹ Throughout this paper, we will be using notation pertaining to images defined over continuous domains, albeit it is implicitly understood that numerical implementations will always make use of appropriate commonplace discretization and quantization.

term (prior) in our variational model to determine the “optimal” K from data, instead. These modes are called intrinsic mode functions (IMF) and can be seen as amplitude- and frequency-modulated (AM–FM) n -D signals, which essentially approximate “plane” waves. Such a mode can have limited spatial support, its local (instantaneous) frequency and amplitude vary smoothly, several modes can overlap in space, and together, the ensemble of modes should reconstruct the given input image up to noise and singular features.

Many fields use signal decomposition as a fundamental tool for quantitative and technical analysis. In remote sensing, decomposing images based on frequency content and signal priors, such as housing lattices and terrain structures, is useful for segmentation, identification, and classification [16]. In oceanography, a combination of baroclinic modes helps model density profiles of seasonal cycles and other geophysical phenomena such as thermal or solar variation [23, 64]. Similarly, in seismology, modes with differing frequency components help highlight different geological and stratigraphic information [35]. In holography, mode decomposition allows reducing speckle [46]. In the fields of energy and power engineering, mode decompositions are used for vibration analysis and fault detection, e.g., [27, 65]. Multivariate mode decomposition and mode entropy analysis are useful tools in neural data analysis [41].

In crystallography, because the crystal lattice exhibits multiple spatial periodicities, interpretable as a superposition of multiple different cosine waves, we wish to couple several “submodes” into a single phase. This coupled-mode decomposition enables robust estimates of mesoscopic properties such as crystal defects, rotations, and grain boundaries. Recent work in crystal orientation detection includes variational methods based on tensor maps in conjunction with a regularization scheme [19] and 2D synchrosqueezed transforms [75]. In nanoscale imaging, segmentation enables analyses and comparisons of surface regions of different structures as well as directed measurements of function, spectra, and dynamics [63, 67]. Ultimately, efficient segmentation will enable directed data acquisition and parsing acquisition time between different modalities to assemble and to converge complementary structural, functional, and other information.

Independent of the scientific discipline, sparse signal decomposition provides expansive utility and a more advanced podium from which to elucidate greater understanding.

1.1 Recent and Related Work

The problem is inspired by the one-dimensional empirical mode decomposition (EMD) algorithm [43] and its more recent derivatives, such as [24, 38–40, 48, 54–56, 61, 66, 71]. We are interested in the two-dimensional (2D) and higher n -

dimensional analogs and extensions of such decomposition problems. The 2D extension of EMD [53] similarly uses recursive sifting of 2D spatial signals by means of interpolating upper and lower envelopes, median envelopes, and thus extracting image components in different “frequency” bands. This 2D-EMD, however, suffers from the same drawbacks in robustness as the original EMD in extremal point finding, interpolation of envelopes, and stopping criteria imposed. More recent work, such as the Prony–Huang transform [59], has only partially improved on some of these drawbacks using modern variational and transform methods.

Classical decomposition methods include the discrete Fourier transform (DFT) and the continuous wavelet transform (CWT), where a fixed basis can yield a sparse representation for appropriate signals. Using more general bases or frames, extended methods such as matching pursuit decomposition (MP), method of frames, best orthogonal basis (BOB), and basis pursuit (BP) are more robust and, in principle, decompose a signal into a superposition of dictionary elements. Though these methods have had success with simple signals, they are still not fully robust to non-stationary waves and may require a large, redundant dictionary of elements.

More specific methods for directional image decomposition work by mostly rigid frames, decomposing the Fourier spectrum into fixed, mostly or strictly disjoint, (quasi-)orthogonal basis elements. Examples include Gabor filters [45], wavelets [13, 15, 47], curvelets [5], or shearlets [34, 44]. More recently, Riesz–Laplace wavelets have been proposed for multiresolution monogenic signal analysis [69]. All these methods, however, are not adaptive relative to the signal, and can attribute principal components of the image to different bands, as well as contain several different image components in the same band. Adaptivity and tuned sparsity concerns have been addressed through synchrosqueezed wavelet transforms [9, 14, 72, 74], where unimportant wavelet coefficients are removed by thresholding based on energy content. In pursuit of the same goal, the 2D empirical wavelet transform (EWT) [29, 30] decomposes an image by creating a more adaptive wavelet basis.

In previous work [17], Dragomiretskiy and Zosso defined a fully variational model for mode decomposition of 1D signals. The so-called *variational mode decomposition* (VMD) in 1D is essentially based on well-established concepts such as the 1D Hilbert transform and the analytic signal, heterodyne demodulation, and Wiener filtering. The goal of 1D-VMD is to decompose an input signal into a discrete number of subsignals (modes), where each mode has limited bandwidth in the spectral domain. In other words, one requires each mode $u_k : \mathbb{R} \rightarrow \mathbb{R}$ to be mostly compact (spectrally concentrated) around a center pulsation ω_k , which is to be determined along with the decomposition. In order to assess the bandwidth of a mode (its degree of spectral

concentration), the following scheme was proposed [17]: (1) For each mode u_k , compute the associated analytic signal by means of the Hilbert transform in order to obtain a unilateral frequency spectrum. (2) For each mode, shift the mode's frequency spectrum to "baseband," by mixing with an exponential tuned to the respective estimated center frequency. (3) Estimate the bandwidth through the H^1 smoothness (Dirichlet energy) of the demodulated signal. The resulting constrained variational problem is the following:

$$\begin{aligned} \min_{u_k: \mathbb{R} \rightarrow \mathbb{R}, \omega_k} & \left\{ \sum_k \alpha_k \left\| \partial_t \left[\left(\delta(\cdot) + \frac{j}{\pi} \right) * u_k(\cdot) \right] (t) e^{-j\omega_k t} \right\|_2^2 \right\} \\ \text{s.t. } \forall t \in \mathbb{R}: & \sum_k u_k(t) = f(t), \end{aligned} \quad (1)$$

where the norm is understood as

$$\| \cdot \|_2^2 := \int_{\mathbb{R}} | \cdot |^2 dt$$

and $*$ denotes convolution.

In [17], it was shown that this variational model can be minimized efficiently and it outperforms empirical mode decomposition algorithms in various respects, most notably regarding noise robustness and mode cleanliness.

1.2 Proposed Method

In this paper, we propose a natural two-dimensional (or higher-dimensional) extension of the (1D) variational mode decomposition algorithm [17] in the context of image segmentation and directional decomposition. The 2D-VMD algorithm is a non-recursive, fully adaptive, variational method that sparsely decomposes images in a mathematically well-founded manner.

Here, we are interested in making the advantages of the variational model accessible for the 2D case (and higher dimensions equally so). The first order of business is thus to generalize the 1D-VMD model to the multidimensional case, as sketched in [18]. Second, we want to address an intrinsic conflict of the VMD model, namely the inverse relation between spatial and frequency support: in 1D VMD, it was noted that the algorithm had difficulties whenever signals exhibited sudden onset and amplitude changes, since these effectively represent a violation of the assumptions of Bedrosian's theorem, a key element of the VMD model. In this work, we address this issue by further introducing a separate amplitude function that masks the underlying mode spatially, which allows decoupling spatial from spectral support. In 2D, this approach allows extraction of modes with sharp boundaries. We then introduce various priors on the shape of the amplitude function. Requiring the amplitude

function to be binary and penalizing its total variation regularize the mode boundaries. Restricting the ensemble of amplitude functions associated with the various modes to the probability simplex at each pixel leads to non-overlapping modes, effectively segmenting the image. Coupling several modes to share a single support function further allows extraction of multi-wave textures, such as lattice patterns.

The remainder of this paper is organized as follows. In Sect. 2, we provide a short definition and description of the 1D Hilbert transform and the directional Hilbert transform, which is one of its generalizations to higher dimensions. We formulate our proposed 2D-VMD model and present a strategy to solve it numerically. We then introduce a separate term for compact spatial support in Sect. 3, by defining binary support functions, and provide a detailed algorithm for this second, compact decomposition model with potentially overlapping modes. In Sect. 4, we can then restrict the support of the modes to form a partition of the image domain, resulting in spectrum-based image segmentation. We provide a complete, third algorithm for this modified problem. Then, in Sect. 5, we couple several submodes together (joint support) to model domains with non-trivial spectral distributions. Finally, we include an artifact detection term to eliminate outlier pixels, as described in Sect. 6. Decomposition and segmentation results on synthetic and real data are provided in Sect. 7, and we discuss the implications of and prospects for this work in Sect. 8.

2 Two-Dimensional Variational Mode Decomposition

We design the 2D model analogously to its 1D predecessor, minimizing the constituent subsignals' bandwidth while maintaining data fidelity. While derivatives in higher dimensions are simply generalized by gradients, and modulation is also straightforward, the generalization of the analytic signal is less obvious. To complete the analogy, we must first define the appropriate "analytic signal"-equivalent in the n -D context.

2.1 n -D Hilbert Transform /Analytic Signal

In the 1D time domain, the analytic signal is achieved by adding the Hilbert transformed copy of the original signal $f: \mathbb{R} \rightarrow \mathbb{R}$ as imaginary part [25]:

$$\begin{aligned} f_{AS}: \mathbb{R} & \rightarrow \mathbb{C} \\ t & \mapsto f(t) + j\mathcal{H}\{f\}(t), \end{aligned} \quad (2)$$

where $j^2 = -1$, and the 1D Hilbert transform is defined as:

$$\mathcal{H}\{f\}(t) := \left\{ \frac{1}{\pi \cdot} * f(\cdot) \right\} (t) = \frac{1}{\pi} \text{p.v.} \int_{\mathbb{R}} \frac{f(s)}{t-s} ds, \tag{3}$$

where $*$ denotes convolution. We note that the real signal is recovered simply by taking the real component of the analytic signal.

In the spectral domain, this definition of analytic signal corresponds to suppressing the negative frequencies, thus giving it a unilateral spectrum:

$$\hat{f}_{AS} : \mathbb{R} \rightarrow \mathbb{C} \\ \omega \mapsto \begin{cases} 2\hat{f}(\omega), & \text{if } \omega > 0, \\ \hat{f}(\omega), & \text{if } \omega = 0, \\ 0, & \text{if } \omega < 0, \end{cases} \tag{4}$$

where

$$\hat{f}(\omega) := \mathcal{F}\{f(\cdot)\}(\omega) = 1/\sqrt{2\pi} \int_{\mathbb{R}} f(t)e^{-j\omega t} dt$$

is the unitary Fourier transform in 1D.

Single-sidedness of the analytic signal spectrum was the key property motivating its use in the 1D case, since this property allowed for easy frequency shifting to baseband by complex exponential mixing. Therefore, to mimic this spectral property in 2D, one half-plane of the frequency domain must effectively be set to zero²; this half-plane is chosen relative to a vector, in our case to ω_k . Thus, the 2D analytic signal of interest can first be defined in the frequency domain by generalizing the concept of half-space spectrum suppression:

$$\hat{f}_{AS} : \mathbb{R}^n \rightarrow \mathbb{C} \\ \omega \mapsto \begin{cases} 2\hat{f}(\omega), & \text{if } \langle \omega, \omega_k \rangle > 0, \\ \hat{f}(\omega), & \text{if } \langle \omega, \omega_k \rangle = 0, \\ 0, & \text{if } \langle \omega, \omega_k \rangle < 0, \end{cases} \tag{5}$$

or, shorthand,

$$\hat{f}_{AS}(\omega) = (1 + \text{sgn}(\langle \omega, \omega_k \rangle))\hat{f}(\omega), \tag{6}$$

where the n -D Fourier transform is defined as

$$\hat{f}(\omega) := \mathcal{F}\{f(\cdot)\}(\omega) = (2\pi)^{-n/2} \int_{\mathbb{R}^n} f(\mathbf{x})e^{-j\langle \omega, \mathbf{x} \rangle} d\mathbf{x}.$$

The 2D analytic signal in the time domain with the aforementioned Fourier property is given in [4]. In particular, we note that the underlying higher-dimensional generalization of Hilbert transform employed here is known as the directional

Hilbert transform [33]. It is easy to see how the generalized analytic signal reduces to the classical definition in 1D.

2.2 n -D VMD Functional

We are now able to put all the generalized VMD ingredients together to define the two-dimensional extension of variational mode decomposition. The functional to be minimized, stemming from this definition of n -D analytic signal, is:

$$\min_{u_k : \mathbb{R}^n \rightarrow \mathbb{R}, \omega_k \in \mathbb{R}^n} \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\mathbf{x})e^{-j\langle \omega_k, \mathbf{x} \rangle} \right] \right\|_2^2 \right\} \\ \text{s.t. } \forall \mathbf{x} \in \mathbb{R}^n : \sum_k u_k(\mathbf{x}) = f(\mathbf{x}), \tag{7}$$

where $u_{AS,k}$ denotes the generalized analytic signal obtained from the mode u_k according to (5) using its associated center frequency ω_k . The coefficient α_k allows to convey different weights to the k modes, a degree of freedom that we do not currently explore, yet. Here, the squared norm is understood as

$$\| \cdot \|_2^2 := \int_{\mathbb{R}^n} | \cdot (\mathbf{x}) |^2 d\mathbf{x},$$

a generalization to higher dimensions of the previous usage in 1D. Note that the argument can be complex and/or vector-valued, in which case the term $| \cdot (\mathbf{x}) |^2$ refers to the squared magnitude.

The model thus minimizes the Dirichlet energy of the modes after half-space spectrum suppression ($u_k \rightarrow u_{AS,k}$) and demodulation to baseband ($e^{-j\langle \omega_k, \mathbf{x} \rangle}$), subject to collective signal fidelity. This model specifically includes the desired two-dimensional case $n = 2$, and reduces to the earlier 1D-VMD for $n = 1$.

Analogous to the 1D-VMD model, the reconstruction constraint is addressed through the introduction of a quadratic penalty and Lagrangian multiplier (the *augmented Lagrangian*, AL, method), and we proceed by a scheme reminiscent of alternate direction minimization (ADMM) for optimization [2, 17, 52]. It is to note that while the model is convex in each argument $u_k : \mathbb{R}^n \rightarrow \mathbb{R}$ and $\omega_k \in \mathbb{R}^n$, separately, the objective is not convex, jointly.³

2.3 Augmented Lagrangian and Optimization

To render the constrained minimization problem (7) unconstrained, we include both a quadratic penalty and a Lagrangian multiplier to enforce the fidelity constraint. We thus define the augmented Lagrangian:

² Similarly, in higher dimensions, a half-space of the frequency domain needs to be suppressed.

³ To be more precise, the objective is convex in ω_k if we consider the analytic signal construction for $u_{AS,k}$ fixed while optimizing for ω_k .

$$\mathcal{L}(\{u_k\}, \{\omega_k\}, \lambda) := \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\mathbf{x}) e^{-j\langle \omega_k, \mathbf{x} \rangle} \right] \right\|_2^2 + \left\| f(\mathbf{x}) - \sum_k u_k(\mathbf{x}) \right\|_2^2 + \left\langle \lambda(\mathbf{x}), f(\mathbf{x}) - \sum_k u_k(\mathbf{x}) \right\rangle. \tag{8}$$

where $\lambda: \mathbb{R}^n \rightarrow \mathbb{R}$ is the Lagrangian multiplier. We can now consider the unconstrained saddle point problem instead of (7):

$$\min_{u_k: \mathbb{R}^n \rightarrow \mathbb{R}, \omega_k \in \mathbb{R}^n} \max_{\lambda: \mathbb{R}^n \rightarrow \mathbb{R}} \mathcal{L}(\{u_k\}, \{\omega_k\}, \lambda). \tag{9}$$

The solution thereof is now found as the saddle point of the augmented Lagrangian \mathcal{L} in a sequence of iterative suboptimizations similar to Gauss–Seidel and the alternate direction method of multipliers (ADMM) for convex problems [2,37,57]. The idea is to iterate the following sequence of variable updates:

$$u_k^{t+1} \leftarrow \arg \min_{u_k: \mathbb{R}^n \rightarrow \mathbb{R}} \mathcal{L}(\{u_i^{t+1}\}, u_k, \{u_i^t\}, \{\omega_i^t\}, \lambda^t) \tag{10a}$$

$$\omega_k^{t+1} \leftarrow \arg \min_{\omega_k \in \mathbb{R}^n} \mathcal{L}(\{u_i^{t+1}\}, \{\omega_i^{t+1}\}, \omega_k, \{\omega_i^t\}, \lambda^t) \tag{10b}$$

$$\lambda^{t+1} \leftarrow \lambda^t + \tau \left(f - \sum_k u_k^{t+1} \right) \tag{10c}$$

for some $\tau \geq 0$.

For simplified notation while considering the subminimization problems (10a) and (10b) in the following paragraphs, we incorporate the Lagrangian multiplier term λ into the quadratic penalty term. Simply by completing the square, we rewrite the objective expression equivalently as follows:

$$\begin{aligned} \mathcal{L}(\{u_k\}, \{\omega_k\}, \lambda) &= \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\mathbf{x}) e^{-j\langle \omega_k, \mathbf{x} \rangle} \right] \right\|_2^2 \\ &+ \left\| f(\mathbf{x}) - \sum_k u_k(\mathbf{x}) + \frac{\lambda(\mathbf{x})}{2} \right\|_2^2 \\ &- \left\| \frac{\lambda(\mathbf{x})}{2} \right\|_2^2 \end{aligned} \tag{11}$$

Since updates happen sequentially, we note that when updating u_k preceding modes $i < k$ will have been updated to u_i^{t+1} already, while subsequent modes $i > k$ still only have u_i^t available. We will omit iteration superscripts in the following to unclutter the notation, and simply refer to u_i , for $i \neq k$, instead; but it is understood that we always use the most recent estimate of a variable.

2.4 Minimization w.r.t. the Modes u_k

The relevant update problem derived from (10a) and (11) is

$$u_k^{t+1} = \arg \min_{u_k: \mathbb{R}^n \rightarrow \mathbb{R}} \left\{ \alpha_k \left\| \nabla \left[u_{AS,k}(\mathbf{x}) e^{-j\langle \omega_k, \mathbf{x} \rangle} \right] \right\|_2^2 + \left\| f(\mathbf{x}) - \sum_i u_i(\mathbf{x}) + \frac{\lambda(\mathbf{x})}{2} \right\|_2^2 \right\} \tag{12}$$

Since we are dealing with L^2 -norms, we can make use of the L^2 Fourier isometry and rewrite the subminimization problem in spectral domain (thus implicitly assuming periodic boundary conditions):

$$\hat{u}_k^{t+1} = \arg \min_{\hat{u}_k | u_k: \mathbb{R}^n \rightarrow \mathbb{R}} \left\{ \alpha_k \left\| j \omega \left[\hat{u}_{AS,k}(\omega + \omega_k) \right] \right\|_2^2 + \left\| \hat{f}(\omega) - \sum_i \hat{u}_i(\omega) + \frac{\hat{\lambda}(\omega)}{2} \right\|_2^2 \right\}. \tag{13}$$

Here, $\hat{u}_k | u_k: \mathbb{R}^n \rightarrow \mathbb{R}$ restricts the minimization to Hermitian spectra, $\hat{u}_k(\omega) = -\hat{u}_k(-\omega)$, corresponding to real-valued modes u_k . The ω_k term in the spectrum of the analytic signal is due to the modulation with the complex exponential and justified by the well-known transform pair:

$$f(\mathbf{x}) e^{-j\langle \omega_0, \mathbf{x} \rangle} \xleftrightarrow{\mathcal{F}} \hat{f}(\omega) * \delta(\omega + \omega_0) = \hat{f}(\omega + \omega_0), \tag{14}$$

where δ is the Dirac distribution and $*$ denotes convolution. Thus, multiplying an analytic signal with a pure exponential results in simple frequency shifting. Further, we can push the frequency shift out of the analytic signal spectrum through a change of variables, to obtain:

$$\hat{u}_k^{t+1} = \arg \min_{\hat{u}_k | u_k: \mathbb{R}^n \rightarrow \mathbb{R}} \left\{ \alpha_k \left\| j(\omega - \omega_k) \left[\hat{u}_{AS,k}(\omega) \right] \right\|_2^2 + \left\| \hat{f}(\omega) - \sum_i \hat{u}_i(\omega) + \frac{\hat{\lambda}(\omega)}{2} \right\|_2^2 \right\}. \tag{15}$$

We now plug in the spectral definition of the n -D analytic signal (5),

$$\hat{u}_{AS,k}(\omega) = (1 + \text{sgn}(\langle \omega, \omega_k \rangle)) \hat{u}_k(\omega).$$

Also, the second term contains spectra of real signals only, which makes it Hermitian and lets us split the norm into two equal parts. To this end, let

$$\Omega_k \subset \mathbb{R}^n: \Omega_k := \{\omega \mid \langle \omega, \omega_k \rangle \geq 0\}$$

denote the frequency domain half-space to which the n -D analytic signal is restricted. We rewrite both terms as integrals

over these frequency domain half-spaces:

$$\hat{u}_k^{t+1} = \arg \min_{\hat{u}_k | u_k: \mathbb{R}^n \rightarrow \mathbb{R}} \left\{ 2\alpha_k \int_{\Omega_k} |\omega - \omega_k|^2 |\hat{u}_k(\omega)|^2 d\omega + \int_{\Omega_k} \left| \hat{f}(\omega) - \sum_i \hat{u}_i(\omega) + \frac{\hat{\lambda}(\omega)}{2} \right|^2 d\omega \right\}. \tag{16}$$

This subminimization problem is now solved by letting the first variation w.r.t. \hat{u}_k vanish.⁴ The optimal mode spectrum thus satisfies:

$$0 = 2\alpha_k |\omega - \omega_k|^2 \hat{u}_k - \left(\hat{f}(\omega) - \sum_i \hat{u}_i(\omega) + \frac{\hat{\lambda}(\omega)}{2} \right), \quad \forall \omega \in \Omega_k. \tag{17}$$

With this optimality condition, solving for \hat{u}_k yields the following Wiener filter update:

$$\hat{u}_k^{t+1}(\omega) = \left(\hat{f}(\omega) - \sum_{i \neq k} \hat{u}_i(\omega) + \frac{\hat{\lambda}(\omega)}{2} \right) \frac{1}{1 + 2\alpha_k |\omega - \omega_k|^2}, \quad \forall \omega \in \Omega_k. \tag{18}$$

The full spectrum \hat{u}_k^{t+1} can then be obtained by symmetric (Hermitian) completion. Equivalently, we can decide to update the half-space analytic signal of the mode, $\hat{u}_{AS,k}^{t+1}$, on the entire frequency domain, instead:

$$\hat{u}_{AS,k}^{t+1}(\omega) = \left(\hat{f}(\omega) - \sum_{i \neq k} \hat{u}_i(\omega) + \frac{\hat{\lambda}(\omega)}{2} \right) \frac{1 + \text{sgn}(\langle \omega, \omega_k \rangle)}{1 + 2\alpha_k |\omega - \omega_k|^2}, \quad \forall \omega \in \mathbb{R}^n, \tag{19}$$

from which the actual mode estimate is recovered as the real part after inverse Fourier transform. The term in parentheses is the signal’s k -th residual, where $\hat{f}(\omega) - \sum_{i \neq k} \hat{u}_i(\omega)$ is the explicit current residual, and $\hat{\lambda}$ accumulates the reconstruction error over iterations (see below). The second term is identified as a frequency filter tuned to the current estimate of the mode’s center pulsation, ω_k , whose bandwidth is controlled by the parameter α_k .

2.5 Minimization w.r.t. the Center Frequencies ω_k

Optimizing for ω_k is even simpler. Indeed, the respective update goal derived from (10b) and (11) is

$$\omega_k^{t+1} = \arg \min_{\omega_k \in \mathbb{R}^n} \left\{ \alpha_k \left\| \nabla \left[u_{AS,k}(\mathbf{x}) e^{-j(\omega_k, \mathbf{x})} \right] \right\|_2^2 \right\}. \tag{20}$$

⁴ Note that the spectrum of u_k is complex-valued, so the process of “taking the first variation” is not self-evident. However, the functional is analytic in \hat{u}_k and complex-valued equivalents to the standard derivatives do indeed apply.

Again we may consider the equivalent problem in the Fourier domain:

$$\omega_k^{t+1} = \arg \min_{\omega_k \in \mathbb{R}^n} \left\{ \alpha_k \left\| j(\omega - \omega_k)(1 + \text{sgn}(\langle \omega_k, \omega \rangle)) \hat{u}_k(\omega) \right\|_2^2 \right\} = \arg \min_{\omega_k \in \mathbb{R}^n} \left\{ 4\alpha_k \int_{\Omega_k} |\omega - \omega_k|^2 |\hat{u}_k(\omega)|^2 d\omega \right\}. \tag{21}$$

The minimization is solved by letting the first variation w.r.t. ω_k vanish, leading to:

$$\int_{\Omega_k} (\omega - \omega_k^{t+1}) |\hat{u}_k(\omega)|^2 d\omega = 0. \tag{22}$$

The resulting solutions are the centers of gravity of the modes’ power spectra, $|\hat{u}_k(\omega)|^2$, restricted to the half-space Ω_k :

$$\omega_k^{t+1} = \frac{\int_{\Omega_k} \omega |\hat{u}_k(\omega)|^2 d\omega}{\int_{\Omega_k} |\hat{u}_k(\omega)|^2 d\omega} = \frac{\int_{\mathbb{R}^n} \omega |\hat{u}_{AS,k}(\omega)|^2 d\omega}{\int_{\mathbb{R}^n} |\hat{u}_{AS,k}(\omega)|^2 d\omega}, \tag{23}$$

where the second form is given for implementation purposes, based on the analytic signal spectrum and involving the entire frequency domain.

2.5.1 Maximization w.r.t. the Lagrangian Multiplier λ

Maximizing the λ is the simplest step in the algorithm. The first variation for λ is just the data reconstruction error, $f(\omega) - \sum_k u_k^{t+1}(\omega)$. We use a standard gradient ascent with some fixed time step $\tau \geq 0$ to achieve this maximization:

$$\lambda^{t+1}(\mathbf{x}) = \lambda^t(\mathbf{x}) + \tau \left(f(\mathbf{x}) - \sum_k u_k^{t+1}(\mathbf{x}) \right). \tag{24}$$

It is important to note that choosing $\tau = 0$ eliminates the Lagrangian update and thus reduces the algorithm to the penalty method for data-fidelity purposes. Doing so is useful when exact data fidelity is not appropriate, such as in (high) noise scenarios, where reconstruction error actually allows capturing noise separately. Effectively the proposed model is capable of denoising.

Note also that the linearity of the Euler–Lagrange equation allows an impartial choice in which space to update the Lagrangian multiplier, either in the time domain or in the frequency domain. In our implementation, we perform our dual ascent update in the frequency domain, since the other appearance of the Lagrangian multiplier in (19) is in spectral terms, as well. Thus:

$$\hat{\lambda}^{t+1}(\omega) = \hat{\lambda}^t(\omega) + \tau \left(\hat{f}(\omega) - \sum_k \hat{u}_k^{t+1}(\omega) \right). \tag{25}$$

2.5.2 Complete 2D VMD Algorithm

The entire proposed algorithm for the 2D-VMD functional optimization problem (7) is summarized in Algorithm 1. Variables are trivially initialized at 0, except for the center frequencies, ω_k , for which smart initialization is of higher importance; initial ω_k^0 can, e.g., be spread randomly, radially uniform, or initialized by user input. Further, we choose to assess convergence in terms of the normalized rate of change of the modes. Typical thresholds $\epsilon > 0$ range in orders of magnitude from 10^{-4} (fast) down to 10^{-7} (very accurate).

Algorithm 1 2D-VMD

Input: signal $f(\mathbf{x})$, number of modes K , parameters α_k, τ, ϵ .

Output: modes $u_k(\mathbf{x})$, center frequencies ω_k .

Initialize $\{\omega_k^0\}, \{\hat{u}_k^0\} \leftarrow 0, \hat{\lambda}^0 \leftarrow 0, n \leftarrow 0$

repeat

$n \leftarrow n + 1$

for $k = 1 : K$ **do**

Create 2D mask for analytic signal Fourier multiplier:

$$\mathcal{H}_k^{t+1}(\omega) \leftarrow 1 + \text{sgn}(\langle \omega_k^t, \omega \rangle)$$

Update $\hat{u}_{AS,k}$:

$$\hat{u}_{AS,k}^{t+1}(\omega) \leftarrow \mathcal{H}_k^{t+1}(\omega) \left[\frac{\hat{f}(\omega) - \sum_{i < k} \hat{u}_i^{t+1}(\omega) - \sum_{i > k} \hat{u}_i^t(\omega) + \frac{\hat{\lambda}^t(\omega)}{2}}{1 + 2\alpha_k |\omega - \omega_k^t|^2} \right]$$

Update ω_k :

$$\omega_k^{t+1} \leftarrow \frac{\int_{\mathbb{R}^2} \omega |\hat{u}_{AS,k}^{t+1}(\omega)|^2 d\omega}{\int_{\mathbb{R}^2} |\hat{u}_{AS,k}^{t+1}(\omega)|^2 d\omega}$$

Retrieve u_k :

$$u_k^{t+1}(\mathbf{x}) \leftarrow \Re \left(\mathcal{F}^{-1} \left\{ \hat{u}_{AS,k}^{t+1}(\omega) \right\} \right)$$

end for

Dual ascent (optional):

$$\hat{\lambda}^{t+1}(\omega) \leftarrow \hat{\lambda}^t(\omega) + \tau \left(\hat{f}(\omega) - \sum_k \hat{u}_k^{t+1}(\omega) \right)$$

until convergence: $\sum_k \|\hat{u}_k^{t+1} - \hat{u}_k^t\|_2^2 / \|\hat{u}_k^t\|_2^2 < \epsilon$.

An example of image decomposition achieved with 2D VMD according to Algorithm 1 is shown in Fig. 1.

3 VMD with Compact Spatial Support

A main assumption regarding the intrinsic mode functions considered so far is that their amplitude (spatially) varies much more slowly than the wavelength of the carrier. Indeed, IMFs can be defined as signals (in time or space) that are both amplitude- and frequency-modulated [14]. In [17], we have

defined the *total practical IMF bandwidth* of such an AM-FM signal, as an extension to Carson’s rule for FM signal bandwidth [6]:

$$BW_{AM-FM} := 2(\Delta f + f_{FM} + f_{AM}), \tag{26}$$

where Δf and f_{FM} represent the frequency swing and modulation bandwidth, respectively, of the FM part, while f_{AM} denotes the bandwidth of the amplitude modulation. The last, AM bandwidth, conflicts with signals composed of modes having sudden signal onset, in particular those with compact spatial support. Indeed, this inverse relation between spatial and spectral compactness is well known and stated by the Heisenberg uncertainty principle.

3.1 Introducing Binary Support Functions A_k

To make our “modes have limited bandwidth” prior compatible with signals of limited spatial support, it is thus necessary to deal with the spatial and spectral compactness of the modes, separately. To this end, we introduce a binary support function for each mode, in order to capture the signal onset and offset disconnected from the smooth AM-FM modulations.

We consider signals and modes $f, u_k : \mathbb{R}^n \rightarrow \mathbb{R}$ (thus including both the 1D-VMD and higher-dimensional signals such as 2D-VMD stated above). Let

$$A_k : \mathbb{R}^n \rightarrow \{0, 1\}$$

denote the binary support functions for each mode u_k . The mode decomposition problem can then formally be stated as

$$\text{find } u_k, A_k \text{ s.t. } f = \sum_k A_k \cdot u_k,$$

i.e., we want the modes u_k , now masked by their binary support function A_k , to reproduce collectively the given input signal. Note that the modes u_k can extend arbitrarily into their inactive regions where $A_k = 0$; in particular, they can decay smoothly or oscillate *ad infinitum*, thus keeping small spectral bandwidth.

3.2 Sparsity Promoting VMD Functional

It is important to introduce sparsity promoting regularity constraints on the support function to achieve reasonable compact local support. Here, we consider both total variation (TV) and L^1 penalties on A_k , thus effectively penalizing support area and boundary length (through the co-area formula).

We incorporate the binary support functions A_k and their regularizers in the n -D VMD functional as follows:




Fig. 1 Synthetic overlapping texture **a** Input image f . **b** 2D-VMD reconstruction $\sum_k u_k$. **c** Compactly supported 2D-TV-VMD reconstruction $\sum_k A_k u_k$. **d** Support boundaries overlaid onto original image.

e 2D-VMD modes u_k . **f** 2D-TV-VMD modes u_k . **g** Detected supports A_k . **h** Masked modes $A_k u_k$. See Sect. 7.1 in text for explanation and discussion

$$\begin{aligned}
 & \min_{u_k: \mathbb{R}^n \rightarrow \mathbb{R}, A_k: \mathbb{R}^n \rightarrow \{0,1\}, \omega_k \in \mathbb{R}^n} \\
 & \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\mathbf{x}) e^{-j(\omega_k, \mathbf{x})} \right] \right\|_2^2 \right. \\
 & \left. + \beta_k \|A_k\|_1 + \gamma_k \text{TV}(A_k) \right\} \\
 & \text{s.t. } \forall \mathbf{x} \in \mathbb{R}^n: \sum_k A_k(\mathbf{x}) u_k(\mathbf{x}) = f(\mathbf{x}), \tag{27}
 \end{aligned}$$

where total variation is defined as

$$\text{TV}(A_k) := \sup_{|\phi| \leq 1} \langle A_k, \text{div} \phi \rangle, \quad \text{for } \phi: \mathbb{R}^n \rightarrow \mathbb{R}^n,$$

and the L^1 -norm is understood as

$$\|A_k\|_1 := \int_{\mathbb{R}^n} |A_k(\mathbf{x})| d\mathbf{x}.$$

The L^1 and TV penalties on A_k ensure that an individual mode is only active in places where it is “sufficiently justified” (i.e., the increased data fidelity outweighs the incurred friction cost), and represent the prior that modes have limited spatial support and regular outlines.

3.3 Model “Relaxation”

Due to the introduction of the binary support functions A_k in the fidelity constraint, and the L^1 - and TV-based prior terms, the functional is no longer directly translatable to the spectral domain. We want to restore spectral solvability of the modes u_k . Currently, the masks A_k prevent this, since in the quadratic penalty addressing the reconstruction constraint, the spatial multiplication translates to spectral convolution. To avoid this, we introduce a splitting of the modes $u_k = v_k$, and apply spectral bandwidth penalty and reconstruction over the separate copies:

$$\begin{aligned}
 & \min_{u_k: \mathbb{R}^n \rightarrow \mathbb{R}, A_k: \mathbb{R}^n \rightarrow \{0,1\}, \omega_k \in \mathbb{R}^n} \\
 & \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\mathbf{x}) e^{-j(\omega_k, \mathbf{x})} \right] \right\|_2^2 \right. \\
 & \left. + \beta_k \|A_k\|_1 + \gamma_k \text{TV}(A_k) \right\} \\
 & \text{s.t. } \forall \mathbf{x} \in \mathbb{R}^n: \begin{cases} u_k(\mathbf{x}) = v_k(\mathbf{x}), \\ \sum_k A_k(\mathbf{x}) v_k(\mathbf{x}) = f(\mathbf{x}). \end{cases} \tag{28}
 \end{aligned}$$

While (28) is equivalent to (27), the important difference for our purposes will appear in the coordinate-wise optimization steps. Similar splitting techniques [11, 12, 31] have been applied to L^1 -based and related optimization problems with great success, such as [22, 32]. The splitting constraint can

be addressed, e.g., with a quadratic penalty (proximal splitting, [12]), or using an augmented Lagrangian [31]. The full saddle point functional (augmented Lagrangian) incorporating both equality constraints through quadratic penalty and Lagrangian multipliers, in analogy to (8), is as follows:

$$\begin{aligned}
 & \mathcal{L}(\{u_k\}, \{v_k\}, \{A_k\}, \{\omega_k\}, \lambda, \{\lambda_k\}) \\
 & := \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\mathbf{x}) e^{-j(\omega_k, \mathbf{x})} \right] \right\|_2^2 \right. \\
 & \quad + \beta_k \|A_k\|_1 + \gamma_k \text{TV}(A_k) \\
 & \quad + \rho \left\| f(\mathbf{x}) - \sum_k A_k(\mathbf{x}) v_k(\mathbf{x}) \right\|_2^2 \\
 & \quad + \langle \lambda(\mathbf{x}), f(\mathbf{x}) - \sum_k A_k(\mathbf{x}) v_k(\mathbf{x}) \rangle \\
 & \quad + \sum_k \rho_k \|u_k(\mathbf{x}) - v_k(\mathbf{x})\|_2^2 \\
 & \quad \left. + \langle \lambda_k(\mathbf{x}), u_k(\mathbf{x}) - v_k(\mathbf{x}) \rangle \right\}, \tag{29}
 \end{aligned}$$

where λ_k are the Lagrangian multipliers associated with the K equality constraints $u_k = v_k$, and ρ, ρ_k are parameters weighting the different quadratic penalties.

3.4 n-D-TV-VMD Minimization

We now proceed to solve the constrained, sparsity promoting n -D VMD functional (27) through its augmented Lagrangian (29). Consider the following saddle point problem:

$$\min_{u_k, v_k: \mathbb{R}^n \rightarrow \mathbb{R}, A_k: \mathbb{R}^n \rightarrow \{0,1\}, \omega_k \in \mathbb{R}^n} \max_{\lambda, \lambda_k: \mathbb{R}^n \rightarrow \mathbb{R}} \left\{ \mathcal{L}(\{u_k\}, \{v_k\}, \{A_k\}, \{\omega_k\}, \lambda, \{\lambda_k\}) \right\}. \tag{30}$$

This saddle point problem is an extended version of the 2D VMD saddle point problem (9) (there, without spatial sparsity promoting terms) and is again efficiently solved through alternate direction minimization and dual ascent scheme similar to ADMM:

$$\begin{aligned}
 u_k^{t+1} & \leftarrow \arg \min_{u_k: \mathbb{R}^n \rightarrow \mathbb{R}} \\
 & \mathcal{L} \left(\left\{ u_{i < k}^{t+1} \right\}, u_k, \left\{ u_{i > k}^t \right\}, \left\{ v_i^t \right\}, \left\{ A_i^t \right\}, \right. \\
 & \left. \left\{ \omega_i^t \right\}, \lambda^t, \left\{ \lambda_i^t \right\} \right) \tag{31a}
 \end{aligned}$$

$$\begin{aligned}
 v_k^{t+1} & \leftarrow \arg \min_{v_k: \mathbb{R}^n \rightarrow \mathbb{R}} \\
 & \mathcal{L} \left(\left\{ u_i^{t+1} \right\}, \left\{ v_{i < k}^{t+1} \right\}, v_k, \left\{ v_{i > k}^t \right\}, \left\{ A_i^t \right\}, \right. \\
 & \left. \left\{ \omega_i^t \right\}, \lambda^t, \left\{ \lambda_i^t \right\} \right) \tag{31b}
 \end{aligned}$$

$$A_k^{t+1} \leftarrow \arg \min_{A_k: \mathbb{R}^n \rightarrow \{0,1\}}$$

$$\mathcal{L} \left(\left\{ u_i^{t+1} \right\}, \left\{ v_i^{t+1} \right\}, \left\{ A_{i < k}^{t+1} \right\}, A_k, \left\{ A_{i > k}^t \right\}, \left\{ \omega_i^t \right\}, \lambda^t, \left\{ \lambda_i^t \right\} \right) \tag{31c}$$

$$\omega_k^{t+1} \leftarrow \arg \min_{\omega_k \in \mathbb{R}^n} \mathcal{L} \left(\left\{ u_i^{t+1} \right\}, \left\{ v_i^{t+1} \right\}, \left\{ A_{i < k}^{t+1} \right\}, \left\{ \omega_{i < k}^{t+1} \right\}, \omega_k, \left\{ \omega_{i > k}^t \right\}, \lambda^t, \left\{ \lambda_i^t \right\} \right) \tag{31d}$$

$$\lambda^{t+1} \leftarrow \lambda^t + \tau \left(f - \sum A_k^{t+1} v_k^{t+1} \right) \tag{31e}$$

$$\lambda_k^{t+1} \leftarrow \lambda_k^t + \tau_k \left(u_k^{t+1} - v_k^{t+1} \right) \tag{31f}$$

where $\tau_k \geq 0$ is the update step size for the Lagrange multipliers associated with the newly introduced mode splitting constraint. We provide details on the individual subminimization problems in the following paragraphs, some of which closely relate to the steps given in Sect. 2.

3.4.1 n-D-TV-VMD Subminimization w.r.t. u_k

The relevant minimization problem (31a) with respect to the modes u_k reads

$$u_k^{t+1} = \arg \min_{u_k: \mathbb{R}^n \rightarrow \mathbb{R}} \left\{ \alpha_k \left\| \nabla \left[u_{AS,k}(\mathbf{x}) e^{-j(\omega_k, \mathbf{x})} \right] \right\|_2^2 + \rho_k \left\| u_k(\mathbf{x}) - v_k(\mathbf{x}) + \frac{\lambda_k(\mathbf{x})}{\rho_k} \right\|_2^2 \right\}. \tag{32}$$

In full analogy to the problem without spatial sparsity terms, (12), the update is most easily computed in spectral domain, like (18). Unsurprisingly, the update rule on the frequency half-space $\Omega_k = \{ \omega \mid \langle \omega, \omega_k \rangle \geq 0 \}$ is found to be:

$$\hat{u}_k^{t+1}(\omega) = (\rho_k \hat{v}_k - \hat{\lambda}_k) \frac{1}{\rho_k + 2\alpha_k |\omega - \omega_k|^2}, \quad \forall \omega \in \Omega_k. \tag{33}$$

From this half-space update, the full spectrum can again be obtained by Hermitian completion, or by updating the mode’s half-space analytic signal instead:

$$\hat{u}_{AS,k}^{t+1}(\omega) = (\rho_k \hat{v}_k - \hat{\lambda}_k) \frac{1 + \text{sgn}(\langle \omega - \omega_k \rangle)}{\rho_k + 2\alpha_k |\omega - \omega_k|^2}. \tag{34}$$

3.4.2 n-D-TV-VMD Subminimization w.r.t. v_k

The update (31b) of v_k reduces to the following minimization problem:

$$v_k^{t+1} = \arg \min_{v_k: \mathbb{R}^n \rightarrow \mathbb{R}} \left\{ \rho \left\| f(\mathbf{x}) - \sum A_i(\mathbf{x}) v_i(\mathbf{x}) + \frac{\lambda(\mathbf{x})}{\rho} \right\|_2^2 + \rho_k \left\| u_k(\mathbf{x}) - v_k(\mathbf{x}) + \frac{\lambda_k(\mathbf{x})}{\rho_k} \right\|_2^2 \right\} \tag{35}$$

This problem admits the following pointwise optimality conditions:

$$-\rho A_k(\mathbf{x}) \left(f(\mathbf{x}) - \sum A_i(\mathbf{x}) v_i(\mathbf{x}) + \frac{\lambda(\mathbf{x})}{\rho} \right) - \rho_k \left(u_k(\mathbf{x}) - v_k(\mathbf{x}) + \frac{\lambda_k(\mathbf{x})}{\rho_k} \right) = 0, \quad \forall \mathbf{x} \in \mathbb{R}^n \tag{36}$$

yielding the simple update rule

$$v_k^{t+1}(\mathbf{x}) = \frac{\rho A_k(\mathbf{x}) \left(f(\mathbf{x}) - \sum_{i \neq k} A_i(\mathbf{x}) v_i(\mathbf{x}) + \frac{\lambda(\mathbf{x})}{\rho} \right) + \rho_k u_k(\mathbf{x}) + \lambda_k(\mathbf{x})}{\rho A_k(\mathbf{x})^2 + \rho_k}, \quad \forall \mathbf{x} \in \mathbb{R}^n. \tag{37}$$

This update is interpreted as a balance between fidelity to the split mode u_k (enforced through Lagrangian multiplier λ_k), and the reconstruction-fidelity constraint where A_k is active (enforced through λ).

3.4.3 n-D-TV-VMD Subminimization w.r.t. A_k

As outlined above, the minimization problem with respect to the binary support functions A_k involves the L^1 - and TV-based priors⁵:

$$A_k^{t+1} = \arg \min_{A_k: \mathbb{R}^n \rightarrow \{0,1\}} \left\{ \beta_k \|A_k\|_1 + \gamma_k \text{TV}(A_k) + \rho \left\| f(\mathbf{x}) - \sum A_i(\mathbf{x}) v_i(\mathbf{x}) + \frac{\lambda(\mathbf{x})}{\rho} \right\|_2^2 \right\}. \tag{38}$$

The first variation associated with the TV term typically involves $\text{div}(\nabla \cdot / |\nabla \cdot|)$. One must expect difficulties with this term when used with binary functions such as A_k : at edges the gradient is not defined, and in flat regions $|\nabla A_k| = 0$. Moreover, if such a gradient descent PDE were integrated explicitly, then the time step is also heavily limited by the stiffness constraint [62]. Further, since A_k is binary we opt for schemes other than split Bregman/shrinkage or dual minimization [7, 32, 77]. We refrain from convex relaxation and instead are inspired by the Merriman–Bence–Osher (MBO) scheme for motion by mean curvature of interfaces [50]. This choice performed most promisingly in our experiments.

The fundamental idea is to reproduce the motion by mean curvature due to the boundary length term $\text{TV}(A_k)$ by more efficient means than direct gradient descent, namely the alternation between heat diffusion and thresholding. For recent comprehensive work on this type of threshold dynamics, see

⁵ Again we omit iteration superscripts for A_i , but it is understood that we always use the most recent estimate of a variable, A_i^{t+1} for $i < k$ and A_i^t for $i > k$.

[20]. These MBO-type schemes have already been successfully integrated with imaging data terms, such as [21, 78], where in addition to the heat diffusion and thresholding steps, a data-driven gradient descent step is included in the iterations. We propose a similar structure here, to account for the balloon force and reconstruction-fidelity term contributions to the A_k minimization.

To this end, we devise a threefold time split gradient descent-emulating iteration: the first step is gradient descent based on the support area and reconstruction-fidelity penalty (all terms involving A_k but the $\text{TV}(A_k)$ term). The second step is diffusion by the heat equation, followed by thresholding as a third step, to deal with the boundary length term $\text{TV}(A_k)$ and the projection onto the admissible set $\{0, 1\}$.

Since A_k is non-negative, it is safe to drop the absolute value and relax the L^1 -area term to $\beta_k \int_{\mathbb{R}^n} A_k$, making the functional smoothly differentiable in the area and reconstruction term.

We thus propose to update the binary support functions A_k^{t+1} in MBO-like fashion by iterating over the following three evolution equations resulting from time splitting⁶:

1. Area penalty and reconstruction-fidelity ODE:

$$\frac{\partial A_k(\mathbf{x})}{\partial t} = -\beta_k + 2\rho v_k(\mathbf{x}) \left(f(\mathbf{x}) - \sum A_i(\mathbf{x})v_i(\mathbf{x}) + \frac{\lambda(\mathbf{x})}{\rho} \right), \tag{39}$$

2. Heat equation PDE for diffusion:

$$\frac{\partial A_k(\mathbf{x})}{\partial t} = \gamma_k \nabla^2 A_k(\mathbf{x}), \tag{40}$$

3. Rectification by thresholding:

$$A_k(\mathbf{x}) = \begin{cases} 0 & \text{if } A_k(\mathbf{x}) \leq \frac{1}{2} \\ 1 & \text{if } A_k(\mathbf{x}) > \frac{1}{2} \end{cases} \quad \forall \mathbf{x} \in \mathbb{R}^n. \tag{41}$$

The differential equations are each propagated forward in time for a total time T . Note that to this end, the ODE problem can be addressed through an implicit (backward) Euler scheme and the heat equation is efficiently solved, e.g., based on convolution or spectral transforms [58].

3.4.4 n -D-TV-VMD Subminimization w.r.t. ω_k

The last, remaining subproblem of the saddle point problem (30) is the update of the mode’s central frequency, ω_k . The

⁶ Here, t is understood as an artificial time introduced for the sole purpose of differential equation notation, but quantized into the discrete iterates of the scheme.

relevant portion of the functional (29) is identical to the non-sparse 2D-VMD model (8). Therefore, the corresponding subminimization problem here is identical to (20), and thus, the update is equally given by (23).

The complete algorithm for the alternate direction (Gauss–Seidel-like) optimization of the 2D-TV-VMD model is shown in Algorithm 2, and illustrative examples of its use are given in Figs. 1 and 4.

4 Spectral Image Segmentation

Up to now, we have considered modes whose spatial support was mutually independent. In particular, this means that VMD and TV-VMD modes can be spatially overlapping and, conversely, that not all parts of a signal are necessarily covered by an active mode. Here, we want to consider the case where modes are restricted to be non-overlapping while covering the entire signal domain. In other words, the modes’ support functions A_k form a partition of the signal domain. For example, such a model includes the image segmentation problem.

4.1 Introducing the Partition Constraint

In terms of the binary support functions, $A_k: \mathbb{R}^n \rightarrow \{0, 1\}$, this means imposing the following constraint:

$$\sum_k A_k(\mathbf{x}) = 1, \quad \forall \mathbf{x} \in \mathbb{R}^n. \tag{42}$$

In return, the area penalty $\beta_k \|A_k\|_1$ can become obsolete, of course, unless not all modes incur the same area penalty due to different size priors, corresponding to $\beta_i \neq \beta_j$ for at least some $i \neq j \in \{1, \dots, K\}$ (this is a degree of freedom that we have not exploited in the examples in this manuscript, however).

We propose the following spatially disjoint n -D-SEG-VMD model, as a modification of (27):

$$\begin{aligned} & \min_{u_k: \mathbb{R}^n \rightarrow \mathbb{R}, A_k: \mathbb{R}^n \rightarrow \{0,1\}, \omega_k \in \mathbb{R}^n} \\ & \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\mathbf{x}) e^{-j(\omega_k, \mathbf{x})} \right] \right\|_2^2 \right. \\ & \left. + \beta_k \|A_k\|_1 + \gamma_k \text{TV}(A_k) \right\} \\ & \text{s.t. } \forall \mathbf{x} \in \mathbb{R}^n: \begin{cases} \sum_k A_k(\mathbf{x})u_k(\mathbf{x}) = f(\mathbf{x}), \\ \sum_k A_k(\mathbf{x}) = 1. \end{cases} \end{aligned} \tag{43}$$

It is important to note that the extra constraint only affects the updates of A_k , while all other steps of the algorithm remain unchanged.

Algorithm 2 2D-TV-VMD (sparsity promoting)

Input: signal $f(\mathbf{x})$, number of modes K , parameters $\alpha_k, \beta_k, \gamma_k, \rho, \rho_k, T, \tau, \tau_k, \epsilon$.

Output: modes $u_k(\mathbf{x})$, support functions $A_k(\mathbf{x})$, center frequencies ω_k .

Initialize $\{\omega_k\}, \{u_k^0\} \leftarrow 0, \{v_k^0\} \leftarrow 0, \{A_k^0\} \leftarrow 1, \{\lambda_k^0\} \leftarrow 0, \lambda^0 \leftarrow 0, n \leftarrow 0$

repeat

$n \leftarrow n + 1$

for $k = 1 : K$ **do**

Create 2D mask for analytic signal Fourier multiplier:

$$\mathcal{H}_k^{t+1}(\omega) \leftarrow 1 + \text{sgn}(\langle \omega_k^t, \omega \rangle)$$

Update $\hat{u}_{AS,k}$:

$$\hat{u}_{AS,k}^{t+1}(\omega) \leftarrow \mathcal{H}_k^{t+1}(\omega) \left[\frac{\rho_k \hat{v}_k^t(\omega) - \hat{\lambda}_k^t(\omega)}{\rho_k + 2\alpha_k |\omega - \omega_k^t|^2} \right]$$

Retrieve u_k :

$$u_k^{t+1}(\mathbf{x}) \leftarrow \Re \left(\mathcal{F}^{-1} \left\{ \hat{u}_{AS,k}^{t+1}(\omega) \right\} \right)$$

Update v_k :

$$v_k^{t+1}(\mathbf{x}) \leftarrow \frac{\rho A_k^t(\mathbf{x}) \left(f(\mathbf{x}) - \sum_{i < k} A_i^t(\mathbf{x}) v_i^{t+1}(\mathbf{x}) - \sum_{i > k} A_i^t(\mathbf{x}) v_i^t(\mathbf{x}) + \frac{\lambda^t(\mathbf{x})}{\rho} \right) + \rho_k u_k^{t+1}(\mathbf{x}) + \lambda_k^t(\mathbf{x})}{\rho A_k^t(\mathbf{x})^2 + \rho_k}$$

Update A_k through modified MBO:

$$A_k^{t+1/3}(\mathbf{x}) \leftarrow \frac{A_k^t(\mathbf{x}) + T \left(-\beta_k + 2\rho v_k^{t+1}(\mathbf{x}) \left(f(\mathbf{x}) - \sum_{i < k} A_i^{t+1}(\mathbf{x}) v_i^{t+1}(\mathbf{x}) - \sum_{i > k} A_i^t(\mathbf{x}) v_i^t(\mathbf{x}) + \frac{\lambda^t(\mathbf{x})}{\rho} \right) \right)}{1 + 2T\rho(v_k^{t+1}(\mathbf{x}))^2}$$

$$\hat{A}_k^{t+2/3}(\omega) \leftarrow \frac{\hat{A}_k^{t+1/3}(\omega)}{1 + T\gamma_k |\omega|^2}$$

$$A_k^{t+1}(\mathbf{x}) \leftarrow \begin{cases} 0 & \text{if } A_k^{t+2/3}(\mathbf{x}) \leq \frac{1}{2} \\ 1 & \text{if } A_k^{t+2/3}(\mathbf{x}) > \frac{1}{2} \end{cases}$$

Update ω_k :

$$\omega_k^{t+1} \leftarrow \frac{\int_{\mathbb{R}^2} \omega |\hat{u}_{AS,k}^{t+1}(\omega)|^2 d\omega}{\int_{\mathbb{R}^2} |\hat{u}_{AS,k}^{t+1}(\omega)|^2 d\omega}$$

Dual ascent u - v coupling:

$$\lambda_k^{t+1}(\mathbf{x}) \leftarrow \lambda_k^t(\mathbf{x}) + \tau_k \left(u_k^{t+1}(\mathbf{x}) - v_k^{t+1}(\mathbf{x}) \right)$$

end for

Dual ascent data fidelity:

$$\lambda^{t+1}(\mathbf{x}) \leftarrow \lambda^t(\mathbf{x}) + \tau \left(f(\mathbf{x}) - \sum_k A_k^{t+1}(\mathbf{x}) v_k^{t+1}(\mathbf{x}) \right)$$

until convergence

4.2 Multiphase MBO and Rearrangement

While the extra partitioning constraint could be addressed through another augmented Lagrangian to be included in the saddle point problem, here we seek to impose the constraint by modifying the current rectification step included in the MBO-like diffusion and threshold dynamics. Indeed, the partitioning problem corresponds to a multiphase interface problem. The fundamental idea is to propagate the data ODE (39) and the heat diffusion PDE (40) on each support function A_k individually, but to replace the individual thresholding step (41) by a single, common “winner-takes-it-all” rectification. This idea has been discussed more rigorously in

[20]. For an application of a similar strategy to graph-based image processing, see [26, 42, 49].

The projection-based partitioning update for A_k becomes:

1. Area penalty and reconstruction-fidelity ODE propagation for each mode k , according to (39).
2. Heat diffusion PDE for each mode k according to (40).
3. “Winner-takes-it-all” rectification; Projection of the intermediate A_k onto the feasible set $A_k \in \{0, 1\} \cap \sum A_k = 1$:

$$A_k^{t+1} = \begin{cases} 1 & \text{if } k = \arg \max_i A_i, \\ 0 & \text{otherwise.} \end{cases} \tag{44}$$

Algorithm 3 2D-TV-VMD with segmentation constraint

Input: signal $f(\mathbf{x})$, number of modes K , parameters $\alpha_k, \beta_k, \gamma_k, \rho, \rho_k, T, \tau, \tau_k, \epsilon$.

Output: modes $u_k(\mathbf{x})$, domain partitioning support functions $A_k(\mathbf{x})$, center frequencies ω_k .

Initialize $\{\omega_k\}, \{u_k^0\} \leftarrow 0, \{v_k^0\} \leftarrow 0, \{A_k^0\} \leftarrow 1, \{\lambda_k\}^0 \leftarrow 0, \lambda^0 \leftarrow 0, n \leftarrow 0$

repeat

$n \leftarrow n + 1$

for $k = 1 : K$ **do**

Create 2D mask for analytic signal Fourier multiplier:

$$\mathcal{H}_k^{t+1}(\omega) \leftarrow 1 + \text{sgn}(\langle \omega_k^t, \omega \rangle)$$

Update $\hat{u}_{AS,k}$:

$$\hat{u}_{AS,k}^{t+1}(\omega) \leftarrow \mathcal{H}_k^{t+1}(\omega) \left[\frac{\rho_k \hat{v}_k^t(\omega) - \hat{\lambda}_k^t(\omega)}{\rho_k + 2\alpha_k |\omega - \omega_k^t|^2} \right]$$

Retrieve u_k :

$$u_k^{t+1}(\mathbf{x}) \leftarrow \Re \left(\mathcal{F}^{-1} \left\{ \hat{u}_{AS,k}^{t+1}(\omega) \right\} \right)$$

Update v_k :

$$v_k^{t+1}(\mathbf{x}) \leftarrow \frac{\rho A_k^t(\mathbf{x}) \left(f(\mathbf{x}) - \sum_{i < k} A_i^t(\mathbf{x}) v_i^{t+1}(\mathbf{x}) - \sum_{i > k} A_i^t(\mathbf{x}) v_i^t(\mathbf{x}) + \frac{\lambda^t(\mathbf{x})}{\rho} \right) + \rho_k u_k^{t+1}(\mathbf{x}) + \lambda_k^t(\mathbf{x})}{\rho A_k^t(\mathbf{x})^2 + \rho_k}$$

Update ω_k :

$$\omega_k^{t+1} \leftarrow \frac{\int_{\mathbb{R}^2} \omega |\hat{u}_{AS,k}^{t+1}(\omega)|^2 d\omega}{\int_{\mathbb{R}^2} |\hat{u}_{AS,k}^{t+1}(\omega)|^2 d\omega}$$

Dual ascent $u-v$ coupling:

$$\lambda_k^{t+1}(\mathbf{x}) \leftarrow \lambda_k^t(\mathbf{x}) + \tau_k \left(u_k^{t+1}(\mathbf{x}) - v_k^{t+1}(\mathbf{x}) \right)$$

end for

for $k = 1 : K$ **do**

Update A_k through time split ODE and PDE propagation:

$$A_k^{t+1/3}(\mathbf{x}) \leftarrow \frac{A_k^t(\mathbf{x}) + T \left(-\beta_k + 2\rho v_k^{t+1}(\mathbf{x}) \left(f(\mathbf{x}) - \sum_{i < k} A_i^{t+2/3}(\mathbf{x}) v_i^{t+1}(\mathbf{x}) - \sum_{i > k} A_i^t(\mathbf{x}) v_i^{t+1}(\mathbf{x}) + \frac{\lambda^t}{\rho} \right) \right)}{1 + 2T\rho(v_k^{t+1}(\mathbf{x}))^2}$$

$$\hat{A}_k^{t+2/3}(\omega) \leftarrow \frac{\hat{A}_k^{t+1/3}(\omega)}{1 + T\gamma_k |\omega|^2}$$

end for

for $k = 1 : K$ **do**

Rectify A_k through winner-takes-it-all:

$$A_k^{t+1}(\mathbf{x}) = \begin{cases} 1 & \text{if } k = \arg \max_i A_i^{t+2/3}(\mathbf{x}) \\ 0 & \text{otherwise} \end{cases}$$

end for

Dual ascent data fidelity:

$$\lambda^{t+1}(\mathbf{x}) \leftarrow \lambda^t(\mathbf{x}) + \tau \left(f(\mathbf{x}) - \sum_k A_k^{t+1}(\mathbf{x}) v_k^{t+1}(\mathbf{x}) \right)$$

until convergence

The modified, 2D-SEG-VMD algorithm with segmentation constraint is given in Algorithm 3, while illustrative examples are shown in Fig. 5 *et seqq.*

5 Lattice Segmentation

Until now, our decomposition associates one spatial characteristic support function, A_k , with only one intrinsic

mode function, u_k . This results in a simple decomposition where each spatial region has exactly one simple oscillation. Let us now consider a case where the image is composed of regions not corresponding to plane waves, but combinations of simple oscillatory patterns, such as a checkerboard or hexagonal pattern. Microscopy of single-molecule layers, colloids, and crystal grains have such patterns. In biochemistry and nanoscience, the decomposition of such microscopy images into regions of homogene-

ity provides a necessary mechanic for further downstream analyses.

In microscopy, a crystal image contains different mesoscopic grains, where each grain typically can be a homogeneous, lattice region. Each grain has different spatial periodicities, depending on the crystal lattice structure. These structures are modeled by Bravais lattices, which, depending on the 2D crystalline arrangement, come in five forms: oblique, rectangular, centered rectangular, hexagonal, and square. Thus, a grain’s Fourier spectrum has several distinct peaks, associated with the various cosine waves that constitute the pattern, which share a common spatial support (function). For example, a grain in a homogeneously hexagonal lattice patch would have three coupled peaks in the spectral half-space. Grains differ by orientation, so it is interesting to find the grain supports, their boundaries and defects, and the Fourier peaks associated with each grain. A crystal image composed of such grains can be considered as an assemblage of 2D general intrinsic mode type functions with non-overlapping supports, specified propagating directions and smoothly varying local wave vectors. A recent state-of-the-art method uses 2D synchrosqueezed transforms together with slow-oscillating, global-structure providing functions, known as shape functions, in order to model atomic crystal images [75]. In general, knowing the Bravais lattice structure yields strong priors on the relative positions of the frequency peaks; here, however, we only make use of the known number of peaks, but not their relative positions.

To accommodate such regions, our spectral image segmentation needs to be adapted to allow for multiple single-Fourier-peak modes to be joined together through a single binary support function. Let $\{u_{kj}\}_j$ denote the set of modes associated with the single binary support function A_k . Each of these modes needs to be individually of small bandwidth, but they contribute to the signal reconstruction jointly through their single support function A_k . This simple modification allows us to *segment* signals into meaningful pieces.

To this end, we modify the spatially disjoint n -D-SEG-VMD model (43) as follows:

$$\begin{aligned}
 & \min_{u_{ki} : \mathbb{R}^n \rightarrow \mathbb{R}, A_k : \mathbb{R}^n \rightarrow \{0,1\}, \omega_{ki} \in \mathbb{R}^n} \\
 & \left\{ \sum_{k,i} \alpha_{ki} \left\| \nabla \left[u_{AS,ki}(\mathbf{x}) e^{-j(\omega_{ki}, \mathbf{x})} \right] \right\|_2^2 \right. \\
 & \left. + \sum_k \beta_k \|A_k\|_1 + \sum_k \gamma_k \text{TV}(A_k) \right\} \\
 & \text{s.t. } \forall \mathbf{x} \in \mathbb{R}^n : \begin{cases} \sum_k A_k(\mathbf{x}) \sum_i u_{ki}(\mathbf{x}) = f(\mathbf{x}), \\ \sum_k A_k(\mathbf{x}) = 1. \end{cases} \quad (45)
 \end{aligned}$$

We call this the n -D-SEG-VMD *lattice* segmentation model. The model can be optimized in much the same way as the

simpler model (43). The only significant difference is in the ODE propagation step of the A_k update: here, all associated modes u_{ki} (resp. their copies v_{ki}) jointly influence the update of the single A_k . Indeed, (39) now becomes:

$$\begin{aligned}
 \frac{\partial A_k(\mathbf{x})}{\partial t} = & -\beta_k \\
 & + 2\rho \left(\sum_i v_{ki}(\mathbf{x}) \right) \left(f(\mathbf{x}) - \sum_l A_l(\mathbf{x}) \sum_j v_{lj}(\mathbf{x}) + \frac{\lambda(\mathbf{x})}{\rho} \right). \quad (46)
 \end{aligned}$$

Explicitly modifying the previous algorithms to incorporate this submode coupling is fairly straightforward and left as an exercise to the reader. Examples of image decomposition with submode coupling are shown in Figs. 9, 10, 12 and 12.

6 Outlier Detection: Artifact Detection and Inpainting

As a final complication regarding crystallography images, we now wish to deal with image features that cannot be explained by the VMD model thus far, such as defects and artifacts. While artifacts can be due to acquisition noise or sample impurities (accidental or intended), defects are irregularities in the regular crystal structure, within crystal grains, or more frequently at the grain boundaries. In imaging terms, these are characterized by a stark deviation from the regular spatial pattern modeled by the band-limited modes of the VMD model. In the presence of imaging noise, one naturally relaxes the data-fidelity constraint by just a quadratic penalty, i.e., not making use of a Lagrangian multiplier. Therefore, unless otherwise accounted for, such defects and artifacts appear in the data-fidelity residual, but due to their non-Gaussian nature as strong outliers will also affect and deteriorate the mode decomposition. It is imperative, therefore, to address these features more specifically beyond making Gaussian noise assumptions.

6.1 Artifact Indicator Function

Recently, a dynamic artifact detection model was introduced in the framework of classical Chan–Vese image segmentation [78]. There, individual pixels were eliminated from the region-based segmentation terms to prevent skewing and misleading the segmentation. This method is related to similar approaches in occlusion detection in optical flow [1] and salt-and-pepper denoising [73]. Here, the goal is to isolate defects and artifacts from interfering with the regular modes. We note that these image defects and artifacts can be the result of important physical structures [8, 36, 67, 68].

We introduce an artifact indicator function,

$$\chi : \mathbb{R}^n \rightarrow \{0, 1\},$$

where for each pixel a 1 denotes an artifact, and 0 absence thereof. We use this artifact indicator function to limit the data-fidelity constraint to non-artifact regions, only, e.g.,

$$\forall \mathbf{x} \in \mathbb{R}^n \mid \chi(\mathbf{x}) = 0: \sum_k A_k(\mathbf{x})u_k(\mathbf{x}) = f(\mathbf{x}). \tag{47}$$

This is equivalent to

$$\forall \mathbf{x} \in \mathbb{R}^n: \sum_k (1 - \chi(\mathbf{x}))A_k(\mathbf{x})u_k(\mathbf{x}) = (1 - \chi(\mathbf{x}))f(\mathbf{x}), \tag{48}$$

where $(1 - \chi(\mathbf{x})) = 1$ in regions not classified as artifacts, which is where data fidelity is to be enforced. A similar modification can be made to all data-fidelity constraints of the previous models.

6.2 Defect and Artifact Detection and Inpainting

We have not described, so far, how the values of the binary defect and artifact indicator function χ are to be determined, in the first place. While there are reasonable grounds to believe that these defect and artifact locations could be heuristically identified from images in preprocessing, we want to integrate this detection process into the very same decomposition model.

At this point, we do not have a concise and simple characterization of the shape and appearance of defects and artifacts, and for the general case, we even want to avoid including too many such priors. Instead, we characterize lattice defects and image artifact locations by what they are not; indeed, at these locations the image simply fails to be sufficiently well modeled by the band-limited modes extracted nearby. We thus decide to classify a certain pixel $f(\mathbf{x})$ as an artifact or defect, $\chi(\mathbf{x}) = 1$, if the incurred data-fidelity cost would be too large, locally, otherwise. This is most simply achieved by including an L^1 term on χ .

We modify the constrained n -D-TV-VMD cost functional (27) to become the n -D-TV-XVMD (with artifact detection) functional as follows:

$$\min_{u_k : \mathbb{R}^n \rightarrow \mathbb{R}, A_k, \chi : \mathbb{R}^n \rightarrow \{0, 1\}, \omega_k \in \mathbb{R}^n} \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\mathbf{x})e^{-j(\omega_k, \mathbf{x})} \right] \right\|_2^2 + \beta_k \|A_k\|_1 + \gamma_k \text{TV}(A_k) + \delta \|\chi\|_1 \right\}$$

$$\begin{aligned} \text{s.t. } \forall \mathbf{x} \in \mathbb{R}^n: & \sum_k (1 - \chi(\mathbf{x}))A_k(\mathbf{x})u_k(\mathbf{x}) \\ & = (1 - \chi(\mathbf{x}))f(\mathbf{x}). \end{aligned} \tag{49}$$

The corresponding unconstrained saddle point problem (without Lagrange multiplier on the data fidelity) then becomes:

$$\begin{aligned} \mathcal{L}(\{u_k\}, \{v_k\}, \{A_k\}, \{\omega_k\}, \chi, \{\lambda_k\}) & := \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\mathbf{x})e^{-j(\omega_k, \mathbf{x})} \right] \right\|_2^2 \right. \\ & + \beta_k \|A_k\|_1 + \gamma_k \text{TV}(A_k) \\ & + \delta \|\chi\|_1 + \rho \left\| (1 - \chi(\mathbf{x})) \left(f(\mathbf{x}) - \sum_k A_k(\mathbf{x})v_k(\mathbf{x}) \right) \right\|_2^2 \\ & \left. + \sum_k \rho_k \|u_k(\mathbf{x}) - v_k(\mathbf{x})\|_2^2 + \langle \lambda_k(\mathbf{x}), u_k(\mathbf{x}) - v_k(\mathbf{x}) \rangle \right\}. \end{aligned} \tag{50}$$

Note that the masking only impacts the data-fidelity evaluation domain, while all other terms are not affected. Indeed, only two subminimization steps will be altered by the introduction of the $(1 - \chi)$ term:

1. The area penalty and reconstruction-fidelity ODE (39) will collapse to just $\partial_t A_k(\mathbf{x}) = -\beta_k$ whenever $\chi(\mathbf{x}) = 1$ (and remain unchanged, otherwise). In particular, the TV- and L^1 terms on the binary support functions A_k will now exclusively drive the evolution of the latter whenever a location is marked as artifact, since the data-fidelity constraint is the only link between modes and support functions.
2. Similarly, the update (37) of v_k collapses to $v_k^{t+1}(\mathbf{x}) = u_k(\mathbf{x}) + \lambda_k(\mathbf{x})/\rho_k$ when $\chi(\mathbf{x}) = 1$, which effectively unlinks the local mode estimate from the observed data and simply *inpaints the artifact regions* by Fourier interpolation of the modes.

On the other hand, the estimation of the artifact indicator function χ itself also leads to a straightforward optimization step. The binary optimization can be carried out independently for each pixel, and the optimal $\chi^*(\mathbf{x})$ chooses between paying data-fidelity penalty versus artifact cost δ , as follows:

$$\chi^*(\mathbf{x}) = \begin{cases} 0 & \text{if } \rho(f(\mathbf{x}) - \sum_k A_k(\mathbf{x})v_k(\mathbf{x}))^2 \leq \delta \\ 1 & \text{otherwise} \end{cases} \tag{51}$$

This thresholding scheme has an immediate interpretation from a hypothesis-testing perspective. Indeed, if we consider the data-fidelity weight ρ to be the precision of the implicitly assumed Gaussian noise distribution, then the expression $\rho(f(\mathbf{x}) - \sum_k A_k(\mathbf{x})v_k(\mathbf{x}))^2$ represents the squared z -score

(standard score) of the local image intensity under such a noise distribution. This squared z -score is compared against the threshold δ . The artifact classification is effectively a concealed statistical hypothesis z -test of the pixel intensity with a Gaussian distribution

$$p(f(\mathbf{x})) = \mathcal{N}(f(\mathbf{x}) \mid \sum A_k(\mathbf{x})v_k(\mathbf{x}), \rho^{-1})$$

as null hypothesis H_0 , and a pixel is classified as an artifact (H_1) if the z -score of its intensity is more extreme than $\sqrt{\delta}$. The model parameter δ is thus intimately related to the level of statistical significance attached to the artifact classification and its expected false positives rate.

Again, in the interest of conciseness, we leave the modification of the algorithms to include the artifact detection and inpainting terms as an exercise for the reader. An inpainting example is illustrated in Fig. 8.

7 Experiments and Results

We have implemented the three Algorithms 1–3, including the submode coupling of Sect. 5 and the artifacts detection and inpainting (Sect. 6) extensions, in MATLAB®. The algorithms can be implemented in a single code file, because they are mostly generalizations of each other.

In the implementation, we make two deliberate choices that have not been discussed, so far. The first choice is with respect to initialization of the center frequencies, where we include four options:

1. initialization of frequencies uniformly spread on a circle (deterministic),
2. random initialization on the positive half-space,
3. user selection through graphical user interface, and
4. user input as parameters.

Unless otherwise noted, all the examples shown below make use of the deterministic radial frequency initialization scheme.

The second particularity is with respect to model selection 2D-VMD, 2D-TV-VMD, and 2D-SEG-VMD. Indeed it is useful in practice to initialize the TV-VMD model by some iterations of unrestricted 2D-VMD, in order to settle the center frequencies close to the optimal location; and similarly, the segmentation model is best initialized based on the outcome of 2D-TV-VMD optimization. We will thus always start optimizing in 2D-VMD mode, and over the iterations, switch to the two more complicated models at user-defined time-points (which may be set to infinity, thereby producing results of simpler models as final output).

These simple initialization strategies turn out to achieve robustness effectively. Note that more sophisticated initial-

ization strategies and enhancements for increased robustness that are known from popular models like k -means, such as repeated runs from random initialization, could also be applied to the proposed models.

Finally, we remark that the examples below typically converge within less than 100–150 iterations (or even much faster for the simplest 2D-VMD model).

Our implementation is publicly available for download at <http://www.math.montana.edu/dzosso/code> and at MATLAB Central.

7.1 Synthetic Overlapping Texture Decomposition

The first, synthetic image is a composition of spatially overlapping basic shapes, more precisely six ellipses and a rectangle, with frequency patterns varying in both periodicity and direction, courtesy of Gilles [28]. The spectrum is ideal for segmentation due to modes being deliberately both well spectrally isolated and narrow-banded. The resolution of the synthetic image is 256×256 .

We feed the synthetic image to our models and show the resulting decompositions for both 2D-VMD and 2D-TV-VMD models in Fig. 1. The parameters are⁷:

K	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	T
5	1000	0.5	500	∞	10	10	2.5	2.5	1.5

In addition, the center frequency of the first mode is held fixed at $\omega_1 = 0$ to account for the DC component of the image. As a result, the first mode contains the solid ellipse and rectangle, while the four remaining decompositions in Fig. 1 show clear separation of the patterned ellipses.

In the simple 2D-VMD model of Fig. 1e, due to the solid pieces having sharp edges, their spectra are not band-limited and only smoothed versions are recovered. This is naturally paired with the two lower frequency modes absorbing residual boundary artifacts of the DC component, and ghost contours appearing in these modes.

The spatially compact 2D-TV-VMD model, Fig. 1f–h, however, can handle sharp boundaries through the support functions A_k , while the modes u_k can smoothly decay. The resulting masked modes, $A_k u_k$, are thus clean and sharp.

These results should be compared against the exhaustive 2D empirical wavelet transform results in [30].

Further, in Fig. 2 we provide an illustration of the *multiresolution monogenic signal analysis by Riesz–Laplace*

⁷ Of course, the simpler 2D-VMD model only uses a subset of these parameters, for the support functions are fixed at $A_k = 1$ uniformly.

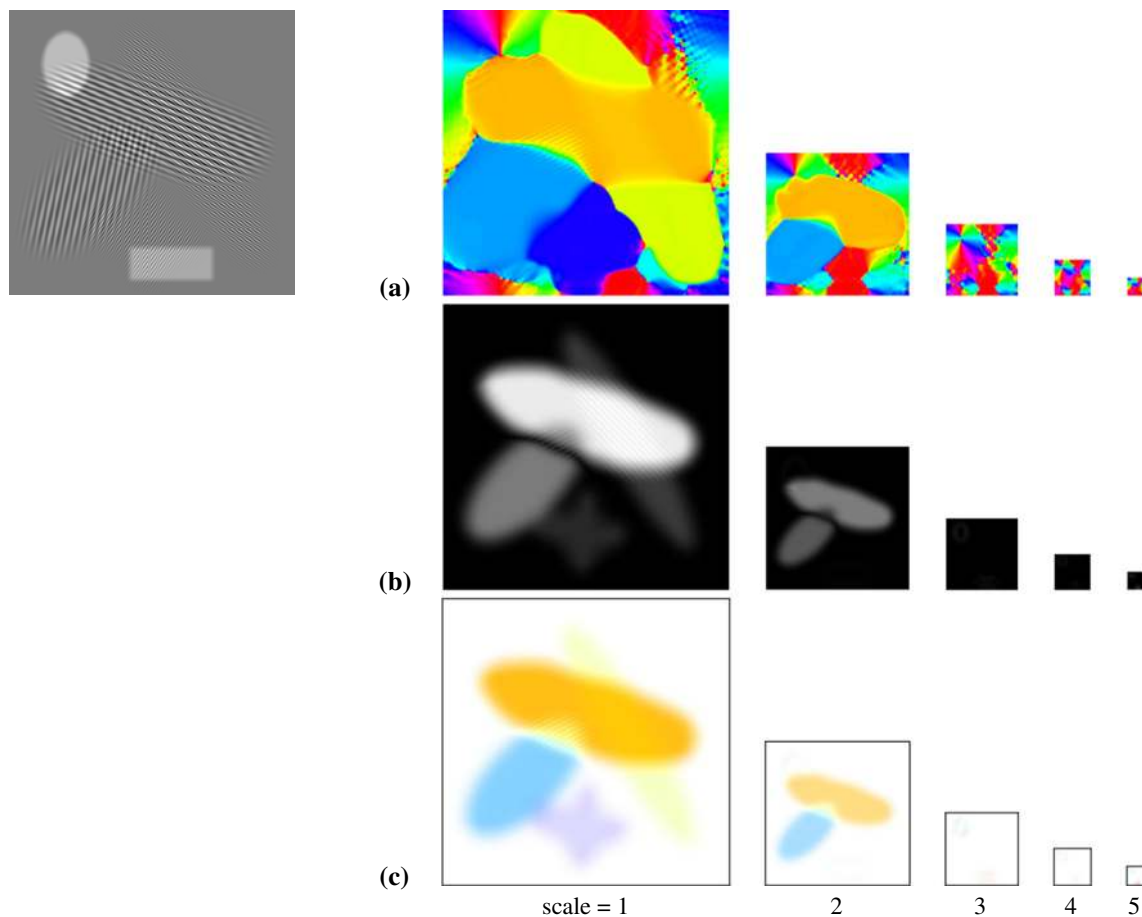


Fig. 2 Synthetic overlapping texture: multiresolution monogenic signal analysis using Riesz–Laplace wavelets (five scales, $\sigma = 1.5$) [69, 70]. **a** The local orientation is color-coded between $-\frac{\pi}{2}$ and $\frac{\pi}{2}$. **b** The coherency (degree of directionality): isotropic features are dark and

directional features are bright. **c** Composite angle (hue) and coherency (saturation). At each scale, the method picks the dominant wave and cannot resolve overlapping signals of same scale. Some waves appear at multiple scales

wavelets [69, 70].⁸ At each (dyadic) scale, the Riesz–Laplace wavelets are able to pick up the dominant local orientation and compute a degree of coherency, reminiscent of our proposed support functions. However, that method cannot properly deal with overlapping signals and due to the rigid construction of the wavelet scales, multiple signals belong to the same scale while a signal may appear at multiple scales.

Finally, our results should also be compared against the results obtained with the *2D Prony–Huang transform* reported in [60],⁹ and partially reproduced in Fig. 3. This latter method performs EMD-based mode decomposition before computing local orientation and coherency information (not shown here). The mode decomposition is rather slow and produces a very different decomposition.

7.2 Overlapping Chirps

The second example problem is still synthetic, but the modes have non-trivial Fourier support. More precisely, the synthetic image is a superposition of three compactly supported yet spatially overlapping 2D chirps (see Fig. 4). Starting from radial initialization, we let our algorithm determine the correct support and appropriate center frequencies for this problem, based on the following parameters:

K	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	T
3	2000	1	1000	∞	7	10	1	1	1

The resulting decomposition is accurate with only little error on the true support functions. The modes are spectrally clean. It is interesting to observe how our model extrapolates the modes outside their rectangular domain boundaries. Note

⁸ MATLAB code available at <http://bigwww.epfl.ch/demo/steerable-wavelets/>.

⁹ Obfuscated MATLAB p-code available at http://perso.ens-lyon.fr/nelly.pustelnik/Software/Toolbox_PHT_2D_v1.0.zip.



Fig. 3 Synthetic overlapping texture: 2D Prony–Huang decomposition (default parameters) [60]. Decomposition into trend and 4 modes; spectral analysis not shown




Fig. 4 Chirp decomposition. **a** Input signal f . **b** 2D-TV-VMD modes u_k . **c** Fourier spectrum \hat{f} . **d** Determined supports A_k . See Sect. 7.2 (Color figure online)

that the decay distance correlates with the wavelength of the mode.

7.3 Textural Segmentation for Denoising

The two examples encountered so far were noise-free and perfect reconstruction was possible through the use of Lagrange multipliers ($\tau, \tau_k > 0$). In the presence of noise, however, enforcing strict data fidelity may be inappropriate, and instead relying on just the quadratic penalty to promote data fidelity is the proper setting. This approach is easily achieved by preventing the Lagrangian multipliers from updating: $\tau, \tau_k = 0$. As a result, the noise can be handled with a residual slack between the splitting variables. In particular, the quadratic penalty term corresponds to a Gaussian noise assumption, where the penalty coefficients ρ, ρ_k relate to the noise precision.

Here, we explore the idea of using the slack in the absence of Lagrangian multipliers for denoising based on spectral sparsity. To this end, we construct a four-quadrant,

non-overlapping unit-amplitude cosine texture image with different levels of noise, shown in Fig. 5. Because the quadrants are non-overlapping, we are interested in the output of the 2D-SEG-VMD model using the following parameters:

K	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	T
4	3500	1.5	750	∞	7	10	0	0	1

Without the Lagrangian multipliers active, it is important to realize that the two copies of the modes, u_k and v_k , may be different and that u_k is the potentially cleaner copy of the two.

In Fig. 5, we can see that even for important noise levels, the partition is recovered with good precision (red contours). In addition, the recovered composite of the four masked modes is very clean, seemingly irrespective of the degrading noise level.




Fig. 5 Denoising. *Top* noisy f (noise SD σ and SNR) with detected phase borders (red). *Bottom* denoised signal $\sum_k A_k u_k$, SNR. See Sect. 7.3 (Color figure online)

7.4 Segmentation of Peptide β -Sheets

The next test case are two atomic force microscopy (AFM) images of peptide β -sheets bonding on a graphite base, courtesy of the Weiss group at the California NanoSystems Institute (CNSI) at UCLA [10,76]. The peptide sheets grow in regions of directional homogeneity and form natural spatial boundaries where the regions meet. It is important to scientists to have accurate segmentation for their dual interests in complementary analysis of the homogeneous regions and their boundaries. Identifying regions of homogeneity enables the subsequent study of isolated peptide sheets of one particular bonding class. For these types of scans, manually finding the boundaries is a tedious problem that demands the attention of a skilled scientist on a rote task. In addition to speed and automation, the proposed 2D-VMD is superior in accuracy to manual boundary identification due to regions potentially having very similar patterns, of which the orientation differs by only a few degrees, difficult to discern by eye.

Nanoscale images such as these are a useful testbed since data are often oversampled relative to the smallest observable features, atoms and molecular parts. Also, segmentation in one imaging modality can be used to guide segmentation or data acquisition in a complementary imaging mode [3,8,36,51,67].

The first example, shown in Fig. 6, is a 512×512 false-color image, of which we only consider the average intensity across color channels as a proxy, *in lieu* of the actual raw data produced by the microscope. Also, as classical preprocessing step, we apply a *Laplacian of Gaussians* (LoG) band-pass filter to the image in order to remove both some noise and the DC component. Expert inspection suggests that there are

six different grain orientations represented in this image. We perform 2D-VMD, 2D-TV-VMD, and 2D-SEG-VMD using these parameters:

K	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	T
6	2000	1	250	∞	7	10	0	0	2.5

The recovered modes are shown in Fig. 6d–f. The unconstrained 2D-VMD model produces overly smooth modes without clear boundaries. The compactly supported 2D-TV-VMD model yields modes with sharp delineation. As can be seen from the grain boundaries overlaid to the input image, in Fig. 6b, the modes are not overlapping, but do not cover the entire image domain, leaving unaccounted space at the grain boundaries. This problem is effectively addressed by the addition of the segmentation constraint, as seen by the boundaries in Fig. 6c.

The second example, shown in Fig. 7, is believed to consist of only three main grain orientations. This 512×512 image is of the same type as the previous example and preprocessed in the same way. The image exhibits strong singular spots due to additional material deposition on the sample surface. In order to address these outliers, we make use of the artifact detection and inpainting extension, for δ finite:

K	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	T
3	2000	1	75	3.5	7	10	0	0	2.5

While the singular deposits (“artifacts”) negatively impact the mode purity for both 2D-VMD and 2D-TV-VMD




Fig. 6 Atomic force microscopy (AFM) image of peptide β -sheets, 512×512 pixels, $500 \text{ nm} \times 500 \text{ nm}$ (I). **a** Input f . **b** 2D-TV-VMD boundaries (red). **c** 2D-SEG-VMD partition (red). **d** 2D-VMD modes u_k . **e** 2D-TV-VMD modes $A_k u_k$. **f** 2D-SEG-VMD modes $A_k u_k$. See Sect. 7.4 in the text for details and discussion (Color figure online)




Fig. 7 Atomic force microscopy (AFM) image of peptide β -sheets, 512×512 (II). **a** Input f . **b** 2D-SEG-VMD partition (red). **c** Partition (red) with enabled artifact detection (cyan). **d** 2D-VMD modes u_k . **e** 2D-SEG-VMD modes $A_k u_k$. **f** Modes obtained with artifact detection enabled. See Sect. 7.4 for details (Color figure online)

(Fig. 7d–e), this effect is partially alleviated by the automatic detection and inpainting capability of the artifact extension (Fig. 7f).¹⁰ In addition to the outlined grain boundaries (red), the location of the detected artifacts is highlighted in cyan, in Fig. 7c. Note that the artifact detection also allows spotting at least some of the grain defects, in addition to the deposits.

7.5 Inpainting

Here, we are interested in exploiting the model’s capability of intrinsically inpainting the modes (and therefore the input image) in regions that are labeled as artifacts/outliers. To this end, we construct a simple checkerboard image, which essentially corresponds to a superposition of two cosine waves with full support each. In addition, portions of the image are corrupted by “pencil-scribble,” as shown in Fig. 8a. We set up the model as a two-modes 2D-VMD image decomposition

¹⁰ Lower artifact threshold δ and higher TV weight γ_k might increase the mode cleanliness even further.




Fig. 8 2D-VMD inpainting. **a** Input image f . **b** Fourier spectrum \hat{f} . **c** Recovered modes $\sum_k u_k$. **d** Detected artifacts χ . See Sect. 7.5

problem, with a finite artifact detection threshold. The data-fidelity Lagrangian is inactive in order to allow some slack (Gaussian noise assumption) and artifact detection, while we maintain an active Lagrangian multiplier on the u – v splitting:

K	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	T
2	1500	n/a	n/a	30	150	20	0	1	n/a

As can be seen in Fig. 8c, d, the model succeeds well in detecting the scribble as outliers. In the artifact-labeled image portions, the submodes are inpainted by intrinsic Fourier interpolation, and as a result, a full checkerboard can be recovered from the decomposition.

7.6 Textural Segmentation: Lattices

We finally turn our attention to the segmentation of images with lattice texture, as observed, for example in crystallography and microscopy images of crystalloid samples. The fundamental assumed property of such images is that they consist of K different domains (grains) forming a partition of the image, such that each grain has a distinct lattice texture composed of a superposition of M different essentially wave-like subbands. As seen earlier, a checkerboard lattice would consist of a superposition of $M = 2$ orthogonal cosine waves, while a hexagonal lattice consist of $M = 3$ modes differing




Fig. 9 Lattice decomposition. **a** Input f . **b** Fourier spectrum \hat{f} . **c, d** Recovered phases $\sum_i A_k u_{ki}$. **e, f** Submodes u_{ki} . See Sect. 7.6.1

by 60° rotation. Our model allows for multiple submodes u_{ki} to share a common support function A_k and thus be spatially coupled.

7.6.1 Checkerboard: 2 Phases with 2 Submodes

As a first simple example, we consider the composite of two checkerboard halves, of which one is slightly rotated, as shown in Fig. 9a. The goal is to find the support of two phases, partitioning the 256×256 image domain, and the




Fig. 10 Three phases–three modes. **a** Input f . **b** Fourier spectrum \hat{f} . **c** 2D-SEG-VMD partition (red). **d** Phases $A_k \sum_i u_{ki}$. See Sect. 7.6.2 (Color figure online)

respective two submodes for each such grain. We run the 2D-SEG-VMD model with the following parameters:

K	M	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	T
2	2	2000	1	250	∞	7	10	0	0	2.5

The resulting decomposition into the two checkerboard phases, $A_k \sum_i u_{ki}$, is shown in Fig. 9c, d, while the constituting two submodes per phase, u_{ki} , are illustrated in Fig. 9e, f.

7.6.2 Hexagonal Lattice: 3 Phases with 3 Submodes

A slightly more complicated problem is illustrated in Fig. 10. We start with a tripartite 256×256 image, where each domain consists of an artificial hexagonal lattice pattern, obtained by superposing three cosine waves rotated by 60° against each other. Each domain has a slightly different lattice orientation ($0^\circ, 15^\circ, 45^\circ$). Like the previous example, this is a 2D-SEG-VMD problem, this time with three phases and three submodes, each. The other parameters remain unchanged:

K	M	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	T
3	3	2000	1	250	∞	7	10	0	0	2.5

As can be seen in Fig. 10c, d, the recovered phases and their boundaries are very precise. Note that this decomposition involves the identification of nine center frequencies and associated wave functions, and the delineation of three support functions partitioning the image domain.

7.6.3 Simulated Hexagonal Crystal

The three-phase three-wave hexagonal lattice image of the previous subsection was an idealized synthetic version of what real-world acquired images of hexagonally arranged crystal structures might look like. In an attempt to make the problem more realistic, we created a more complicated synthetic lattice image as follows: We predefine a 5-partition of the 256×256 image domain. In each domain, individual pixels corresponding to approximate “bubble locations” of the crystal lattice are activated. The exact center position is affected by discretization noise (the pixel locations are obviously limited to the Cartesian grid) as well as additional, controllable jitter. The resulting “nail board” is then convoluted with a circular point spread function designed to mimic the approximate appearance of an individual lattice element, and Gaussian white noise is added. An example is shown in Fig. 11a. Due to this construction the grain boundaries exhibit very irregular defects. All of these complications make the resulting image much more interesting and challenging to segment.

In a first, simple attempt, we configure the 2D-SEG-VMD algorithm as follows:

K	M	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	T
5	3	2000	1	250	∞	7	10	1	1	2.5

In contrast to the actually noise-free preceding examples, here, we enforce data fidelity strictly by picking $\tau = \tau_k = 1$, so as to make sure the phases and modes pick up the relevant center frequencies and do not lazily get stuck in local minima (see a discussion in [17] for the role of the Lagrangian multipliers in low-noise regimes). The model is thus obliged to over-explain all image noise (jitter and Gaussian noise) in terms of mode decomposition. As a result, the obtained partition captures the five phases largely, but suffers from strong noise, as shown in the middle of Fig. 11. Most importantly, though, this procedure found the correct 5×3 center frequencies.

These correctly identified center frequencies can now be used as a very strong prior when running the 2D-SEG-VMD model a second time, in a different regime with inactive Lagrangian multipliers to allow noise slack. To this end, we use the obtained center frequencies as user initialization for a second run, with parameters as follows:

K	M	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	T
5	3	2e4	1	500	∞	7	10	0	0	2.5

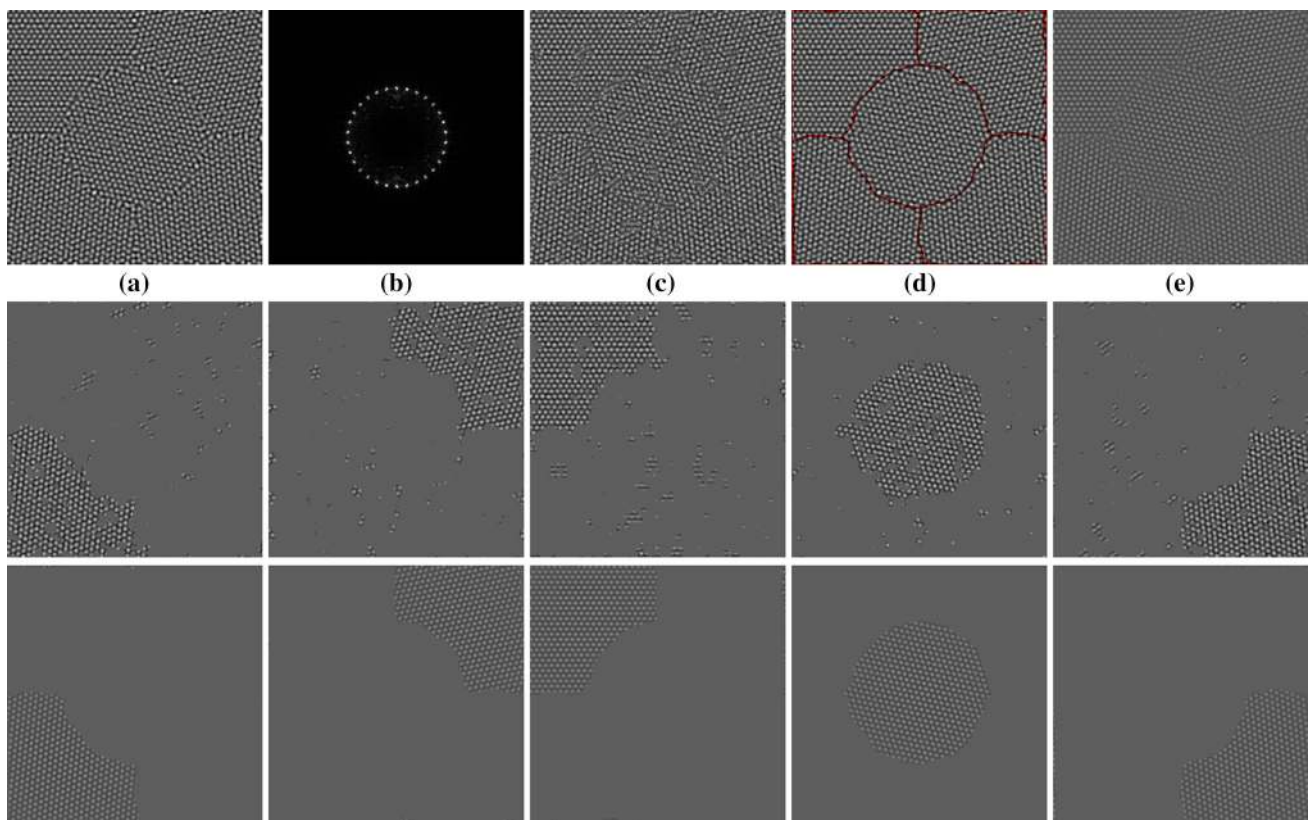


Fig. 11 Simulated crystal lattice. 2D-SEG-VMD decomposition in two runs, first with, then without Lagrangian multipliers. See Sect. 7.6.3 in text for details and discussion. **a** Input image f . **b** Fourier spectrum \hat{f} . **c** First run reconstruction $\sum_{k,i} A_k u_{ki}$. **d** Partition (red) of second

run. **e** Reconstruction of second run. *Middle row* phases obtained in first run with $\tau, \tau_k > 0$ to find correct $\omega_k i$. *Bottom row* clean phases $A_k \sum_i u_{ik}$ of second run with $\tau = \tau_k = 0$ and well-initialized $\omega_k i$ (Color figure online)

Now, the increased α_k renders the modes more pure, and also keeps the center frequencies from drifting too much, while the partition regularity is regularized slightly stronger (increased γ_k). The main difference are the inactivated Lagrangian multipliers, relaxing the data-fidelity constraint considerably. The resulting decomposition is shown in Fig. 11. In the correctly initialized denoising regime, we obtain a very accurate partition and much cleaner crystal grain estimates.

7.6.4 Colloidal Image

As a last example problem, we consider a bright-field light microscopy image of 10 μm -sized spherical glass particles suspended in water.¹¹ These glass particles form a collection of small 2D colloidal crystals with grain boundaries between them. These grains have a hexagonal lattice structure similar to the previously considered examples. For our purposes, the original image is cropped, band-pass filtered with a LoG

filter, and downsampled to a final dimension of 256×256 . The effective input image is shown in Fig. 12a.

Visual inspection of the Fourier spectrum suggests that there are probably four different grain orientations to be found in the image (see Fig. 12b). We thus configure the 2D-SEG-VMD model with the following parameter choice:

K	M	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	T
4	3	2000	1	250	∞	10	50	0.1	0.1	2.5

The resulting grain boundaries shown in Fig. 12c should be compared to computationally determined lattice irregularities (grain boundaries, defects) in Fig. 12d.¹²

8 Conclusions and Outlook

In this paper, we have presented a variational method for decomposing a multidimensional signal, $f: \mathbb{R}^n \rightarrow \mathbb{R}$,

¹¹ Image used with permission, courtesy by Richard Wheeler, Sir William Dunn School of Pathology, University of Oxford, UK.

¹² *Ibid.*




Fig. 12 Bright-field microscopy image of colloidal crystal and its segmentation. Individual beads are $10\ \mu\text{m}$ in diameter. See Sect. 7.6.4. **a** Cropped, LoG-filtered, and downsampled input image f . **b** Fourier spectrum \hat{f} . **c** 2D-SEG-VMD 4-partition (*red*) overlaid on input image. **d** Colloidal connectivity graph for comparison: *white edges* indicate hexagonal alignment (six equally spaced neighbors) and that a particle is therefore part of a crystalline domain (grain), *white colored edges* indicate grain boundaries and defects (Color figure online)

(images for $n = 2$) into ensembles of constituent modes, $u_k: \mathbb{R}^n \rightarrow \mathbb{R}$, intrinsic mode functions that have specific directional and oscillatory characteristics. This multidimensional extension of the variational mode decomposition (VMD) method [17] yields a sparse representation with band-limited modes around a center frequency ω_k , which reconstructs the initial signal, exactly or approximately.

In addition to generalizing the 1D VMD model to higher dimensions, we introduce a binary support function $A_k: \mathbb{R}^n \rightarrow \{0, 1\}$ for each mode u_k , such that the signal decomposition obeys $f \approx \sum_k A_k \cdot u_k$. In order to encourage compact spatial support, an L^1 and a TV penalty term on A_k are introduced. After appropriate variable splitting, we present an ADMM scheme for efficient optimization of this model. In particular, this includes MBO-like threshold dynamics to tackle the motion by mean curvature stemming from the support function regularizing TV term.

In this general setting, our model allows for spatially compact modes that may be spatially overlapping. By restricting the support functions on the probability simplex, $\sum_k A_k = 1$, the modes have mutually exclusive spatial support and actually form a partition of the signal domain. In this fashion, we obtain an image segmentation model that can be seen as a Chan–Vese-like region-based model, where the homo-

geneity is assessed through spectral bandwidth. Our variable splitting and the handling of region boundaries through the binary support functions elegantly overcomes the usual trade-off between spatial and spectral compactness/bandwidth.

In order to deal with images of crystal grains, each region being more complicated than a simple cosine wave, we introduce the coupling of submodes with a single binary support function. This allows the segmentation of crystal grain images, e.g., from microscopy, into respective grains of different lattice orientation. Further, non-Gaussian image noise, outliers, and lattice defects are efficiently addressed by the introduction of an artifact indicator function, $\chi: \mathbb{R}^n \rightarrow \{0, 1\}$.

In summary, the models and algorithms allow decomposing a signal/image into modes that may:

- have smooth or sharp boundaries (with or without TV/L^1 terms on A_k),
- overlap or form a partition of the domain (image segmentation),
- be essentially wavelike (single mode) or crystalline (coupled submodes),
- reconstruct the input image exactly or up to Gaussian noise,
- identify outlier pixels/regions and inpaint them.

Acknowledgements The authors are grateful to Jérôme Gilles for providing the synthetic texture image for Fig. 1; to Diana Yugay for providing peptide β -sheet images used in Figs. 6 and 7; and to Richard Wheeler for permitting the use of the colloidal crystal image of Fig. 12.

References

1. Ayvaci, A., Raptis, M., Soatto, S.: Sparse occlusion detection with optical flow. *Int. J. Comput. Vis.* **97**(3), 322–338 (2011). doi:[10.1007/s11263-011-0490-7](https://doi.org/10.1007/s11263-011-0490-7)
2. Bertsekas, D.P.: Multiplier methods: a survey. *Automatica* **12**(2), 133–145 (1976)
3. Bonnell, D.A., Basov, D.N., Bode, M., Diebold, U., Kalinin, S.V., Madhavan, V., Novotny, L., Salmeron, M., Schwarz, U.D., Weiss, P.S.: Imaging physical phenomena with local probes: from electrons to photons. *Rev. Mod. Phys.* **84**(3), 1343–1381 (2012). doi:[10.1103/RevModPhys.84.1343](https://doi.org/10.1103/RevModPhys.84.1343)
4. Bülow, T., Sommer, G.: A novel approach to the 2D analytic signal. In: *Computer Analysis of Images and Patterns*, pp. 25–32 (1999)
5. Candes, E.J., Donoho, D.L.: Curvelets: A surprisingly effective nonadaptive representation for objects with edges. In: *Curve and Surface Fitting*, pp. 105–120 (1999)
6. Carson, J.: Notes on the theory of modulation. *Proc. IRE* **10**(1), 57–64 (1922). doi:[10.1109/JRPROC.1922.219793](https://doi.org/10.1109/JRPROC.1922.219793)
7. Chambolle, A.: An algorithm for total variation minimization and applications. *J. Math. Imaging Vis.* **20**, 89–97 (2004)
8. Claridge, S.A., Schwartz, J.J., Weiss, P.S.: Electrons, photons, and force: quantitative single-molecule measurements from physics to biology. *ACS Nano* **5**(2), 693–729 (2011). doi:[10.1021/nn103298x](https://doi.org/10.1021/nn103298x)

9. Clausel, M., Oberlin, T., Perrier, V.: The monogenic synchrosqueezed wavelet transform: a tool for the decomposition/demodulation of AM-FM images (2012). [arxiv:1211.5082](https://arxiv.org/abs/1211.5082)
10. Claridge, S.A., Thomas, J.C., Silverman, M.A., Schwartz, J.J., Yang, Y., Wang, C., Weiss, P.S.: Differentiating amino acid residues and side chain orientations in peptides using scanning tunneling microscopy. *J. Am. Chem. Soc.* **135**(49), 18528–18535 (2013). doi:[10.1021/ja408550a](https://doi.org/10.1021/ja408550a)
11. Cohen, L.D.: Auxiliary variables and two-step iterative algorithms in computer vision problems. *J. Math. Imaging Vis.* **6**(1), 59–83 (1996). doi:[10.1007/BF00127375](https://doi.org/10.1007/BF00127375)
12. Combettes, P.L., Pesquet, J.C.: Proximal splitting methods in signal processing. In: Bauschke, H.H., Burachik, R.S., Combettes, P.L., Elser, V., Luke, D.R., Wolkowicz, H. (eds.) *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, Springer Optimization and Its Applications, vol. 49, pp. 185–212. Springer, New York, NY (2011). doi:[10.1007/978-1-4419-9569-8](https://doi.org/10.1007/978-1-4419-9569-8)
13. Daubechies, I.: Orthonormal bases of compactly supported wavelets. *Commun. Pure Appl. Math.* **41**(7), 909–996 (1988). doi:[10.1002/cpa.3160410705](https://doi.org/10.1002/cpa.3160410705)
14. Daubechies, I., Lu, J., Wu, H.T.: Synchrosqueezed wavelet transforms: an empirical mode decomposition-like tool. *Appl. Comput. Harmonic Anal.* **30**(2), 243–261 (2011). doi:[10.1016/j.acha.2010.08.002](https://doi.org/10.1016/j.acha.2010.08.002)
15. Do, M., Vetterli, M.: Pyramidal directional filter banks and curvelets. In: *Proceedings 2001 International Conference on Image Processing*, vol. 2, pp. 158–161. IEEE (2001). doi:[10.1109/ICIP.2001.958075](https://doi.org/10.1109/ICIP.2001.958075)
16. Dong, W., Li, X., Lin, X., Li, Z.: A bidimensional empirical mode decomposition method for fusion of multispectral and panchromatic remote sensing images. *Remote Sens.* **6**(9), 8446–8467 (2014). doi:[10.3390/rs6098446](https://doi.org/10.3390/rs6098446)
17. Dragomiretskiy, K., Zosso, D.: Variational mode decomposition. *IEEE Trans. Signal Process.* **62**(3), 531–544 (2014). doi:[10.1109/TSP.2013.2288675](https://doi.org/10.1109/TSP.2013.2288675)
18. Dragomiretskiy, K., Zosso, D.: Two-dimensional variational mode decomposition. In: Tai, X.C., Bae, E., Chan, T.F., Lysaker, M. (eds.) *EMMCVPR 2015. Lecture Notes in Computer Science*, vol. 8932, pp. 197–208. Springer, Berlin (2015). doi:[10.1007/978-3-319-14612-6_15](https://doi.org/10.1007/978-3-319-14612-6_15)
19. Elsey, M., Wirth, B.: Fast automated detection of crystal distortion and crystal defects in polycrystal images. *Multiscale Model. Simul.* **12**(1), 1–24. doi:[10.1137/130916515](https://doi.org/10.1137/130916515)
20. Esedoglu, S., Otto, F.: Threshold dynamics for networks with arbitrary surface tensions. *Commun. Pure Appl. Math.* **68**(5), 808–864 (2015). doi:[10.1002/cpa.21527](https://doi.org/10.1002/cpa.21527)
21. Esedoglu, S., Tsai, Y.H.R.: Threshold dynamics for the piecewise constant Mumford–Shah functional. *J. Comput. Phys.* **211**, 367–384 (2006). doi:[10.1016/j.jcp.2005.05.027](https://doi.org/10.1016/j.jcp.2005.05.027)
22. Estellers, V., Zosso, D., Bresson, X., Thiran, J.P.: Harmonic active contours. *IEEE Trans. Image Process.* **23**(1), 69–82 (2014). doi:[10.1109/TIP.2013.2286326](https://doi.org/10.1109/TIP.2013.2286326)
23. Fauchereau, N., Pegram, G.G.S., Sinclair, S.: Empirical mode decomposition on the sphere: application to the spatial scales of surface temperature variations. *Hydrol. Earth Syst. Sci.* **12**(3), 933–941 (2008). doi:[10.5194/hess-12-933-2008](https://doi.org/10.5194/hess-12-933-2008)
24. Flandrin, P., Gonçalves, P., Rilling, G.: EMD equivalent filter banks, from interpretation to applications. In: *Hilbert–Huang Transform and Its Applications*, pp. 57–74 (2005)
25. Gabor, D.: Theory of communication. *J. Inst. Electr. Eng. Part III Radio Commun. Eng.* **93**(26), 429–457 (1946)
26. Garcia-Cardona, C., Merkurjev, E., Bertozzi, A.L., Flenner, A., Percus, A.G.: Multiclass data segmentation using diffuse interface methods on graphs. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**(8), 1600–1613 (2014). doi:[10.1109/TPAMI.2014.2300478](https://doi.org/10.1109/TPAMI.2014.2300478)
27. Georgoulas, G., Loutas, T., Stylios, C.D., Kostopoulos, V.: Bearing fault detection based on hybrid ensemble detector and empirical mode decomposition. *Mech. Syst. Signal Process.* **41**(1–2), 510–525 (2013). doi:[10.1016/j.ymssp.2013.02.020](https://doi.org/10.1016/j.ymssp.2013.02.020)
28. Gilles, J.: Multiscale texture separation. *Multiscale Model. Simul.* **10**(4), 1409–1427 (2012). doi:[10.1137/120881579](https://doi.org/10.1137/120881579)
29. Gilles, J.: Empirical wavelet transform. *IEEE Trans. Signal Process.* **61**(16), 3999–4010 (2013). doi:[10.1109/TSP.2013.2265222](https://doi.org/10.1109/TSP.2013.2265222)
30. Gilles, J., Tran, G., Osher, S.: 2D empirical transforms. Wavelets, ridgelets, and curvelets revisited. *SIAM J. Imaging Sci.* **7**(1), 157–186 (2014). doi:[10.1137/130923774](https://doi.org/10.1137/130923774)
31. Glowinski, R., Le Tallec, P.: *Augmented Lagrangian and Operator-Splitting Methods in Nonlinear Mechanics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia (1989)
32. Goldstein, T., Osher, S.: The split Bregman method for L1-regularized problems. *SIAM J. Imaging Sci.* **2**(2), 323–343 (2009). doi:[10.1137/080725891](https://doi.org/10.1137/080725891)
33. Grafakos, L.: *Classical Fourier Analysis*, Graduate Texts in Mathematics, vol. 249. Springer, New York, NY (2009). doi:[10.1007/978-0-387-09432-8](https://doi.org/10.1007/978-0-387-09432-8)
34. Guo, K., Labate, D.: Optimally sparse multidimensional representation using shearlets. *SIAM J. Math. Anal.* **39**(1), 298–318 (2007). doi:[10.1137/060649781](https://doi.org/10.1137/060649781)
35. Han, J., van der Baan, M.: Empirical mode decomposition for seismic time-frequency analysis. *Geophysics* **78**(2), O9–O19 (2013). doi:[10.1190/geo2012-0199.1](https://doi.org/10.1190/geo2012-0199.1)
36. Han, P., Kurland, A.R., Giordano, A.N., Nanayakkara, S.U., Blake, M.M., Pochas, C.M., Weiss, P.S.: Heads and tails: simultaneous exposed and buried interface imaging of monolayers. *ACS Nano* **3**(10), 3115–3121 (2009). doi:[10.1021/nn901030x](https://doi.org/10.1021/nn901030x)
37. Hestenes, M.R.: Multiplier and gradient methods. *J. Optim. Theory Appl.* **4**(5), 303–320 (1969)
38. Hou, T.Y., Shi, Z.: Adaptive data analysis via sparse time-frequency representation. *Adv. Adapt. Data Anal.* **03**(1 & 2), 1–28 (2011). doi:[10.1142/S1793536911000647](https://doi.org/10.1142/S1793536911000647)
39. Hou, T.Y., Shi, Z.: Data-driven time frequency analysis. *Appl. Comput. Harmonic Anal.* **35**(2), 284–308 (2013). doi:[10.1016/j.acha.2012.10.001](https://doi.org/10.1016/j.acha.2012.10.001)
40. Hou, T.Y., Shi, Z.: Sparse time-frequency decomposition for multiple signals with same frequencies (2015). [arxiv:1507.02037](https://arxiv.org/abs/1507.02037)
41. Hu, M., Liang, H.: Intrinsic mode entropy based on multivariate empirical mode decomposition and its application to neural data analysis. *Cogn. Neurodyn.* **5**(3), 277–284 (2011). doi:[10.1007/s11571-011-9159-8](https://doi.org/10.1007/s11571-011-9159-8)
42. Hu, H., Sunu, J., Bertozzi, A.L.: Multi-class graph Mumford–Shah model for plume detection using the MBO scheme. In: Tai, X.C., Bae, E., Chan, T.F., Lysaker, M. (eds.) *Energy Minimization Methods in Computer Vision and Pattern Recognition, Lecture Notes in Computer Science*, vol. 8932, pp. 209–222. Springer, Berlin (2015). doi:[10.1007/978-3-319-14612-6_16](https://doi.org/10.1007/978-3-319-14612-6_16)
43. Huang, N.E., Shen, Z., Long, S.R., Wu, M.C., Shih, H.H., Zheng, Q., Yen, N.C., Tung, C.C., Liu, H.H.: The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **454**(1971), 903–995 (1998). doi:[10.1098/rspa.1998.0193](https://doi.org/10.1098/rspa.1998.0193)
44. Labate, D., Lim, W.Q., Kutyniok, G., Weiss, G.: Sparse multidimensional representation using shearlets. In: Papadakis, M., Laine, A.F., Unser, M.A. (eds.) *Optics and Photonics 2005*, pp. 1–9. International Society for Optics and Photonics, Bellingham (2005). doi:[10.1117/12.613494](https://doi.org/10.1117/12.613494)
45. Lee, T.S.: Image representation using 2D Gabor wavelets. *IEEE Trans. Pattern Anal. Mach. Intell.* **18**(10), 959–971 (1996). doi:[10.1109/34.541406](https://doi.org/10.1109/34.541406)
46. Leo, M., Piccolo, R., Distanto, C., Memmolo, P., Paturzo, M., Ferraro, P.: Multilevel bidimensional empirical mode decomposition:

- a new speckle reduction method in digital holography. *Opt. Eng.* **53**(11), 112,314 (2014). doi:[10.1117/1.OE.53.11.112314](https://doi.org/10.1117/1.OE.53.11.112314)
47. Mallat, S.: A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. Pattern Anal. Mach. Intell.* **11**(7), 674–693 (1989). doi:[10.1109/34.192463](https://doi.org/10.1109/34.192463)
 48. Meignen, S., Perrier, V.: A new formulation for empirical mode decomposition based on constrained optimization. *IEEE Signal Process. Lett.* **14**(12), 932–935 (2007). doi:[10.1109/LSP.2007.904706](https://doi.org/10.1109/LSP.2007.904706)
 49. Merkurjev, E., Garcia-Cardona, C., Bertozzi, A.L., Flenner, A., Percus, A.G.: Diffuse interface methods for multiclass segmentation of high-dimensional data. *Appl. Math. Lett.* **33**, 29–34 (2014). doi:[10.1016/j.aml.2014.02.008](https://doi.org/10.1016/j.aml.2014.02.008)
 50. Merriman, B., Bence, J.K., Osher, S.J.: Motion of multiple junctions: a level set approach. *J. Comput. Phys.* **112**(2), 334–363 (1994). doi:[10.1006/jcph.1994.1105](https://doi.org/10.1006/jcph.1994.1105)
 51. Moore, A.M., Yeganeh, S., Yao, Y., Claridge, S.A., Tour, J.M., Ratner, M.A., Weiss, P.S.: Polarizabilities of adsorbed and assembled molecules: measuring the conductance through buried contacts. *ACS Nano* **4**(12), 7630–7636 (2010). doi:[10.1021/nn102371z](https://doi.org/10.1021/nn102371z)
 52. Nocedal, J., Wright, S.J.: *Numerical Optimization*, 2nd edn. Springer, Berlin (2006)
 53. Nunes, J., Bouaoune, Y., Delechelle, E., Niang, O., Bunel, P.: Image analysis by bidimensional empirical mode decomposition. *Image Vis. Comput.* **21**(12), 1019–1026 (2003). doi:[10.1016/S0262-8856\(03\)00094-5](https://doi.org/10.1016/S0262-8856(03)00094-5)
 54. Rilling, G., Flandrin, P.: One or two frequencies? The empirical mode decomposition answers. *IEEE Trans. Signal Process.* **56**(1), 85–95 (2008). doi:[10.1109/TSP.2007.906771](https://doi.org/10.1109/TSP.2007.906771)
 55. Rilling, G., Flandrin, P.: Sampling effects on the empirical mode decomposition. *Adv. Adapt. Data Anal.* **01**(01), 43–59 (2009). doi:[10.1142/S1793536909000023](https://doi.org/10.1142/S1793536909000023)
 56. Rilling, G., Flandrin, P., Gonçalves, P.: On empirical mode decomposition and its algorithms. In: *IEEE-EURASIP Workshop on Nonlinear Signal and Image Processing* (2003)
 57. Rockafellar, R.T.: A dual approach to solving nonlinear programming problems by unconstrained optimization. *Math. Program.* **5**(1), 354–373 (1973). doi:[10.1007/BF01580138](https://doi.org/10.1007/BF01580138)
 58. Ruuth, S.: Efficient algorithms for diffusion-generated motion by mean curvature. *J. Comput. Phys.* **625**, 603–625 (1998). doi:[10.1006/jcph.1998.6025](https://doi.org/10.1006/jcph.1998.6025)
 59. Schmitt, J., Pustelnik, N., Borgnat, P., Flandrin, P., Condat, L.: 2-D Prony-Huang Transform: A New Tool for 2-D Spectral Analysis, p. 24 (2014). [arxiv:1404.7680](https://arxiv.org/abs/1404.7680)
 60. Schmitt, J., Pustelnik, N., Borgnat, P., Flandrin, P., Condat, L.: 2D Prony-Huang transform: a new tool for 2D spectral analysis. *IEEE Trans. Image Process.* **23**(12), 5233–5248 (2014). doi:[10.1109/TIP.2014.2363000](https://doi.org/10.1109/TIP.2014.2363000)
 61. Sharpley, R.C., Vatchev, V.: Analysis of the intrinsic mode functions. *Constr. Approx.* **24**(1), 17–47 (2005). doi:[10.1007/s00365-005-0603-z](https://doi.org/10.1007/s00365-005-0603-z)
 62. Strikwerda, J.C.: *Finite Difference Schemes and Partial Differential Equations*, 2nd edn. SIAM, Philadelphia (2004)
 63. Sykes, E.C.H., Mantooth, B.A., Han, P., Donhauser, Z.J., Weiss, P.S.: Substrate-mediated intermolecular interactions: a quantitative single molecule analysis. *J. Am. Chem. Soc.* **127**(19), 7255–7260 (2005). doi:[10.1021/ja0472331](https://doi.org/10.1021/ja0472331)
 64. Szuts, Z.B., Blundell, J.R., Chidichimo, M.P., Marotzke, J.: A vertical-mode decomposition to investigate low-frequency internal motion across the Atlantic at 26° N. *Ocean Sci.* **8**(3), 345–367 (2012). doi:[10.5194/os-8-345-2012](https://doi.org/10.5194/os-8-345-2012)
 65. Tang, J., Zhao, L., Yue, H., Yu, W., Chai, T.: Vibration analysis based on empirical mode decomposition and partial least square. *Proc. Eng.* **16**, 646–652 (2011). doi:[10.1016/j.proeng.2011.08.1136](https://doi.org/10.1016/j.proeng.2011.08.1136)
 66. Tavallali, P., Hou, T.Y., Shi, Z.: Extraction of intrawave signals using the sparse time-frequency representation method. *Multiscale Model. Simul.* **12**(4), 1458–1493 (2014). doi:[10.1137/140957767](https://doi.org/10.1137/140957767)
 67. Thomas, J.C., Schwartz, J.J., Hohman, J.N., Claridge, S.A., Auluck, H.S., Serino, A.C., Spokoyny, A.M., Tran, G., Kelly, K.F., Mirkin, C.A., Gilles, J., Osher, S.J., Weiss, P.S.: Defect-tolerant aligned dipoles within two-dimensional plastic lattices. *ACS Nano* **9**(5), 4734–4742 (2015). doi:[10.1021/acsnano.5b01329](https://doi.org/10.1021/acsnano.5b01329)
 68. Thomas, J.C., Goronzy, D.P., Dragomiretskiy, K., Zosso, D., Gilles, J., Osher, S.J., Bertozzi, A.L., Weiss, P.S.: Mapping buried hydrogen-bonding networks. *ACS Nano* **10**(5), 5446–5451 (2016). doi:[10.1021/acsnano.6b01717](https://doi.org/10.1021/acsnano.6b01717)
 69. Unser, M., Sage, D., Van De Ville, D.: Multiresolution monogenic signal analysis using the Riesz–Laplace wavelet transform. *IEEE Trans. Image Process.* **18**(11), 2402–2418 (2009). doi:[10.1109/TIP.2009.2027628](https://doi.org/10.1109/TIP.2009.2027628)
 70. Unser, M., Chenouard, N., Van De Ville, D.: Steerable pyramids and tight wavelet frames in $L_2(\mathbb{R}^d)$. *IEEE Trans. Image Process.* **20**(10), 2705–2721 (2011). doi:[10.1109/TIP.2011.2138147](https://doi.org/10.1109/TIP.2011.2138147)
 71. Wu, Z., Huang, N.E.: Ensemble empirical mode decomposition: a noise-assisted data analysis method. *Adv. Adapt. Data Anal.* **01**(01), 1–41 (2009). doi:[10.1142/S1793536909000047](https://doi.org/10.1142/S1793536909000047)
 72. Wu, H.T., Flandrin, P., Daubechies, I.: One or two frequencies? The synchrosqueezing answers. *Adv. Adapt. Data Anal.* **03**(01n02), 29–39 (2011). doi:[10.1142/S179353691100074X](https://doi.org/10.1142/S179353691100074X)
 73. Yan, M.: Restoration of images corrupted by impulse noise and mixed Gaussian impulse noise using blind inpainting. *SIAM J. Imaging Sci.* **6**(3), 1227–1245 (2013). doi:[10.1137/12087178X](https://doi.org/10.1137/12087178X)
 74. Yang, H., Ying, L.: Synchrosqueezed wave packet transform for 2D mode decomposition. *SIAM J. Imaging Sci.* **6**(4), 1979–2009 (2013). doi:[10.1137/120891113](https://doi.org/10.1137/120891113)
 75. Yang, H., Lu, J., Ying, L.: Crystal image analysis using 2D synchrosqueezed transforms, p. 27 (2014). [arxiv:1402.1262](https://arxiv.org/abs/1402.1262)
 76. Yugay, D., Goronzy, D.P., Kawakami, L.M., Claridge, S.A., Song, T.B., Yan, Z., Xie, Y.H., Gilles, J., Yang, Y., Weiss, P.S.: Copper ion binding site in β -amyloid peptide. *Nano Lett.* **16**(10), 6282–6289 (2016). doi:[10.1021/acs.nanolett.6b02590](https://doi.org/10.1021/acs.nanolett.6b02590)
 77. Zhu, M., Wright, S.J., Chan, T.F.: Duality-based algorithms for total-variation-regularized image restoration. *Comput. Optim. Appl.* **47**(3), 377–400 (2010). doi:[10.1007/s10589-008-9225-2](https://doi.org/10.1007/s10589-008-9225-2)
 78. Zosso, D., An, J., Stevick, J., Takaki, N., Weiss, M., Slaughter, L.S., Cao, H.H., Weiss, P.S., Bertozzi, A.L.: Image segmentation with dynamic artifacts detection and bias correction. *AIMS Journal of Inverse Problems Imaging*, p. 24 (2017) (accepted)



Dominique Zosso received his M.Sc. and Ph.D. degrees in Electrical and Electronics Engineering from Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland, in 2006 and 2011, respectively. From 2012 to 2016, he was a postdoctoral fellow and CAM Assistant Adjunct Professor in the Department of Mathematics at the University of California, Los Angeles (UCLA), partially supported by fellowships from the Swiss National Science Foundation. In 2016, he joined the Department of Mathematical Sciences of Montana State University, Bozeman, as an Assistant Professor in Mathematics. Dr. Zosso's

research interests are variational and PDE methods, PDE on graphs, minimal surfaces, and efficient algorithms to solve inverse problems in imaging, computer vision, data science, and related machine learning applications.



Konstantin Dragomiretskiy received the B.S. (Hons.) degree in Mathematics and the B.A. degree in Economics from the University of California, San Diego, CA, USA, in 2010. Afterward, he received his Ph.D. degree in Applied Mathematics from the University of California, Los Angeles, CA, USA, in 2015. He was a Research Assistant at the Department of Mathematics at UCLA working with Prof. Andrea L. Bertozzi. His research interests include variational and

PDE-based methods applied to signal decomposition and image processing problems.



Andrea L. Bertozzi received the BA, MA, and Ph.D. degrees in Mathematics all from Princeton University, New Jersey, in 1987, 1988, and 1991, respectively. She was on the faculty of the University of Chicago, Illinois, from 1991 to 1995 and Duke University, Durham, North Carolina, from 1995 to 2004. During 1995–1996, she was the Maria Goeppert Mayer distinguished scholar at Argonne National Laboratory. Since 2003, she has been with the University of California,

Los Angeles, as a professor of Mathematics, and currently serves as the director of applied Mathematics. In 2012, she was appointed the Betsy Wood Knapp chair for Innovation and Creativity. Her research interests include machine learning, image processing, cooperative control

of robotic vehicles, swarming, fluid interfaces, and crime modeling. She is a fellow of both the Society for Industrial and Applied Mathematics and the American Mathematical Society; she is a member of the American Physical Society. She has served as a plenary/ distinguished lecturer for both SIAM and AMS and is an associate editor for the SIAM journals *Multiscale Modelling and Simulation*, *Imaging Sciences*, and *Mathematical Analysis*. She also serves on the editorial board of *Interfaces and Free Boundaries*, *Nonlinearity*, *Applied Mathematics Letters*, *Journal of the American Mathematical Society*, *Journal of Nonlinear Science*, *Journal of Statistical Physics*, and *Communications in Mathematical Sciences*. Her past honors include a Sloan Foundation Research Fellowship, the Presidential Career Award for Scientists and Engineers, and the SIAM Kovalevsky Prize in 2009.



Paul S. Weiss is a UC Presidential Chair and Distinguished Professor of Chemistry and Biochemistry and of Materials Science and Engineering at UCLA. He received his S.B. and S.M. from MIT in 1980 and his Ph.D. from UC Berkeley in 1986, all in chemistry. He was a postdoc at AT&T Bell Labs and IBM Almaden Research Center. After 20 years on the faculty at Penn State, he moved to UCLA in 2009. He has held visiting faculty appointments at Caltech, Har-

vard University, Kyoto University, the University of Washington, and elsewhere. He and his interdisciplinary research group focus on exploring the ultimate limits of miniaturization; understanding and controlling single-molecule/assembly structure, function, and spectra; precisely assembled materials; chemical patterning at all scales; nanobiosensor arrays for neuroscience, the microbiome, and medicine; and Developing and applying nanoscale analysis tools. He works closely with Mathematicians to apply ideas from compressive sensing and sparsity to enhancing and enriching multimodal data acquisition and analyses. He is the founding editor-in-chief of *ACS Nano*.