TWO-DIMENSIONAL COMPACT VARIATIONAL MODE DECOMPOSITION

DOMINIQUE ZOSSO, KONSTANTIN DRAGOMIRETSKIY, ANDREA L. BERTOZZI, AND PAUL S. WEISS

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, and/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 re-composition. 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 (MBO,[48]) 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.

1. Introduction

In this paper, we are interested in decomposing images $f\colon \mathbb{R}^n \to \mathbb{R}$ into ensembles of constituent modes (components) that have specific directional and oscillatory characteristics. Simply put, the goal is to retrieve a small number K of modes $u_k\colon \mathbb{R}^n \to \mathbb{R}$, that each have a very limited bandwidth around their characteristic center frequency ω_k. These modes are called intrinsic mode functions (IMF) and can be seen as amplitude- and frequency-modulated (AM-FM) n-D signals ("plane"-waves). Such a mode can have limited spatial support, its local (instantaneous)

Date: November 8, 2015.

2010 *Mathematics Subject Classification.* Primary: 68U10.

Key words and phrases. Image decomposition, Image segmentation, Spatio-spectral decomposition, Microscopy, Crystal grains, Artifact detection, Threshold dynamics, Variational methods, Sparse time-frequency analysis.

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.

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 [34]. In holography, mode decomposition allows reducing speckle [44]. In the fields of energy and power engineering, mode decompositions are used for vibration analysis and fault detection, e.g., [27, 66]. 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 "sub-modes" 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 [72]. 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, 68]. 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 [42] and its more recent derivates, such as [24, 37, 38, 39, 46, 55, 56, 57, 61, 67, 70]. We are interested in the two-dimensional (2D) analogs and extensions of such decomposition problems. The 2D extension of EMD [53] similarily 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 [60], 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 be used to find a sparse representation. 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 an "optimal" superposition of dictionary elements. Though these methods have had success with simple signals, they are still not fully robust to non-stationary waves and require a large, redundant dictionary of elements, which are not reflective of the specifics of the given signal.

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 [65], wavelets [13, 15, 45], curvelets [5], or shearlets [33, 43]. These methods are not adaptive relative to the signal, and can attribute principle 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 [10, 14, 69, 73], 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 Wiener filtering, the 1D Hilbert transform and the analytic signal, and heterodyne demodulation. The goal of 1D-VMD is to decompose an input signal into a discrete number of sub-signals (modes), where each mode has limited bandwidth in the spectral domain. In other words, one requires each mode $u_k\colon \mathbb{R} \to \mathbb{R}$ to be mostly compact around a center pulsation ω_k, which is to be determined along with the decomposition. In order to assess the bandwidth of a mode, 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) The bandwidth is now estimated through the H^1 smoothness (Dirichlet energy) of the demodulated signal. The resulting constrained variational problem is the following:

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

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 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 regularizes 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 hexagonal lattice patterns.

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

2. Two-Dimensional Variational Mode Decomposition

We design the 2D model analogously to its 1D predecessor, minimizing the constituent sub-signals 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 $f : \mathbb{R} \to \mathbb{R}$ as imaginary part [25]:

$$\begin{array}{rrcl} f_{AS}: & \mathbb{R} & \to & \mathbb{C} \\ & f_{AS}(t) & \mapsto & f(t) + j\mathcal{H}\{f\}(t), \end{array} \tag{2}$$

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

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

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:

$$\hat{f}_{AS}(\omega) = \begin{cases} 2\hat{f}(\omega), & \text{if} \quad \omega > 0, \\ \hat{f}(\omega), & \text{if} \quad \omega = 0, \\ 0, & \text{if} \quad \omega < 0, \end{cases} \tag{4}$$

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 zero;[1] this half-plane is chosen relative to a vector, in our case to $\vec{\omega}_k$. Thus the 2D analytic signal of interest can first be defined in the frequency domain by generalizing the concept of half-space spectrum suppression:

$$\begin{aligned} \hat{f}_{AS}(\vec{\omega}) &= \begin{cases} 2\hat{f}(\omega), & \text{if} \quad \langle \vec{\omega}, \vec{\omega}_k \rangle > 0, \\ \hat{f}(\omega), & \text{if} \quad \langle \vec{\omega}, \vec{\omega}_k \rangle = 0, \\ 0, & \text{if} \quad \langle \vec{\omega}, \vec{\omega}_k \rangle < 0, \end{cases} \\ &= (1 + \operatorname{sgn}(\langle \vec{\omega}, \vec{\omega}_k \rangle)) \hat{f}(\vec{\omega}) \end{aligned} \tag{5}$$

where the n-D Fourier transform is defined as

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

The 2D analytic signal in the time domain with the aforementioned Fourier property is given in [4]. 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:

$$\min_{u_k : \, \mathbb{R}^n \to \mathbb{R}, \, \vec{\omega}_k \in \mathbb{R}^n} \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle \vec{\omega}_k, \vec{x} \rangle} \right] \right\|_2^2 \right\} \tag{6}$$
$$\text{s.t.} \quad \forall \vec{x} \in \mathbb{R}^n : \sum_k u_k(\vec{x}) = f(\vec{x}),$$

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. We thus minimize the Dirichlet energy of the modes after half-space spectrum suppression ($u_k \to u_{AS,k}$) and demodulation to baseband ($e^{-j\langle \vec{\omega}_k, \vec{x} \rangle}$), 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

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

augmented Lagrangian, AL, method), and we proceed by alternate direction minimization (ADMM) for optimization [2, 17, 52].

2.3. **Augmented Lagrangian and ADMM Optimization.** To render the constrained minimization problem (6) unconstrained, we include both a quadratic penalty and a Lagrangian multiplier to enforce the fidelity constraint. We thus define the augmented Lagrangian:

$$\begin{aligned}(7)\quad \mathcal{L}(\{u_k\},\{\omega_k\},\lambda) := \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle \vec{\omega}_k, \vec{x}\rangle} \right] \right\|_2^2 \\ + \left\| f(\vec{x}) - \sum u_k(\vec{x}) \right\|_2^2 + \left\langle \lambda(\vec{x}), f(\vec{x}) - \sum u_k(\vec{x}) \right\rangle .\end{aligned}$$

where $\lambda\colon \mathbb{R}^n \to \mathbb{R}$ is the Lagrangian multiplier. We can now solve the unconstrained saddle point problem instead of (6):

$$(8)\qquad \min_{u_k\colon \mathbb{R}^n\to\mathbb{R},\ \vec{\omega}_k\in\mathbb{R}^n} \ \max_{\lambda\colon \mathbb{R}^n\to\mathbb{R}} \mathcal{L}(\{u_k\},\{\omega_k\},\lambda)$$

The solution to the original constrained minimization problem (6) is now found as the saddle point of the augmented Lagrangian $\mathcal{L}$ in a sequence of iterative sub-optimizations called alternate direction method of multipliers (ADMM) [2, 36, 58]. The idea is to iterate the following sequence of variable updates:

$$(9a)\qquad u_k^{t+1} \leftarrow \operatorname*{arg\,min}_{u_k\colon \mathbb{R}^n\to\mathbb{R}} \mathcal{L}\left(\left\{u_{i<k}^{t+1}\right\}, u_k, \left\{u_{i>k}^{t}\right\}, \left\{\omega_i^t\right\}, \lambda^t\right)$$

$$(9b)\qquad \vec{\omega}_k^{t+1} \leftarrow \operatorname*{arg\,min}_{\vec{\omega}_k\in\mathbb{R}^n} \mathcal{L}\left(\left\{u_i^{t+1}\right\}, \left\{\vec{\omega}_{i<k}^{t+1}\right\}, \vec{\omega}_k, \left\{\vec{\omega}_{i>k}^{t}\right\}, \lambda^t\right)$$

$$(9c)\qquad \lambda^{t+1} \leftarrow \lambda^t + \tau\left(f - \sum u_k^{t+1}\right)$$

for $1 > \tau \geq 0$. For simplified notation while considering the subminimization problems (9a) and (9b) in the following paragraphs, we incorporate the Lagrangian multiplier term λ into the quadratic penalty term, and rewrite the objective expression slightly different:

$$\begin{aligned}(10)\quad \mathcal{L}(\{u_k\},\{\omega_k\},\lambda) = \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle \vec{\omega}_k, \vec{x}\rangle} \right] \right\|_2^2 \\ + \left\| f(\vec{x}) - \sum u_k(\vec{x}) + \frac{\lambda(\vec{x})}{2} \right\|_2^2 - \left\| \frac{\lambda(\vec{x})^2}{4} \right\|_2^2\end{aligned}$$

2.4. **Minimization w.r.t. the modes u_k.** The relevant update problem derived from (10) is

$$\begin{aligned}(11)\quad u_k^{n+1} = \operatorname*{arg\,min}_{u_k\colon \mathbb{R}^n\to\mathbb{R}} \Bigg\{ \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle \vec{\omega}_k, \vec{x}\rangle} \right] \right\|_2^2 \\ + \left\| f(\vec{x}) - \sum_i u_i(\vec{x}) + \frac{\lambda(\vec{x})}{2} \right\|_2^2 \Bigg\}\end{aligned}$$

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):

$$(12)\quad \hat{u}_k^{n+1} = \underset{\hat{u}_k | u_k:\, \mathbb{R}^n \to \mathbb{R}}{\arg\min} \Bigg\{ \alpha_k \left\| j\vec{\omega} \left[\hat{u}_{AS,k}(\vec{\omega} + \vec{\omega}_k) \right] \right\|_2^2 + \left\| \hat{f}(\vec{\omega}) - \sum_i \hat{u}_i(\vec{\omega}) + \frac{\hat{\lambda}(\vec{\omega})}{2} \right\|_2^2 \Bigg\}.$$

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:

$$(13)\qquad f(\vec{x}) e^{-j\langle \vec{\omega}_0, \vec{x} \rangle} \xleftrightarrow{\mathcal{F}} \hat{f}(\vec{\omega}) * \delta(\vec{\omega} + \vec{\omega}_0) = \hat{f}(\vec{\omega} + \vec{\omega}_0),$$

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:

$$(14)\quad \hat{u}_k^{n+1} = \underset{\hat{u}_k | u_k:\, \mathbb{R}^n \to \mathbb{R}}{\arg\min} \Bigg\{ \alpha_k \left\| j(\vec{\omega} - \vec{\omega}_k) \left[\hat{u}_{AS,k}(\vec{\omega}) \right] \right\|_2^2 + \left\| \hat{f}(\vec{\omega}) - \sum_i \hat{u}_i(\vec{\omega}) + \frac{\hat{\lambda}(\vec{\omega})}{2} \right\|_2^2 \Bigg\}.$$

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

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

Also, the spectra in the second term have Hermitian symmetry, since they correspond to real signals. Let

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

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:

$$(15)\quad \hat{u}_k^{n+1} = \underset{\hat{u}_k | u_k:\, \mathbb{R}^n \to \mathbb{R}}{\arg\min} \Bigg\{ 2\alpha_k \int_{\Omega_k} |\vec{\omega} - \vec{\omega}_k|^2 |\hat{u}_k(\vec{\omega})|^2 d\vec{\omega} + \int_{\Omega_k} \left| \hat{f}(\vec{\omega}) - \sum_i \hat{u}_i(\vec{\omega}) + \frac{\hat{\lambda}(\vec{\omega})}{2} \right|^2 d\vec{\omega} \Bigg\}.$$

This subminimization problem is now solved by letting the first variation w.r.t. $\hat{u}_k$ vanish[2]. The optimal mode spectrum thus satisfies:

$$(16)\qquad 0 = 2\alpha_k |\vec{\omega} - \vec{\omega}_k|^2 \hat{u}_k - \left(\hat{f}(\vec{\omega}) - \sum_i \hat{u}_i(\vec{\omega}) + \frac{\hat{\lambda}(\vec{\omega})}{2} \right), \quad \forall \vec{\omega} \in \Omega_k.$$

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

$$(17)\qquad \hat{u}_k^{n+1}(\vec{\omega}) = \left(\hat{f}(\vec{\omega}) - \sum_{i \neq k} \hat{u}_i(\vec{\omega}) + \frac{\hat{\lambda}(\vec{\omega})}{2} \right) \frac{1}{1 + 2\alpha_k |\vec{\omega} - \vec{\omega}_k|^2}, \quad \forall \vec{\omega} \in \Omega_k.$$

[2]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.

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

$$\hat{u}_{AS,k}^{n+1}(\vec{\omega}) = \Big(\hat{f}(\vec{\omega}) - \sum_{i \neq k} \hat{u}_i(\vec{\omega}) + \frac{\hat{\lambda}(\vec{\omega})}{2}\Big) \frac{1 + \operatorname{sgn}(\langle \vec{\omega}, \vec{\omega}_k \rangle)}{1 + 2\alpha_k |\vec{\omega} - \vec{\omega}_k|^2}, \quad \forall \vec{\omega} \in \mathbb{R}^n, \tag{18}$$

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 $\hat{f}(\vec{\omega}) - \sum_{i \neq k} \hat{u}_i(\vec{\omega})$ 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, $\vec{\omega}_k$, and whose bandwidth is controlled by the parameter α_k.

2.5. **Minimization w.r.t. the center frequencies $\vec{\omega}_k$.** Optimizing for $\vec{\omega}_k$ is even simpler. Indeed, the respective update goal derived from (10) is

$$\vec{\omega}_k^{n+1} = \underset{\vec{\omega}_k \in \mathbb{R}^n}{\arg\min} \left\{ \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j \langle \vec{\omega}_k, \vec{x} \rangle} \right] \right\|_2^2 \right\}. \tag{19}$$

Or, again we may consider the equivalent problem in the Fourier domain:

$$\begin{aligned} \vec{\omega}_k^{n+1} &= \underset{\vec{\omega}_k \in \mathbb{R}^n}{\arg\min} \left\{ \alpha_k \left\| j(\vec{\omega} - \vec{\omega}_k)(1 + \operatorname{sgn}(\langle \vec{\omega}_k, \vec{\omega} \rangle)) \hat{u}_k(\vec{\omega}) \right\|_2^2 \right\} \\ &= \underset{\vec{\omega}_k \in \mathbb{R}^n}{\arg\min} \left\{ 4\alpha_k \int_{\Omega_k} |\vec{\omega} - \vec{\omega}_k|^2 \left| \hat{u}_k(\vec{\omega}) \right|^2 d\vec{\omega} \right\}. \end{aligned} \tag{20}$$

The minimization is solved by letting the first variation w.r.t. $\vec{\omega}_k$ vanish, leading to:

$$\int_{\Omega_k} (\vec{\omega} - \vec{\omega}_k^{n+1}) \left| \hat{u}_k(\vec{\omega}) \right|^2 d\vec{\omega} = 0. \tag{21}$$

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

$$\vec{\omega}_k^{n+1} = \frac{\int_{\Omega_k} \vec{\omega} |\hat{u}_k(\vec{\omega})|^2 d\vec{\omega}}{\int_{\Omega_k} |\hat{u}_k(\vec{\omega})|^2 d\vec{\omega}} = \frac{\int_{\mathbb{R}^n} \vec{\omega} |\hat{u}_{AS,k}(\vec{\omega})|^2 d\vec{\omega}}{\int_{\mathbb{R}^n} |\hat{u}_{AS,k}(\vec{\omega})|^2 d\vec{\omega}}, \tag{22}$$

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, $f(\vec{\omega}) - \sum_k u_k^{n+1}(\vec{\omega})$. We use a standard gradient ascent with fixed time step $1 > \tau \geq 0$ to achieve this maximization:

$$\lambda^{n+1}(\vec{x}) = \lambda^n(\vec{x}) + \tau \left(f(\vec{x}) - \sum_k u_k^{n+1}(\vec{x}) \right). \tag{23}$$

It is important to note that choosing $\tau = 0$ effectively 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, we reconstruction error actually allows capturing noise separately.

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 (18) is in spectral terms, as well. Thus:

$$\hat{\lambda}^{n+1}(\vec{\omega}) = \hat{\lambda}^{n}(\vec{\omega}) + \tau \left(\hat{f}(\vec{\omega}) - \sum_k \hat{u}_k^{n+1}(\vec{\omega}) \right). \tag{24}$$

2.5.2. *Complete 2D VMD algorithm.* The entire proposed algorithm for the 2D-VMD functional optimization problem (6) is summarized in algorithm 1. Variables are trivially initialized at 0, except for the center frequencies, $\vec{\omega}_k$, for which smart initialization is of higher importance; initial $\vec{\omega}_k^0$ 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 $f(\vec{x})$, number of modes K, parameters α_k, τ, ϵ.
Output: modes $u_k(\vec{x})$, center frequencies $\vec{\omega}_k$.

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:

$$\mathcal{H}_k^{n+1}(\vec{\omega}) \leftarrow 1 + \operatorname{sgn}(\langle \vec{\omega}_k^n, \vec{\omega} \rangle)$$

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

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

    Update $\vec{\omega}_k$:

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

    Retrieve u_k:

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

  end for
  Dual ascent (optional):

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

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

An example of image decomposition achieved with 2D VMD according to algorithm 1 is shown in figure 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]:

$$BW_{\text{AM-FM}} := 2(\Delta f + f_{\text{FM}} + f_{\text{AM}}), \tag{25}$$

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

$$A_k \colon \mathbb{R}^n \to \{0, 1\}$$

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

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

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:

$$\begin{aligned} &\min_{u_k \colon \mathbb{R}^n \to \mathbb{R},\, A_k \colon \mathbb{R}^n \to \{0,1\},\, \vec{\omega}_k \in \mathbb{R}^n} \\ &\left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle \vec{\omega}_k, \vec{x} \rangle} \right] \right\|_2^2 + \beta_k \|A_k\|_1 + \gamma_k \|\nabla A_k\|_1 \right\} \\ &\qquad \text{s.t.} \quad \forall \vec{x} \in \mathbb{R}^n \colon \sum_k A_k(\vec{x}) u_k(\vec{x}) = f(\vec{x}). \end{aligned} \tag{26}$$

The L^1 penalties on A_k and ∇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 on modes to 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-based prior terms, the functional is no longer directly translatable to the spectral domain. Moreover, the L^1-terms do not lend themselves to standard calculus of variations methods, directly. Instead, we propose an ensemble of splitting techniques [11, 12, 31] that have been applied to L^1-based and related optimization problems with great success, such as [22, 32, 76].

First, we would like 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. Spectral solvability for u_k is restored by introducing a splitting of the modes $u_k = v_k$, and applying spectral bandwidth penalty and reconstruction over the separate copies:

$$\begin{aligned}(27)\qquad &\min_{u_k:\,\mathbb{R}^n\to\mathbb{R},\ A_k:\,\mathbb{R}^n\to\{0,1\},\ \vec{\omega}_k\in\mathbb{R}^n} \\ &\left\{\sum_k \alpha_k \left\|\nabla\left[u_{AS,k}(\vec{x})e^{-j\langle\vec{\omega}_k,\vec{x}\rangle}\right]\right\|_2^2 + \beta_k\|A_k\|_1 + \gamma_k\|\nabla A_k\|_1\right\} \\ &\qquad\qquad \text{s.t.}\quad \forall\vec{x}\in\mathbb{R}^n:\ \begin{cases} u_k(\vec{x}) = v_k(\vec{x}), \\ \sum_k A_k(\vec{x})v_k(\vec{x}) = f(\vec{x}).\end{cases}\end{aligned}$$

The splitting constraint can be addressed with a quadratic penalty (proximal splitting, [12]), or using an augmented Lagrangian [31]. As an intermediate illustration, and since the latter includes the former, we give the full saddle-point functional (augmented Lagrangian) incorporating both equality constraints through quadratic penalty and Lagrangian multipliers, in analogy to (7):

$$\begin{aligned}(28)\quad &\mathcal{L}(\{u_k\},\{v_k\},\{A_k\},\{\omega_k\},\lambda,\{\lambda_k\}) := \\ &\left\{\sum_k \alpha_k \left\|\nabla\left[u_{AS,k}(\vec{x})e^{-j\langle\vec{\omega}_k,\vec{x}\rangle}\right]\right\|_2^2 + \beta_k\|A_k\|_1 + \gamma_k\|\nabla A_k\|_1\right. \\ &+ \rho\left\|f(\vec{x}) - \sum A_k(\vec{x})v_k(\vec{x})\right\|_2^2 + \left\langle \lambda(\vec{x}), f(\vec{x}) - \sum A_k(\vec{x})v_k(\vec{x})\right\rangle \\ &\qquad \left. + \sum_k \rho_k \left\|u_k(\vec{x}) - v_k(\vec{x})\right\|_2^2 + \langle\lambda_k(\vec{x}), u_k(\vec{x}) - v_k(\vec{x})\rangle\right\},\end{aligned}$$

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. All terms involving u_k translate nicely into the spectral domain, while all terms in v_k lend themselves to efficient point-wise optimization in time domain. Before actually looking at the specific subminimization problems, we want to study the L^1-terms further by recognizing them as essentially balloon and motion-by-mean-curvature forces acting on the binary support functions A_k.

3.4. **Excursion on MBO.** The first variation associated with the TV-term is proportional to $\operatorname{div}(\nabla A_k/|\nabla A_k|)$. One can expect difficulties with this term, for example in flat regions where $|\nabla A_k| \to 0$. Moreover, if the gradient descent PDE is integrated explicitly, then the time step is also heavily limited by the stiffness constraint [62].

An important contribution stems from the diffusion-threshold scheme for approximating motion by mean curvature proposed by Merriman, Bence, and Osher (MBO) [49]. The fundamental idea is to reproduce the motion by mean curvature due to the boundary-length term $TV(A_k)$ by more efficient means than direct gradient descent.

Since A_k is binary we opt for alternative schemes other than split-Bregman/shrinkage or dual minimization [7, 32, 74]. As a preliminary, motivational step, let us replace the total variation of the support function A_k, by the real Ginzburg-Landau (GL, also known as Allen-Cahn) functional [50]:

$$E^{\epsilon}_{\mathrm{GL}}(A_k) := \epsilon \int_{\Omega} |\nabla A_k(\vec{x})|^2 d\vec{x} + \frac{1}{\epsilon}\int_{\Omega} W(A_k(\vec{x}))d\vec{x}, \quad \epsilon > 0, \tag{29}$$

where $W(s)$ is a double-well potential with two equal minima at $s = 0$ and $s = 1$, for example $W(s) := s^2(1-s)^2$. Minimizing this functional yields a phase field that is smooth and tends to be binary. In particular, it has been shown [50] that the GL-functional Γ-converges to the total variation functional of binary phase-fields $A_k \in \{0, 1\}$ as $\epsilon \to 0$:

$$E^{0}_{\mathrm{GL}}(A_k) = \sigma(W) \int_{\Omega} |\nabla A_k|, \tag{30}$$

where $\sigma(W)$ is a surface tension term depending on the double well potential. The minimizing flow of this functional for $\epsilon \to 0^+$ produces motion by mean curvature of the interface, which is exactly what one needs in the spatially sparse VMD model minimization. However, now, the PDE associated with the GL-functional minimization is

$$\frac{\partial A_k}{\partial t} = 2\epsilon\nabla^2 A_k - \frac{1}{\epsilon}W'(A_k), \tag{31}$$

and this PDE is conveniently solved in a discrete-time two step time-splitting approach:

(1) Propagate A_k according to the heat equation,

$$\frac{\partial A_k}{\partial t} = 2\epsilon\nabla^2 A_k$$

(2) Propagate A_k according to the double well potential gradient descent,

$$\frac{\partial A_k}{\partial t} = -\frac{1}{\epsilon}W'(A_k).$$

The heat equation is efficiently solved, e.g., based on convolution or spectral transforms [59].

Now, the MBO-scheme [49] improves on this time-split GL-optimization in that the ODE is recognized as essentially performing thresholding. While the first step is reduced to propagation according to the standard heat equation, the second step in MBO is actual thresholding (projection onto the binary set $\{0, 1\}$):

(1) Propagate A_k according to the heat equation,

$$\frac{\partial A_k}{\partial t} = \nabla^2 A_k$$

(2) Rectify A_k by thresholding:

$$A_k(\vec{x}) = \begin{cases} 0 & \text{if } A_k(\vec{x}) \le \frac{1}{2} \\ 1 & \text{if } A_k(\vec{x}) > \frac{1}{2} \end{cases} \qquad \forall \vec{x} \in \mathbb{R}^n$$

These MBO threshold dynamics have already been successfully integrated with imaging data terms, such as [21, 75], 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.

3.5. ***n*-D-TV-VMD Minimization.** Based on the preparatory steps of the preceding sections, we now propose to solve the constraint, sparsity promoting n-D VMD functional (26) through its augmented Lagrangian (28). Consider the following saddle point problem:

$$\min_{u_k,v_k\,:\,\mathbb{R}^n\to\mathbb{R},\ A_k\,:\,\mathbb{R}^n\to\{0,1\},\ \vec{\omega}_k\in\mathbb{R}^n} \quad \max_{\lambda,\lambda_k\,:\,\mathbb{R}^n\to\mathbb{R}} \left\{ \mathcal{L}(\{u_k\},\{v_k\},\{A_k\},\{\omega_k\},\lambda,\{\lambda_k\}) \right\}. \tag{32}$$

This saddle point problem is an extended version of the 2D VMD saddle point problem (8) (without spatial sparsity promoting terms), and is again efficiently solved through alternate direction minimization and dual ascent (ADMM):

$$u_k^{t+1} \leftarrow \operatorname*{arg\,min}_{u_k\,:\,\mathbb{R}^n\to\mathbb{R}} \mathcal{L}\left(\left\{u_{i<k}^{t+1}\right\}, u_k, \left\{u_{i>k}^t\right\}, \left\{v_i^t\right\}, \left\{A_i^t\right\}, \left\{\omega_i^t\right\}, \lambda^t, \left\{\lambda_i^t\right\}\right) \tag{33a}$$

$$v_k^{t+1} \leftarrow \operatorname*{arg\,min}_{v_k\,:\,\mathbb{R}^n\to\mathbb{R}} \mathcal{L}\left(\left\{u_i^{t+1}\right\}, \left\{v_{i<k}^{t+1}\right\}, v_k, \left\{v_{i>k}^t\right\}, \left\{A_i^t\right\}, \left\{\omega_i^t\right\}, \lambda^t, \left\{\lambda_i^t\right\}\right) \tag{33b}$$

$$A_k^{t+1} \leftarrow \operatorname*{arg\,min}_{A_k\,:\,\mathbb{R}^n\to\{0,1\}} \mathcal{L}\left(\left\{u_i^{t+1}\right\}, \left\{v_i^{t+1}\right\}, \left\{A_{i<k}^{t+1}\right\}, A_k, \left\{A_{i>k}^t\right\}, \left\{\omega_i^t\right\}, \lambda^t, \left\{\lambda_i^t\right\}\right) \tag{33c}$$

$$\vec{\omega}_k^{t+1} \leftarrow \operatorname*{arg\,min}_{\vec{\omega}_k\in\mathbb{R}^n} \mathcal{L}\left(\left\{u_i^{t+1}\right\}, \left\{v_i^{t+1}\right\}, \left\{A_i^{t+1}\right\}, \left\{\vec{\omega}_{i<k}^{t+1}\right\}, \vec{\omega}_k, \left\{\vec{\omega}_{i>k}^t\right\}, \lambda^t, \left\{\lambda_i^t\right\}\right) \tag{33d}$$

$$\lambda^{t+1} \leftarrow \lambda^t + \tau\left(f - \sum A_k^{t+1} v_k^{t+1}\right) \tag{33e}$$

$$\lambda_k^{t+1} \leftarrow \lambda_k^t + \tau_k\left(u_k^{t+1} - v_k^{t+1}\right) \tag{33f}$$

We provide details on the individual sub-minimization problems in the following paragraphs. The complete algorithm for n-D-TV-VMD functional (with spatial sparsity promoting terms) is then easily derived in analogy to algorithm 1.

3.5.1. *n-D-TV-VMD Subminimization w.r.t. u_k.* The relevant minimization problem (33a) with respect to the modes u_k reads

$$u_k^{t+1} = \operatorname*{arg\,min}_{u_k\,:\,\mathbb{R}^n\to\mathbb{R}} \left\{ \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle\vec{\omega}_k,\vec{x}\rangle} \right] \right\|_2^2 + \rho_k \left\| u_k(\vec{x}) - v_k(\vec{x}) + \frac{\lambda_k(\vec{x})}{\rho_k} \right\|_2^2 \right\}. \tag{34}$$

In full analogy to the problem without spatial sparsity terms, (11), the update is most easily computed in spectral domain, like (17). Unsurprisingly, the update rule on the frequency halfspace $\Omega_k = \{\omega \mid \langle\omega,\omega_k\rangle \geq 0\}$ is found to be:

$$\hat{u}_k^{t+1}(\vec{\omega}) = (\rho_k \hat{v}_k - \hat{\lambda}_k)\frac{1}{\rho_k + 2\alpha_k|\omega-\omega_k|^2}, \quad \forall\omega\in\Omega_k. \tag{35}$$

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:

$$\hat{u}_{AS,k}^{t+1}(\vec{\omega}) = (\rho_k \hat{v}_k - \hat{\lambda}_k)\frac{1+\operatorname{sgn}(\langle\omega,\omega_k\rangle)}{\rho_k + 2\alpha_k|\omega-\omega_k|^2}. \tag{36}$$

3.5.2. *n-D-TV-VMD Subminimization w.r.t. v_k.* The update (33b) of v_k reduces to the following minimization problem:

$$v_k^{t+1} = \operatorname*{arg\,min}_{v_k\,:\,\mathbb{R}^n\to\mathbb{R}} \left\{ \rho \left\| f(\vec{x}) - \sum A_i(\vec{x}) v_i(\vec{x}) + \frac{\lambda(\vec{x})}{\rho} \right\|_2^2 + \rho_k \left\| u_k(\vec{x}) - v_k(\vec{x}) + \frac{\lambda_k(\vec{x})}{\rho_k} \right\|_2^2 \right\} \tag{37}$$

This problem admits the following pointwise Euler-Lagrange equations:

$$-\rho A_k(\vec{x})\left(f(\vec{x}) - \sum A_i(\vec{x}) v_i(\vec{x}) + \frac{\lambda(\vec{x})}{\rho} \right) - \rho_k \left(u_k(\vec{x}) - v_k(\vec{x}) + \frac{\lambda_k(\vec{x})}{\rho_k} \right) = 0, \quad \forall\vec{x}\in\mathbb{R}^n \tag{38}$$

yielding the simple update rule

$$v_k^{t+1}(\vec{x}) = \frac{\rho A_k(\vec{x})\left(f(\vec{x}) - \sum_{i\neq k} A_i(\vec{x}) v_i(\vec{x}) + \frac{\lambda(\vec{x})}{\rho} \right) + \rho_k u_k(\vec{x}) + \lambda_k(\vec{x})}{\rho A_k(\vec{x})^2 + \rho_k}, \quad \forall\vec{x}\in\mathbb{R}^n. \tag{39}$$

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).

3.5.3. *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-based priors:

$$A_k^{t+1} = \operatorname*{arg\,min}_{A_k\,:\,\mathbb{R}^n\to\{0,1\}} \Bigg\{ \beta_k \|A_k\|_1 + \gamma_k \|\nabla A_k\|_1 + \rho \left\| f(\vec{x}) - \sum A_i(\vec{x}) v_i(\vec{x}) + \frac{\lambda(\vec{x})}{\rho} \right\|_2^2 \Bigg\}. \tag{40}$$

Motivated by successful implementation for image segmentation problems, for example, we want to employ diffusion and threshold dynamics for the efficient solution of this problem. In analogy to the image segmentation scheme, we devise a threefold time-split gradient descent iteration: The first step is gradient descent based on the support area and reconstruction-fidelity penalty. The second step is diffusion by the heat equation, followed by thresholding, to deal with the boundary length term and the projection on 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 $\beta_k \int_{\mathbb{R}^n} A_k$. This makes 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 [21, 49, 75] by iterating over the following three evolution equations:

(1) Area penalty and reconstruction fidelity ODE:

$$\frac{\partial A_k(\vec{x})}{\partial t} = -\beta_k + 2\rho v_k(\vec{x}) \left(f(\vec{x}) - \sum A_i(\vec{x}) v_i(\vec{x}) + \frac{\lambda(\vec{x})}{\rho} \right), \tag{41}$$

(2) Heat equation PDE for diffusion:

$$\frac{\partial A_k(\vec{x})}{\partial t} = \gamma_k \nabla^2 A_k(\vec{x}), \tag{42}$$

(3) Rectification by thresholding:

$$A_k(\vec{x}) = \begin{cases} 0 & \text{if } A_k(\vec{x}) \le \frac{1}{2} \\ 1 & \text{if } A_k(\vec{x}) > \frac{1}{2} \end{cases} \qquad \forall \vec{x} \in \mathbb{R}^n. \tag{43}$$

Note that the ODE problem can be addressed through an implicit (backward) Euler scheme, and the heat equation PDE is efficiently solved spectrally.

3.5.4. *n-D-TV-VMD Subminimization w.r.t.* ω_k. The last, remaining sub-problem of the saddle-point problem (32) is the update of the mode's central frequency, ω_k. The relevant portion of the functional (28) is identical to the non-sparse 2D-VMD model (7). Therefore the corresponding subminimization problem here is identical to (19), and thus the update is equally given by (22).

The complete algorithm for the ADMM optimization of the 2D-TV-VMD model is shown in algorithm 2, and illustrative examples of its use are given in figures 1 and 2.

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

Input: signal $f(\vec{x})$, number of modes K, parameters α_k, β_k, γ_k, ρ, ρ_k, t, τ, τ_k, ϵ.

Output: modes $u_k(\vec{x})$, support functions $A_k(\vec{x})$, center frequencies $\vec{\omega}_k$.

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

repeat

$n \leftarrow n+1$

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

Create 2D mask for analytic signal Fourier multiplier:

$$\mathcal{H}_k^{n+1}(\vec{\omega}) \leftarrow 1 + \operatorname{sgn}(\langle \vec{\omega}_k^n, \vec{\omega} \rangle)$$

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

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

Retrieve u_k:

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

Update v_k:

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

Update A_k through modified MBO:

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

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

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

Update $\vec{\omega}_k$:

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

Dual ascent u-v coupling:

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

end for

Dual ascent data fidelity:

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

until convergence

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 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.

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

$$\sum_k A_k(\vec{x}) = 1, \qquad \forall \vec{x} \in \mathbb{R}^n. \tag{44}$$

In return, the area penalty $\beta_k \|A_k\|_1$ is 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 $(i,j) \in \{1,\dots,K\}^2$.

We propose the following spatially disjoint n-D-TV-VMD model, as a modification of (26):

$$\begin{aligned} &\min_{u_k\colon \mathbb{R}^n \to \mathbb{R},\ A_k\colon \mathbb{R}^n \to \{0,1\},\ \vec{\omega}_k \in \mathbb{R}^n} \\ &\qquad \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle \vec{\omega}_k, \vec{x}\rangle} \right] \right\|_2^2 + \beta_k \|A_k\|_1 + \gamma_k \|\nabla A_k\|_1 \right\} \\ &\qquad\qquad \text{s.t.} \quad \forall \vec{x} \in \mathbb{R}^n\colon \begin{cases} \sum_k A_k(\vec{x}) u_k(\vec{x}) = f(\vec{x}), \\ \sum_k A_k(\vec{x}) = 1. \end{cases} \end{aligned} \tag{45}$$

Next we outline two different strategies to accommodate this extra constraint on the support functions in the minimization scheme. The first strategy incorporates the partitioning constraint through another augmented Lagrangian to be included in the saddle point problem. The second model deals with the restricted solution space through projection, more precisely by modifying the current rectification step included in the MBO-like diffusion and threshold-dynamics.

4.1. Augmented Lagrangian method.

In the first approach, we incorporate the segmentation constraint as a third augmented Lagrangian term. Based on the AL (28) of the spatially overlapping compact VMD functional (26), we write:

$$\begin{aligned} &\mathcal{L}(\{u_k\},\{v_k\},\{A_k\},\{\omega_k\},\lambda,\{\lambda_k\}) := \\ &\quad \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle \vec{\omega}_k, \vec{x}\rangle} \right] \right\|_2^2 + \beta_k \|A_k\|_1 + \gamma_k \|\nabla A_k\|_1 \right. \\ &\quad + \rho \left\| f(\vec{x}) - \sum A_k(\vec{x}) v_k(\vec{x}) \right\|_2^2 + \left\langle \lambda(\vec{x}), f(\vec{x}) - \sum A_k(\vec{x}) v_k(\vec{x}) \right\rangle \\ &\quad + \sum \rho_k \left\| u_k(\vec{x}) - v_k(\vec{x}) \right\|_2^2 + \sum \left\langle \lambda_k(\vec{x}), u_k(\vec{x}) - v_k(\vec{x}) \right\rangle \\ &\quad \left. + \rho' \left\| \sum A_k(\vec{x}) - 1 \right\|_2^2 + \left\langle \lambda'(\vec{x}), \sum A_k(\vec{x}) - 1 \right\rangle \right\}, \end{aligned} \tag{46}$$

where $\lambda'\colon \mathbb{R}^n \to \mathbb{R}$ is the newly introduced Lagrangian multiplier, and ρ' the weight of the corresponding quadratic penalty term. Sticking to the alternate direction gradient descent and dual ascent scheme (32) for optimization, we realize that all sub-optimization problems remain unchanged, except for the A_k update and an additional dual ascent step.

The heat diffusion and thresholding steps are not affected by the extra terms in the functional. Instead, the corresponding first variation is incorporated in the first, ODE step (41), as follows:

$$\begin{aligned}\frac{\partial A_k(\vec{x})}{\partial t} = -\beta + 2\rho v_k(\vec{x}) \left(f(\vec{x}) - \sum A_i(\vec{x}) v_i(\vec{x}) + \frac{\lambda(\vec{x})}{\rho} \right) \\ - 2\rho'(\sum A_k(\vec{x}) - 1) - \lambda'(\vec{x}).\end{aligned} \tag{47}$$

4.2. **Projection: Multiphase MBO and rearrangement.** Instead of the additional penalty and Lagrangian multiplier term, the partitioning constraint can be dealt with by the rectification step in the MBO-like part. Indeed, the partitioning problem corresponds to a multiphase interface problem. The fundamental idea is to propagate the data-ODE (41) and the heat-diffusion PDE (42) on each support function A_k individually, but to replace the individual thresholding step (43) by a single, common "winner-takes-it-all" rectification. This idea has been discussed more rigorously in [20], and is related to the rearrangement algorithm for the discrete graph partitioning problem [54].

The projection-based partitioning update for A_k becomes:

(1) Area penalty and reconstruction fidelity ODE propagation for each mode k, according to (41).
(2) Heat diffusion PDE for each mode k according to (42).
(3) "Winner-takes-it-all" rectification; Projection of the intermediate A_k on the feasible set $A_k \in \{0,1\} \cap \sum A_k = 1$:

$$A_k^{t+1} = \begin{cases} 1 & \text{if} \quad k = \arg\max_i A_i, \\ 0 & \text{otherwise.} \end{cases} \tag{48}$$

For an application of the same strategy to graph-based image processing see [26, 40, 47]. The modified 2D-TV-VMD algorithm with segmentation constraint is given in algorithm 3, while illustrative examples are shown in figures 3 *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 homogeneity 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 modelled by Bravais lattices, which, depending on the 2D crystalline arrangement, come in five forms: oblique, rectangular, centered rectangular, hexagonal,

Algorithm 3 2D-TV-VMD with segmentation constraint

Input: signal $f(\vec{x})$, number of modes K, parameters α_k, β_k, γ_k, ρ, ρ_k, t, τ, τ_k, ϵ.

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

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

repeat

$n \leftarrow n+1$

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

Create 2D mask for analytic signal Fourier multiplier:

$$\mathcal{H}_k^{n+1}(\vec{\omega}) \leftarrow 1 + \operatorname{sgn}(\langle \vec{\omega}_k^n, \vec{\omega} \rangle)$$

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

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

Retrieve u_k:

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

Update v_k:

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

Update $\vec{\omega}_k$:

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

Dual ascent u-v coupling:

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

end for

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

Update A_k through time split ODE and PDE propagation:

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

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

end for

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

Rectify A_k through winner-takes-it-all:

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

end for

Dual ascent data fidelity:

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

until convergence

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 [72]. 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-TV-VMD model (45) as follows:

$$\min_{u_{ki}\,:\,\mathbb{R}^n\to\mathbb{R},\ A_k\,:\,\mathbb{R}^n\to\{0,1\},\ \vec{\omega}_{ki}\in\mathbb{R}^n} \left\{ \sum_{k,i} \alpha_{ki} \left\| \nabla \left[u_{AS,ki}(\vec{x}) e^{-j\langle \vec{\omega}_{ki}, \vec{x} \rangle} \right] \right\|_2^2 + \sum_k \beta_k \|A_k\|_1 + \sum_k \gamma_k \|\nabla A_k\|_1 \right\}$$
$$\text{s.t.} \quad \forall \vec{x} \in \mathbb{R}^n : \begin{cases} \sum_k A_k(\vec{x}) \sum_i u_{ki}(\vec{x}) = f(\vec{x}), \\ \sum_k A_k(\vec{x}) = 1. \end{cases} \tag{49}$$

We call this the n-D-TV-VMD lattice segmentation model. The model can be optimized in much the same way as the simpler model (45). 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, (41) now becomes:

$$\frac{\partial A_k(\vec{x})}{\partial t} = -\beta + 2\rho \left(\sum_i v_{ki}(\vec{x}) \right) \left(f(\vec{x}) - \sum_l A_l(\vec{x}) \sum_j v_{lj}(\vec{x}) + \frac{\lambda(\vec{x})}{\rho} \right). \tag{50}$$

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 figures 7–10.

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 [75]. 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 [71]. Here, the goal is to isolate defects and artifacts from interfering with the regular modes.

We introduce an artifact indicator function,

$$\chi\colon \mathbb{R}^n \to \{0,1\},$$

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.,

$$\forall \vec{x} \in \mathbb{R}^n \mid \chi(\vec{x}) = 0\colon \sum_k A_k(\vec{x})u_k(\vec{x}) = f(\vec{x}). \tag{51}$$

This is equivalent to

$$\forall \vec{x} \in \mathbb{R}^n\colon \sum_k (1-\chi(\vec{x}))A_k(\vec{x})u_k(\vec{x}) = (1-\chi(\vec{x}))f(\vec{x}), \tag{52}$$

where $(1-\chi(\vec{x})) = 1$ 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 $f(\vec{x})$ as an artifact or defect, $\chi(\vec{x}) = 1$, 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 (26) to become the n-D-TV-XVMD (with artifact detection) functional as follows:

$$\begin{aligned}(53)\quad &\min_{u_k:\,\mathbb{R}^n\to\mathbb{R},\ A_k,\chi:\,\mathbb{R}^n\to\{0,1\},\ \vec{\omega}_k\in\mathbb{R}^n} \\ &\left\{\sum_k \alpha_k \left\|\nabla\left[u_{AS,k}(\vec{x})e^{-j\langle\vec{\omega}_k,\vec{x}\rangle}\right]\right\|_2^2 + \beta_k\|A_k\|_1 + \gamma_k\|\nabla A_k\|_1 + \delta\|\chi\|_1\right\} \\ &\qquad\qquad \text{s.t.}\quad \forall\vec{x}\in\mathbb{R}^n:\ \sum_k(1-\chi(\vec{x}))A_k(\vec{x})u_k(\vec{x}) = (1-\chi(\vec{x}))f(\vec{x}).\end{aligned}$$

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

$$\begin{aligned}(54)\quad &\mathcal{L}(\{u_k\},\{v_k\},\{A_k\},\{\omega_k\},\chi,\{\lambda_k\}) := \\ &\left\{\sum_k \alpha_k \left\|\nabla\left[u_{AS,k}(\vec{x})e^{-j\langle\vec{\omega}_k,\vec{x}\rangle}\right]\right\|_2^2 + \beta_k\|A_k\|_1 + \gamma_k\|\nabla A_k\|_1\right. \\ &\quad + \delta\|\chi\|_1 + \rho\left\|(1-\chi(\vec{x}))(f(\vec{x}) - \sum A_k(\vec{x})v_k(\vec{x}))\right\|_2^2 \\ &\quad\left. + \sum_k \rho_k\left\|u_k(\vec{x}) - v_k(\vec{x})\right\|_2^2 + \langle\lambda_k(\vec{x}), u_k(\vec{x}) - v_k(\vec{x})\rangle\right\}.\end{aligned}$$

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

(1) the area penalty and reconstruction fidelity ODE (41) will collapse to just $\partial_t A_k(\vec{x}) = -\beta_k$ whenever $\chi(\vec{x}) = 1$ (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 (39) of v_k collapses to $v_k^{t+1}(\vec{x}) = u_k(\vec{x}) + \lambda_k(\vec{x})/\rho_k$ when $\chi(\vec{x}) = 1$, which effectively unlinks the local mode estimate from the observed data and simply *in-paints 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 $\chi^*(\vec{x})$ chooses between paying data-fidelity penalty versus artifact cost δ, as follows:

$$\chi^*(\vec{x}) = \begin{cases} 0 & \text{if} \quad \rho(f(\vec{x}) - \sum A_k(\vec{x})v_k(\vec{x}))^2 \le \delta \\ 1 & \text{otherwise}\end{cases} \tag{55}$$

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 $\rho(f(\vec{x}) - \sum A_k(\vec{x})v_k(\vec{x}))^2$ 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

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

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 figure 6.

7. Experiments and Results

We have implemented the above three algorithms 1–3, including the submode coupling of section §5 and the artifacts detection and inpainting (§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).

Our implementation is publicly available for download at `http://www.math.ucla.edu/~zosso/code.html`, and on 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 J. 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 figure 1. The parameters are[3]:

[3] 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.

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 figure 1 show clear separation of the patterned ellipses.

In the simple 2D-VMD model of figure 1(e), 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, figure 1(f)–(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.

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 figure 2). 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 that the decay distance correlates with the wave-length 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 way to go. This 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 figure 3. 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

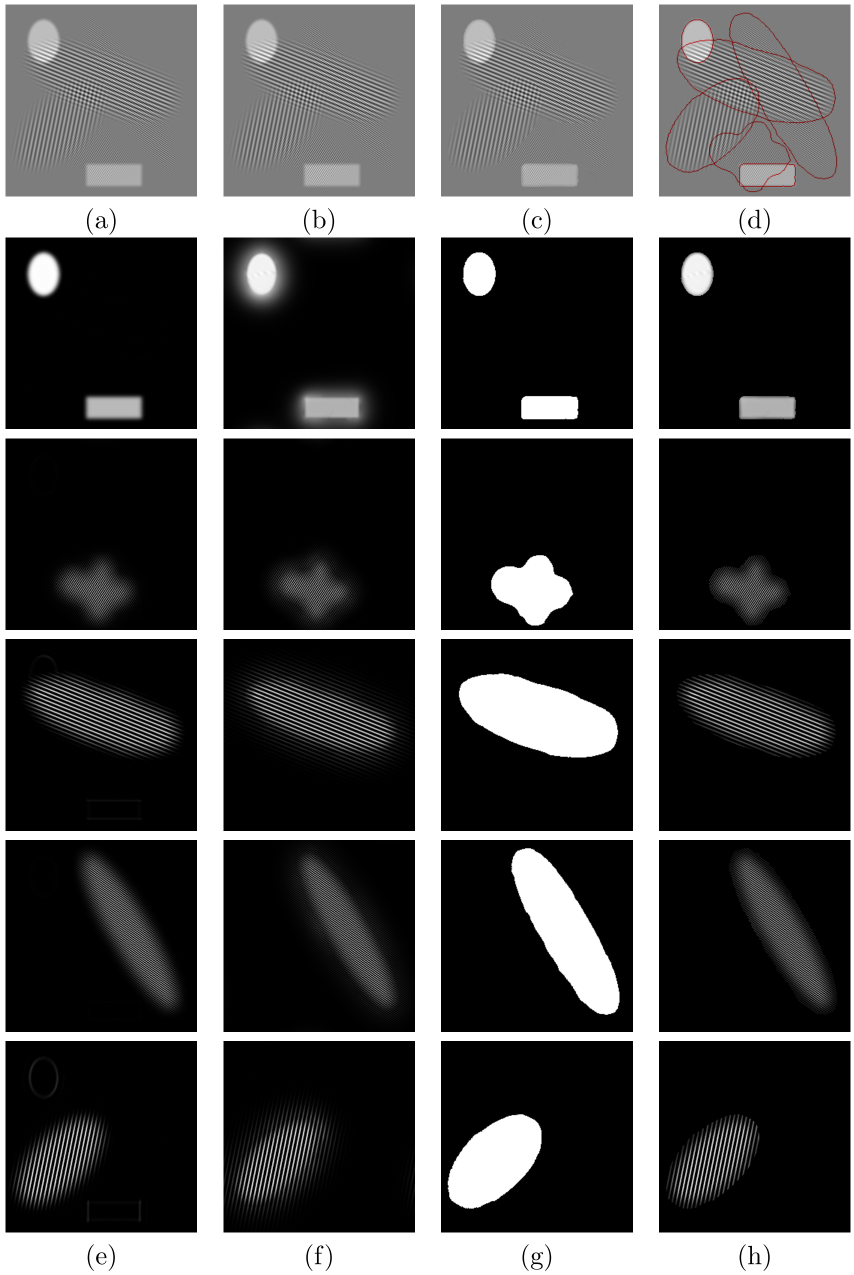

FIGURE 1. Synthetic overlapping texture (a) Input image f. (b) 2D-VMD reconstruction $\sum_k u_k$. (c) Compactly supported 2D-TV-VMD reconstruction $\sum_k A_k u_k$. (d) Support boundaries overlaid onto original image. (e) 2D-VMD modes u_k. (f) 2D-TV-VMD modes u_k. (g) Detected supports A_k. (h) Masked modes $A_k u_k$. See §7.1 in text for explanation and discussion.

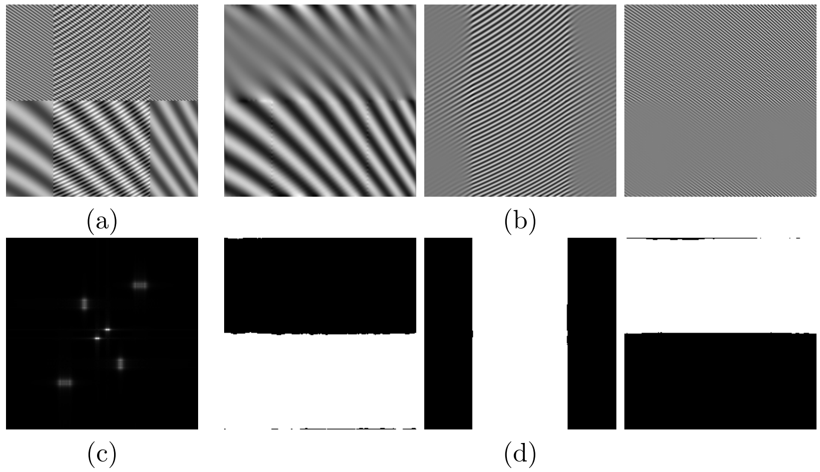

FIGURE 2. Chirp decomposition. (a) Input signal f. (b) 2D-TV-VMD modes u_k. (c) Fourier spectrum $\hat{f}$. (d) Determined supports A_k. See §7.2.

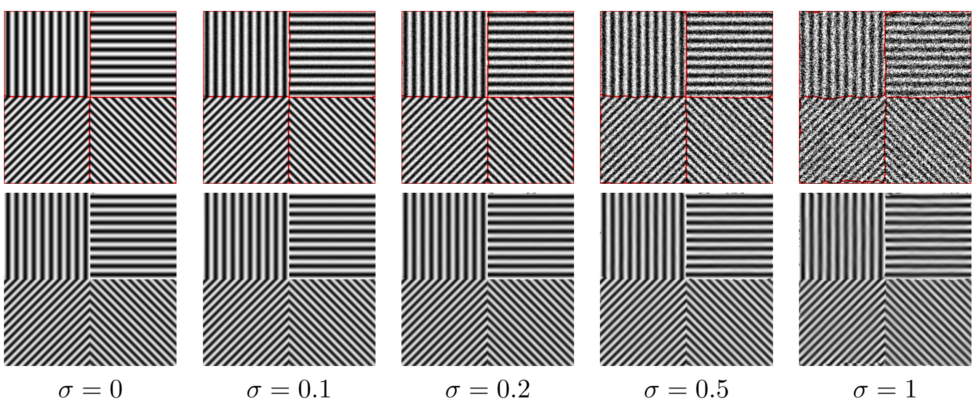

FIGURE 3. Denoising. Noise standard deviation σ. Top: noisy f with detected phase borders (red). Bottom: denoised signal $\sum_k A_k u_k$. See §7.3.

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 figure 3, 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.

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, [9]. The peptide sheets grow in regions of directional homogeneity and form natural

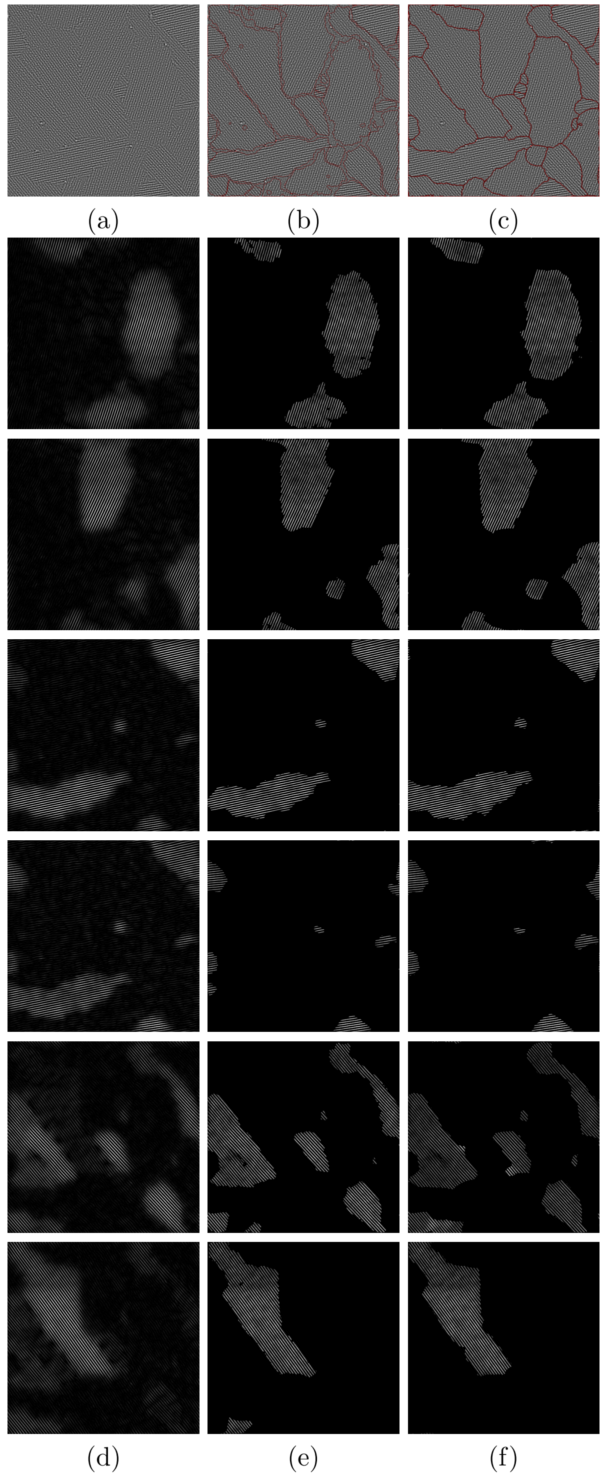

FIGURE 4. Atomic force microscopy (AFM) image of peptide β-sheets, 512×512 pixels, 500 nm $\times$ 500 nm (I). (a) Input f. (b) 2D-TV-VMD boundaries (red). (c) 2D-SEG-VMD partition (red). (d) 2D-VMD modes u_k. (e) 2D-TV-VMD modes $A_k u_k$. (f) 2D-SEG-VMD modes $A_k u_k$. See §7.4 in the text for details and discussion.

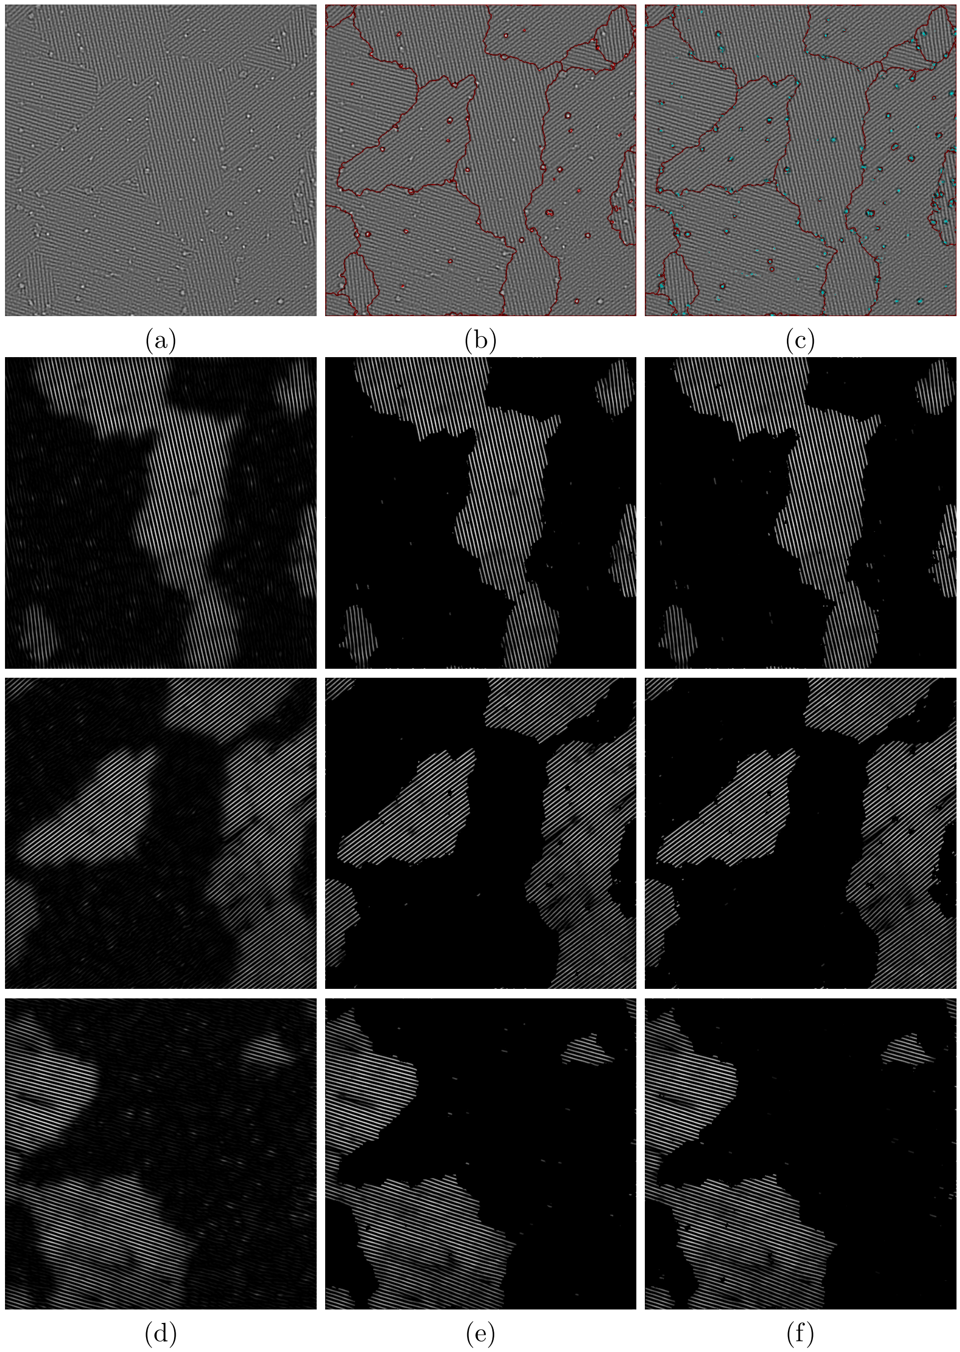

Figure 5. Atomic force microscopy (AFM) image of peptide β-sheets, 512×512 (II). (a) Input f. (b) 2D-SEG-VMD partition (red). (c) Partition (red) with enabled artifact detection (cyan). (d) 2D-VMD modes u_k. (e) 2D-SEG-VMD modes $A_k u_k$. (f) Modes obtained with artifact detection enabled. See §7.4 for details.

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 over-sampled 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, 35, 51, 68].

The first example, shown in figure 4, 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 pre-processing 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 figure 4(d)–(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 figure 4(b), 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 4(c).

The second example, shown in figure 5, is believed to consist of only three main grain orientations. This 512×512 image is of the same type as the previous example and pre-processed 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 (figure 5(d)–(e)), this effect is partially alleviated by the automatic detection and inpainting capability of the artifacts-extension (figure 5(f))[4]. In addition to the outlined grain boundaries (red), the location of the detected artifacts is highlighted in cyan, in figure 5(c). Note that the artifact detection also allows spotting at least some of the grain defects, in addition to the deposits.

[4]Lower artifact threshold δ and higher TV-weight γ_k might increase the mode cleanliness even further.

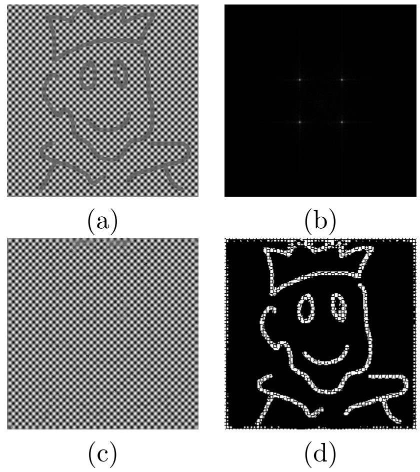

FIGURE 6. 2D-VMD inpainting. (a) Input image f. (b) Fourier spectrum $\hat{f}$. (c) Recovered modes $\sum_k u_k$. (d) Detected artifacts χ. See §7.5.

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 figure 6(a). We set up the model as a two-modes 2D-VMD image decomposition 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 figure 6(c)–(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 wavelike sub-bands. 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 by 60° rotation. Our model allows for multiple sub-modes u_{ki} to share a common support function A_k, and thus be spatially coupled.

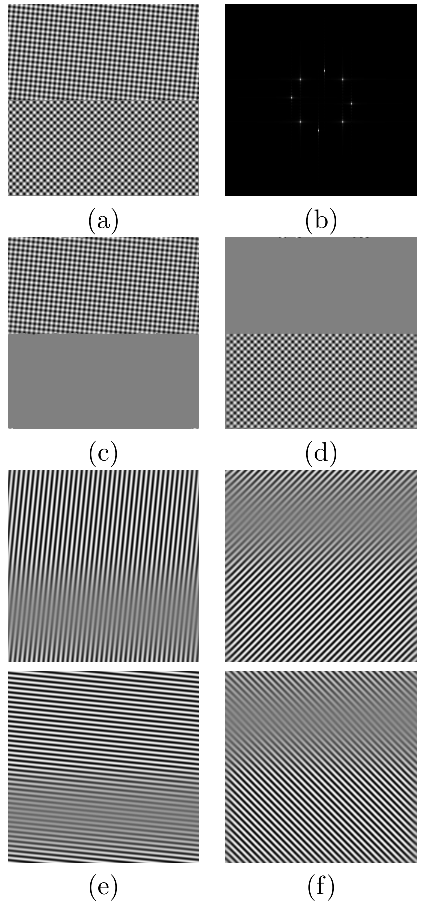

FIGURE 7. Lattice decomposition. (a) Input f. (b) Fourier spectrum $\hat{f}$. (c)–(d) Recovered phases $\sum_i A_k u_{ki}$. (e)–(f) Submodes u_{ki}. See §7.6.1.

7.6.1. *Checkerboard: 2 phases with 2 sub-modes.* As a first simple example, we consider the composite of two checkerboard halves, of which one is slightly rotated, as shown in figure 7(a). The goal is to find the support of two phases, partitioning the 256×256 image domain, and the respective two sub-modes 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 figure 7(c)–(d), while the constituting two sub-modes per phase, u_{ki}, are illustrated in figure 7(e)–(f).

7.6.2. *Hexagonal lattice: 3 phases with 3 sub-modes.* A slightly more complicated problem is illustrated in figure 8. We start with a tripartite 256×256 image, where each domain consists of an artificial hexagonal lattice pattern, obtained by

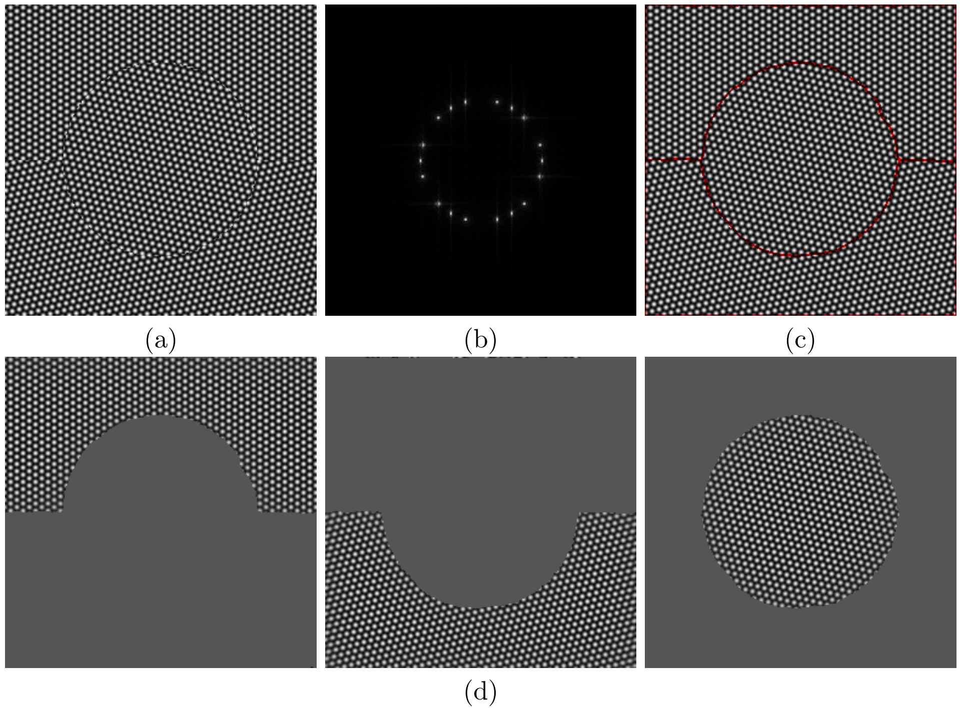

FIGURE 8. 3 phase 3 modes. (a) Input f. (b) Fourier spectrum $\hat{f}$. (c) 2D-SEG-VMD partition (red). (d) Phases $A_k \sum_i u_{ki}$. See §7.6.2.

superposing three cosine waves rotated by $60\,°$ against each other. Each domain has a slightly different lattice orientation ($0\,°, 15\,°, 45\,°$). Like the previous example, this is a 2D-SEG-VMD problem, this time with three phases and three sub-modes, 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 figure 8(c)–(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 3-phase-3-waves hexagonal lattice image of the previous sub-section 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

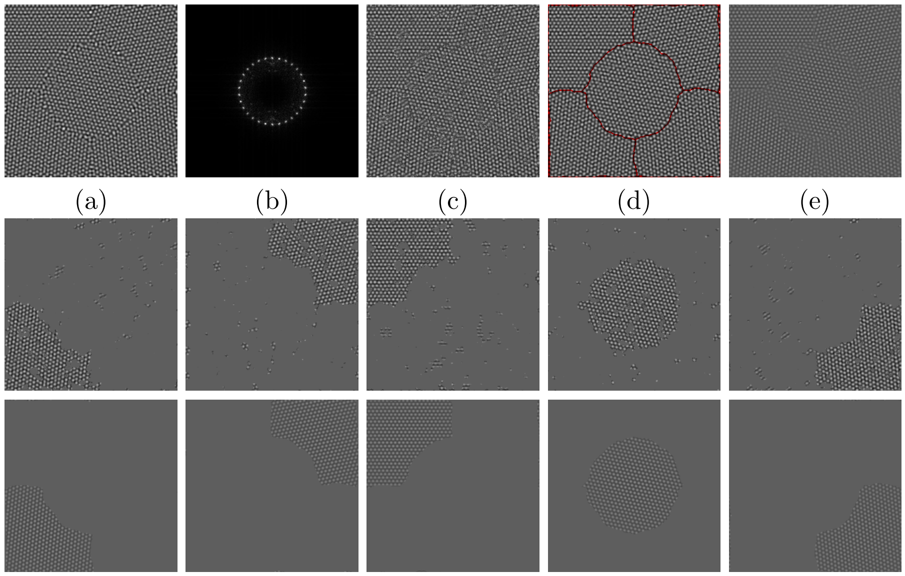

FIGURE 9. Simluated crystal lattice. 2D-SEG-VMD decomposition in two runs, first with, then without Lagrangian multipliers. See §7.6.3 in text for details and discussion. (a) Input image f. (b) Fourier spectrum $\hat{f}$. (c) First run reconstruction $\sum_{k,i} A_k u_{ki}$. (d) Partition (red) of second run. (e) Reconstruction of second run. Middle row: Phases obtained in first run with with $\tau, \tau_k > 0$ to find correct ω_{ki}. Bottom row: Clean phases $A_k \sum_i u_{ik}$ of second run with $\tau = \tau_k = 0$ and well-initialized ω_{ki}.

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 figure 9(a). 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 figure 9. 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

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 figure 9. 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[5]. 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 figure 10(a).

Visual inspection of the Fourier spectrum suggests that there are probably four different grain orientations to be found in the image (see figure 10(b)). 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 figure 10(c) should be compared to computationally determined lattice irregularities (grain boundaries, defects) in figure 10(d)[6].

8. Conclusions and Outlook

In this paper, we have presented a variational method for decomposing a multidimensional signal, $f\colon \mathbb{R}^n \to \mathbb{R}$, (images for $n = 2$) into ensembles of constituent modes, $u_k\colon \mathbb{R}^n \to \mathbb{R}$, intrinsic mode functions which 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 $A_k\colon \mathbb{R}^n \to \{0, 1\}$ 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

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

[6] *Ibid.*

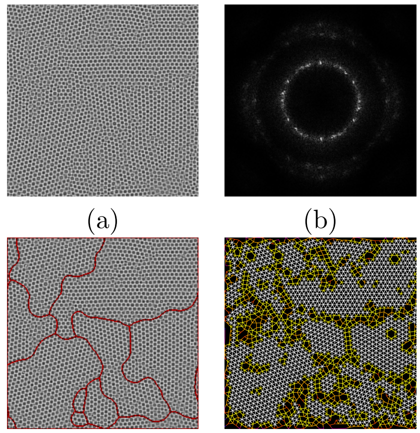
(a) (b) (c) (d)

FIGURE 10. Bright-field microscopy image of colloidal crystal and its segmentation. Individual beads are 10 μm in diameter. See §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), while colored edges indicate grain boundaries and defects.

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 homogeneity is assessed through spectral bandwidth. Our variable splitting and the handling of region boundaries through the binary support functions elegantly overcomes the usual tradeoff 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 sub-modes 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, $\chi\colon \mathbb{R}^n \to \{0,1\}$.

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 sub-modes),
- 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 figure 1; to Diana Yugay for providing peptide β-sheet images used in figures 4&5; and to Richard Wheeler for permitting the use of the colloidal crystal image of figure 10.

References

[1] Alper Ayvaci, Michalis Raptis, and Stefano Soatto. Sparse occlusion detection with optical flow. *International Journal of Computer Vision*, 97(3):322–338, October 2011.
[2] Dimitri P Bertsekas. Multiplier methods: A survey. *Automatica*, 12(2):133–145, 1976.
[3] Dawn A. Bonnell, D. N. Basov, Matthias Bode, Ulrike Diebold, Sergei V. Kalinin, Vidya Madhavan, Lukas Novotny, Miquel Salmeron, Udo D. Schwarz, and Paul S. Weiss. Imaging physical phenomena with local probes: from electrons to photons. *Reviews of Modern Physics*, 84(3):1343–1381, September 2012.
[4] Thomas Bülow and Gerald Sommer. A novel approach to the 2D analytic signal. In *Computer Analysis of Images and Patterns*, pages 25–32, 1999.
[5] Emmanuel J. Candes and David L. Donoho. Curvelets: A surprisingly effective nonadaptive representation for objects with edges. In *Curve and Surface Fitting*, pages 105–120, 1999.
[6] J.R. Carson. Notes on the theory of modulation. *Proceedings of the IRE*, 10(1):57–64, February 1922.
[7] Antonin Chambolle. An algorithm for total variation minimization and applications. *Journal of Mathematical Imaging and Vision*, 20:89–97, 2004.
[8] Shelley A. Claridge, Jeffrey J. Schwartz, and Paul S. Weiss. Electrons, photons, and force: quantitative single-molecule measurements from physics to biology. *ACS Nano*, 5(2):693–729, February 2011.
[9] Shelley A Claridge, John Christopher Thomas, Miles A Silverman, Jeffrey J Schwartz, Yanlian Yang, Chen Wang, and Paul S Weiss. Differentiating amino acid residues and side chain orientations in peptides using scanning tunneling microscopy. *Journal of the American Chemical Society*, 135(49):18528–18535, November 2013.
[10] Marianne Clausel, Thomas Oberlin, and Valérie Perrier. The monogenic synchrosqueezed wavelet transform: A tool for the decomposition/demodulation of AM-FM images, November 2012.
[11] Laurent D. Cohen. Auxiliary variables and two-step iterative algorithms in computer vision problems. *J. Math. Imaging Vis.*, 6(1):59–83, January 1996.
[12] Patrick L. Combettes and Jean-Christophe Pesquet. Proximal splitting methods in signal processing. In Heinz H. Bauschke, Regina S. Burachik, Patrick L. Combettes, Veit Elser, D. Russell Luke, and Henry Wolkowicz, editors, *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, volume 49 of *Springer Optimization and Its Applications*, pages 185–212. Springer New York, New York, NY, 2011.
[13] Ingrid Daubechies. Orthonormal bases of compactly supported wavelets. *Communications on Pure and Applied Mathematics*, 41(7):909–996, October 1988.
[14] Ingrid Daubechies, Jianfeng Lu, and Hau-Tieng Wu. Synchrosqueezed wavelet transforms: An empirical mode decomposition-like tool. *Applied and Computational Harmonic Analysis*, 30(2):243–261, March 2011.
[15] M.N. Do and M. Vetterli. Pyramidal directional filter banks and curvelets. In *Proceedings 2001 International Conference on Image Processing*, volume 2, pages 158–161. IEEE, 2001.
[16] Weihua Dong, Xian'en Li, Xiangguo Lin, and Zhilin Li. A bidimensional empirical mode decomposition method for fusion of multispectral and panchromatic remote sensing images. *Remote Sensing*, 6(9):8446–8467, September 2014.
[17] Konstantin Dragomiretskiy and Dominique Zosso. Variational mode decomposition. *IEEE Transactions on Signal Processing*, 62(3):531–544, February 2014.
[18] Konstantin Dragomiretskiy and Dominique Zosso. Two-dimensional variational mode decomposition. In Xue-Cheng Tai, Egil Bae, Tony F. Chan, and Marius Lysaker, editors, *Energy Minimization Methods in Computer Vision and Pattern Recognition*, volume 8932 of *Lecture Notes in Computer Science*, pages 197–208. Springer International Publishing, 2015.

[19] Matt Elsey and Benedikt Wirth. Fast automated detection of crystal distortion and crystal defects in polycrystal images. *Multiscale Modeling & Simulation*, January 2014.
[20] Selim Esedoglu and Felix Otto. Threshold dynamics for networks with arbitrary surface tensions. *Communications on Pure and Applied Mathematics*, 68(5):808–864, 2015.
[21] Selim Esedoglu and Yen Hsi Richard Tsai. Threshold dynamics for the piecewise constant Mumford-Shah functional. *Journal of Computational Physics*, 211:367–384, 2006.
[22] Virginia Estellers, Dominique Zosso, Xavier Bresson, and Jean-Philippe Thiran. Harmonic active contours. *IEEE Transactions on Image Processing*, 23(1):69–82, January 2014.
[23] N. Fauchereau, G. G. S. Pegram, and S. Sinclair. Empirical mode decomposition on the sphere: application to the spatial scales of surface temperature variations. *Hydrology and Earth System Sciences*, 12(3):933–941, June 2008.
[24] Patrick Flandrin, Paulo Gonçalvès, and Gabriel Rilling. EMD equivalent filter banks, from interpretation to applications. In *Hilbert-Huang Transform and Its Applications*, pages 57–74. 2005.
[25] D. Gabor. Theory of communication. *Journal of the Institution of Electrical Engineers - Part III: Radio and Communication Engineering*, 93(26):429–457, 1946.
[26] Cristina Garcia-Cardona, Ekaterina Merkurjev, Andrea L Bertozzi, Arjuna Flenner, and Allon G. Percus. Multiclass data segmentation using diffuse interface methods on graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(8):1600–1613, 2014.
[27] George Georgoulas, Theodore Loutas, Chrysostomos D. Stylios, and Vassilis Kostopoulos. Bearing fault detection based on hybrid ensemble detector and empirical mode decomposition. *Mechanical Systems and Signal Processing*, 41(1-2):510–525, December 2013.
[28] Jérôme Gilles. Multiscale texture separation. *Multiscale Modeling & Simulation*, 10(4):1409–1427, December 2012.
[29] Jérôme Gilles. Empirical wavelet transform. *IEEE Transactions on Signal Processing*, 61(16):3999–4010, August 2013.
[30] Jérôme Gilles, Giang Tran, and Stanley Osher. 2D empirical transforms. Wavelets, ridgelets, and curvelets revisited. *SIAM Journal on Imaging Sciences*, 7(1):157–186, January 2014.
[31] Roland Glowinski and Patrick Le Tallec. *Augmented Lagrangian and operator-splitting methods in nonlinear mechanics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, 1989.
[32] Tom Goldstein and Stanley Osher. The split Bregman method for L1-regularized problems. *SIAM Journal on Imaging Sciences*, 2(2):323–343, January 2009.
[33] Kanghui Guo and Demetrio Labate. Optimally sparse multidimensional representation using shearlets. *SIAM Journal on Mathematical Analysis*, 39(1):298–318, January 2007.
[34] Jiajun Han and Mirko van der Baan. Empirical mode decomposition for seismic time-frequency analysis. *GEOPHYSICS*, 78(2):O9–O19, February 2013.
[35] Patrick Han, Adam R. Kurland, Andrea N. Giordano, Sanjini U. Nanayakkara, Meaghan M. Blake, Chris M. Pochas, and Paul S. Weiss. Heads and tails: simultaneous exposed and buried interface imaging of monolayers. *ACS Nano*, 3(10):3115–3121, October 2009.
[36] Magnus R. Hestenes. Multiplier and gradient methods. *Journal of Optimization Theory and Applications*, 4(5):303–320, 1969.
[37] Thomas Y. Hou and Zuoqiang Shi. Adaptive data analysis via sparse time-frequency representation. *Advances in Adaptive Data Analysis*, 03(1 & 2):1–28, April 2011.
[38] Thomas Y. Hou and Zuoqiang Shi. Data-driven timefrequency analysis. *Applied and Computational Harmonic Analysis*, 35(2):284–308, September 2013.
[39] Thomas Y. Hou and Zuoqiang Shi. Sparse time-frequency decomposition for multiple signals with same frequencies. July 2015.
[40] Huiyi Hu, Justin Sunu, and Andrea L. Bertozzi. Multi-class Graph Mumford-Shah Model for Plume Detection Using the MBO scheme. volume 8932 of *Lecture Notes in Computer Science*, pages 209–222. Springer International Publishing, 2015.
[41] Meng Hu and Hualou Liang. Intrinsic mode entropy based on multivariate empirical mode decomposition and its application to neural data analysis. *Cognitive Neurodynamics*, 5(3):277–284, September 2011.
[42] N. E. Huang, Z. Shen, S. R. Long, M. C. Wu, H. H. Shih, Q. Zheng, N.-C. Yen, C. C. Tung, and H. H. Liu. The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 454(1971):903–995, March 1998.

[43] Demetrio Labate, Wang-Q Lim, Gitta Kutyniok, and Guido Weiss. Sparse multidimensional representation using shearlets. In Manos Papadakis, Andrew F. Laine, and Michael A. Unser, editors, *Optics & Photonics 2005*, pages 1–9. International Society for Optics and Photonics, August 2005.

[44] Marco Leo, Roberta Piccolo, Cosimo Distante, Pasquale Memmolo, Melania Paturzo, and Pietro Ferraro. Multilevel bidimensional empirical mode decomposition: a new speckle reduction method in digital holography. *Optical Engineering*, 53(11):112314, June 2014.

[45] S.G. Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):674–693, July 1989.

[46] S. Meignen and V. Perrier. A new formulation for empirical mode decomposition based on constrained optimization. *IEEE Signal Processing Letters*, 14(12):932–935, December 2007.

[47] Ekaterina Merkurjev, Cristina Garcia-Cardona, Andrea L. Bertozzi, Arjuna Flenner, and Allon G. Percus. Diffuse interface methods for multiclass segmentation of high-dimensional data. *Applied Mathematics Letters*, 33:29–34, July 2014.

[48] B Merriman, J K Bence, and S Osher. Diffusion generated motion by mean curvature. Technical report, UCLA CAM Report 92-18, 1992.

[49] Barry Merriman, James K Bence, and Stanley J Osher. Motion of multiple junctions: a level set approach. *Journal of Computational Physics*, 112(2):334–363, June 1994.

[50] Luciano Modica. The gradient theory of phase transitions and the minimal interface criterion. *Archive for Rational Mechanics and Analysis*, 98:123–142, 1987.

[51] Amanda M. Moore, Sina Yeganeh, Yuxing Yao, Shelley A. Claridge, James M. Tour, Mark A. Ratner, and Paul S. Weiss. Polarizabilities of adsorbed and assembled molecules: measuring the conductance through buried contacts. *ACS Nano*, 4(12):7630–7636, December 2010.

[52] J Nocedal and S J Wright. *Numerical optimization*. Springer, Berlin, 2nd edition, 2006.

[53] J.C Nunes, Y Bouaoune, E Delechelle, O Niang, and Ph Bunel. Image analysis by bidimensional empirical mode decomposition. *Image and Vision Computing*, 21(12):1019–1026, November 2003.

[54] Braxton Osting, Chris D. White, and Édouard Oudet. Minimal Dirichlet energy partitions for graphs. *SIAM Journal on Scientific Computing*, 36(4):A1635–A1651, August 2014.

[55] G. Rilling and P. Flandrin. One or two frequencies? The empirical mode decomposition answers. *IEEE Transactions on Signal Processing*, 56(1):85–95, January 2008.

[56] Gabriel Rilling and Patrick Flandrin. Sampling effects on the empirical mode decomposition. *Advances in Adaptive Data Analysis*, 01(01):43–59, January 2009.

[57] Gabriel Rilling, Patrick Flandrin, and Paulo Gonçalvès. On empirical mode decomposition and its algorithms. In *IEEE-EURASIP workshop on Nonlinear Signal and Image Processing*, 2003.

[58] R. Tyrrell Rockafellar. A dual approach to solving nonlinear programming problems by unconstrained optimization. *Mathematical Programming*, 5(1):354–373, December 1973.

[59] SJ Ruuth. Efficient algorithms for diffusion-generated motion by mean curvature. *Journal of Computational Physics*, 625:603–625, 1998.

[60] Jérémy Schmitt, Nelly Pustelnik, Pierre Borgnat, Patrick Flandrin, and Laurent Condat. 2-D Prony-Huang transform: A new tool for 2-D spectral analysis. page 24, April 2014.

[61] Robert C. Sharpley and Vesselin Vatchev. Analysis of the intrinsic mode functions. *Constructive Approximation*, 24(1):17–47, August 2005.

[62] John C. Strikwerda. *Finite Difference Schemes and Partial Differential Equations*. SIAM, second edition, 2004.

[63] E. Charles H. Sykes, Brent A. Mantooth, Patrick Han, Zachary J. Donhauser, and Paul S. Weiss. Substrate-mediated intermolecular interactions: a quantitative single molecule analysis. *Journal of the American Chemical Society*, 127(19):7255–7260, May 2005.

[64] Z. B. Szuts, J. R. Blundell, M. P. Chidichimo, and J. Marotzke. A vertical-mode decomposition to investigate low-frequency internal motion across the Atlantic at 26° N. *Ocean Science*, 8(3):345–367, June 2012.

[65] Tai Sing Lee. Image representation using 2D Gabor wavelets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(10):959–971, 1996.

[66] Jian Tang, Lijie Zhao, Heng Yue, Wen Yu, and Tianyou Chai. Vibration analysis based on empirical mode decomposition and partial least square. *Procedia Engineering*, 16:646–652, 2011.

[67] Peyman Tavallali, Thomas Y. Hou, and Zuoqiang Shi. Extraction of intrawave signals using the sparse time-frequency representation method. *Multiscale Modeling & Simulation*, 12(4):1458–1493, October 2014.

[68] John C. Thomas, Jeffrey J. Schwartz, J. Nathan Hohman, Shelley A. Claridge, Harsharn S. Auluck, Andrew C. Serino, Alexander M. Spokoyny, Giang Tran, Kevin F. Kelly, Chad A. Mirkin, Jerome Gilles, Stanley J. Osher, and Paul S. Weiss. Defect-tolerant aligned dipoles within two-dimensional plastic lattices. *ACS Nano*, 9(5):4734–4742, May 2015.

[69] Hau-Tieng Wu, Patrick Flandrin, and Ingrid Daubechies. One or two frequencies? The synchrosqueezing answers. *Advances in Adaptive Data Analysis*, 03(01n02):29–39, April 2011.

[70] Zhaohua Wu and Norden E. Huang. Ensemble empirical mode decomposition: A noise-assisted data analysis method. *Advances in Adaptive Data Analysis*, 01(01):1–41, January 2009.

[71] Ming Yan. Restoration of images corrupted by impulse noise and mixed Gaussian impulse noise using blind inpainting. *SIAM Journal on Imaging Sciences*, 6(3):1227–1245, July 2013.

[72] Haizhao Yang, Jianfeng Lu, and Lexing Ying. Crystal image analysis using 2D synchrosqueezed transforms. page 27, February 2014.

[73] Haizhao Yang and Lexing Ying. Synchrosqueezed wave packet transform for 2D mode decomposition. *SIAM Journal on Imaging Sciences*, 6(4):1979–2009, January 2013.

[74] Mingqiang Zhu, Stephen J. Wright, and Tony F. Chan. Duality-based algorithms for total-variation-regularized image restoration. *Computational Optimization and Applications*, 47(3):377–400, December 2010.

[75] Dominique Zosso, Jing An, James Stevick, Nicholas Takaki, Morgan Weiss, Liane S. Slaughter, Huan H. Cao, Paul S. Weiss, and Andrea L. Bertozzi. Image segmentation with dynamic artifacts detection and bias correction. *submitted to: AIMS J. Inverse Problems and Imaging*, page 24, 2015.

[76] Dominique Zosso, Xavier Bresson, and Jean-Philippe Thiran. Fast geodesic active fields for image registration based on splitting and augmented Lagrangian approaches. *IEEE Transactions on Image Processing*, 23(2):673–683, March 2014.

(D. Zosso, K. Dragomiretskiy, A. L. Bertozzi) Department of Mathematics, University of California, Los Angeles, Los Angeles, CA 90095-1555, USA

E-mail address, Corresponding author: `zosso@math.ucla.edu`

E-mail address: `konstantin@math.ucla.edu`

E-mail address: `bertozzi@math.ucla.edu`

(P. S. Weiss) California NanoSystems Institute, Department of Chemistry and Biochemistry, Department of Materials Science and Engineering, University of California, Los Angeles, Los Angeles, CA 90095, USA

E-mail address: `psw@cnsi.ucla.edu`