

Semi-supervised Learning for Aggregated Multilayer Graphs Using Diffuse Interface Methods and Fast Matrix-Vector Products*

Kai Bergermann[†], Martin Stoll[†], and Toni Volkmer[†]

Abstract. We generalize a graph-based multiclass semi-supervised classification technique based on diffuse interface methods to multilayer graphs. Besides the treatment of various applications with an inherent multilayer structure, we present a very flexible approach that interprets high-dimensional data in a low-dimensional multilayer graph representation. Highly efficient numerical methods involving the spectral decomposition of the corresponding differential graph operators as well as fast matrix-vector products based on the nonequispaced fast Fourier transform enable the rapid treatment of large and high-dimensional data sets. We perform various numerical tests putting a special focus on image segmentation. In particular, we test the performance of our method on data sets with up to 10 million nodes per layer as well as up to 104 dimensions, resulting in graphs with up to 52 layers. While all presented numerical experiments can be run on an average laptop computer, the linear dependence per iteration step of the runtime on the network size in all stages of our algorithm makes it scalable to even larger and higher-dimensional problems.

Key words. power mean Laplacian, multiclass semi-supervised learning, graph Laplacian, fast eigenpair computation, nonequispaced fast Fourier transform, diffuse interface methods, feature grouping

AMS subject classifications. 68R10, 05C50, 65F15, 65T50, 68T05, 62H30

DOI. 10.1137/20M1352028

1. Introduction. Complex networks have become an indispensable tool in the modeling of phenomena ranging from neurobiology to statistical physics [58]. As most sets of items interact in a variety of relationships, multilayer graphs have emerged as a flexible tool to reflect these complex interactions; see also [32, 8]. The power of these networks to model phenomena from social interactions to energy networks has greatly fueled research for a better understanding of the network properties and also to tailor numerical methods to incorporate their mathematical structures.

In this paper, we propose a technique for semi-supervised learning (cf. [70, 71] for an introduction) on multilayer graphs. Graph-based semi-supervised learning has recently risen to prominence with the introduction of graph convolutional neural networks [31]. In other semi-supervised settings, the graph Laplacian enters via a regularization term [24]. Besides that, the unsupervised case was also considered; e.g., the detection of communities without any labeled data in multiplex networks based on Laplacian dynamics was discussed in [28, 50].

*Received by the editors July 10, 2020; accepted for publication (in revised form) March 11, 2021; published electronically June 24, 2021.

<https://doi.org/10.1137/20M1352028>

Funding: The work of the third author was partially funded by the Sächsische Aufbaubank – Förderbank - (SAB) 100378180.

[†]Chair of Scientific Computing, Technische Universität Chemnitz, Department of Mathematics, 09107 Chemnitz, Germany (kai.bergermann@math.tu-chemnitz.de, martin.stoll@mathematik.tu-chemnitz.de, toni.volkmer@mathematik.tu-chemnitz.de).

Furthermore, [65] serves as a survey about spectral clustering.

Our proposed method represents significant progress over existing methods that rely on spectral information of fully connected networks where the adjacency matrices are formed by evaluating kernel functions on the feature vectors representing the graph nodes. This is achieved by the incorporation of fast matrix-vector products with the graph Laplacian. In this paper, we discuss the extension of a recently introduced acceleration technique [1] to multilayer graphs. We find that the matrix-vector products with the graph Laplacian scale linearly not only in the number of graph nodes, but also in the feature space dimension when employing a feature grouping approach introduced later in this paper.

In particular we focus on semi-supervised learning where only around 0.5% to 5% of the data is prelabeled. In order to classify the remaining unlabeled graph nodes, we rely on a diffuse interface approach which was first introduced in [6]. This method by Bertozzi and Flenner borrows from well-known results that have mainly been studied in the context of phase separation phenomena in materials science [51, 68, 69, 5]. The crucial formulation on a graph then requires the use of a discrete differential operator, namely, the graph Laplacian [65, 12]. Based on its properties and additional terms in the loss function, this method has shown great potential for different applications. This technique has recently been extended to various different scenarios, including multiclass segmentation [20], the use of a Merriman–Bence–Osher (MBO) scheme [38, 64] and of nonsmooth potentials [9], application to signed networks [43], and many more; see also [7] for an overview. Its extension and efficient implementation for the multilayer case is at the heart of this paper.

Of course, we will require a corresponding differential operator, and we rely on the formulation of graph Laplacians for multilayer graphs. There exist many different representations for multilayer networks as well as different aggregation approaches [32]. The use of an aggregated Laplacian that effectively combines the crucial information of the single-layer graphs will be essential for our method, especially when considering artificially created multilayer structures using the feature grouping approach proposed in section 7. Encouraged by the results presented in [44], we focus on the power mean Laplacian in this work. The use of this differential graph operator for semi-supervised learning was introduced in [45], and we compare the numerical results of our scheme to those of that method. For general data with an inherent multilayer structure, different tools, such as, e.g., supra-Laplacians with interlayer edges (cf. [32, section 2.3] for the definition), can, from a numerical perspective, be used in combination with the methods introduced in this paper. Although methods for a separate treatment of inter- and intralayer edges might be required from an application point of view, this approach would in principle allow for labels to change over different layers, i.e., time in temporal multiplex networks [32]. In this paper, however, we only consider classification problems where nodes are not allowed to change labels across different layers.

We demonstrate the performance of our method on synthetic as well as real data sets with a focus on the application of graph-based image segmentation, where the resulting graph is often fully connected and the sparsity of the network cannot be exploited. Additionally, the dimensionality of the feature space is often vast, making it a crucial task to be able to work with the resulting matrices in an efficient manner [57]. Many techniques in machine learning rely on the spectral information of the graph Laplacian in question; see also [10, 33, 41, 47, 48] in addition to the ones mentioned before. The diffuse interface method considered in

this paper combines the favorable properties of the eigeninformation of the graph Laplacian with a nonlinear function pushing the graph nodes into their corresponding classes. The computation of both eigenvalues and eigenvectors heavily relies on the efficiency of the matrix-vector products with the graph Laplacian. For moderate dimensions of the feature space, in particular $d \leq 3$, and when the weight function of the graph is a smooth kernel function like a Gaussian kernel function, fast summation techniques [30, 54, 55, 1, 49] show great potential to implicitly realize the matrix-vector multiplication in $\mathcal{O}(n)$, where n is the number of nodes in the graph. These techniques are often based on arguments from Fourier analysis.

In order to also take advantage of fast summation techniques for medium- to high-dimensional data, e.g., ranging from 4 to more than 100 spatial dimensions, we present a feature grouping approach that splits the feature space into several low-dimensional subspaces. We interpret each feature subspace as a layer in a multilayer graph giving rise to a novel class of artificially created multilayer graphs which are then recombined again using our aggregated graph Laplacian. Each feature subspace can then take advantage of the fast matrix-vector products described before. This highly scalable technique not only enables us to efficiently classify large data sets like large images with several megapixels that would normally produce enormous graph Laplacian matrices but also allows for a high feature space dimensionality, which can be found in various applications, including, for example, hyperspectral imaging.

We achieve the outlined tasks in this paper as follows. First, we give the necessary definitions for both graphs and multilayer graphs, including the discrete differential operators in section 2, and we comment on eigenpairs computation approaches in section 3. The fast summation technique based on the nonequispaced fast Fourier transform (NFFT), which is utilized in this paper, is introduced in section 4. We review the graph Allen–Cahn equation for (single-layer) graphs in section 5 focusing on the multiclass case. Its extension to the multilayer case is given in section 6. In section 7 we propose the reformulation of a graph-based problem with a high-dimensional feature space as a multilayer graph via a feature grouping technique, allowing for the application of the NFFT-based fast summation introduced in section 4. Finally, the derived methods are applied to the classification of various data sets. Several real-world data sets with an inherent multilayer structure are considered in subsection 8.1, and large parts of the results obtained in [45] are further improved. Subsection 8.2 then presents the segmentation of a 10 megapixel image, taking full advantage of fast matrix-vector products, and shows that our method generalizes very well to similar unseen images. Finally, in subsection 8.3 all methods presented in this paper join forces in order to efficiently treat the both large and high-dimensional hyperspectral Pavia center data set [52] with 148 152 pixels and 102 frequency bands, achieving excellent classification accuracies while working directly on the unfiltered raw data without requiring problem-tailored hardware or software architectures. In particular, all numerical experiments can be run on a laptop computer. Additionally, in the accompanying supplementary material, SMM135202.pdf ([local/web 349KB]; cf. section SM3), some persuasive features of the power mean Laplacian are illustrated by classifying data sets generated by the stochastic block model [26].

2. Graphs and multilayer graphs. First, we briefly introduce the notation of graphs and graph Laplacians. For more details and properties, we refer the reader to [47, 48, 12].

A graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ consists of vertices $x_i \in \mathcal{V}$, $|\mathcal{V}| = n$, and edges $e \in \mathcal{E} \subset \mathcal{V} \times \mathcal{V}$, where

an edge e connects any pair of vertices $x_i, x_j \in \mathcal{V}$. In this paper, we do not allow self-edges, i.e., we require $(x_i, x_i) \notin \mathcal{E} \forall i$. In particular, we use weighted graphs \mathcal{G} with a weight function $w: \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$. A value $w(x_i, x_j) > 0$ indicates that two vertices $x_i, x_j \in \mathcal{V}$ are connected by an edge and $w(x_i, x_j) = 0$ means $(x_i, x_j) \notin \mathcal{E}$. The weight matrix $\mathbf{W} := (w(x_i, x_j))_{i,j=1}^n \in \mathbb{R}_{\geq 0}^{n \times n}$ collects the weight information. Here, we only consider undirected graphs, which yields $w(x_i, x_j) = w(x_j, x_i)$, leading to a symmetric weight matrix \mathbf{W} . Since self-edges are not allowed, the diagonal of \mathbf{W} only contains zeros.

Based on the weight function w , the degree $\deg(x_i)$ of a node x_i can be defined as

$$\deg(x_i) := \sum_{x_j \in \mathcal{V}} w(x_i, x_j)$$

and the degree matrix $\mathbf{D} \in \mathbb{R}^{n \times n}$ as the diagonal matrix

$$\mathbf{D} := \text{diag}(\deg(x_1), \dots, \deg(x_n)) = \text{diag}(\mathbf{W}\mathbf{1}), \quad \mathbf{1} := (1, 1, \dots, 1)^\top \in \mathbb{R}^n.$$

Then the (unnormalized symmetric) graph Laplacian $\mathbf{L} \in \mathbb{R}^{n \times n}$ is defined as $\mathbf{L} := \mathbf{D} - \mathbf{W}$ and the symmetric normalized graph Laplacian as

$$\mathbf{L}_{\text{sym}} := \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}.$$

The eigeninformation of these Laplacians plays a key role in many graph-based learning techniques, especially in classification tasks. For the efficient computation of this eigeninformation, repeated matrix-vector multiplications with the graph Laplacian, particularly with the weight matrix \mathbf{W} , are required.

In order to be able to perform these tasks within a reasonable time frame for a large number n of nodes, the computation time for such matrix-vector products has to be reasonable. This can be achieved, in particular, if the weight matrix \mathbf{W} fulfills one of the following properties:

- (i) The weight matrix $\mathbf{W} \in \mathbb{R}_{\geq 0}^{n \times n}$ is dense and has no exploitable structure, leading to $\mathcal{O}(n^2)$ computation time for each matrix-vector multiplication with \mathbf{W} , but the number n of nodes is not too large. \mathbf{W} can be either built once and stored, or it can be assembled on the fly for each matrix-vector multiplication.
- (ii) The weight matrix \mathbf{W} is sparse, e.g., containing only $\mathcal{O}(n)$ nonzero entries and resulting in $\mathcal{O}(n)$ computation time for a matrix-vector multiplication, when \mathbf{W} is stored appropriately.
- (iii) There exists a rank- r factorization of \mathbf{W} with low rank $r \ll n$, resulting in $\mathcal{O}(rn)$ computation time for each matrix-vector multiplication with \mathbf{W} .
- (iv) A feature vector $\mathbf{x}_i \in \mathbb{R}^d$ is associated with each node $x_i \in \mathcal{V}$ of the graph \mathcal{G} , and the weight function w is given by a suitable kernel function, $w(x_i, x_j) = K(\mathbf{x}_i - \mathbf{x}_j)$, such that the matrix-vector multiplication with \mathbf{W} can still be realized in $\mathcal{O}(n)$ computation time via a highly efficient algorithm without the need to explicitly store \mathbf{W} , although \mathbf{W} may be densely populated. This case will be discussed in more detail in [section 4](#).




Figure 1. Example of a multilayer network with 3 layers and 4 nodes.

As many applications produce data graphs with an inherent multilayer structure, describing, e.g., different types of interactions between nodes, time series data, or data sets combining data from independent sources [32], we next consider multilayer graphs, which consist of $T \in \mathbb{N}$ graph layers; see Figure 1 for an example. Now, each layer $\mathcal{G}^{(t)}$, $t = 1, \dots, T$, is a graph based on the same vertex set \mathcal{V} , $|\mathcal{V}| = n$. The edge sets $\mathcal{E}^{(t)} \subset \mathcal{V} \times \mathcal{V}$, however, are typically different across the layers and correspondingly also the weight matrices $\mathbf{W}^{(t)} \in \mathbb{R}_{\geq 0}^{n \times n}$. In this paper, we do not allow interlayer edges.

Following [44], the symmetric normalized graph Laplacian \mathbf{L}_{sym} of a graph \mathcal{G} can then be generalized to multilayer graphs. \mathbf{L}_{sym} is now defined for each layer t separately as $\mathbf{L}_{\text{sym}}^{(t)}$. To merge the information of all graph layers into one Laplacian, the power mean Laplacian \mathbf{L}_p , which is defined as

$$(2.1) \quad \mathbf{L}_p := \left(\frac{1}{T} \sum_{t=1}^T (\mathbf{L}_{\text{sym}}^{(t)})^p \right)^{1/p},$$

was introduced in [44], where $p \in \mathbb{R} \setminus \{0\}$ denotes matrix powers and is not meant elementwise. Similarly, for a positive definite matrix \mathbf{A} , the matrix p th root $\mathbf{A}^{1/p}$ is the unique positive definite solution of the matrix equation $\mathbf{X}^p = \mathbf{A}$; cf. [25]. For $p > 0$, the definition (2.1) can be directly applied to the symmetric positive semidefinite graph Laplacians $\mathbf{L}_{\text{sym}}^{(t)}$, where all eigenvalues are ≥ 0 and at least one eigenvalue is zero; cf. [65] for this property. Note that for the $p = 1$ case, similar constructions have already been used; see, e.g., [61, 27].

For $p < 0$, however, (2.1) conflicts with the noninvertibility of $\mathbf{L}_{\text{sym}}^{(t)}$ due to the occurrence of at least one zero eigenvalue. To circumvent this issue, [44, section 2.2] proposes applying a diagonal shift of $\delta \geq 0$ to each symmetric graph Laplacian to obtain a (strictly) positive definite version

$$\mathbf{L}_{\text{sym},\delta}^{(t)} = \mathbf{L}_{\text{sym}}^{(t)} + \delta \mathbf{I},$$

where the choice $\delta = \log(1 + |p|)$ was suggested¹ in [44] for $p < 0$. For $p > 0$ we set $\delta = 0$. For the special case $p = 0$, we refer the reader to [44]. Combining the $\mathbf{L}_{\text{sym},\delta}^{(t)}$ for the different

¹While, from a theoretical perspective, an infinitesimal shift would suffice to guarantee the invertibility of $\mathbf{L}_{\text{sym},\delta}^{(t)}$, the dependency on p provides numerical stability as it prevents negative powers of the eigenvalues of $\mathbf{L}_{\text{sym},\delta}^{(t)}$ from becoming large.

layers yields the shifted power mean Laplacian

$$(2.2) \quad \mathbf{L}_{p,\delta} := \left(\frac{1}{T} \sum_{t=1}^T (\mathbf{L}_{\text{sym},\delta}^{(t)})^p \right)^{1/p}.$$

The benefit of considering negative powers p in (2.2) can be motivated by [44], where the convergence of $\mathbf{L}_{p,\delta}$ to the AND- and OR-operators on the presence of a common clustering structure across the layers is proven for $p \rightarrow \infty$ and $p \rightarrow -\infty$, respectively, in certain stochastic block model situations. With this interpretation, the power mean Laplacian tends to capture the underlying clustering structure even if it is only present in some graph layers for $p < 0$, while the $p > 0$ case requires a more uniform distribution of the clustering information across the layers. Thus, $\mathbf{L}_{p,\delta}$ should be more robust with respect to noisy layers for negative p . We illustrate this with an example in section SM3.

Note that for real symmetric matrices $\mathbf{A} \in \mathbb{R}^{n \times n}$ with eigendecomposition $\mathbf{A} = \mathbf{\Phi} \mathbf{\Lambda} \mathbf{\Phi}^\top$, the result of a matrix function $f(\mathbf{A})$, defined as in [25, Definition 1.1], can be obtained via $f(\mathbf{A}) = \mathbf{\Phi} f(\mathbf{\Lambda}) \mathbf{\Phi}^\top$, where $f(\mathbf{\Lambda}) = \text{diag}(f(\lambda_i)_{i=1}^n)$. This means f only acts on the eigenvalues. In particular, λ^p is an eigenvalue of the matrix power \mathbf{A}^p if λ is an eigenvalue of \mathbf{A} and $1 - \lambda$ is an eigenvalue of $\mathbf{I} - \mathbf{A}$; cf., e.g., [25].

3. Computation of eigenpairs and the polynomial Krylov subspace method. It is well known that for many classification tasks involving graph Laplacians, the computation of its k smallest eigenpairs is of particular interest as they provide optimal solutions to different graph cut problems; cf., e.g., [65]. To this end, when using the power mean Laplacian \mathbf{L}_1 or the shifted version $\mathbf{L}_{p,\delta}$, $p < 0$, one needs to perform matrix-vector products with this matrix or the p th power

$$(3.1) \quad \mathbf{L}_{p,\delta}^p = \frac{1}{T} \sum_{t=1}^T (\mathbf{L}_{\text{sym},\delta}^{(t)})^p.$$

In the case $p = 1$, one has $\mathbf{L}_1 = \frac{1}{T} \sum_{t=1}^T \mathbf{L}_{\text{sym}}^{(t)}$, and the relevant eigeninformation of \mathbf{L}_1 can be obtained by computing the eigenpairs for the k largest eigenvalues of

$$(3.2) \quad \mathbf{I} - \mathbf{L}_1 = \frac{1}{T} \sum_{t=1}^T (\mathbf{D}^{(t)})^{-1/2} \mathbf{W}^{(t)} (\mathbf{D}^{(t)})^{-1/2}$$

via the Lanczos method [21]. As mentioned in section 2, the resulting eigenvectors ϕ are identical to those of \mathbf{L}_1 and the eigenvalues λ of (3.2) correspond to $1 - \lambda$ of \mathbf{L}_1 .

Likewise, in the case of general p for a given eigenvalue λ and eigenvector ϕ of the real symmetric matrix $\mathbf{L}_{p,\delta}$, the corresponding eigenvalue of $\mathbf{L}_{p,\delta}^p$ is λ^p and the eigenvector remains ϕ ; cf. section 2. Since the function $f(\lambda) = \lambda^p$ is order reversing for $p < 0$, it holds for $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ that $\lambda_1^p \geq \lambda_2^p \geq \dots \geq \lambda_n^p$. In order to obtain the first k smallest eigenvalues $\lambda_1, \dots, \lambda_k$ and corresponding eigenvectors ϕ_1, \dots, ϕ_k in the case $p < 0$, it is sufficient to compute the k largest eigenvalues $\lambda_1^p, \dots, \lambda_k^p$ with its eigenvectors of $\mathbf{L}_{p,\delta}^p$. For this, we propose utilizing the Lanczos method, which requires matrix-vector multiplications of the

graph Laplacian matrices $(\mathbf{L}_{\text{sym},\delta}^{(t)})^p$; cf. (3.1). The latter can be realized by using the polynomial Krylov subspace method (PKSM) [44]. For this we again rely on the Lanczos method for $\mathbf{L}_{\text{sym},\delta}^{(t)}$; cf. [25, Chapter 13].

As with any Krylov subspace method, the main algorithmic cost comes from the matrix-vector products, i.e., multiplication with $\mathbf{L}_{\text{sym},\delta}^{(t)}$. The matrix function is then approximated using a projected matrix of drastically reduced dimension for which well-established methods from dense linear algebra can be employed.

When the weight functions $w^{(t)}$ and, correspondingly, the weight matrices $\mathbf{W}^{(t)}$ of the layers $t \in \{1, \dots, T\}$ have a special structure, the matrix-vector multiplications and eigenpair computations can be accelerated considerably, as we discuss in the next section.

4. NFFT-based fast summation for fast matrix-vector multiplications with the weight matrix. For a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, the nodes $x_i \in \mathcal{V}$ are identified with feature vectors $\mathbf{x}_i \in \mathbb{R}^d$, and the weight function $w: \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$ is associated with a kernel function $K: \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, $w(x_i, x_j) = K(\mathbf{x}_i - \mathbf{x}_j)$ for $\mathbf{x}_i \neq \mathbf{x}_j$. Since self-loops are excluded, we have $w(x_i, x_i) = 0$ due to $(x_i, x_i) \notin \mathcal{E}$. The matrix-vector multiplications with the weight matrix $\mathbf{W} \in \mathbb{R}_{\geq 0}^{n \times n}$ can then be sped up dramatically, especially if the matrix is nonsparse, e.g., when K can be well approximated by a trigonometric polynomial achieving a desired tolerance. For instance, K could be the Gaussian radial basis function (RBF) kernel $K(\mathbf{y}) = \exp(-\|\mathbf{y}\|^2/\sigma^2)$ or Laplacian RBF kernel $K(\mathbf{y}) = \exp(-\|\mathbf{y}\|/\sigma)$. For dimensions $d \leq 3$, one very efficient method is the NFFT-based fast summation [54, 55], which achieves a runtime complexity of $\mathcal{O}(n)$ for fixed accuracy. Subsequently, we briefly describe the general ideas and give a fast algorithm. For more details, we refer the reader to [1, section 3]. In case of higher dimensions $d \geq 4$,² we propose a feature grouping approach later in section 7.

For technical reasons,³ see also [1, section 3], we consider the modified weight function $\tilde{w}(x_i, x_j) = K(\mathbf{x}_i - \mathbf{x}_j) \forall x_i, x_j \in \mathcal{V}$ with associated matrix

$$\tilde{\mathbf{W}} := \mathbf{W} + K(\mathbf{0}) \mathbf{I},$$

and we have $\mathbf{W} = \tilde{\mathbf{W}} - K(\mathbf{0}) \mathbf{I}$. Then the matrix-vector multiplication of \mathbf{W} with an arbitrary vector $\mathbf{v} \in \mathbb{C}^n$ can be written as $\mathbf{W}\mathbf{v} = \tilde{\mathbf{W}}\mathbf{v} - K(\mathbf{0})\mathbf{v}$. The last part $K(\mathbf{0})\mathbf{v}$ is simply a multiplication of a scalar value with a vector, and we will compute the first part $\tilde{\mathbf{W}}\mathbf{v}$ in a fast way using the NFFT-based fast summation. Each entry of the result $\tilde{\mathbf{W}}\mathbf{v}$ reads as

$$(4.1) \quad (\tilde{\mathbf{W}}\mathbf{v})_i = f(\mathbf{x}_i) := \sum_{j=1}^n v_j K(\mathbf{x}_i - \mathbf{x}_j).$$

The key idea for the efficient computation of (4.1) uses methods from Fourier analysis [53]. In particular, the kernel function K is approximated by a d -variate trigonometric polynomial

²Here, the curse of dimensionality makes computations burdensome as the runtime depends exponentially on the feature space dimension, as will be discussed later in this section. Note that the threshold $d \geq 4$ may shift to larger values of d in the future when increased computing power becomes available.

³On the one hand, one has $w(x_i, x_i) = 0$ since $(x_i, x_i) \notin \mathcal{E}$. On the other hand, $K(\mathbf{x}_i - \mathbf{x}_i) = K(\mathbf{0}) \neq 0$ in general. Modifying the kernel function (only) at the origin is not an option since K should be smooth (or at least continuous).

K_{RF} , which allows us to separate the computations involving the nodes x_i and x_j in (4.3).⁴ This will be one main ingredient for reducing the computational complexity from $\mathcal{O}(n^2)$ to $\mathcal{O}(n)$. Assuming we have a suitable approximation of K , e.g., up to a tolerance of 10^{-4} with respect to the L^∞ norm,⁵ given by

$$(4.2) \quad K(\mathbf{y}) \approx K_{\text{RF}}(\mathbf{y}) := \sum_{\mathbf{l} \in I_N} \hat{b}_{\mathbf{l}} e^{2\pi i \mathbf{l} \mathbf{y}}, \quad I_N := \{-N/2, -N/2 + 1, \dots, N/2 - 1\}^d,$$

with bandwidth $N \in 2\mathbb{N}$ and Fourier coefficients $\hat{b}_{\mathbf{l}}$, we replace K by K_{RF} in (4.1) and obtain

$$(4.3) \quad \begin{aligned} (\tilde{\mathbf{W}}\mathbf{v})_i &= f(\mathbf{x}_i) \approx f_{\text{RF}}(\mathbf{x}_i) := \sum_{j=1}^n v_j K_{\text{RF}}(\mathbf{x}_i - \mathbf{x}_j) = \sum_{j=1}^n v_j \sum_{\mathbf{l} \in I_N} \hat{b}_{\mathbf{l}} e^{2\pi i \mathbf{l}(\mathbf{x}_i - \mathbf{x}_j)} \\ &= \sum_{\mathbf{l} \in I_N} \hat{b}_{\mathbf{l}} \underbrace{\left(\sum_{j=1}^n v_j e^{-2\pi i \mathbf{l} \mathbf{x}_j} \right)}_{=: \hat{f}_{\mathbf{l}}} e^{2\pi i \mathbf{l} \mathbf{x}_i} \quad \forall i = 1, \dots, n. \end{aligned}$$

Comparing (4.3) with the initial problem of evaluating (4.1), the Fourier approximation K_{RF} of the kernel function K has so far only introduced an additional sum over I_N . In situations of large n (starting at ten thousands ranging to millions and billions), however, the NFFT [29] manages to substantially speed up the evaluation of the inner sums $\hat{f}_{\mathbf{l}} := (\sum_{j=1}^n v_j e^{-2\pi i \mathbf{l} \mathbf{x}_j})$, $\mathbf{l} \in I_N$, as well as the computation of the outer sums $f_{\text{RF}}(\mathbf{x}_i) := \sum_{\mathbf{l} \in I_N} \hat{b}_{\mathbf{l}} \hat{f}_{\mathbf{l}}$, $i = 1, \dots, n$. We remark that one cannot apply standard (equispaced) fast Fourier transform (FFT) in most cases since the feature vectors \mathbf{x}_i are not located on an equispaced grid.

Therefore, the NFFT is the second main ingredient for the reduction in the computational complexity from $\mathcal{O}(n^2)$, when using the direct summation (4.1), to $\mathcal{O}((\mathfrak{m}_{\text{NFFT}})^d n + d 2^d N^d \log N)$ for computing (4.3) via the NFFT-based fast summation, where $\mathfrak{m}_{\text{NFFT}}$ represents an internal window cut-off parameter of the NFFT controlling the desired precision,⁶ cf. [34, section 3]. In situations where the number n of nodes of the graph is large and the feature space dimension d is not too big, this represents a crucial gain in computational complexity.

The accuracy of the Fourier approximation K_{RF} of the kernel function K depends on the decay of the Fourier coefficients of K , which are influenced by smoothness properties of the kernel function. For instance, smooth rotational invariant kernel functions are particularly

⁴We remark that in K_{RF} , the subscript R stands for regularized kernel function (smooth 1-periodic) and F for the Fourier approximation by a trigonometric polynomial. The main trick—from a linear algebra point of view—is that one diagonalizes the (modified) weight matrix $\tilde{\mathbf{W}}$ (in a certain sense) and computes $\tilde{\mathbf{W}}\mathbf{v} \approx \mathbf{A}(\text{diag}(\hat{b}_{\mathbf{l}})_{\mathbf{l} \in I_N})\mathbf{A}^*\mathbf{v}$ (cf. (4.2) for the definitions of $\hat{b}_{\mathbf{l}}$ and I_N). Here $\mathbf{A} := (e^{2\pi i \mathbf{l} \mathbf{x}_j})_{j=1, \dots, n; \mathbf{l} \in I_N}$ is the nonequispaced Fourier matrix, which can be applied to a vector using $\mathcal{O}((\mathfrak{m}_{\text{NFFT}})^d n + d 2^d N^d \log N)$ arithmetic operations. Please note that, assuming one evaluation of K has complexity $\mathcal{O}(d)$, one can evaluate K at all N^d nodes \mathbf{k}/N , $\mathbf{k} \in I_N$, and then compute by d -dimensional fast Fourier transform the Fourier coefficients $\hat{b}_{\mathbf{l}} := N^{-d} \sum_{\mathbf{k} \in I_N} K(\mathbf{k}/N) e^{-2\pi i \mathbf{l} \mathbf{k}/N}$, $\mathbf{l} \in I_N$, in $\mathcal{O}(d N^d \log N)$ runtime complexity.

⁵For the formulation of the tolerance, the kernel function K is assumed to be normalized, i.e., $|K(\mathbf{x})| \leq 1$. For more details, we refer the reader to [54] and [1, section 3].

⁶In practice, $\mathfrak{m}_{\text{NFFT}}$ is independent of n and typically chosen between 2 and 6. $\mathfrak{m}_{\text{NFFT}} = 8$ yields errors close to machine precision for IEEE double precision; see also [30, section 5.2].

suited, e.g., the already mentioned Gaussian RBF kernel $K(\mathbf{y}) = \exp(-\|\mathbf{y}\|^2/\sigma^2)$, Laplacian RBF kernel $K(\mathbf{y}) = \exp(-\|\mathbf{y}\|/\sigma)$, and many others. The Fourier coefficients \hat{b}_l of the Fourier approximation K_{RF} can be computed easily by sampling the kernel function K or a regularized version of K on an equispaced grid and applying an FFT which takes $\mathcal{O}(d N^d \log N)$ arithmetic operations, assuming that K can be evaluated in $\mathcal{O}(d)$ arithmetic operations; see also [1, section 3] for more details. Since the parameter m_{NFFT} and the bandwidth N only depend on the desired accuracy, we have a computational complexity of $\mathcal{O}(n)$ for fixed d and fixed accuracy. For further technical details, we refer the reader to [54] and [1, section 3].

In total, we have a fast approximate algorithm for the matrix-vector multiplication $\tilde{\mathbf{W}}\mathbf{v}$ of complexity $\mathcal{O}(n)$ available; cf. Algorithm 4.1. This algorithm is implemented as `applications/fastsum` and `matlab/fastsum` in C and MATLAB within the NFFT3 software library and freely available; see [29]. Then one easily computes $\mathbf{W}\mathbf{v}$ from $\tilde{\mathbf{W}}\mathbf{v}$ by subtracting the vector $K(\mathbf{0})\mathbf{v}$.

Algorithm 4.1 ([1, Algorithm 1]). Fast approximate matrix-vector multiplication $\tilde{\mathbf{W}}\mathbf{v}$ using NFFT-based fast summation, $(\tilde{\mathbf{W}}\mathbf{v})_i = \sum_{j=1}^n v_j K(\mathbf{x}_i - \mathbf{x}_j) \forall i = 1, \dots, n$, e.g., $K(\mathbf{x}_i - \mathbf{x}_j) = \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|^2/\sigma^2)$.

Input:	$(\hat{b}_l)_{l \in I_N}$	Fourier coefficients of trigonometric polynomial K_{RF} which approximates K ,
	$\{\mathbf{x}_i\}_{i=1}^n$	nodes, $\mathbf{x}_i \in \mathbb{R}^d$, $\ \mathbf{x}_i\ \in [-1/4, 1/4]^d$,
	$\mathbf{v} = [v_1, v_2, \dots, v_n]^T$	vector $\in \mathbb{R}^n$.
1. Apply d -dimensional adjoint NFFT (see also [30, section 2.3]) on \mathbf{v} and obtain $\hat{v}_l \approx \sum_{j=1}^n v_j e^{-2\pi i l \mathbf{x}_j} \forall l \in I_N$.		
2. Multiply result by Fourier coefficients $(\hat{b}_l)_{l \in I_N}$ of K_{RF} and obtain $\hat{f}_l := \hat{b}_l \hat{v}_l \forall l \in I_N$.		
3. Apply d -dimensional NFFT on $(\hat{f}_l)_{l \in I_N}$ and obtain output $\tilde{f}_{\text{RF}}(\mathbf{x}_i) \approx \sum_{l \in I_N} \hat{f}_l e^{2\pi i l \mathbf{x}_i} \forall i = 1, \dots, n$.		
Output:	$\left[\tilde{f}_{\text{RF}}(\mathbf{x}_i)\right]_{i=1, \dots, n}$	$\tilde{f}_{\text{RF}}(\mathbf{x}_i) \approx (\tilde{\mathbf{W}}\mathbf{v})_i \forall i = 1, \dots, n$.
Complexity:	$\mathcal{O}(n)$ for fixed accuracy.	

Note that Algorithm 4.1 can be used for accelerating the eigenpair computation of the power mean Laplacian \mathbf{L}_1 and of the shifted power mean Laplacian $\mathbf{L}_{p,\delta}^p$, $p < 0$. In the latter case, this means applying Algorithm 4.1 inside the PKSM; see also section 3.

In the next section, the eigeninformation of \mathbf{L}_1 and $\mathbf{L}_{p,\delta}^p$, $p < 0$, is used in order to perform semi-supervised learning based on multilayer graphs.

5. Graph Allen–Cahn for multiclass problems. Diffuse interface methods are heavily used in materials science and beyond; see, e.g., [51, 68, 69, 5] and the references therein. They offer an efficient and flexible way to model phase separation, with one of the most

prominent models being the Allen–Cahn equation [2], which originally describes the evolution of a binary liquid over time. The Allen–Cahn equation is derived as the L^2 gradient flow of the Ginzburg–Landau energy functional

$$(5.1) \quad E(u) = \int \frac{\epsilon}{2} |\nabla u|^2 d\mathbf{x} + \int \frac{1}{\epsilon} \psi(u) d\mathbf{x},$$

where we obtain

$$(5.2) \quad \frac{\partial u}{\partial t} = -\nabla E(u) = \epsilon \Delta u - \frac{1}{\epsilon} \psi'(u).$$

Here, u and $-u$ are defined on a physical domain Ω and describe the two (liquid) components. Furthermore, t is the time variable and $\epsilon > 0$ is a (typically small) parameter which influences the width of the interface regions where the transition from one pure phase to the other happens. The gradient term represents the Dirichlet energy that penalizes the length of the interface and ψ is a suitable potential function with minima at the pure phases. Standard solution methods for partial differential equations, such as the finite element method or the finite difference method, can be used to solve (5.2), given suitable initial conditions as well as boundary conditions.⁷

The Allen–Cahn equation has been adapted to binary semi-supervised classification on graphs; see [6, 10]. Most importantly, the spatial domain Ω is replaced by the graph domain, and, in particular, this means that Ω is identified with the set of vertices $x_i \in \mathcal{V}$ of a graph \mathcal{G} . Moreover, a data fidelity term is added to the Ginzburg–Landau energy functional (5.1) to include the correct classification of the a priori labeled data as an additional objective.

Furthermore, the method has been extended to the multiclass case; see, e.g., [20, 9, 43] as well as Figures 2 and 3 for an image segmentation example. For the general case of $m \geq 2$ classes, the vector $u \in \mathbb{R}^n$ of the phase-field description on the finite set of vertices $\mathcal{V} = \{x_i\}_{i=1}^n$ is identified with a matrix $\mathbf{U} \in \mathbb{R}^{n \times m}$ containing one row of m entries per vertex x_i . The goal is to identify vertices x_i with the j th class whenever the j th entry in the corresponding row \mathbf{u}_i^\top of \mathbf{U} is largest for the class ID $j \in \{1, \dots, m\}$. In order to incorporate the known class information of a priori labeled data, the data fidelity term

$$\frac{1}{2} \sum_{i=1}^n \omega(x_i) \|\mathbf{f}_i - \mathbf{u}_i\|_{\ell^2}^2$$

is added to the Ginzburg–Landau energy functional (5.1) in the multiclass case, where $\omega(x_i)$ is a penalty parameter that is equal to the constant $\omega_0 \gg 0$ for labeled vertices x_i and 0 for unlabeled vertices. For the class information of the known labels, one-hot encoding is used, i.e., the vectors \mathbf{f}_i are set to the j th unit vector $\mathbf{e}_j \in \mathbb{R}^m$ for labeled data,

$$(5.3) \quad \mathbf{f}_i := \begin{cases} \mathbf{e}_j & \text{if node } x_i \text{ is of class } j, \\ \mathbf{0} & \text{if class of node } x_i \text{ is unknown.} \end{cases}$$

⁷For details on (5.2) as well as binary semi-supervised Allen–Cahn classification on graphs we refer the reader to section SM1 in the accompanying supplementary material, SMM135202.pdf ([local/web 349KB]).




Figure 2. Example of the evolution of a phase-field simulation for image segmentation based on graph Allen–Cahn with four classes, where the four colors red, green, blue, and yellow indicate different class affiliations with score values higher than 0.66. The image in this example is a resized version (520×293 pixels) of Figure 3.



Figure 3. Left: Test image 1 (4160×2340 pixels) for image segmentation. Right: Prelabeled pixels.

The choice of ω_0 controls the trade-off between the classical Ginzburg–Landau energy and the least squares data fidelity term. Choosing ω_0 too small bears the risk of underfitting, while choosing it too large may cause overfitting to the preclassified data.

In addition, the Dirichlet energy term $\int \frac{\epsilon}{2} |\nabla u|^2 d\mathbf{x}$ in the Ginzburg–Landau energy functional (5.1) is replaced by $\frac{\epsilon}{2} \text{trace}(\mathbf{U}^\top \mathbf{L}_{\text{sym}} \mathbf{U})$, which is motivated by [65, section 5.2]. More-

over, the multiwell potential

$$(5.4) \quad \psi(\mathbf{u}_i) := \prod_{j=1}^m \frac{1}{4} \|\mathbf{u}_i - \mathbf{e}_j\|_{\ell^1}^2 = \prod_{j=1}^m \frac{1}{4} \left(\sum_{l=1}^n |\mathbf{u}_{il} - \delta_{jl}| \right)^2$$

is used with minima of 0 in the corners of the Gibbs simplex

$$(5.5) \quad \Sigma^m := \left\{ (s_1, \dots, s_m) \in [0, 1]^m : \sum_{j=1}^m s_j = 1 \right\}$$

(cf. [20]), and the integral is replaced by a finite sum over the vertices x_i , $i = 1, \dots, n$. The locations of these minima correspond to the one-hot encoding (5.3) of the known class labels \mathbf{f}_i and model the goal that each row \mathbf{u}_i^\top of the solution \mathbf{U} should be close to one of the unit vectors \mathbf{e}_j^\top , $j \in \{1, \dots, m\}$. Note that for the binary case $m = 2$, the multiwell potential (5.4) corresponds to the double-well potential (SM1.1).

With the modifications discussed above, the discretized Ginzburg–Landau functional for the multiclass case becomes

$$(5.6) \quad \tilde{E}(\mathbf{U}) = \frac{\epsilon}{2} \text{trace}(\mathbf{U}^\top \mathbf{L}_{\text{sym}} \mathbf{U}) + \frac{1}{2\epsilon} \sum_{i=1}^n \left(\prod_{l=1}^m \frac{1}{4} \|\mathbf{u}_i - \mathbf{e}_l\|_{\ell^1}^2 \right) + \frac{1}{2} \sum_{i=1}^n \omega(x_i) \|\mathbf{f}_i - \mathbf{u}_i\|_{\ell^2}^2$$

and forms the basis for the semi-supervised classification technique considered in this work, as it can be viewed as its loss function. The potential term in this case slightly varies in contrast to the binary case as the scaling changes from $\frac{1}{\epsilon}$ to $\frac{1}{2\epsilon}$, and we stick to this formulation as this is used throughout the literature for the multiclass case.

For the solution of the Allen–Cahn equation (5.2), a numerical scheme called convexity splitting [17, 16] is commonly applied (see, e.g., [43, 9, 6, 20]), where $E(u)$ is split up into a convex part E_1 and a concave part $-E_2$ so that $E(u) = E_1(u) - E_2(u)$. The convex part E_1 is then treated implicitly to allow for numerical stability, while the concave part $-E_2$ is treated explicitly. For the general multiclass case $m \geq 2$, one possible splitting $\tilde{E}(\mathbf{U}) = \tilde{E}_1(\mathbf{U}) - \tilde{E}_2(\mathbf{U})$ reads

$$(5.7) \quad \tilde{E}_1(\mathbf{U}) = \frac{\epsilon}{2} \text{trace}(\mathbf{U}^\top \mathbf{L}_{\text{sym}} \mathbf{U}) + \frac{c}{2} \text{trace}(\mathbf{U}^\top \mathbf{U}),$$

$$(5.8) \quad \begin{aligned} \tilde{E}_2(\mathbf{U}) &= \frac{c}{2} \text{trace}(\mathbf{U}^\top \mathbf{U}) - \frac{1}{2\epsilon} \sum_{i=1}^n \left(\frac{1}{4} \prod_{l=1}^m \|\mathbf{u}_i - \mathbf{e}_l\|_{\ell^1}^2 \right) - \frac{1}{2} \sum_{i=1}^n \omega(x_i) \|\mathbf{f}_i - \mathbf{u}_i\|_{\ell^2}^2 \\ &= \sum_{i=1}^n \frac{c}{2} (\mathbf{u}_i^\top \mathbf{u}_i) - \frac{1}{2\epsilon} \left(\frac{1}{4} \prod_{l=1}^m \|\mathbf{u}_i - \mathbf{e}_l\|_{\ell^1}^2 \right) - \frac{1}{2} \omega(x_i) \|\mathbf{f}_i - \mathbf{u}_i\|_{\ell^2}^2, \end{aligned}$$

where a productive zero is inserted into $\tilde{E}(\mathbf{U})$ by adding and subtracting the convex function $\frac{c}{2} \text{trace}(\mathbf{U}^\top \mathbf{U})$ with a constant $c > 0$ that ensures the strict convexity of \tilde{E}_2 ; cf. [20]. The resulting scheme is then given by

$$(5.9) \quad \frac{\mathbf{U}^{l+1} - \mathbf{U}^l}{\Delta t} = -\epsilon \mathbf{L}_{\text{sym}} \mathbf{U}^{l+1} - c \mathbf{U}^{l+1} + c \mathbf{U}^l - \frac{1}{2\epsilon} \mathcal{T}^l - \omega(\mathbf{U}^l - \mathbf{F}).$$

Here the derivative of the term $\frac{1}{2\epsilon} \sum_{i=1}^n (\prod_{l=1}^m \frac{1}{4} \|\mathbf{u}_i - \mathbf{e}_l\|_{\ell^1}^2)$ is given by the matrix $\mathcal{T}(\mathbf{U}) \in \mathbb{R}^{n \times m}$ with the entries

$$(5.10) \quad \mathcal{T}_{ij}(\mathbf{U}) := \sum_{q=1}^m \frac{1}{2} (1 - 2\delta_{jq}) \|\mathbf{u}_i - \mathbf{e}_q\|_{\ell^1} \prod_{\substack{p=1 \\ p \neq q}}^m \frac{1}{4} \|\mathbf{u}_i - \mathbf{e}_p\|_{\ell^1}^2,$$

$\mathcal{T}^l := \mathcal{T}(\mathbf{U}^l)$, $\boldsymbol{\omega} := \text{diag}(\omega(x_i)_{i=1}^n)$, and $\mathbf{F} := (\mathbf{f}_i^\top)_{i=1}^n \in \mathbb{R}^{n \times m}$, which contains the known label information; cf. (5.3). All of this leads to solving the equation

$$(5.11) \quad \mathbf{U}^{l+1} = \underbrace{[(1 + c(\Delta t))\mathbf{I} + \epsilon(\Delta t)\mathbf{L}_{\text{sym}}]}_{:=\mathbf{B}}^{-1} \left((1 + c(\Delta t))\mathbf{U}^l - \frac{\Delta t}{2\epsilon} \mathcal{T}^l - (\Delta t) \boldsymbol{\omega} (\mathbf{U}^l - \mathbf{F}) \right),$$

which is made computationally efficient using a projection onto the dominating eigenspace of \mathbf{L}_{sym} ; see section SM2 for details.

For the initialization of \mathbf{U}^0 , we set $\mathbf{u}_i := \mathbf{e}_j \in \mathbb{R}^m$ for the prelabeled vertices x_i , where j is the corresponding class. For the vertices with unknown labels, [20, Algorithm 1] suggests using randomized initial conditions \mathbf{U}^0 , which are then scaled to the Gibbs simplex. Especially in unsupervised classification problems, this initialization strategy can be beneficial; cf. [62, section 5.5] for a detailed discussion. However, interpreting the vector entries of $\mathbf{u}_i \in \Sigma^m$ as empirical probabilities for class affiliation, a random initialization bears the risk of initially assigning high probabilities to wrong classes, which potentially reduces the classification accuracy of the whole method. Having no a priori information about the unlabeled nodes, we therefore propose to initialize the respective rows of \mathbf{U}^0 with uniform empirical probabilities, i.e., $\mathbf{u}_i := \frac{1}{m} \mathbf{1} \in \mathbb{R}^m$. Numerical experiments with our semi-supervised classifier (not included in section 8) underpin this heuristic as classification results consistently degrade with an increase of randomization in the initialization.

Note that the presented method requires the choice of several hyperparameters. While the Allen–Cahn parameters in the ranges $\epsilon \in [5 \cdot 10^{-3}, 0.5]$, $\omega_0 \in [10^3, 10^4]$, $c \in [\omega_0 + 1/\epsilon, \omega_0 + 3/\epsilon]$, and $\Delta t \in [10^{-2}, 1]$ were experimentally found not to have a major influence on classification results, the number k of eigenpairs of the Laplacian operator and the scaling parameter σ in the Gaussian kernel have to be carefully chosen. We employed a parameter-grid search for both parameters for each data set presented in the numerical experiments in section 8.

A related technique involving fewer hyperparameters is based on the MBO scheme, where the nonlinear term of the PDE is dropped and a thresholding procedure is added to the method; cf., e.g., [46, 20, 7] and references therein. This approach could potentially reduce the computation time per iteration as the nonlinearity does not need to be evaluated. A naive parameter selection did not yield improved performance for this scheme compared to our approach when applied on the WebKB data set from subsection 8.1. We leave a detailed investigation to future research. Moreover, we point out that all techniques suggested in sections 3, 4, and 7 should also be applicable for the MBO scheme.

6. Graph Allen–Cahn on multilayer graphs for multiclass problems. In this section, we extend the graph-based multiclass Allen–Cahn approach from section 5 to multilayer graphs.

For this, we employ the concept of the power mean Laplacian in order to combine the information of all graph layers; see [section 2](#). Formally, we replace the graph Laplacian matrix \mathbf{L}_{sym} in (5.6) by the power mean Laplacian \mathbf{L}_p from (2.1) for $p > 0$ (corresponding to a shift $\delta = 0$) and by the shifted version $\mathbf{L}_{p,\delta}$ from (2.2) with $\delta = \log(1 + |p|)$ for $p < 0$. This causes changes in the term $\tilde{E}_1(\mathbf{U})$ in (5.7), and, correspondingly, \mathbf{L}_{sym} is replaced in the iteration formula (5.11), yielding

$$\mathbf{B} = \begin{cases} (1 + c(\Delta t))\mathbf{I} + \epsilon(\Delta t)\mathbf{L}_p & \text{for } p > 0, \\ (1 + c(\Delta t))\mathbf{I} + \epsilon(\Delta t)\mathbf{L}_{p,\delta} & \text{for } p < 0. \end{cases}$$

Analogously to the single-layer case, we utilize an eigendecomposition $\Phi\Lambda\Phi^\top$ of \mathbf{L}_p and $\mathbf{L}_{p,\delta}$ for $p > 0$ and $p < 0$, respectively, where $\Lambda \in \mathbb{R}^{n \times n}$ and $\Phi \in \mathbb{R}^{n \times n}$. The eigendecomposition will then be approximated by a truncated version $\Lambda_k \in \mathbb{R}^{k \times k}$, $\Phi_k \in \mathbb{R}^{n \times k}$ using the k smallest eigenvalues. In order to compute this eigeninformation, we use the Lanczos method, which relies on matrix-vector products with \mathbf{L}_p and $\mathbf{L}_{p,\delta}^p$. For the case $p = 1$, which is the computationally most efficient one as it does not involve matrix powers, we compute the eigenpairs belonging to the k largest eigenvalues of (3.2) as described in detail in [section 3](#). Otherwise, for the case $p < 0$, we use the relation

$$\mathbf{L}_{p,\delta}^p = \frac{1}{T} \sum_{t=1}^T (\mathbf{L}_{\text{sym},\delta}^{(t)})^p$$

to compute the eigenpairs belonging to the k largest eigenvalues $\lambda_i^p, i \in \{1, \dots, k\}$, of the p th power of the shifted power mean Laplacian $\mathbf{L}_{p,\delta}^p$ again using the Lanczos method. Within the Lanczos process, we employ the PKSM [44] in order to approximate the matrix powers of the single-layer graph Laplacians $(\mathbf{L}_{\text{sym},\delta}^{(t)})^p$ as described in [section 3](#).

The graph Allen–Cahn equation can then be solved using the iteration⁸

$$\mathbf{V}^{l+1} = [(1 + c(\Delta t))\mathbf{I} + \epsilon(\Delta t)\Lambda_k]^{-1} \Phi_k^\top \left((1 + c(\Delta t))\mathbf{U}^l - \frac{\Delta t}{2\epsilon} \mathcal{T}^l - (\Delta t) \omega(\mathbf{U}^l - \mathbf{F}) \right),$$

since it only depends on the eigendecomposition of the graph Laplacian, the known label information \mathbf{F} , and the previous iterate \mathbf{U}^l (\mathcal{T}^l is a function of \mathbf{U}^l). As before, we project each row of the result $\tilde{\mathbf{U}}^{l+1} := \Phi_k \mathbf{V}^{l+1} \in \mathbb{R}^{n \times m}$ of each iteration l back to the Gibbs simplex Σ^m by the method [11] and use these projected values as the input \mathbf{U}^{l+1} for the next iteration $l + 1$.

Finally, we obtain the method summarized in [Algorithm 6.1](#). The algorithm takes the multilayer graph Laplacian and the known label information as inputs. [Algorithm 6.1](#) outputs the scores for the class affiliations of each node $x_i \in \mathcal{V}$ of the multilayer graph. Based on these scores, we predict the class by a majority vote, i.e., we take the rowwise maximum of the output matrix \mathbf{U}^l , where ties are broken by prioritizing the class with the lowest class ID.

Note that [Algorithm 6.1](#) does not employ a fast technique for the case $1 \neq p > 0$. The reason is that this would require the smallest, i.e., informative eigenvalues, of the power mean

⁸This can be obtained from (5.11), as shown in the detailed derivation in [section SM2](#).

Laplacian, which due to the nature of typical iterative eigenvalue solvers need to be computed using the inverse of the matrix; cf., e.g., [21, section 7.6.1]. This would then add another Lanczos iteration on top of the existing ones, which would further increase the algorithmic complexity. It would be possible to investigate the use of the Rayleigh–Chebyshev iteration [3] in combination with the power mean Laplacian, but the results presented in [44] as well as our experiments in section SM3 indicate that positive p 's tend to achieve worse classification results compared to the case $p < 0$, for which we present efficient numerical methods.

Algorithm 6.1 Computation of the multiclass scores for the class affiliations of each node of the multilayer graph using a graph Allen–Cahn-type method.

Input: $\mathbf{L}_{\text{sym}}^{(t)} \in \mathbb{R}^{n \times n}$, graph Laplacian matrix for each layer or function
 $t = 1, \dots, T$ realizing matrix-vector multiplication of graph
 Laplacian matrix with a vector,
 $\mathbf{F} := (\mathbf{f}_i^\top)_{i=1}^n \in \{0, 1\}^{n \times m}$ known label information as per (5.3).

Parameters: $p, k, \epsilon, \omega_0, c, \Delta t, \text{tolerance}, \text{max_iter}$

1. If $p = 1$, then compute eigenpairs $\tilde{\mathbf{\Lambda}}_k, \Phi_k$ of $\mathbf{I} - \mathbf{L}_1$ belonging to the k largest eigenvalues, using the Lanczos method. Set $\mathbf{\Lambda}_k := \mathbf{I} - \tilde{\mathbf{\Lambda}}_k$.
 Otherwise, for $p < 0$, using the Lanczos method in combination with the PKSM [44], compute the k largest eigenvalues $\lambda_1^p, \dots, \lambda_k^p$ of $\mathbf{L}_{p,\delta}^p$ with its eigenvectors, using $\delta := \log(1 + |p|)$. Collect the eigenvectors in Φ_k . Obtain $\mathbf{\Lambda}_k := \text{diag}(\lambda_1, \dots, \lambda_k)$.
2. Compute $\mathbf{Z} := [(1 + c(\Delta t))\mathbf{I} + \epsilon(\Delta t)\mathbf{\Lambda}_k]^{-1}\Phi_k^\top$.
3. Initialize $\mathbf{U}_0 := ((\mathbf{u}_i^0)^\top)_{i=1}^n \in \mathbb{R}^{n \times m}$ with $\mathbf{u}_i^0 := \mathbf{e}_j \in \mathbb{R}^m$ for the prelabeled vertices x_i , where j is the corresponding class. Otherwise, use $\mathbf{u}_i^0 := \frac{1}{m}\mathbf{1} \in \mathbb{R}^m$ for unlabeled x_i . Initialize $l := 0$.
4. **do**
 - (a) Compute the matrix $\mathcal{T}^l = \mathcal{T}(\mathbf{U}^l)$ with entries (5.10).
 - (b) Compute $\mathbf{V}^{l+1} := \mathbf{Z}((1 + c(\Delta t))\mathbf{U}^l - \frac{\Delta t}{2\epsilon}\mathcal{T}^l - (\Delta t)\omega(\mathbf{U}^l - \mathbf{F}))$.
 - (c) Project rows of $\tilde{\mathbf{U}}^{l+1} := \Phi_k \mathbf{V}^{l+1} \in \mathbb{R}^{n \times m}$ to Gibbs simplex Σ^m using [11] and obtain new iterate \mathbf{U}^{l+1} .
 - (d) Calculate $\text{relative_change} := \max_{i=1, \dots, n} \|\mathbf{u}_i^{l+1} - \mathbf{u}_i^l\|^2 / \max_{i=1, \dots, n} \|\mathbf{u}_i^{l+1}\|^2$.
 - (e) $l \leftarrow l + 1$.

while $\text{relative_change} > \text{tolerance}$ and $l < \text{max_iter}$

Output: $\mathbf{U}^l := ((\mathbf{u}_i^l)^\top)_{i=1}^n \in [0, 1]^{n \times m}$ scores for class affiliation of each node x_i , $\mathbf{u}_i^l \in \Sigma^m$.

When the weight matrix $\mathbf{W}^{(t)}$, which encodes node similarities in each layer $t \in \{1, \dots, T\}$, is associated with a suitable⁹ kernel function $K^{(t)}: \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ as discussed in section 4, we can substantially accelerate the computation of the required eigenpairs of the power mean

⁹For example, when K is rotationally invariant and smooth.

Laplacian by using the NFFT-based fast summation in the case $d \leq 3$.¹⁰ This is true even when each weight matrix $\mathbf{W}^{(t)}$ is densely populated, as the computational complexity of one matrix-vector multiplication with $\mathbf{W}^{(t)}$ is reduced from $\mathcal{O}(n^2)$ to $\mathcal{O}(n)$ with the help of Algorithm 4.1.

7. Higher-dimensional data and feature grouping. In this section, we consider the case $d \geq 4$ for which Algorithm 4.1 from section 4 would be computationally too expensive and thus would no longer be directly applicable within a reasonable time frame. This section introduces a framework that makes the NFFT-based methods efficient again. For these higher-dimensional data sets, the formation of the weight matrix \mathbf{W} can be accomplished with a computational complexity of $\mathcal{O}(dn^2)$ when the evaluation of the kernel function is $\mathcal{O}(d)$. The memory requirement of this approach is at most n^2 , and each matrix-vector product can be realized in a computational complexity of $\mathcal{O}(n^2)$. When \mathbf{W} cannot be stored explicitly, each row of \mathbf{W} can also be assembled on-the-fly whenever a matrix-vector multiplication with \mathbf{W} is performed, increasing the computational complexity from $\mathcal{O}(n^2)$ to $\mathcal{O}(dn^2)$ (still assuming $\mathcal{O}(d)$ for each evaluation of the kernel function). For large n , however, the explicit storage or on-the-fly construction of the matrices quickly becomes infeasible for practical computations, and an alternative approach is required.

Incorporating the NFFT-based fast summation from section 4 adapts the computational complexity from $\mathcal{O}(n^2)$ to $\mathcal{O}((m_{\text{NFFT}})^d n + d 2^d N^d \log N)$. This linear dependence on n enables the treatment of large data sets. In order to decrease the exponential dependence of this complexity on the spatial dimension d , we propose a feature grouping approach for $d \geq 4$. The idea is to decompose the d -dimensional feature space into T subspaces of (not necessarily equal) dimensions $d^{(t)} \leq 3$, so that

$$d = \sum_{t=1}^T d^{(t)}.$$

We interpret the data $\mathbf{X}^{(t)}$ in each subspace as one graph layer $\mathcal{G}^{(t)}$ in a multilayer graph as introduced in section 2. Defining a weight function $w^{(t)}$ for each layer enables the formation of the weight matrices $\mathbf{W}^{(t)}$, the degree matrices $\mathbf{D}^{(t)}$, and finally the individual layer graph Laplacians $\mathbf{L}_{\text{sym}}^{(t)}$.

While there is no theory on an optimal feature grouping yet, numerical experiments show that grouping “suitable” features together leads to much higher accuracies in classification tasks. Considering the example of an RGB image, the best image segmentation results can be obtained using a very reasonable grouping of the 5-dimensional feature space consisting of the RGB values and the x - and y -pixel coordinates into one layer containing the RGB color information and a second layer containing the spatial xy information of the pixel;¹¹ cf. [60]. This approach combined with the method presented in section 7 will allow us to consider large image data sets in our numerical tests in subsection 8.2, applying Algorithm 6.1 to a 10-megapixel image which then corresponds to a two-layer graph with 10 million nodes per layer.

¹⁰The case $d \geq 4$ is discussed in the next section.

¹¹We ran several numerical tests with different groupings. The results are not included in this manuscript, but can be reproduced using our code linked in section 8.

The feature grouping approach thus gives rise to a new class of multilayer graphs, where the information from the different graph layers $\mathcal{G}^{(t)}$ is merged together again using the individual layer graph Laplacians $\mathbf{L}_{\text{sym}}^{(t)}$ in the power mean Laplacian defined in (2.2). The combination of this approach with the ideas from section 3 allows a reduction of the computational complexity of matrix-vector products with the power mean Laplacian $\mathbf{L}_{p,\delta}$ of large high-dimensional data sets to approximately $\mathcal{O}(dn)$. The power mean Laplacian is then used in the Allen–Cahn classification scheme¹² as described in section 6.

8. Numerical experiments. In this section, we present the numerical results for the various methods and approaches discussed in the previous sections.¹³ The corresponding MATLAB code is available at <https://github.com/KBergermann/MultilayerSSL>. All runtime measurements were performed on a laptop computer with 16 GB RAM and an Intel Core i5-8265U CPU with $4 \times 1.60\text{--}3.90$ GHz cores. We use the shift $\delta = \log(1 + |p|)$ for $p < 0$ and $\delta = 0$ for $p > 0$ for the power mean Laplacian defined in (2.2) throughout this section.

8.1. Small multilayer data sets. We start by classifying some small real-world example data sets with an inherent multilayer structure that still permit the explicit formation of the power mean Laplacian. Here we consider the data sets from [45, section 6], where the three data sets Citeseer [37], Cora [40], and WebKB [13] have one layer, the three data sets 3sources [36], BBCS [23], and Wikipedia [56] have three layers, the data set BBC [22] has four layers, and the data set UCI [63] six layers. For each multilayer data set, the feature space dimensionality varies across the graph layers ranging between $d^{(t)} = 6$ and $d^{(t)} = 4684$. As n is sufficiently small in all examples, we are in situation (i) of the enumeration in section 2 and explicitly assemble and store all matrices and do not apply feature grouping and the NFFT.

We chose these data sets as they are readily available, and we can compare our new results with the multilayer semi-supervised learning (SSL) approach presented in [45, section 3]. This approach uses the power mean Laplacian for a generalized Tikhonov-type regularization, also called ridge regression, solving the optimization problem

$$(8.1) \quad (u_{i,j})_{i=1}^n = \arg \min_{\mathbf{y}_j \in \mathbb{R}^n} \|\mathbf{y}_j - (f_{i,j})_{i=1}^n\|^2 + \lambda \mathbf{y}_j^\top \mathbf{L} \mathbf{y}_j$$

for each class $j \in \{1, \dots, m\}$, where \mathbf{L} is either \mathbf{L}_1 or $\mathbf{L}_{p,\delta}$ with $p < 0$, and the label of a node v_i is determined by $\arg \max\{u_{i,j}\}_{j=1}^m$. In particular, the test code for [45, section 6] including the known label information for the data sets can be obtained from [42], and we were able to exactly reproduce the corresponding results in [45, Table 2], which presents the mean misclassification rates of 10 test runs with different random known label information. We denote these results by the prefix “SSL” in Table 1. In contrast to the MBO scheme we here solve a linear system and then only once perform a thresholding to obtain the classes.

We employ the multiclass multilayer Allen–Cahn classification scheme (Algorithm 6.1 from section 6) on those data sets using the power mean Laplacians \mathbf{L}_1 , $\mathbf{L}_{-1,\delta}$, and $\mathbf{L}_{-10,\delta}$, where we set the parameters $\epsilon = 5 \cdot 10^{-3}$, $\omega_0 = 1000$, $c = \omega_0 + 3/\epsilon$, $\Delta t = 0.01$, **tolerance** = 10^{-6} , and **max_iter** = 300. The resulting mean misclassification rates are shown in Table 1 with the

¹²This is another $\mathcal{O}(n)$ operation which depends on the termination of the eigeninformation computations.

¹³In addition, section SM3 presents numerical experiments on synthetic stochastic block model data.

Table 1

Mean misclassification rate in percent of the multiclass Allen–Cahn scheme (*Algorithm 6.1* denoted by “AC”) using power mean Laplacians with different parameters p in comparison with results from [45] (denoted by “SSL”) using the same power mean Laplacians. For *Algorithm 6.1*, the parameters $\epsilon = 5 \cdot 10^{-3}$, $\omega_0 = 1000$, $c = \omega_0 + 3/\epsilon$, $\Delta t = 0.01$, **tolerance** = 10^{-6} , and **max_iter** = 300 are used.

3sources [36] ($n = 169$, $T = 3$, $m = 6$; AC: $\sigma = 5$, $k = 16$)

Known labels	1%	5%	10%	15%	20%	25%
SSL L_1	33.5	23.9	23.4	20.1	15.6	14.6
SSL $L_{-1,\delta}$	28.4	20.0	21.8	22.0	17.2	17.9
SSL $L_{-10,\delta}$	40.9	29.1	21.9	19.3	14.8	14.7
AC L_1	39.9	25.1	17.3	15.6	10.6	10.2
AC $L_{-1,\delta}$	40.2	25.3	17.3	15.3	10.6	10.4
AC $L_{-10,\delta}$	38.5	25.1	17.6	15.4	10.7	10.6

BBC [22] ($n = 685$, $T = 4$, $m = 5$; AC: $\sigma = 6$, $k = 31$)

Known labels	1%	5%	10%	15%	20%	25%
SSL L_1	31.3	22.8	17.4	13.5	10.2	8.9
SSL $L_{-1,\delta}$	31.0	17.0	11.5	10.5	9.2	8.7
SSL $L_{-10,\delta}$	51.6	26.9	16.6	12.8	10.3	9.5
AC L_1	41.9	12.9	8.9	7.6	7.0	6.1
AC $L_{-1,\delta}$	41.9	13.0	8.8	7.5	6.9	6.1
AC $L_{-10,\delta}$	42.5	13.2	8.7	7.5	6.9	6.2

BBCS [23] ($n = 544$, $T = 2$, $m = 5$; AC: $\sigma = 2$, $k = 62$)

Known labels	1%	5%	10%	15%	20%	25%
SSL L_1	29.9	15.0	13.5	10.6	8.7	7.2
SSL $L_{-1,\delta}$	23.8	11.6	8.7	6.3	5.8	5.1
SSL $L_{-10,\delta}$	48.7	22.5	14.2	9.1	7.8	6.1
AC L_1	52.4	14.5	8.5	6.5	5.8	5.0
AC $L_{-1,\delta}$	52.3	14.7	8.6	6.5	5.8	5.0
AC $L_{-10,\delta}$	52.1	14.2	8.6	6.4	6.0	5.0

Wikipedia [56] ($n = 693$, $T = 2$, $m = 10$; AC: $\sigma = 2$, $k = 74$)

Known labels	1%	5%	10%	15%	20%	25%
SSL L_1	68.2	61.1	53.6	48.3	44.1	42.3
SSL $L_{-1,\delta}$	59.1	52.3	40.2	36.3	35.1	34.1
SSL $L_{-10,\delta}$	66.9	57.2	43.2	38.7	36.3	34.9
AC L_1	48.4	39.4	31.4	30.1	28.7	27.9
AC $L_{-1,\delta}$	48.4	39.4	31.4	30.0	28.7	27.9
AC $L_{-10,\delta}$	48.4	39.4	31.4	30.1	28.8	27.9

UCI [63] ($n = 2000$, $T = 6$, $m = 10$; AC: $\sigma = 10$, $k = 98$)

Known labels	1%	5%	10%	15%	20%	25%
SSL L_1	31.3	23.8	18.7	15.6	14.4	13.2
SSL $L_{-1,\delta}$	30.5	17.1	13.8	12.6	12.3	11.9
SSL $L_{-10,\delta}$	57.0	33.8	23.7	17.6	15.3	13.4
AC L_1	39.6	11.6	8.7	7.9	7.1	6.5
AC $L_{-1,\delta}$	44.5	11.5	8.6	8.0	6.9	6.6
AC $L_{-10,\delta}$	47.4	12.3	9.1	8.1	7.2	6.9

Citeseer [37] ($n = 3312$, $T = 1$, $m = 6$; AC: $\sigma = 2$, $k = 130$)

Known labels	1%	5%	10%	15%	20%	25%
SSL L_1	56.3	44.1	41.2	38.5	36.1	34.7
SSL $L_{-1,\delta}$	52.4	39.0	35.6	32.6	30.9	29.5
SSL $L_{-10,\delta}$	68.6	54.6	48.5	43.0	39.7	37.2
AC L_1	57.9	34.3	32.1	31.3	30.0	29.4
AC $L_{-1,\delta}$	57.9	34.3	32.1	31.3	30.0	29.4
AC $L_{-10,\delta}$	57.9	34.3	32.1	31.3	30.0	29.4

Cora [40] ($n = 2708$, $T = 1$, $m = 7$; AC: $\sigma = 2$, $k = 85$)

Known labels	1%	5%	10%	15%	20%	25%
SSL L_1	50.7	38.2	33.4	31.2	28.2	25.6
SSL $L_{-1,\delta}$	43.2	31.8	24.5	21.1	18.8	17.2
SSL $L_{-10,\delta}$	62.0	46.3	35.4	29.4	25.2	22.3
AC L_1	62.3	43.1	34.6	32.1	30.3	29.3
AC $L_{-1,\delta}$	62.3	43.1	34.6	32.1	30.3	29.3
AC $L_{-10,\delta}$	62.3	43.1	34.6	32.1	30.3	29.3

WebKB [13] ($n = 187$, $T = 1$, $m = 5$; AC: $\sigma = 2$, $k = 15$)

Known labels	1%	5%	10%	15%	20%	25%
SSL L_1	58.5	49.0	44.8	44.3	44.5	44.4
SSL $L_{-1,\delta}$	49.9	45.5	40.7	39.5	39.9	40.3
SSL $L_{-10,\delta}$	52.3	41.9	38.0	38.1	36.8	39.5
AC L_1	43.7	33.2	23.3	19.6	15.5	14.8
AC $L_{-1,\delta}$	43.7	33.2	23.3	19.6	15.5	14.8
AC $L_{-10,\delta}$	43.7	33.2	23.3	19.6	15.5	14.8

prefix “AC”. We use the Gaussian kernel for the weight matrices \mathbf{W} with scaling parameter σ as mentioned in the table for each data set.

In general, we observe that *Algorithm 6.1* has a lower misclassification rate than the SSL classification scheme [45] for a ratio of known labels $\geq 5\%$, while the latter often performs better for 1% known labels. Moreover, the value of p seems to have a higher influence on the misclassification rate for [45], whereas the influence of p is distinctly smaller in the case of *Algorithm 6.1*. Interestingly, the misclassification rates when applying the SSL classification scheme from [45] also depend on p for the single-layer data sets Citeseer, Cora, and WebKB,

Table 2

Mean runtimes (over 10 runs) in seconds for BBC and UCI data set from Table 1.

BBC							UCI						
Known labels	1%	5%	10%	15%	20%	25%	Known labels	1%	5%	10%	15%	20%	25%
SSL \mathbf{L}_1	0.28	0.27	0.27	0.28	0.30	0.29	SSL \mathbf{L}_1	11.2	11.5	11.3	11.4	11.5	11.3
SSL $\mathbf{L}_{-1,\delta}$	0.30	0.30	0.32	0.31	0.33	0.31	SSL $\mathbf{L}_{-1,\delta}$	11.8	11.4	11.6	11.4	11.2	11.5
SSL $\mathbf{L}_{-10,\delta}$	0.32	0.34	0.34	0.34	0.34	0.34	SSL $\mathbf{L}_{-10,\delta}$	11.2	11.2	11.2	11.2	11.4	11.3
AC \mathbf{L}_1	0.30	0.27	0.21	0.18	0.17	0.16	AC \mathbf{L}_1	3.26	3.16	3.03	3.08	2.83	2.78
AC $\mathbf{L}_{-1,\delta}$	0.39	0.33	0.27	0.24	0.23	0.22	AC $\mathbf{L}_{-1,\delta}$	4.57	4.49	4.21	4.17	4.21	3.96
AC $\mathbf{L}_{-10,\delta}$	0.78	0.73	0.65	0.63	0.63	0.61	AC $\mathbf{L}_{-10,\delta}$	14.1	14.0	14.0	13.8	13.7	13.6

which is not the case for Algorithm 6.1. In addition, the results for the method [45] are better in the case of the Cora data set, whereas the results for Algorithm 6.1 in the case of the WebKB data set are distinctly improved.

One likely explanation for the different behavior of the two considered methods is as follows. Since the modified Ginzburg–Landau energy functional \tilde{E} in (5.6) can be viewed as the loss function of our approach, one main difference between method [45] and Algorithm 6.1 is the additional potential term $\frac{1}{2\epsilon} \sum_{i=1}^n \left(\prod_{l=1}^m \frac{1}{4} \|\mathbf{u}_i - \mathbf{e}_l\|_{\ell^1}^2 \right)$ in (5.6), which promotes more distinct class affiliations.

The observed runtimes behave differently depending on the number of nodes. For instance, for the BBC data set, the SSL classification scheme [45] requires approximately 0.3 seconds (average over 10 test runs) for each considered p and known label ratio, while Algorithm 6.1 requires between 0.16 seconds and 0.78 seconds; cf. Table 2. Moreover, for the UCI data set, the SSL classification scheme [45] has a runtime between 11.2 and 11.8 seconds, whereas Algorithm 6.1 takes between approximately 2.8 and 14.1 seconds.

8.2. Image data. Image segmentation tasks have been a key application for machine learning techniques for many years. While many methods rely on convolutions, image data can also be represented as a graph. A typical approach is to interpret each pixel of an image as a graph node, which is represented by its 3-dimensional 8-bit color channel (i.e., RGB) values, which range in the interval $[0, 255] \cap \mathbb{N}_0$.

One possible way to form a dense weight matrix \mathbf{W} which also takes the spatial relation between the pixels into account is adding each pixel's location in the horizontal direction x as well as in the vertical direction y to the feature space; cf. [60] for a related discussion in the computer vision context. This way, we can again operate on a fully connected graph with feature space dimension $d = 5$.

We employ our feature grouping approach from section 7 and divide the 5-dimensional feature space into two separate graph layers of a multilayer graph with $\mathcal{G}^{(1)}$ containing the 3-dimensional color information and $\mathcal{G}^{(2)}$ the 2-dimensional spatial information of the pixel locations. The information of the two layers is then recombined using the power mean Laplacian introduced in (2.1). This feature grouping enables the fast computation of the first eigenpairs of \mathbf{L}_1 belonging to the smallest eigenvalues by applying the NFFT-based fast summation to both low-dimensional graph layers. Furthermore, numerical experiments revealed that this grouping of “similar” features achieves better image segmentation results than the single-layer graph Laplacian on the full 5-dimensional feature space.

We test that approach on the image of Figure 3 in section 5, which has about 9.7 megapixels. We vectorize the image data, which leads to the two data matrices $\mathbf{X}^{(1)} \in \mathbb{R}^{9734400 \times 3}$ and $\mathbf{X}^{(2)} \in \mathbb{R}^{9734400 \times 2}$, which are centered and then scaled to the boxes $[-1, 1]^3$ and $[-1, 1]^2$, respectively. We apply a multiclass approach with $m = 4$ classes in order to segment the image into the four regions “tree,” “beach,” “sea,” and “sky.” The right image in Figure 3 marks the classified pixels (approximately 4%), which are read out for the initialization of the Allen–Cahn multiclass method in order to act as the known label matrix \mathbf{F} . The rows of the initial solution matrix \mathbf{U}_0 are again initialized with the corresponding unit vector \mathbf{e}_j^\top , where available, and $\frac{1}{m}\mathbf{1}^\top$ otherwise. For this example we use the scaling parameters $\sigma^{(1)} = 1$ and $\sigma^{(2)} = 4$ in the Gaussian kernel as well as compute $k = 12$ eigenpairs. We set the Allen–Cahn parameters $\epsilon = 0.005$, $\omega_0 = 1000$, $c = \omega_0 + 3/\epsilon$, $\Delta t = 0.01$, `max_iter` = 500, and `tolerance` = 10^{-6} . Moreover, for the NFFT-based fast summation, we choose the NFFT parameters bandwidth $N = 64$, window cutoff parameter `m_NFFT` = 5, regularization length $\epsilon_B = 1/16$, and regularization degree `p_NFFT` = 5.

For \mathbf{L}_1 , i.e., the case $p = 1$, the computation of the $k = 12$ eigenpairs using the Lanczos method and Algorithm 4.1 requires approximately 980 seconds. The Allen–Cahn scheme (Algorithm 6.1) reaches its tolerance and terminates after 260 iterations, which require approximately 950 seconds. Note that all runtimes scale almost linearly with respect to the number of pixels n , which makes the segmentation of larger images possible.

We then take the rowwise maximum of the output matrix \mathbf{U} in order to make our prediction, to which class each pixel most likely belongs. Figure 4 shows the original pixel color for pixels, which are identified as belonging to the respective class, and white pixels otherwise. Apart from minor systematic misclassifications for objects on the sea which do not possess their own class, the method segments the image very well. When we consider $\mathbf{L}_{p,\delta}$ with $p = -10$ instead of \mathbf{L}_1 , the classification results improve slightly, but the runtime for computing the eigenpairs increases to approximately 7090 seconds since we have to use the PKSM in addition. Algorithm 6.1 requires again approximately 950 seconds.

In a second step, we consider a downscaled version of the same image as well as a similar image of the same size and concatenate both images to one image with $n = 304720$ pixels. By keeping only the prelabeled nodes from the first image we aim to transfer the known-label information not only to the image itself but also to an unseen image consisting of the same classes. We use the same feature grouping and power mean Laplacian \mathbf{L}_1 . In the second graph layer $\mathcal{G}^{(2)}$ containing the 2-dimensional spatial information of the pixel locations, we now use two connected components, one for each partial image, and the corresponding graph Laplacian $\mathbf{L}_{\text{sym}}^{(2)}$ is block diagonal with two large densely populated blocks belonging to each partial image. For the Gaussian kernel, the NFFT-based fast summation, and the Allen–Cahn scheme, we use the same parameters as for the single image. Moreover, we increase the number of eigenpairs to $k = 35$. Now, we only have around 2% of prelabeled nodes and obtain the classification results displayed in Figure 5 showing obvious misclassifications in Figures 5(c) and 5(d). The computation of the $k = 35$ eigenpairs requires approximately 3870 seconds and the Allen–Cahn scheme additionally approximately 74 seconds. The runtime for the eigenpairs computation can be drastically reduced by requesting fewer values, e.g., to approximately 55 seconds for $k = 27$ while achieving visually similar segmentation results as in Figure 5. The

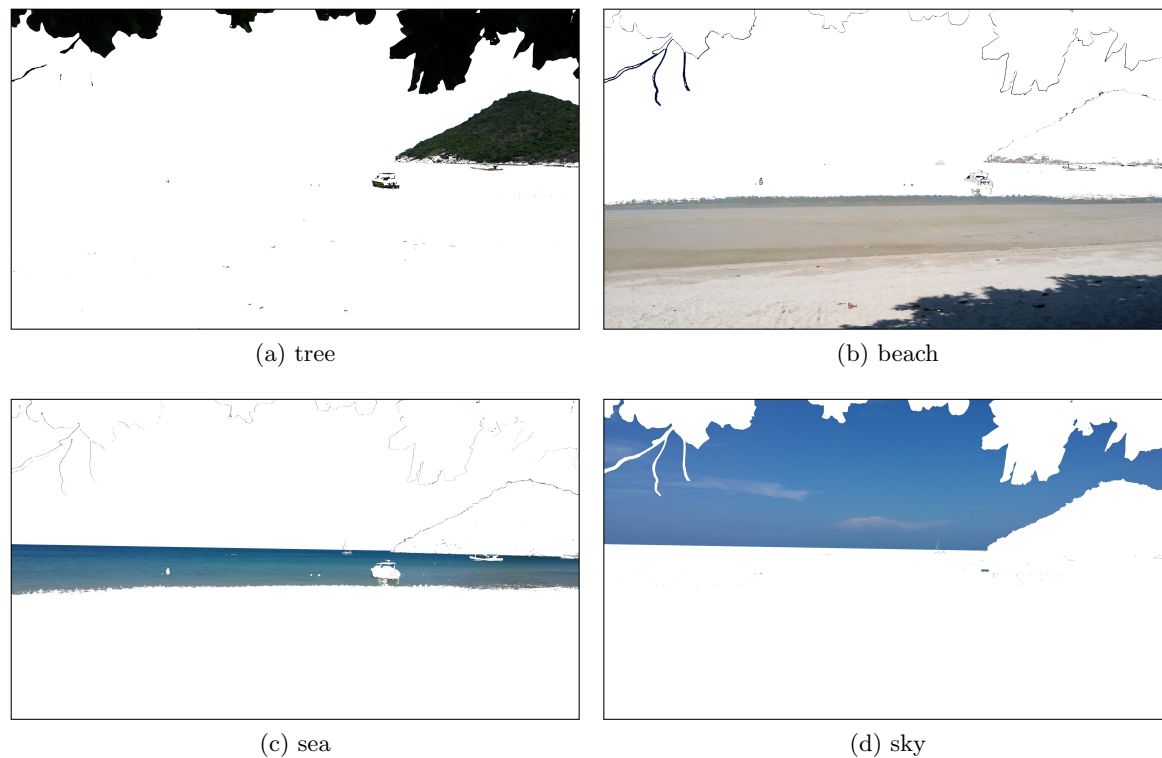


Figure 4. 4-class image segmentation result for test image from [Figure 3](#) using the two-layer power mean Laplacian L_1 .



Figure 5. 4-class image segmentation result for label transfer to a second image using the two-layer power mean Laplacian L_1 .

reason for the larger runtime in case of $k = 35$ eigenpairs is that the used Krylov–Schur method almost stagnates at 27 eigenpairs.

Similar to the observations made in [section SM3](#), the result can be enhanced using a



Figure 6. 4-class image segmentation result for label transfer to a second image using the two-layer power mean Laplacian $L_{-10,\delta}$.

negative power $p < 0$ in the power mean Laplacian. Here in our case, using $p = -10$ leads to much better results, as displayed in Figure 6. The correct detection of the different areas in the second image is a strong result given the varying colors in the two images, especially in the sea and sky regions in Figures 6(c) and 6(d), respectively. The computation of the $k = 35$ eigenpairs requires approximately 585 seconds and the Allen–Cahn scheme approximately 58 seconds. Reducing k to 30 yields a further reduction of the computation time of the eigenpairs to approximately 438 seconds, while still achieving results similar to those in Figure 6. Setting $k < 30$ worsens the segmentation results similar to the ones in Figure 5. We discuss the influence of the parameter k in more detail in the following subsection.

8.3. Hyperspectral data. Finally, we tackle a problem that necessitates the full arsenal of methods derived in this paper. We consider the urban mapping problem posed by the Pavia center data set [52] as an example for the classification of the vast amounts of earth observation data gathered these days.

The original image size of 1096 by 1096 pixels contains valid ground truth labels for $n = 148\,152$ of those pixels. The considered graph contains these labeled nodes. This together with the considerable number of classes $m = 9$ as well as the high feature space dimension arising from 102 atmospherically corrected hyperspectral frequency bands makes this classification problem a demanding task. The frequently discussed combination of spectral and spatial information in the context of hyperspectral image classification [52, 67, 18] is here accomplished by adding the x - and y -pixel coordinates to the feature space, leading to a total of $d = 104$ feature variables.

Since, depending on the assembling strategy of the weight matrix, the memory requirement and/or the runtime of forming the Laplacian matrices would become infeasible, we enable the efficient numerical treatment of the problem by Algorithm 6.1 using our feature grouping approach from section 7 for the 104-dimensional feature space and the NFFT-based fast summation as discussed in section 4 and given in Algorithm 4.1. Note that we refrain from data preprocessing beyond grouping the features into layers and directly work on the raw

Table 3

Average classification accuracies and standard deviations for Pavia center data set [52] using Algorithm 6.1 in combination with Algorithm 4.1 for 5% known labels per class. In the multilayer cases, $p = -10$ is used. “rand.” means random and “det.” deterministic frequency band selection, and “coord.” means coordinate layer. The Gaussian kernel with scaling parameter $\sigma = 8000$ is used for frequency bands and $\sigma = 2 \cdot 1095$ for the coordinates layer.

Layers	Type	#tests	$k = 20$	$k = 40$	$k = 120$
1	2 bands (rand.)	100	0.760 ± 0.031	—	—
2	2 bands (rand.) + coord.	100	0.850 ± 0.043	0.859 ± 0.062	—
2	2 bands (det.) + coord.	51	0.864 ± 0.036	0.879 ± 0.034	—
52	51×2 bands (det.) + coord.	100	0.928 ± 0.001	0.942 ± 0.001	0.957 ± 0.001
1	3 bands (rand.)	100	0.787 ± 0.028	—	—
2	3 bands (rand.) + coord.	100	0.885 ± 0.040	0.896 ± 0.035	—
2	3 bands (det.) + coord.	34	0.915 ± 0.020	0.920 ± 0.014	—
35	34×3 bands (det.) + coord.	100	0.930 ± 0.001	0.943 ± 0.001	0.959 ± 0.001

hyperspectral data. For the Gaussian kernel, we set the scaling parameter to $\sigma = 8000$ for features involving hyperspectral frequency bands, which take values between 0 and 8000, and $\sigma = 2 \cdot 1095$ for coordinates layers, which have pixel coordinates between 1 and 1096. We set the Allen–Cahn parameters $\epsilon = 0.5$, $\omega_0 = 10\,000$, $c = \omega_0 + 3/\epsilon$, $\Delta t = 0.01$, $\text{max_iter} = 300$, $\text{tolerance} = 10^{-6}$. For the NFFT-based fast summation, we choose the NFFT parameters bandwidth $N = 64$, window cutoff parameter $\mathfrak{m}_{\text{NFFT}} = 3$, regularization length $\varepsilon_B = 1/16$, and regularization degree $\mathfrak{p}_{\text{NFFT}} = 3$. Moreover, we use 5 percent random known labels per class.

The results shown in Table 3 illustrate the classification performance for different data modeling approaches. The lowest accuracies are obtained in the cases where we consider single hyperspectral layers of two and three bands, respectively. In each case, we draw 100 random combinations with replacement and average over the 100 test runs. Following the approach presented in subsection 8.2, these results can be distinctly improved by adding the pixel coordinates as a second layer and using the power mean Laplacian with $p = -10$. Interestingly, a deterministic band selection improves the achieved accuracy further. Here, we compute the results averaged over the 51 combinations $(1, 52)$, $(2, 53)$, \dots , $(51, 102)$ of two bands and the 34 combinations $(1, 35, 69)$, $(2, 36, 70)$, \dots , $(34, 68, 102)$ of three bands respectively. For all experiments with only one hyperspectral layer, we obtain a rather high variance in the classification accuracy. In addition to the random choice of hyperspectral bands, these deviations originate from the different random choices of prelabeled nodes in each test run.

The best results, however, are obtained by utilizing all band information by employing our feature grouping approach from section 7. Here, we use all bands with two bands per layer $((1, 52), (2, 53), \dots, (51, 102))$, resulting in 51 frequency bands and 1 coordinate layer, as well as all bands with three bands per layer $((1, 35, 69), (2, 36, 70), \dots, (34, 68, 102))$, resulting in $34+1$ layers, respectively.

Within the range of the number k of eigenpairs of the respective graph Laplacian we consider in our experiments, a higher number of eigenpairs tends to improve the achieved accuracy at the cost of an increased runtime. Depending of the feature space dimension, a

Table 4

Average accuracies and standard deviations over 100 test runs for Pavia center data set [52] using *Algorithm 6.1* with $p = -10$ and $k = 120$ in combination with *Algorithm 4.1* for varying percentage of known labels per class. “det.” means deterministic frequency band selection and “coord.” means coordinates layer. The Gaussian kernel with scaling parameter $\sigma = 8000/2$ is used for frequency bands and $\sigma = 1095/2$ for the coordinates layer.

Type	Known labels per class			
	0.25%	0.5%	1%	5%
51×2 bands (det.) + coord.	0.965 ± 0.005	0.972 ± 0.003	0.968 ± 0.002	0.975 ± 0.001
34×3 bands (det.) + coord.	0.967 ± 0.005	0.974 ± 0.003	0.972 ± 0.002	0.977 ± 0.001

saturation of the accuracy typically sets in at some point. Choosing k too large, however, bears the risk of including noisy eigenvectors, resulting in worse accuracies or potential convergence issues for the eigenvector computations by the Lanczos algorithm. Similar observations can be made for different data sets.

Furthermore, the results can be distinctly improved in the 52-layer case of two bands per layer by modifying the scaling parameter $\sigma = 8000/2$ for the frequency band layers and $\sigma = 1095/2$ for the coordinate layer. We achieve an average accuracy of 0.958 ± 0.001 when using $k = 40$ eigenvectors and of 0.975 ± 0.001 for $k = 120$. We obtain almost the same numbers using the same scaling parameters in the case of all bands with three bands per hyperspectral layer and the additional coordinate layer. Moreover, varying the percentage of known labels per class has only a relatively small influence on the accuracies; cf. Table 4. In particular, for only 0.5% known labels per class, we achieve average accuracies of 0.972 ± 0.003 and 0.974 ± 0.003 for two frequency bands per layer and three frequency bands per layer, respectively. Here, in the two frequency bands per layer case, the eigenpair computation takes approximately 12 800 seconds as well as approximately 26 900 seconds for the three bands per layer case. These increased runtimes in comparison to earlier examples are caused by the increased number of eigenvalues as well as graph layers. The latter, however, only enters the computational complexity linearly, as discussed in section 7. In both cases, the Allen–Cahn scheme requires approximately 71 seconds on average. Note that the computation time for the eigenpairs could be easily reduced on a many-core computer using parallelization.¹⁴

Similarly, applying the same approach to other hyperspectral data sets like the Pavia university data set, another data set presented in [52], or the Indian pines data set [4] also yields mean classification accuracies above 0.97 given 5% and 10% preknown labels, respectively.

Note that there are many results for different classification approaches of the Pavia cen-

¹⁴For instance, we tested the two frequency bands per layer case from Table 4 using a rather simple parallelization based on the `spmd` statement of the MATLAB *Parallel Computing Toolbox*. On a computer with $4 \times$ Intel Xeon E5-4640 each with 8×2.40 GHz cores (in total 32 cores / 64 threads), we observed a runtime of 849 seconds for the eigenpair computation using 52 MATLAB processes, including the setup time for these 52 processes. For comparison, the single threaded case required 15 770 seconds on the same computer, resulting in a speedup of approximately 18.6 and a parallel efficiency of approximately 36%. Moreover, for the three frequency bands per layer case, we observed a runtime of 3120 seconds for the eigenpair computation using 35 processes as well as 54 417 seconds for the single-threaded case, which yields a speedup of approximately 17.4 and a parallel efficiency of approximately 50%. Please note that the runtimes in this footnote cannot be directly compared to the previously stated runtimes as the measurements were performed on another machine with a different CPU type.

ter data set available in the literature with varying foci (e.g., feature selection, maximizing classification accuracy, minimizing prelabeled data) and classification accuracies. While there are examples for the application of “classical” methods [52, 66], some authors reduce the feature space dimension by feature selection techniques [14, 15], whereas the top classification results are achieved with highly specialized convolutional neural network (CNN) architectures [39, 35]. Our results can compete with most results presented in the literature. Better results have been reported for problem-tailored CNN architectures and certain support vector machine classifiers [19, 59]. However, we emphasize again that our method operates directly on the raw data without feature selection and data preprocessing beyond grouping the features into layers, and that it does not require excessive computational power or specialized hardware. In fact, all numerical experiments presented in this paper can be run on an average laptop computer.

9. Conclusion. We have studied the applicability of the power mean Laplacian in the context of diffuse interface-based semi-supervised learning for multilayer networks. The presented feature grouping approach allowed us to tackle data sets with a feature space of larger dimension than was previously possible with standard NFFT-based fast matrix-vector products.

Acknowledgment. The authors are indebted to the anonymous referees for their helpful comments.

REFERENCES

- [1] D. ALFKE, D. POTTS, M. STOLL, AND T. VOLKMER, *NFFT meets Krylov methods: Fast matrix-vector products for the graph Laplacian of fully connected networks*, *Frontiers Appl. Math. Statist.*, 4 (2018).
- [2] S. M. ALLEN AND J. W. CAHN, *A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsening*, *Acta Metallurgica*, 27 (1979), pp. 1085–1095.
- [3] C. R. ANDERSON, *A Rayleigh–Chebyshev procedure for finding the smallest eigenvalues and associated eigenvectors of large sparse Hermitian matrices*, *J. Comput. Phys.*, 229 (2010), pp. 7477–7487.
- [4] M. F. BAUMGARDNER, L. L. BIEHL, AND D. A. LANDGREBE, *220 Band AVIRIS Hyperspectral Image Data Set: June 12, 1992 Indian Pine Test Site 3*, 2015.
- [5] K. BERGERMANN, *Modeling the morphology evolution of organic solar cells*, *GAMM Arch. Students*, 1 (2019), pp. 18–27.
- [6] A. L. BERTOZZI AND A. FLENNER, *Diffuse interface models on graphs for classification of high dimensional data*, *Multiscale Model. Simul.*, 10 (2012), pp. 1090–1118, <https://doi.org/10.1137/11083109X>.
- [7] A. L. BERTOZZI AND A. FLENNER, *Diffuse interface models on graphs for classification of high dimensional data*, *SIAM Rev.*, 58 (2016), pp. 293–328, <https://doi.org/10.1137/16M1070426>.
- [8] S. BOCCALETTI, G. BIANCONI, R. CRIADO, C. I. DEL GENIO, J. GÓMEZ-GARDENES, M. ROMANCE, I. SENDINA-NADAL, Z. WANG, AND M. ZANIN, *The structure and dynamics of multilayer networks*, *Phys. Rep.*, 544 (2014), pp. 1–122.
- [9] J. BOSCH, S. KLAMT, AND M. STOLL, *Generalizing diffuse interface methods on graphs: Nonsmooth potentials and hypergraphs*, *SIAM J. Appl. Math.*, 78 (2018), pp. 1350–1377, <https://doi.org/10.1137/17M1117835>.
- [10] L. CALATRONI, Y. VAN GENNIP, C.-B. SCHÖNLIEB, H. M. ROWLAND, AND A. FLENNER, *Graph clustering, variational image segmentation methods and Hough transform scale detection for object measurement in images*, *J. Math. Imaging Vision*, 57 (2017), pp. 269–291.
- [11] Y. CHEN AND X. YE, *Projection onto a Simplex*, preprint, <https://arxiv.org/abs/1101.6081>, 2011.
- [12] F. R. K. CHUNG, *Spectral Graph Theory*, CBMS Reg. Conf. Ser. in Math. 92, AMS, Providence, RI, 1997.

- [13] M. CRAVEN, D. DiPASQUO, D. FREITAG, A. MCCALLUM, T. MITCHELL, K. NIGAM, AND S. SLATTERY, *Learning to extract symbolic knowledge from the world wide web*, in Proceedings of the Fifteenth National/Tenth Conference on Artificial Intelligence/Innovative Applications of Artificial Intelligence, (AAAI '98/IAAI '98), AAAI Press, Menlo Park, CA, 1998, pp. 509–516.
- [14] M. DALLA MURA, J. A. BENEDIKTSSON, J. CHANUSSOT, AND L. BRUZZONE, *The evolution of the morphological profile: From panchromatic to hyperspectral images*, in Optical Remote Sensing, Springer, 2011, pp. 123–146.
- [15] A. DAVARI, V. CHRISTLEIN, S. VESAL, A. MAIER, AND C. RIESS, *GMM supervectors for limited training data in hyperspectral remote sensing image classification*, in International Conference on Computer Analysis of Images and Patterns, Springer, 2017, pp. 296–306.
- [16] D. J. EYRE, *An Unconditionally Stable One-Step Scheme for Gradient Systems*, 1997.
- [17] D. J. EYRE, *Unconditionally gradient stable time marching the Cahn-Hilliard equation*, MRS Proc., 529 (1998), 39.
- [18] B. FANG, Y. LI, H. ZHANG, AND J. C.-W. CHAN, *Semi-supervised deep learning classification for hyperspectral image based on dual-strategy sample selection*, Remote Sensing, 10 (2018), 574.
- [19] M. FAUVEL, Y. TARABALKA, J. A. BENEDIKTSSON, J. CHANUSSOT, AND J. C. TILTON, *Advances in spectral-spatial classification of hyperspectral images*, Proc. IEEE, 101 (2012), pp. 652–675.
- [20] C. GARCIA-CARDONA, E. MERKURJEV, A. L. BERTOZZI, A. FLENNER, AND A. G. PERCUS, *Multiclass data segmentation using diffuse interface methods on graphs*, IEEE Trans. Pattern Anal. Mach. Intell., 36 (2014), pp. 1600–1613.
- [21] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, 2012.
- [22] D. GREENE AND P. CUNNINGHAM, *Producing accurate interpretable clusters from high-dimensional data*, in European Conference on Principles of Data Mining and Knowledge Discovery, Springer, 2005, pp. 486–494.
- [23] D. GREENE AND P. CUNNINGHAM, *A matrix factorization approach for integrating multiple data views*, in Joint European Conference on Machine Learning and Knowledge Discovery in Databases, Springer, 2009, pp. 423–438.
- [24] M. HEIN, S. SETZER, L. JOST, AND S. S. RANGAPURAM, *The total variation on hypergraphs-learning on hypergraphs revisited*, in Advances in Neural Information Processing Systems, 2013, pp. 2427–2435.
- [25] N. J. HIGHAM, *Functions of Matrices: Theory and Computation*, SIAM, Philadelphia, 2008, <https://doi.org/10.1137/1.9780898717778>.
- [26] P. W. HOLLAND, K. B. LASKEY, AND S. LEINHARDT, *Stochastic blockmodels: First steps*, Social Networks, 5 (1983), pp. 109–137.
- [27] H.-C. HUANG, Y.-Y. CHUANG, AND C.-S. CHEN, *Affinity aggregation for spectral clustering*, in Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2012, pp. 773–780.
- [28] L. G. S. JEUB, M. W. MAHONEY, P. J. MUCHA, AND M. A. PORTER, *A local perspective on community structure in multilayer networks*, Network Sci., 5 (2017), p. 144–163.
- [29] J. KEINER, S. KUNIS, AND D. POTTS, *NFFT 3.5, C Subroutine Library*, F. Bartel, M. Fenn, T. Görner, M. Kirchels, T. Knopp, M. Quellmalz, M. Schmiscke, T. Volkmer, and A. Vollrath, contributors, <http://www.tu-chemnitz.de/~potts/nfft>, <https://github.com/NFFT/nfft>.
- [30] J. KEINER, S. KUNIS, AND D. POTTS, *Using NFFT3—a software library for various nonequispaced fast Fourier transforms*, ACM Trans. Math. Softw., 36 (2009), 19.
- [31] T. N. KIPF AND M. WELLING, *Semi-supervised classification with graph convolutional networks*, in Proceedings of the 5th International Conference on Learning Representations (ICLR '17), Open-Review.net, 2017.
- [32] M. KIVELÄ, A. ARENAS, M. BARTHELEMY, J. P. GLEESON, Y. MORENO, AND M. A. PORTER, *Multilayer networks*, J. Complex Networks, 2 (2014), pp. 203–271.
- [33] J. KUNEGIS, S. SCHMIDT, A. LOMMATZSCH, J. LERNER, E. W. DE LUCA, AND S. ALBAYRAK, *Spectral analysis of signed graphs for clustering, prediction and visualization*, in Proceedings of the 2010 SIAM International Conference on Data Mining, SIAM, Philadelphia, 2010, pp. 559–570, <https://doi.org/10.1137/1.9781611972801.49>.

- [34] S. KUNIS AND D. POTTS, *Time and memory requirements of the nonequispaced FFT*, *Sampl. Theory Signal Image Process.*, 7 (2008), pp. 77–100.
- [35] L. LIN, C. CHEN, AND T. XU, *Spatial-spectral hyperspectral image classification based on information measurement and CNN*, *EURASIP J. Wireless Commun. Networking*, 2020 (2020), pp. 1–16.
- [36] J. LIU, C. WANG, J. GAO, AND J. HAN, *Multi-view clustering via joint nonnegative matrix factorization*, in *Proceedings of the 2013 SIAM International Conference on Data Mining*, SIAM, Philadelphia, 2013, pp. 252–260, <https://doi.org/10.1137/1.9781611972832.28>.
- [37] Q. LU AND L. GETOOR, *Link-based classification*, in *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, AAAI Press, Menlo Park, CA, 2003, pp. 496–503.
- [38] X. LUO AND A. L. BERTOZZI, *Convergence of the graph Allen–Cahn scheme*, *J. Statist. Phys.*, 167 (2017), pp. 934–958.
- [39] K. MAKANTASIS, K. KARANTZALOS, A. DOULAMIS, AND N. DOULAMIS, *Deep supervised learning for hyperspectral data classification through convolutional neural networks*, in *Proceedings of the 2015 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, IEEE, 2015, pp. 4959–4962.
- [40] A. K. MCCALLUM, K. NIGAM, J. RENNIE, AND K. SEYMORE, *Automating the construction of internet portals with machine learning*, *Inform. Retrieval*, 3 (2000), pp. 127–163.
- [41] Z. MENG, E. MERKURJEV, A. KONIGES, AND A. L. BERTOZZI, *Hyperspectral image classification using graph clustering methods*, *Image Process. On Line*, 7 (2017), pp. 218–245.
- [42] P. MERCADO, *MATLAB Implementation of the Paper: Generalized Matrix Means for Semi-supervised Learning with Multilayer Graphs*, https://github.com/melopeo/PM_SSL.
- [43] P. MERCADO, J. BOSCH, AND M. STOLL, *Node classification for signed social networks using diffuse interface methods*, in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, Springer, 2019, pp. 524–540.
- [44] P. MERCADO, A. GAUTIER, F. TUDISCO, AND M. HEIN, *The power mean Laplacian for multilayer graph clustering*, in *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, A. Storkey and F. Perez-Cruz, eds., *Proc. Mach. Learn. Res.* 84, PMLR, 2018, pp. 1828–1838.
- [45] P. MERCADO, F. TUDISCO, AND M. HEIN, *Generalized matrix means for semi-supervised learning with multilayer graphs*, in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, eds., Curran Associates, 2019, pp. 14877–14886.
- [46] E. MERKURJEV, T. KOSTIĆ, AND A. L. BERTOZZI, *An MBO scheme on graphs for classification and image processing*, *SIAM J. Imaging Sci.*, 6 (2013), pp. 1903–1930, <https://doi.org/10.1137/120886935>.
- [47] B. MOHAR, *The Laplacian spectrum of graphs*, in *Graph Theory, Combinatorics, and Applications*, Wiley, 1991, pp. 871–898.
- [48] B. MOHAR, *Some Applications of Laplace Eigenvalues of Graphs*, Springer, Dordrecht, 1997, pp. 225–275.
- [49] V. I. MORARIU, B. V. SRINIVASAN, V. C. RAYKAR, R. DURAI SWAMI, AND L. S. DAVIS, *Automatic online tuning for fast Gaussian summation*, in *Advances in Neural Information Processing Systems*, Curran Associates, 2009, pp. 1113–1120.
- [50] P. J. MUCHA, T. RICHARDSON, K. MACON, M. A. PORTER, AND J.-P. ONNELA, *Community structure in time-dependent, multiscale, and multiplex networks*, *Science*, 328 (2010), pp. 876–878.
- [51] A. ONUKI, *Phase Transition Dynamics*, Cambridge University Press, 2002.
- [52] A. PLAZA, J. A. BENEDIKTSSON, J. W. BOARDMAN, J. BRAZILE, L. BRUZZONE, G. CAMPS-VALLS, J. CHANUSSOT, M. FAUVEL, P. GAMBA, A. GUALTIERI ET AL., *Recent advances in techniques for hyperspectral image processing*, *Remote Sensing Environ.*, 113 (2009), pp. S110–S122.
- [53] G. PLONKA, D. POTTS, G. STEIDL, AND M. TASCHE, *Numerical Fourier Analysis*, in *Applied and Numerical Harmonic Analysis*, Birkhäuser/Springer, Cham, 2018.
- [54] D. POTTS AND G. STEIDL, *Fast summation at nonequispaced knots by NFFTs*, *SIAM J. Sci. Comput.*, 24 (2003), pp. 2013–2037, <https://doi.org/10.1137/S1064827502400984>.
- [55] D. POTTS, G. STEIDL, AND A. NIESLONY, *Fast convolution with radial kernels at nonequispaced knots*, *Numer. Math.*, 98 (2004), pp. 329–351.

- [56] N. RASIWASIA, J. COSTA PEREIRA, E. COVIELLO, G. DOYLE, G. R. LANCKRIET, R. LEVY, AND N. VASCONCELOS, *A new approach to cross-modal multimedia retrieval*, in Proceedings of the 18th ACM International Conference on Multimedia, ACM, 2010, pp. 251–260.
- [57] M. STOLL, *A literature survey of matrix methods for data science*, GAMM-Mitt., 43 (2020), e202000013.
- [58] S. H. STROGATZ, *Exploring complex networks*, Nature, 410 (2001), pp. 268–276.
- [59] Z. SUN, C. WANG, P. LI, H. WANG, AND J. LI, *Hyperspectral image classification with SVM-based domain adaption classifiers*, in Proceedings of the 2012 International Conference on Computer Vision in Remote Sensing, IEEE, 2012, pp. 268–272.
- [60] C. TOMASI AND R. MANDUCHI, *Bilateral filtering for gray and color images*, in Proceedings of the Sixth International Conference on Computer Vision, IEEE, 1998, pp. 839–846.
- [61] K. TSUDA, H. SHIN, AND B. SCHÖLKOPF, *Fast protein classification with multiple networks*, Bioinform., 21 (2005), pp. 59–65.
- [62] F. TUDISCO, P. MERCADO, AND M. HEIN, *Community detection in networks via nonlinear modularity eigenvectors*, SIAM J. Appl. Math., 78 (2018), pp. 2393–2419, <https://doi.org/10.1137/17M1144143>.
- [63] M. VAN BREUKELN, R. P. W. DUIN, D. M. J. TAX, AND J. E. DEN HARTOG, *Handwritten digit recognition by combined classifiers*, Kybernetika, 34 (1998), pp. 381–386.
- [64] Y. VAN GENNIP, N. GUILLEN, B. OSTING, AND A. L. BERTOZZI, *Mean curvature, threshold dynamics, and phase field theory on finite graphs*, Milan J. Math., 82 (2014), pp. 3–65.
- [65] U. VON LUXBURG, *A tutorial on spectral clustering*, Stat. Comput., 17 (2007), pp. 395–416.
- [66] Q. WANG AND J. ZHANG, *A data transfer fusion method for discriminating similar spectral classes*, Sensors, 16 (2016), 1895.
- [67] Y. WANG, H. SONG, AND Y. ZHANG, *Spectral-spatial classification of hyperspectral images using joint bilateral filter and graph cut based model*, Remote Sensing, 8 (2016), 748.
- [68] A. A. WHEELER, W. J. BOETTINGER, AND G. B. MCFADDEN, *Phase-field model for isothermal phase transitions in binary alloys*, Phys. Rev. A, 45 (1992), pp. 7424–7439.
- [69] O. WODO AND B. GANAPATHYSUBRAMANIAN, *Modeling morphology evolution during solvent-based fabrication of organic solar cells*, Comput. Materials Sci., 55 (2012), pp. 113–126.
- [70] D. ZHOU, O. BOUSQUET, T. N. LAL, J. WESTON, AND B. SCHÖLKOPF, *Learning with local and global consistency*, in Advances in Neural Information Processing Systems, MIT Press, 2004, pp. 321–328.
- [71] X. ZHU AND A. B. GOLDBERG, *Introduction to semi-supervised learning*, Synthesis Lectures Artif. Intell. Mach. Learning, 3 (2009), pp. 1–130.