Antonio G. Marques, Santiago Segarra,
and Gonzalo Mateos

Signal Processing on Directed Graphs

The role of edge directionality when processing and learning from network data

©ISTOCKPHOTO.COM/ALISEFOX

Digital Object Identifier 10.1109/MSP.2020.3014597
Date of current version: 28 October 2020

This article provides an overview of the current landscape of signal processing (SP) on directed graphs (digraphs). Directionality is inherent to many real-world (information, transportation, biological) networks, and it should play an integral role in processing and learning from network data. We thus lay out a comprehensive review of recent advances in SP on digraphs, offering insights through comparisons with results available for undirected graphs, discussing emerging directions, establishing links with related areas in machine learning and causal inference in statistics as well as illustrating their practical relevance to timely applications. To this end, we begin by surveying (orthonormal) signal representations and their graph-frequency interpretations based on novel measurements of signal variation for digraphs. We then move on to filtering, a central component in deriving a comprehensive theory of SP on digraphs. Indeed, through the lens of filter-based generative signal models, we explore a unified framework to study inverse problems (e.g., sampling and deconvolution on networks), the statistical analysis of random signals, and the topology inference of digraphs from nodal observations.

Introduction and motivation

Coping with the panoply of challenges found at the confluence of data and network sciences necessitates fundamental breakthroughs in the modeling, identification, and controllability of networked (complex) system processes—often conceptualized as signals defined on graphs [1]. Graph-supported signals abound in real-world applications, including vehicle-congestion levels over road networks, neurological activity signals supported on brain-connectivity networks, and fake news that diffuse on online social networks. There is, however, an evident mismatch between our scientific understanding of signals defined over regular domains such as time or space and graph signals, due, in part, to the fact that the prevalence of network-related problems and access to quality network data are recent events. To address these problems, machine learning and SP over graphs have emerged as active

areas aimed at making sense of large-scale data sets from a network-centric perspective.

Upon modeling the domain of the information as a graph and the observations at hand as graph signals, the graph SP (GSP) body of work has put forth models that relate the properties of the signals with those of the graph, along with algorithms that fruitfully leverage this relational structure to better process and learn from network data. Most GSP efforts to date assume that the underlying networks are undirected [2]. Said graphs are equivalently represented by symmetrical matrices whose (well-behaved) spectral properties can be used to process the signals associated with the network. The most prominent example is the graph Laplacian, which not only gives rise to a natural definition of signal smoothness but also offers a complete set of orthonormal eigenvectors that serve as a Fourier-type basis for graph signals [3].

Most GSP efforts to date assume that the underlying networks are undirected.

Their scarcer adoption notwithstanding, digraph models are more adequate (and, in fact, more accurate) for a number of applications. Information networks such as scientific citations or the World Wide Web itself are typically directed, and flows in technological (e.g., transportation, power, and communication) networks are often unidirectional. The presence of directionality plays a critical role when the measurements taken in those networks need to be processed to remove noise, outliers, and artifacts, and this requires new tools and algorithms that do not assume that the matrices representing the underlying graphs are symmetrical.

Gene-regulatory networks are highly nonreciprocal, and this lack of reciprocity needs to be accounted for when, for example, the goal is to predict a gene or a protein functionality from a small set of observations obtained from expensive experiments. Pairwise relations among social actors are rarely purely symmetrical [4] and, in fact, when the graph captures some level of influence on a social network, the lack of symmetry is essential to accurately solve inverse problems that aim to separate the leaders from the followers [5]. More abstractly, when the graph encodes (often unknown) relations between observed variables, directionality is vital to identify the nodes representing the cause and those representing the effect [6], calling for fundamental changes in the algorithms that use available signal observations to learn the topology of the underlying graph. Accordingly, a first step to address these and other related questions is to develop judicious models that account for directionality while leading to tractable processing tools and efficient algorithms. That is precisely the goal of this tutorial article, which aims at delineating the analytical background and relevance of innovative tools to analyze and process signals defined over digraphs. Throughout, concepts will be made accessible to SP researchers (including those without a strong background in network science) via a combination of rigorous problem formulations and intuitive reasoning. A recurrent message with important practical ramifications interweaves the narrative—different from the undirected case, where graph spectrum-based tools offer a number of distinct advantages [3], vertex-domain generative graph signal models that rely on nonsymmetric network operators may be preferable when it comes to signal and information processing on directed networks.

GSP preliminaries, frequency analysis, and signal representations

After introducing the necessary graph-theoretic notation and background, this section presents different generalizations of smoothness and total-variation measurements for signals defined on digraphs. This is particularly relevant to the graph Fourier transform (GFT), which decomposes a graph signal into components that describe different modes of variation with respect to the graph topology. Although adopting the real-valued orthonormal eigenvectors of the Laplacian as the frequency basis for undirected graphs is well motivated and widely used in practice [1], extending the GFT framework to digraphs is not a simple pursuit and different alternatives exist, as we explain in the "Digraph Fourier Transforms: Spectral Methods" and "Digraph Fourier Transforms: Orthonormal Transform Learning" sections.

Graph signals and the graph shift operator

Let $\mathcal{G}$ denote a digraph with a set of nodes $\mathcal{N}$ (with cardinality N) and a set of links $\mathcal{E}$. If i is connected to j, then $(i,j) \in \mathcal{E}$. Because $\mathcal{G}$ is directed, local connectivity is captured by the set $\mathcal{N}_i := \{j \,|\, (j,i) \in \mathcal{E}\}$, which stands for the (incoming) neighborhood of i. For any given $\mathcal{G}$, we define the adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ as a sparse matrix with nonzero elements A_{ji} if and only if $(i,j) \in \mathcal{E}$. The value of A_{ji} captures the strength of the connection from i to j and, because the graph is directed, the matrix $\mathbf{A}$ is, in general, nonsymmetric.

The focus of the article is on analyzing and modeling (graph) signals defined on the node set $\mathcal{N}$. These signals can be represented as vectors $\mathbf{x} = [x_1, ..., x_N]^T \in \mathbb{R}^N$, with x_i being the value of the signal at node i. As the vectorial representation does not explicitly account for the structure of the graph, $\mathcal{G}$ can be endowed with the so-called graph shift operator (GSO) $\mathbf{S}$ [7], [8]. The shift $\mathbf{S} \in \mathbb{R}^{N \times N}$ is a matrix whose entry S_{ji} can be nonzero only if $i = j$ or if $(i,j) \in \mathcal{E}$. The sparsity pattern of the matrix $\mathbf{S}$ captures the local structure of $\mathcal{G}$, but we make no specific assumptions on the values of its nonzero entries, which will depend on the application at hand [1].

To justify the adopted graph shift terminology, consider the directed cycle graph whose circulant adjacency matrix $\mathbf{A}_{dc}$ is zero, except for entries $A_{ji} = 1$ whenever $i = \text{mod}_N(j) + 1$, where $\text{mod}_N(x)$ denotes the modulus (remainder) obtained after dividing x by N. Such a graph can be used to represent the domain of discrete-time periodic signals with period N. If $\mathbf{S} = \mathbf{A}_{dc}$, then $\mathbf{Sx}$ implements a circular shift of the entries in $\mathbf{x}$, which corresponds to a one-unit time delay under the aforementioned interpretation [1]. Note though that, in general, $\mathbf{S}$ need neither be invertible nor isometric, an important departure from the shift in discrete-time SP. The intuition behind

$\mathbf{S}$ is to represent a linear transformation that can be computed locally at the nodes of the graph, while it can be more general than the adjacency matrix. More rigorously, if the graph signal $\mathbf{y}$ is defined as $\mathbf{y} = \mathbf{S}\mathbf{x}$, then node i can compute y_i as a linear combination of the signal values x_j at node i's neighbors $j \in \mathcal{N}_i$. The GSO will play a fundamental role in defining the counterpart of the Fourier transform for graph signals, which is discussed in this section, as well as graph filters that are introduced in the "Graph Filters and Nonlinear Graph Signal Operators" section.

Vertex-domain generative graph signal models that rely on nonsymmetric network operators may be preferable when it comes to signal and information processing on directed networks.

Digraph fourier transforms: Spectral methods

An instrumental GSP tool is the GFT, which decomposes a graph signal into orthonormal components describing different modes of variation with respect to the graph topology encoded in an application-dictated GSO $\mathbf{S}$. The GFT allows equivalently representing a graph signal in two different domains—the vertex domain consisting of the nodes in $\mathcal{N}$ and the graph-frequency domain spanned by the spectral basis of $\mathcal{G}$. Therefore, signals can be manipulated in the frequency domain for the purpose of, e.g., denoising, compression, and feature extraction (see also the "Graph Filters and Nonlinear Graph Signal Operators" section). For didactic purposes, it is informative to introduce first the GFT for symmetrical graph Laplacians associated with undirected graphs (see "A Motivating Starting Point: The Graph Fourier Transform for Undirected Graphs"). In the remainder of this section, we show that the GFT can be defined for digraphs where the interpretation of components as different modes of variability is not as clean and, Parseval's identity may not hold, but its value toward yielding parsimonious spectral representations of network processes remains.

The Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{A}$ is not well defined for digraphs because $\mathbf{D}$ is rendered meaningless when edges have directionality. One can instead consider a generic asymmetric GSO $\mathbf{S}$, for instance, the adjacency matrix $\mathbf{A}$ or one of the several generalized Laplacians for digraphs; see, e.g., [9] and [10]. Suppose the GSO is diagonalizable as $\mathbf{S} = \mathbf{V}\text{diag}(\boldsymbol{\lambda})\mathbf{V}^{-1}$, with $\mathbf{V} := [\mathbf{v}_1, ..., \mathbf{v}_N]$ denoting the (nonorthogonal) eigenvectors of $\mathbf{S}$ and $\boldsymbol{\lambda} := [\lambda_1, ..., \lambda_N]^T$ its possibly complex-valued eigenvalues. Then, a widely adopted alternative is to redefine the GFT as $\tilde{\mathbf{x}} = \mathbf{V}^{-1}\mathbf{x}$ [8]. Otherwise, one can resort to the Jordan decomposition of $\mathbf{S}$ and use its generalized eigenvectors as the GFT basis; see also [11] for a careful treatment of the nondiagonalizable case, which relies on oblique spectral projectors to define the GFT. Setting the GFT to $\mathbf{V}^{-1}$ for the directed case is an intuitively pleasing definition because frequency shift operator, as in discrete-time SP. Moreover, allowing for generic GSOs reveals the encompassing nature of the GFT relative to the time-domain discrete Fourier transform (DFT), the multidimensional DFT, and principal component analysis [12]. Toward interpreting graph frequencies, which are defined by the (possibly complex-valued, nonorthogonal) eigenvectors of the nonsymmetric $\mathbf{S}$, consider the total-variation measure

$$\text{TV}_1(\mathbf{x}) := \left\| \mathbf{x} - \tilde{\mathbf{S}}\mathbf{x} \right\|_1, \tag{1}$$

where $\mathbf{S} = \mathbf{S}/|\lambda_{\max}|$, and $\lambda_{\max}$ is the spectral radius of $\mathbf{S}$ [compare (S1)]. Using (1) and following the rationale for undirected

A Motivating Starting Point: The Graph Fourier Transform for Undirected Graphs

Consider an undirected graph $\mathcal{G}$ with combinatorial Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{A}$ chosen as the graph shift operator [3], where $\mathbf{D}$ stands for the diagonal degree matrix. The symmetrical $\mathbf{L}$ can always be decomposed as $\mathbf{L} = \mathbf{V}\text{diag}(\boldsymbol{\lambda})\mathbf{V}^T$, with $\mathbf{V} := [\mathbf{v}_1, ..., \mathbf{v}_N]$ collecting the orthonormal eigenvectors of the Laplacian and $\boldsymbol{\lambda} := [\lambda_1, ..., \lambda_N]^T$ its nonnegative eigenvalues. The graph Fourier transform (GFT) of $\mathbf{x}$ with respect to $\mathbf{L}$ is the signal $\tilde{\mathbf{x}} = [\tilde{x}_1, ..., \tilde{x}_N]^T$ defined as $\tilde{\mathbf{x}} = \mathbf{V}^T\mathbf{x}$. The inverse GFT (iGFT) of $\tilde{\mathbf{x}}$ is given by $\mathbf{x} = \mathbf{V}\tilde{\mathbf{x}}$, which is a proper inverse by the orthogonality of $\mathbf{V}$.

The iGFT formula $\mathbf{x} = \mathbf{V}\tilde{\mathbf{x}} = \sum_{k=1}^{N} \tilde{x}_k \mathbf{v}_k$ allows one to synthesize $\mathbf{x}$ as a sum of orthogonal-frequency components $\mathbf{v}_k$. The contribution of $\mathbf{v}_k$ to the signal $\mathbf{x}$ is the real-valued GFT coefficient $\tilde{x}_k$. The GFT encodes a notion of signal variability over the graph akin to the notion of frequency in the Fourier analysis of temporal signals. To understand this analogy, define the total variation of the graph signal $\mathbf{x}$ with respect to the Laplacian $\mathbf{L}$ (also known as *Dirichlet energy*) as the following quadratic form:

$$\text{TV}_2(\mathbf{x}) := \mathbf{x}^T\mathbf{L}\mathbf{x} = \sum_{i<j} A_{ij}(x_i - x_j)^2. \tag{S1}$$

The total variation $\text{TV}_2(\mathbf{x})$ is a smoothness measure, quantifying how much the signal $\mathbf{x}$ changes with respect to the graph topology encoded in $\mathbf{A}$.

Back to the GFT, consider the total variation of the eigenvectors $\mathbf{v}_k$, which is given by $\text{TV}_2(\mathbf{v}_k) = \mathbf{v}_k^T\mathbf{L}\mathbf{v}_k = \lambda_k$. It follows that the eigenvalues $0 = \lambda_1 \le \lambda_2 \le \ldots \le \lambda_N$ can be viewed as graph frequencies, indicating how the eigenvectors (i.e., frequency components) vary over the graph $\mathcal{G}$. Accordingly, the GFT and iGFT offer a decomposition of the graph signal $\mathbf{x}$ into spectral components that characterize different levels of variability.

graphs, one can define a frequency ordering $\lambda_i \succ \lambda_j$ if $\mathrm{TV}_1(\mathbf{v}_i) > \mathrm{TV}_1(\mathbf{v}_j)$ [8]. Although applicable to signals on digraphs, unlike (S1), the signal-variation measure (1) does not ensure that constant signals have zero variation. In addition, (generalized) eigenvectors of asymmetric GSOs need not be orthonormal, implying that Parseval's identity will not hold and hence, the signal power is not preserved across the vertex and dual domains. In general, this can be an issue for graph-filtering methods operating in the spectral domain, thus motivating this article's overarching theme of relying on vertex-domain operations for extensions to digraphs. From a computational standpoint, obtaining the Jordan decomposition for moderate-sized graphs is expensive and often numerically unstable; see also [11] and the references therein for recent attempts toward mitigating this instability issue. Addressing the uniqueness of the representation is also critical when the GSO (even the combinatorial Laplacian) has repeated eigenvalues because the corresponding eigenspaces exhibit rotational ambiguities that can hinder the interpretability of graph-frequency analyses. To address this (often overlooked shortcoming), [11] puts forth a quasi-coordinate-free GFT definition based on oblique spectral projectors. Other noteworthy GFT approaches rely on projections onto the (nonorthogonal) eigenvectors of a judicious random-walk operator on the digraph [9], [10]; the interested reader is referred to [9, Sec. 7] for a collection of examples involving semisupervised learning and signal modeling on digraphs.

Orthonormal linear transformations excel at separating signals from noise.

Alternatives to the spectral GFT methods described thus far are surveyed in the following section. The focus shifts to orthonormal transform learning approaches, whereby optimization problems are formulated to find a suitable spectral representation basis for graph signals.

Digraph Fourier transforms: Orthonormal transform learning

The history of SP has repeatedly taught us how low frequencies are more meaningful in human speech for the purpose of compression, high frequencies represent borders in images whose identification is key for segmentation, and different principal components offer varying discriminative powers when it comes to face recognition. Although analogous interpretations are not always possible in more advanced representations obtained with modern tools such as learned overcomplete dictionaries and neural networks (NNs), at a basic level, it remains true that orthonormal linear transformations excel at separating signals from noise. Motivated by this general signal representation principle, a fresh look at the GFT for digraphs was put forth in [13] based on the minimization of the convex Lovász extension of the graph cut size (which can be interpreted as a measurement of signal variation on the graph capturing the edges' directionality), subject to orthonormality constraints on the desired basis. The rationale behind the graph cut criterion is that its minimization leads to identifying clusters in $\mathcal{G}$. Accordingly, the learned GFT basis in [13] tends to be constant across clusters of the graph, offering parsimonious spectral representations of signals that are real valued and piecewise-constant over said clusters. The price paid for all of these desirable properties is that the resulting GFT basis may fail to yield atoms capturing different levels of signal variation with respect to $\mathcal{G}$, and the optimization procedure in [13] is computationally expensive due to repeated singular-value decompositions.

A related (optimization-based) approach in [14] searches for an orthonormal digraph Fourier transform (DGFT) basis $\mathbf{U} := [\mathbf{u}_1, \ldots, \mathbf{u}_N] \in \mathbb{R}^{N \times N}$, where $\mathbf{u}_k \in \mathbb{R}^N$ represents the kth frequency component. Toward defining frequencies, a more general notion of signal-directed variation (DV) for digraphs is introduced as $\mathrm{DV}(\mathbf{x}) := \sum_{i \neq j} A_{ji} [x_i - x_j]_+^2$, where $[x]_+ := \max(0, x)$ denotes projection onto the nonnegative reals. To gain insight on DV, consider a graph signal $\mathbf{x}$ on the digraph $\mathcal{G}$, and suppose a directed edge represents the direction of signal flow from a larger value to a smaller one. Thus, an edge from node i to node j (i.e., $A_{ji} > 0$) contributes to $\mathrm{DV}(\mathbf{x})$ only if $x_i > x_j$. Moreover, notice that if $\mathcal{G}$ is undirected, then $\mathrm{DV}(\mathbf{x}) \equiv \mathrm{TV}_2(\mathbf{x})$. Analogous to the GFTs surveyed in the "Digraph Fourier Transforms: Spectral Methods" section, we define the frequency $f_k := \mathrm{DV}(\mathbf{u}_k)$ as the DV of the frequency component $\mathbf{u}_k$. Because for all previous GFT approaches the spacing between frequencies can be highly irregular, the idea in [14] to better capture low, bandpass, and high frequencies is to design a DGFT such that the orthonormal frequency components are as spread as possible in the graph-spectral domain. Beyond offering parsimonious representations of slowly varying signals on digraphs, a DGFT with spread frequency components can facilitate more interpretable frequency analyses and aid filter design in the spectral domain. To this end, a viable approach is to minimize a so-termed spectral-dispersion criterion

$$\mathbf{U}^* = \operatorname*{argmin}_{\mathbf{U}} \sum_{i=1}^{N-1} [\mathrm{DV}(\mathbf{u}_{i+1}) - \mathrm{DV}(\mathbf{u}_i)]^2$$
$$\text{s.t. } \mathbf{U}^T\mathbf{U} = \mathbf{I}_N,\ \mathbf{u}_1 = \frac{\mathbf{1}_N}{\sqrt{N}},\ \mathbf{u}_N = \operatorname*{argmax}_{\|\mathbf{u}\|=1} \mathrm{DV}(\mathbf{u}). \quad (2)$$

The cost function measures how well spread the corresponding frequencies are over $[0, \mathrm{DV}(\mathbf{u}_N)]$. Having fixed the first and last columns of $\mathbf{U}$, the dispersion function is minimized when the free DV values are selected to form an arithmetic sequence over the attainable bandwidth. However, as the variables here are the columns of $\mathbf{U}$, we can expect to obtain only approximately equidistributed frequencies. Finding the global optimum of (2) is challenging due to the nonconvexity arising from the orthonormality (Stiefel manifold) constraints, yet a stationary point can be provably obtained via the algorithm in [14]. Accordingly, the basis $\mathbf{U}^*$ in (2) and its counterpart in [13] may not be unique. In the "Applications" section, we illustrate a graph signal-denoising task whereby the DGFT basis learned from (2) is used to

decompose and then (low- pass) filter temperatures recorded across the United States.

Graph filters and nonlinear graph signal operators

Here we consider operators whose inputs and outputs are signals defined on a digraph [see Figure 1(a) for a pictorial representation]. These operators are not only used to process information defined on digraphs (see also the "Applications" section), but also to postulate generative signal models for network data and solve statistical inference tasks surveyed in the "Inverse Problems on Digraphs," "Statistical Digraph SP," and "Digraph Topology Inference" sections. A key aspect throughout the discussion is how the topology of the digraph impacts the transformation of signals. The section begins by discussing linear graph filters [2, Ch. 11] and then builds on those to describe nonlinear (deep) architectures. After a brief outline of the current filtering landscape for undirected graphs, we focus on recent progress to tackle the challenges faced when extending those operators to the directed case.

Linear graph filters

Several definitions for graph filters coexist in the GSP literature. Early works focused on using the graph Laplacian $\mathbf{L}$ as the GSO and leveraged its eigendecomposition $\mathbf{L} = \mathbf{V}\text{diag}(\boldsymbol{\lambda})\mathbf{V}^T$ (see "A Motivating Starting Point: The Graph Fourier Transform for Undirected Graphs") to define the graph-filtering operation in the spectral domain [3]. Specifically, if $\mathbf{x}$ denotes the input of the graph filter and $\mathbf{y}$ is its output, filtering a graph signal is tantamount to transforming the input signal to the graph Fourier domain as $\tilde{\mathbf{x}} = \mathbf{V}^T\mathbf{x}$, applying a pointwise (diagonal) operator in the spectral domain to generate the output $\tilde{\mathbf{y}}$, and finally, transforming the obtained output back onto the vertex domain as $\mathbf{y} = \mathbf{V}\tilde{\mathbf{y}}$. The pointwise spectral operator can be expressed as the multiplication by a diagonal matrix $\text{diag}(\tilde{\mathbf{g}})$ so that $\tilde{\mathbf{y}} = \text{diag}(\tilde{\mathbf{g}})\tilde{\mathbf{x}}$. Alternatively, one can adopt a scalar kernel function $g : \mathbb{R} \to \mathbb{R}$ applied to the eigenvalues of the Laplacian so that the frequency response of the filter can be obtained as $\text{diag}(\tilde{\mathbf{g}}) = \text{diag}(g(\boldsymbol{\lambda}))$, where $g(\cdot)$ is applied entrywise. Regardless of the particular choice, the input–output relation can be written as

$$\mathbf{y} = \mathbf{V}\text{diag}(g(\boldsymbol{\lambda}))\mathbf{V}^T\mathbf{x} = \mathbf{V}\text{diag}(\tilde{\mathbf{g}})\mathbf{V}^T\mathbf{x}, \tag{3}$$

with the $N \times N$ matrix $\mathbf{V}\text{diag}(g(\boldsymbol{\lambda}))\mathbf{V}^T$ representing the linear transformation in the nodal domain. An alternative definition consists of leveraging the interpretation of $\mathbf{S}$ as a reference graph signal operator and then building more general linear operators of the form [7]

$$\begin{aligned}\mathbf{y} &= h_0\mathbf{x} + h_1\mathbf{S}\mathbf{x} + \cdots + h_{L-1}\mathbf{S}^{L-1}\mathbf{x} \\ &:= \mathbf{H}\mathbf{x}, \text{ with } \mathbf{H} := \sum_{l=0}^{L-1} h_l\mathbf{S}^l,\end{aligned} \tag{4}$$

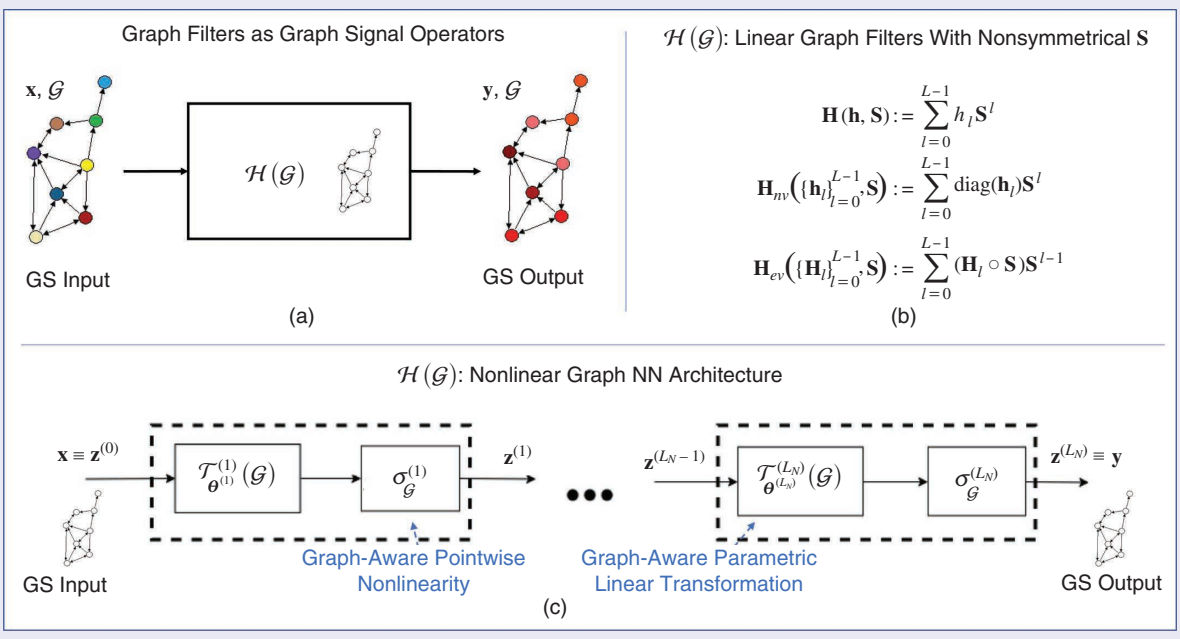

FIGURE 1. (a) Graph filters as generic operators that transform a graph signal (GS) input into an output. The graph filter processes the features of the input, taking into account the topology of the digraph where the signals are defined. (b) The different types of linear graph filters: a regular (shift-invariant) graph filter $\mathbf{H}$, a node-variant graph filter $\mathbf{H}_{nv}$, and an edge-variant graph filter $\mathbf{H}_{ev}$. The number of parameters (coefficients) is L, NL, and $|\mathcal{E}|L$, respectively. Due to their polynomial definition, all of these filters can operate over digraphs (nonsymmetric $\mathbf{S}$). (c) Nonlinear graph signal operators using a (potentially deep) NN with L_N layers. Each layer consists of a parametrized graph-aware linear transformation (given, e.g., by any of the linear graph filters described previously) followed by a pointwise nonlinearity [cf. (6)–(8)].

where the filter coefficients are collected in $\mathbf{h} := [h_0, \ldots, h_{L-1}]^T$, with $L-1$ denoting the filter degree. Upon defining $\mathbf{x}^{(l+1)} := \mathbf{S}\mathbf{x}^{(l)}$ and $\mathbf{x}^{(0)} = \mathbf{x}$, the output $\mathbf{y}$ in (4) can be equivalently written as $\mathbf{y} = \sum_{l=0}^{L-1} h_l \mathbf{x}^{(l)}$. As the application of the GSO $\mathbf{S}$ requires only local exchanges among (one-hop) neighbors, the latter expression reveals that the operators in (4) can be implemented in a distributed fashion with $L-1$ successive exchanges of information among neighbors [15]. This is a key insight (and advantage) of (4) that will be leveraged in subsequent sections. Note that the coefficients $\mathbf{h}$ can be given (e.g., when modeling known network-diffusion dynamics) or designed to accomplish a particular SP task, such as low-pass filtering (see, e.g., [15] for further details on graph-filter designs).

When $L = N$ and $\mathcal{G}$ is symmetrical so that the GSO is guaranteed to be diagonalizable, the two previous definitions can be rendered equivalent. But this is not the case when the graph filter is defined on a digraph. To see why, note that the polynomial definition in (4) is valid regardless of whether the GSO is symmetrical or not. Its interpretation as a local operator also holds true for digraphs provided that the notion of locality is understood, in this case, considering only the neighbors with incoming connections. The generalization of the definition in (3) to the directed case is, however, more intricate.

Because both $\mathbf{H}_{nv}$ and $\mathbf{H}_{ev}$ build on a polynomial definition, they can seamlessly operate over digraphs.

As explained in the "GSP Preliminaries, Frequency Analysis, and Signal Representations" section, different GFTs for digraphs exist. If the iGFT is given by the eigenvectors of the GSO, then one need only replace $\mathbf{V}^T$ with the (nonorthogonal) $\mathbf{V}^{-1}$ in (3). If the GSO is diagonalizable and $\mathbf{V}^{-1}$ is adopted as the GFT, then the polynomial definition in (4) and the updated version of (3) are equivalent. If the GSO is not diagonalizable, the generalization of (3) is unclear, while (4) still holds. On the other hand, if the GFT is not chosen to be $\mathbf{V}^{-1}$ but is one of the orthogonal (graph-smoothness-related) dictionaries $\mathbf{U}$ presented in the "Digraph Fourier Transforms: Orthonormal Transform Learning" section, then the two definitions diverge. Specifically, linear operators of the form $\mathbf{U}\text{diag}(\tilde{\mathbf{g}})\mathbf{U}^T$ will be symmetrical (meaning that the influence of the input at node i on the output at node j will be the same as that of node j on node i), while operators of the form $\sum_{l=0}^{L-1} h_l \mathbf{S}^l$ will not. Equally important is that, although a polynomial filter can always be implemented using local exchanges, there is no guarantee that the symmetrical transformation $\mathbf{U}\text{diag}(\tilde{\mathbf{g}})\mathbf{U}^T$ can be implemented in a distributed fashion [15]. All in all, if the definition in (4) is adopted for the directed case, then graph filters are always well defined, their distributed implementation is still feasible, and the design and interpretation of the filter coefficients $\mathbf{h}$ as weights given to the information obtained after successive local exchanges is preserved. Their interpretation as diagonal spectral operators only holds, however, if $\mathbf{V}^{-1}$ is used as a GFT and the GSO at hand is diagonalizable.

Generalizations of graph filters were introduced within the class of linear graph-aware signal operators. These include node- [15] and edge-variant graph filters [16], whose respective expressions are given by

$$\mathbf{H}_{nv} := \sum_{l=0}^{L-1} \text{diag}(\mathbf{h}_l)\mathbf{S}^l \quad \text{and} \quad \mathbf{H}_{ev} := \sum_{l=0}^{L-1} (\mathbf{H}_l \circ \mathbf{S})\mathbf{S}^{l-1}, \tag{5}$$

where $\circ$ denotes the Hadamard product, $\mathbf{h}_l$ is a vector of dimension N, and $\mathbf{H}_l$ is a sparse matrix with the same support as $\mathbf{S}$. Compared with its (node-invariant) counterpart in (4), we observe that the output generated by a node-variant filter can also be viewed as a linear combination of locally shifted inputs $\mathbf{x}^{(l)} = \mathbf{S}^l\mathbf{x}$, but in this case, each node has the flexibility of using a different set of weights. The flexibility is even larger for edge-variant graph filters because nodes can change the weight they give to each of their neighbors (compare all of their neighbors for node-variant filters). Because both $\mathbf{H}_{nv}$ and $\mathbf{H}_{ev}$ build on a polynomial definition, they can seamlessly operate over digraphs. They thus inherit most of the properties described for the original polynomial graph filters in (4).

Graph NN architectures

Graph filters have also been used to define nonlinear operators that account for the topology of the graph, such as median filters [17] and Volterra graph filters. All of these works build their definitions from the polynomial expression in (4) and hence can handle digraphs, although some of their properties (e.g., the conditions that a signal needs to satisfy to be a root of a median graph filter [17]) require minor modifications. A case of particular interest is that of deep graph NN (GNN) architectures [18], which have attracted significant attention in recent years for tackling machine learning problems involving network data. Traditional (e.g., convolutional) NNs have been remarkably successful in tasks involving images, video, and speech, all of which represent data with an underlying Euclidean domain that is regularly sampled over a grid-like structure. However, said structure one almost takes for granted is missing when it comes to signals defined on graphs. As argued next, GSP offers an ideal framework to fill in this fundamental gap.

The overall idea in GNN architectures is to define an input–output relation by using a concatenation of L_N layers composed of a linear transformation that combines the different signal values and a scalar (pointwise) nonlinear function that increases the expressiveness of the mapping. Mathematically, with $\mathbf{x}$ and $\mathbf{y}$ denoting the input and the output to the overall NN architecture and ℓ being the layer index, we have that

$$\mathbf{z}^{(0)} = \mathbf{x}, \ \text{and} \ \mathbf{y} = \mathbf{z}^{(L_N)}, \ \text{where} \tag{6}$$

$$\hat{\mathbf{z}}^{(\ell)} = \mathcal{T}^{(\ell)}_{\boldsymbol{\theta}^{(\ell)}}\{\mathbf{z}^{(\ell-1)} | \mathcal{G}\}, \ 1 \le \ell \le L_N, \tag{7}$$

$$\mathbf{z}^{(\ell+1)}_{ij} = \sigma^{(\ell)}_{\mathcal{G}}([\hat{\mathbf{z}}^{(\ell)}]_{ij}), \ 1 \le \ell \le L_N, \text{and all } i,j \in \mathcal{N}. \tag{8}$$

In (6)–(8), $\mathbf{z}^{(\ell)}$ is the output of layer ℓ and serves as input to layer $\ell+1$. The transformation $\mathcal{T}^{(\ell)}_{\boldsymbol{\theta}^{(\ell)}}\{\cdot | \mathcal{G}\}$ is the linear operator

implemented at layer ℓ, $\boldsymbol{\theta}^{(\ell)}$ are the learnable parameters that define such a transformation, and $\sigma_{\mathcal{G}}^{(\ell)}: \mathbb{R} \to \mathbb{R}$ is a scalar nonlinear operator (possibly different per layer). When applied to graph signals, the NN architecture in (6)–(8) must account for the topology of the graph and, for that reason, the dependence of both the linear and nonlinear operators on $\mathcal{G}$ was made explicit. In most works, the role of the graph is considered when defining the linear operator in $\mathcal{T}_{\boldsymbol{\theta}^{(\ell)}}^{(\ell)}\{\cdot|\mathcal{G}\}$ with the most widely used approach for (convolutional) GNNs being to replace $\mathcal{T}_{\boldsymbol{\theta}^{(\ell)}}^{(\ell)}\{\cdot|\mathcal{G}\}$ with a graph filter. Precisely in inspiring this approach is where GSP insights and advances have been transformative because basic shift-invariance properties and convolution operations are otherwise not well defined for graph signals [19].

Early contributions following the graph-filtering rationale have emerged from the machine learning community. The spectral approach in [19] relies on the Laplacian eigenvectors $\mathbf{V}$ and parametrizes the transformation $\mathcal{T}_{\boldsymbol{\theta}^{(\ell)}}^{(\ell)}\{\cdot|\mathcal{G}\} = \mathbf{V}\text{diag}(\boldsymbol{\theta}^{(\ell)})\mathbf{V}^T$ via $\tilde{\mathbf{g}} = \boldsymbol{\theta}^{(\ell)}$, the filter's frequency response in (3) that is learned using backpropagation. Although successful in many applications, when dealing with digraphs these approaches suffer from the same limitations as those discussed for their linear counterparts. Moreover, scalability is often an issue due to the computational burden associated with calculating the eigenvectors of large (albeit sparse) graphs. Alternative architectures proposed replacing $\mathcal{T}_{\boldsymbol{\theta}^{(\ell)}}^{(\ell)}\{\mathbf{z}^{(\ell-1)}|\mathcal{G}\}$ with $(\mathbf{I} - \theta^{(\ell)}\mathbf{A})\mathbf{z}^{(\ell-1)}$, where $\mathbf{A}$ is the (possibly nonsymmetric) adjacency matrix of the graph, and $\theta^{(\ell)}$ is a learnable scalar. To increase the number of parameters, some authors have considered learning the nonzero entries of $\mathbf{A}$, assuming that its support is known. A more natural approach is to replace $\mathcal{T}_{\boldsymbol{\theta}^{(\ell)}}^{(\ell)}\{\cdot|\mathcal{G}\}$ with the polynomial filter in (4) and consider the filter taps $\mathbf{h} = \boldsymbol{\theta}^{(\ell)}$ as the parameters to be learned (see [18] and the references therein). Once again, implementing (6)–(8) with $\mathbf{H}^{(\ell)} = \sum_{l=0}^{L_\ell - 1} h_l^{(\ell)} \mathbf{S}^l$ in lieu of $\mathcal{T}_{\boldsymbol{\theta}^{(\ell)}}^{(\ell)}\{\cdot|\mathcal{G}\}$ exhibits a number of advantages because 1) the graph filter is always well defined (even for nondiagonalizable GSOs), 2) the degree of the filter controls the complexity of the architecture (the number of learnable parameters), and 3) the polynomial definition guarantees that the resultant graph filter can be implemented efficiently (via the successive application of sparse matrices), which is essential in scaling to large data sets. As in standard NN architectures, GNN parameters (i.e., the filter coefficients for each of the layers) are learned using stochastic gradient descent. For supervised learning tasks, the goal is to minimize a suitable loss function over a training set of (labeled) examples. The sparsity of $\mathbf{S}$ and the efficient implementation of polynomial graph filters [compare (3)] are cardinal properties to keep the overall computational complexity in check.

As in standard NN architectures, GNN parameters (i.e., the filter coefficients for each of the layers) are learned using stochastic gradient descent.

Beyond convolutional GNNs, the aforementioned findings are also valid for recurrent GNNs. Furthermore, one can also replace the graph filter $\mathbf{H}^{(\ell)} = \sum_{l=0}^{L_\ell - 1} h_l^{(\ell)} \mathbf{S}^l$ either with a set of parallel filters, or with its node-variant $\mathbf{H}_{nv}^{(\ell)}$ or edge-variant $\mathbf{H}_{ev}^{(\ell)}$ counterparts. All of them preserve the distributed implementation of (4) while increasing the number of learnable parameters. As a result, the use of polynomial-based graph filter definitions that operate directly in the nodal domain to design NN architectures for digraphs opens a number of research avenues for deep learning over digraphs (see, e.g., [18]–[20] as well as other relevant articles in this special issue for additional details).

Inverse problems on digraphs

Inverse problems such as sampling and deconvolution have played a central role in the development of GSP. Different modeling assumptions must be considered when addressing these problems for digraphs; a good practice is to leverage the concepts introduced in the "GSP Preliminaries, Frequency Analysis, and Signal Representations" and "Graph Filters and Nonlinear Graph Signal Operators" sections and balancing practical utility with mathematical tractability. For instance, parsimonious signal models based on graph smoothness or bandlimitedness are widely adopted. Alternatively, observations can be modeled as the outputs of graph filters driven by white, sparse, or piecewise constant inputs. This approach is particularly useful in applications dealing with diffusion processes defined over real-world networks with directional links. In this section, we formally introduce a selection of prominent inverse problems, present established approaches for their solution, and identify the main challenges when the signals at hand are defined over digraphs.

Sampling and reconstruction

The sampling of graph signals and their subsequent reconstruction have arguably been the most widely studied problems within GSP [2, Ch. 9]. Broadly speaking, the objective is to infer the value of the signal at every node from the observations at a few nodes by leveraging the structure of the graph. To describe the problem formally, let us introduce the fat, binary, $M \times N$ sampling matrix $\mathbf{C}_{\mathcal{M}}$ and define the sampled signal as $\tilde{\mathbf{x}} = \mathbf{C}_{\mathcal{M}}\mathbf{x}$. Notice that if $\mathcal{M}$ represents the subset of $M < N$ nodes where the signal is sampled, $\mathbf{C}_{\mathcal{M}}$ has exactly one nonzero element per row, and the position of those nonzero elements correspond to the indexes of the nodes in $\mathcal{M}$. Then, the signal $\tilde{\mathbf{x}}$ is indeed a selection of M out of the N elements of $\mathbf{x}$. This raises two fundamental questions: How can we reconstruct $\mathbf{x}$ from $\tilde{\mathbf{x}}$, and how can we design $\mathbf{C}_{\mathcal{M}}$ to facilitate this reconstruction?

Starting with the first question, early works assumed the graph to be undirected and the signal $\mathbf{x}$ to be bandlimited, i.e., to be a linear combination of just a few leading eigenvectors of the GSO. The GSO was typically set to the Laplacian $\mathbf{L}$, with its eigenvectors $\mathbf{V} = [\mathbf{v}_1, ..., \mathbf{v}_N]$ being real valued and orthogonal. That is, the signal was assumed to be expressible as $\mathbf{x} = \sum_{k=1}^{K} \tilde{x}_k \mathbf{v}_k := \mathbf{V}_K \tilde{\mathbf{x}}_K$, where $\tilde{\mathbf{x}}_K \in \mathbb{R}^K$ collects the K active frequency coefficients, and $\mathbf{V}_K$ is a submatrix of the GFT. Indeed, because the leading

eigenvectors in $\mathbf{V}$ are those with the smallest total variation [cf. (S1)], this model was originally motivated by the practical importance of signals that vary smoothly with the underlying graph. Under the band-limited assumption, the sampled signal $\bar{\mathbf{x}}$ is given by $\bar{\mathbf{x}} = \mathbf{C}_{\mathcal{M}}\mathbf{x} = \mathbf{C}_{\mathcal{M}}\mathbf{V}_K\tilde{\mathbf{x}}_K$. Clearly, if the linear transformation represented by matrix $\mathbf{C}_{\mathcal{M}}\mathbf{V}_K \in \mathbb{R}^{M\times K}$ has full column rank (that is, if the rank of $\mathbf{C}_{\mathcal{M}}\mathbf{V}_K$ is equal to K), then $\tilde{\mathbf{x}}_K$ can be recovered from $\bar{\mathbf{x}}$. Once the coefficients $\tilde{\mathbf{x}}_K$ are known, the signal in the original domain can be found as $\mathbf{x} = \mathbf{V}_K\tilde{\mathbf{x}}_K = \mathbf{V}_K(\mathbf{C}_{\mathcal{M}}\mathbf{V}_K)^{\dagger}\bar{\mathbf{x}}$. Hence, the critical factor used to characterize the recovery of $\mathbf{x}$ from $\bar{\mathbf{x}}$ is the invertibility (and conditioning) of matrix $\mathbf{C}_{\mathcal{M}}\mathbf{V}_K$, which is a submatrix of $\mathbf{V}$ formed by the K columns corresponding to the active frequencies and the M rows corresponding to the sampled nodes in $\mathcal{M}$. Note that a key difference with sampling in classical SP is that, designing matrix $\mathbf{C}_{\mathcal{M}}$ as a regular sampler is meaningless in GSP because the node indexing is completely arbitrary.

When dealing with the sampling of real-world signals defined over digraphs, a first step is to decide which type of signal dictionary is going to be used.

Indeed, multiple approaches have been proposed to identify the most informative nodes on a graph for subsequent reconstruction. This is tantamount to leveraging the (spectral) properties of $\mathbf{C}_{\mathcal{M}}\mathbf{V}_K$ to design sampling matrices $\mathbf{C}_{\mathcal{M}}$ that lead to an optimal reconstruction. For example, by maximizing the minimum singular value of $\mathbf{C}_{\mathcal{M}}\mathbf{V}_K$, the sampling set is designed to minimize the effect of noise in a mean-squared-error sense [21]. The fact that the reconstruction matrix is a submatrix of the eigenvectors of the graph has also been exploited to design optimal low-pass graph-filtering operators that can reconstruct the original signal $\mathbf{x}$ by implementing local exchanges [22] as well as efficient algorithms that leverage the sparsity of the graph to compute $\mathbf{V}_K$ efficiently [23].

When dealing with the sampling and reconstruction of signals defined on digraphs, a number of challenges arise. As introduced in the "Digraph Fourier Transforms: Spectral Methods" section, multiple definitions of GFT coexist for digraphs. Some of those are based on generalizations of smoothness and lead to real-valued orthogonal dictionaries. In those cases, the results presented for signals in undirected graphs still hold, but the connections with polynomial low-pass filtering and the ability to find the eigenvectors efficiently are lost. Alternatively, one can use (a subset of) the eigenvectors of the nonsymmetric GSO as the basis for the signal $\mathbf{x}$. The caveats being, in this case, that the GSO needs to be diagonalizable and that the resulting eigenvectors $\mathbf{V}$ are neither orthogonal nor real valued. The latter point implies that the frequency coefficients $\tilde{\mathbf{x}}_K$ are complex valued as well so that the recovery methods for digraphs must be conceived in the complex field. Regarding the loss of orthogonality, this will typically deteriorate the conditioning of the submatrix $\mathbf{C}_{\mathcal{M}}\mathbf{V}_K$, which is critical in regimes where noise is present and M is close to K. Hence, when dealing with the sampling of real-world signals defined over digraphs, a first step is to decide which type of signal dictionary is going to be used. This likely depends on the prior domain knowledge as well as on the properties of the signals at hand. If no prior knowledge exists, schemes considering different dictionaries (at the expense of increasing the sample complexity) may be prudent. Moreover, in the cases where the selected basis is composed of the eigenvectors of the GSO, the recovery problems need to be formulated in the complex domain, and oversampling is likely to be required in scenarios where noise, outliers, or model mismatches are present.

Additional models for the observed signal have been studied in the digraph literature, including the cases where 1) the K dictionary atoms spanning $\mathbf{x}$ are not known a priori (thus leading to a sparse regression problem) [21], 2) the observations do not correspond to values of $\mathbf{x}$ but rather of $\mathbf{S}^i\mathbf{x}$ for varying i (which can be interpreted as sampling an evolving network process as opposed to a static one) [24], 3) total-variation metrics are considered in the form of regularizers or constraints [25], and 4) the signal $\mathbf{x}$ is modeled as the output of a graph filter excited by a structured input [5], [26]. We revisit the two last cases while studying the next collection of inverse problems.

(Blind) deconvolution, system identification, and source localization

We now introduce a family of recovery and reconstruction problems involving signals over digraphs. The common denominator across all of them is the assumption that the generative model $\mathbf{y} = \mathbf{H}\mathbf{x}$ holds, where $\mathbf{y}$ is a (partially) observed graph signal, $\mathbf{H}$ is a linear graph filter, and $\mathbf{x}$ is a potentially unknown and structured input. Building on this model and assuming that we have access to samples of the output $\mathbf{y}$, the supporting digraph, and side information on $\mathbf{H}$ and $\mathbf{x}$, the goal is to recover 1) the graph filter $\mathbf{H}$ (system identification), 2) the values of $\mathbf{x}$ (deconvolution), 3) the support of $\mathbf{x}$ (source localization), and 4) both the graph filter and the values of $\mathbf{x}$ (blind deconvolution).

Because graph filters can be efficiently used to model local diffusion dynamics, the relevance of the aforementioned schemes goes beyond signal reconstruction and permeates to broader domains, such as opinion formation and source identification in social networks, inverse problems of biological signals supported on graphs, and the modeling and estimation of diffusion processes in multiagent networks, all of which are typically directed. In particular, we envision applications in marketing where, e.g., social media advertisers want to identify a small set of influencers so that an online campaign can go viral; in a health-care policy implementing network analytics to infer hidden needle-sharing networks of injecting drug users; or, in environmental monitoring using wireless sensor networks to localize heat or seismic sources. As an encompassing formal framework, consider the following optimization problem:

$$\{\mathbf{x}^*, \mathbf{h}^*, \mathbf{y}^*\} = \underset{\{\mathbf{x},\mathbf{h},\mathbf{y}\}}{\operatorname{argmin}} \mathcal{L}_0\left(\mathbf{y}, \sum_{l=0}^{L-1} h_l \mathbf{S}^l \mathbf{x}\right) + \alpha_x r_x(\mathbf{x}) + \alpha_h r_h(\mathbf{h}) + \alpha_y r_y(\mathbf{y}), \quad \text{s.t.} \quad \mathbf{x} \in \mathcal{X}, \mathbf{h} \in \mathcal{H}, \mathbf{y} \in \mathcal{Y}, \tag{9}$$

where $\mathcal{L}_0$ is a loss function between the observed signal $\mathbf{y}$ and its prediction generated by the chosen $\mathbf{x}$ and $\mathbf{h}$. The regularizers r_x, r_h, and r_y promote desirable features on the optimization variables, and $\mathcal{X}$, $\mathcal{H}$, and $\mathcal{Y}$ represent prespecified feasibility sets. Although for the undirected case the generative filter $\mathbf{H}$ can be either defined in the spectral or in the vertex domain, in (9) the polynomial form has been selected. As pointed out in the "Graph Filters and Nonlinear Graph Signal Operators" section, the reasons for this choice are multiple: polynomial filters are always well defined (even for nondiagonalizable GSOs); the number of parameters is L (in contrast with N for those spectral formulations that do not consider an explicit parametrization), which is beneficial in the context of inverse problems; and the filter can be used to capture distributed diffusion dynamics on directed networks, strengthening the practical value of the formulation in (9). Finally, even though (9) was posited for the generic case where both $\mathbf{x}$ and $\mathbf{h}$ are unknown and $\mathbf{y}$ might be only partially observed, it can readily incorporate perfect knowledge of any of these variables just by fixing its value and dropping the corresponding feasibility constraint and regularization term.

Focusing first on the problem of deconvolution, notice that the nonsymmetric filter $\mathbf{H}$ is completely known because both the GSO $\mathbf{S}$ and the filter coefficients $\mathbf{h}$ are assumed to be given. The goal then is to use incomplete observations of $\mathbf{y}$ to recover the values of $\mathbf{y}$ in the nonobserved nodes and to obtain the seeding values in $\mathbf{x}$. Leveraging the notation introduced in the "Sampling and Reconstruction" section, we denote by $\bar{\mathbf{y}} = \mathbf{C}_{\mathcal{M}}\mathbf{y} = \mathbf{H}_{\mathcal{M}}\mathbf{x}$ the sampled output, with $\mathbf{H}_{\mathcal{M}} = \mathbf{C}_{\mathcal{M}}\mathbf{H}$ being the corresponding M rows of $\mathbf{H}$. As $\mathbf{y}$ and $\mathbf{x}$ are graph signals of the same size, the deconvolution problem is ill posed when $M < N$. Hence, to overcome this, we may assume some structural prior on the input $\mathbf{x}$. A common assumption is that $\mathbf{x}$ is sparse. This corresponds to setups where the observed signal $\mathbf{y}$ can be accurately modeled by a few sources percolating across the entire network. Applications fitting this setup range from social networks where a rumor originated by a small group of people is spread across the network via local opinion exchanges, to brain networks where an epileptic seizure emanating from few regions is later diffused across the entire brain [22]. Formally, (9) reduces to

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \left\| \bar{\mathbf{y}} - \mathbf{H}_{\mathcal{M}}\mathbf{x} \right\|_2^2 + \alpha_x \left\| \mathbf{x} \right\|_1, \tag{10}$$

which is a classical sparse-regression problem with well-established results showing that the recovery performance provably depends on the coherence of the nonsymmetric matrix $\mathbf{H}_{\mathcal{M}}$. The ℓ_1-norm regularizer in (10) acts as a convex surrogate of the sparsity-measuring ℓ_0 pseudonorm. Whenever sparsity is assumed as a structural property of the input and the emphasis is on recovering the support of $\mathbf{x}$, (10) and variations thereof (with imperfect knowledge of $\mathbf{h}$) are referred to as *source localization problems*. In terms of the samples of $\mathbf{y}$ that are observed, the optimal selection (in cases where this selection can be designed) is nontrivial and considerations similar to those discussed in the "Sampling and Reconstruction" section apply here as well. Finally, note that the generative model $\mathbf{y} = \mathbf{H}\mathbf{x}$ can also be used for undirected graphs and, as a result, the formulation in (10) and the associated algorithms can be used in such a case, the main difference being that the theoretical analysis of identifiability and recovery is simpler when the GSO (and hence the filter) is symmetrical.

Moving on to the system identification problem, where the main objective is to find the filter coefficients $\mathbf{h}$, it is crucial to note that $\mathbf{y}$ is a bilinear function in $\mathbf{h}$ and $\mathbf{x}$. Hence, if we assume that $\mathbf{x}$ is given, then the system identification problem is very similar to the deconvolution problem where the roles of $\mathbf{x}$ and $\mathbf{h}$ are interchanged. In terms of structural priors for an unknown $\mathbf{h}$, sparsity can also be employed. More specifically, it is instrumental to consider a weighted ℓ_1-norm regularization $r_h(\mathbf{h}) = \left\| \operatorname{diag}(\omega)\mathbf{h} \right\|_1$, where $\omega \in \mathbb{R}_+^L$ is a weighting vector whose weights increase with $l = 1, \ldots, L$, the entry index. In this way, the coefficients associated with higher powers of $\mathbf{S}$ in the filter specification are more heavily penalized, thus promoting a low-complexity and numerically stable model for explaining the observed data. For undirected graphs, the cost $\mathcal{L}_0$ that enforces the generative graph filter model to hold is often formulated in the spectral domain, bypassing the need of computing the powers of $\mathbf{S}$. Although we advocate working on the nodal domain, when the GSO is diagonalizable, formulating the problem in the spectral domain is also feasible for the directed case. The matrices mapping the unknown $\mathbf{h}$ to the observations $\bar{\mathbf{y}}$ would be complex valued, but the optimization would still be carried over the real-valued vector $\mathbf{h}$. From an algorithmic perspective, the main challenge would be to find the eigenvectors of the nonsymmetric $\mathbf{S}$, while from an analytical point of view, the issue would be the characterization of the conditioning of the (complex-valued) matrix that maps $\mathbf{h}$ to $\bar{\mathbf{y}}$.

The more challenging problem of blind deconvolution arises when both the input $\mathbf{x}$ and the filter coefficients $\mathbf{h}$ are unknown. To formally tackle this problem, we explicitly write the fact that $\mathbf{y}$ is a bilinear function of $\mathbf{h}$ and $\mathbf{x}$ as $\mathbf{y} = \mathcal{A}(\mathbf{x}\mathbf{h}^T)$, where the linear operator $\mathcal{A}$ is a function of the nonsymmetric $\mathbf{S}$ and acts on the outer product of the sought vectors. A direct implementation of the general framework (9) can be employed for the problem of blind identification, where the goodness-of-fit loss $\left\| \mathbf{y} - \mathcal{A}(\mathbf{x}\mathbf{h}^T) \right\|_2^2$ is combined with structure-promoting regularizers for both $\mathbf{x}$ and $\mathbf{h}$. Notice, however, that this leads to a nonconvex optimization problem for which alternating minimization schemes (e.g., a block-coordinate descent method that alternates between $\mathbf{x}$ and $\mathbf{h}$) can be implemented. To derive a convex relaxation, notice that $\mathbf{y}$ is a linear function of the entries of the rank-one matrix $\mathbf{Z} = \mathbf{x}\mathbf{h}^T$. This motivates the statement of the following convex optimization problem:

Problem	Given	Find	Fig
System Identification	$\mathbf{x}$ $\bar{\mathbf{y}}$	$\mathbf{h}$	(a)
Deconvolution	$\mathbf{h}$ $\bar{\mathbf{y}}$	$\mathbf{x}$	(a)
Source Localization	$\mathbf{h}$ $\bar{\mathbf{y}}$	$\text{supp}(\mathbf{x})$	(a)
Blind Deconvolution	$\bar{\mathbf{y}}$	$\mathbf{x}$ $\mathbf{h}$	(a)
MIMO Blind Deconvolution	$\{\mathbf{y}_i\}_{i=1}^P$	$\{\mathbf{x}_i\}_{i=1}^P$ $\mathbf{h}$	(b) Left
SIMO Blind Deconvolution	$\{\mathbf{y}_i\}_{i=1}^P$	$\mathbf{x}$ $\{\mathbf{h}_i\}_{i=1}^P$	(b) Middle
Blind Demixing	$\mathbf{y}$	$\{\mathbf{x}_i\}_{i=1}^P$ $\{\mathbf{h}_i\}_{i=1}^P$	(b) Right

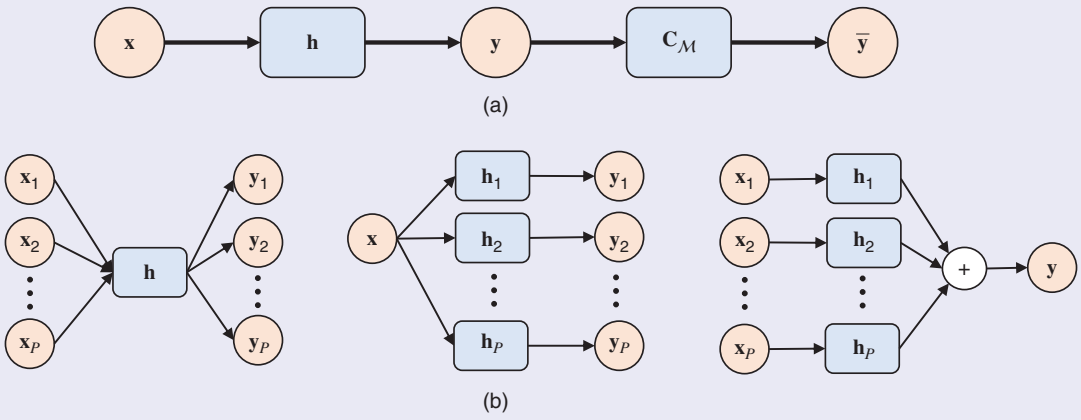

FIGURE 2. A summary of the inverse problems introduced. In the schematic representation, graph signals are depicted as red circles and graph operators as blue rectangles. Notice that filters are functions of the coefficients $\mathbf{h}$ and the GSO $\mathbf{S}$ [cf. (4)]. However, because $\mathbf{S}$ is assumed to be known for every problem considered, we succinctly represent filters by their coefficients, $\mathbf{h}$. The first four problems refer to the single-input, single-output scenario with blind deconvolution being the most challenging because only (a sampled version of) the output is observed. Note that the problem frameworks in (b) can be further extended to the case where the output is partially observed, as in (a). We omit this illustration to minimize redundancy and because these more challenging problems are generally ill posed even in the case where $\mathbf{y}$ is fully observed.

$$\mathbf{Z}^* = \underset{\mathbf{Z}}{\text{argmin}} \| \mathbf{y} - \mathcal{A}(\mathbf{Z}) \|_2^2 + \alpha_1 \| \mathbf{Z} \|_* + \alpha_2 \| \mathbf{Z} \|_{2,1}. \quad (11)$$

The nuclear-norm regularizer $\| \cdot \|_*$ in (11) promotes a low-rank solution because we know that $\mathbf{Z}$ should be the outer product of the true variables of interest, that is, $\mathbf{x}$ and $\mathbf{h}$. On the other hand, the $\ell_{2,1}$ mixed-norm $\| \mathbf{Z} \|_{2,1} = \Sigma_{i=1}^N \| \mathbf{z}_i \|_2$ is the sum of the ℓ_2 norms of the rows of $\mathbf{Z}$, thus promoting a row-sparse structure in $\mathbf{Z}$. This is aligned with a sparse input $\mathbf{x}$ forcing rows of $\mathbf{Z}$ to be entirely zero from the outer product. After solving for $\mathbf{Z}^*$, one may recover $\mathbf{x}$ and $\mathbf{h}$ from, e.g., a rank-one decomposition of $\mathbf{Z}^*$.

Extensions to multiple input, multiple-output (MIMO) pairs (with a common filter) along with theoretical guarantees for the case where the GSO $\mathbf{S}$ is normal (i.e., $\mathbf{S}\mathbf{S}^H = \mathbf{S}^H\mathbf{S}$) can be found in [5]. Interestingly, it was empirically observed and theoretically demonstrated that blind deconvolution in circulant graphs (such as the directed cycle that represents the domain of classic SP) corresponds to the most favorable setting. The related case of a single graph signal as the input to multiple filters (generating multiple outputs) was recently studied in [27], thus providing a generalization of the classical blind multichannel identification problem in digital SP. Moreover, [26] addresses the blind demixing case where a single observation formed by the sum of multiple outputs is available, and it is assumed that these outputs are generated by different sparse inputs diffused through different graph filters. This variation of the problem is severely ill posed, and strong regularization conditions should be assumed to ensure recovery, with the problem being easier if the graph filters are defined over different GSOs $\{\mathbf{S}_i\}_{i=1}^I$. Figure 2 provides an overarching view of the problems mentioned in this section. See also [25] for additional signal-recovery problems that can be written in the encompassing framework of (9).

As previously explained, a key feature that allows using the problem formulations introduced in this section for signals defined on digraphs is that the generative graph filter was incorporated in polynomial form. Unfortunately, many of the theoretical guarantees for solving these problems rely heavily on the spectral analysis of the GSO, thus assuming

symmetry or at least normality of $\mathbf{S}$. One of the main remaining challenges for inverse problems in digraphs is the derivation of recovery guarantees, along with the identification of key performance drivers that can accommodate nondiagonalizable GSOs and generalized (complex) eigenvectors. Another important potential research direction is the incorporation of alternative generative models by replacing graph filters with the more general graph-signal operators presented in the "Graph Filters and Nonlinear Graph Signal Operators" section, such as node- and edge-variant filters (5) or GNNs (6)–(8). Especially in this latter case, system identification and blind deconvolution would become extremely challenging due to the incorporation of nonlinearities, making the convex relaxation in (11) based on the linear operator $\mathcal{A}$ no longer valid.

Statistical digraph SP

Randomness is pervasive in engineering, and graph signals are no exception. For this reason, here we build upon the results presented in the previous sections to discuss recent advances and challenges to develop statistical models for random graph signals defined over digraphs. In the field of statistics, graphs quickly emerged as a convenient intuitive mathematical structure to describe complex statistical dependencies across multidimensional variables. A prominent example is that of Markov random fields (MRFs), which are symmetrical graphical models whose edges capture conditional dependencies across the variables represented by the nodes. Inference over MRFs is computationally affordable and, for the particular case of the signals being Gaussian, the graph describing the MRF can be inferred directly from the precision (inverse covariance) matrix of the data. In parallel, digraphs have been used to capture one-directional conditional dependence (hence causal) relations, with Bayesian networks—which, in addition to being directed, are acyclic—being the most tractable graphical model within this class.

GSP literature has also contributed to the statistical modeling of random graph signals. The first step is to postulate how the graph structure plays a role in shaping the signal's statistical properties and then analyze how the model put forth can be used to tackle inference tasks more effectively. As in the case of graphical models, most existing results focused on undirected graphs.

Arguably, the most relevant line of work has been the generalization of the definition of *weak stationarity* to signals supported on either undirected graphs or on graphs whose GSO is a normal matrix [2, Ch. 12]. Although this latter characterization includes some digraphs (such as circulant and skew-Hermitian), the definitions cannot be applied to a generic nonsymmetric GSO. The key contribution of the articles reviewed in [2, Ch. 12] was to provide a dual definition for stationary graph processes, which was consistent with the vertex and frequency interpretations of graph signals. Specifically, it was stated that a zero-mean random graph signal $\mathbf{x}$ was weakly stationary on a known graph $\mathcal{G}$ if 1) its covariance matrix has the same eigenvectors as those of the GSO; or, equivalently, 2) the process can be modeled as the output of a graph filter excited with a white input. This allowed for establishing parallelisms with the classical definition of weak stationarity for time-varying signals and opened the door to the development of efficient algorithms that estimate the second moment of a graph stationary process using fewer samples. For example, if the eigenvectors of the covariance and the GSO are the same, instead of estimating the N^2 entries of the covariance matrix, one can focus on estimating only its N eigenvalues.

However, from the initial discussion in the "Graph Filters and Nonlinear Graph Signal Operators" section, it follows that this convenient equivalence between the frequency and the vertex domains does not hold for digraphs. Indeed, if the GSO is not a normal matrix, its eigenvectors cannot coincide with those of the covariance matrix, which is guaranteed to be normal. As a result, one must adapt the definitions and sacrifice some of the properties shown for the symmetrical case. To be mathematically precise, let us recall that $\mathbf{x}$ is a zero-mean random process defined on the digraph $\mathcal{G}$ with GSO $\mathbf{S}$, and let us denote by $\mathbf{C}_\mathbf{x} := \mathbb{E}[\mathbf{x}\mathbf{x}^T]$ the $N \times N$ covariance matrix of $\mathbf{x}$. We say that the random graph signal $\mathbf{x}$ is stationary in the nonsymmetric $\mathbf{S}$ if it can be described as

$$\mathbf{x} = \mathbf{H}\mathbf{w}, \text{ with } \mathbf{H} := \sum_{l=0}^{L-1} h_l \mathbf{S}^l \text{ and } \mathbb{E}[\mathbf{w}\mathbf{w}^T] = \mathbf{I}, \tag{12}$$

where $L \leq N$, and $\mathbf{w}$ is a white zero-mean random signal. By adopting the generative model in (12), it follows that the covariance of $\mathbf{x}$ can be written as $\mathbf{C}_\mathbf{x} = \mathbb{E}[\mathbf{x}\mathbf{x}^T] = \mathbf{H}\mathbb{E}[\mathbf{w}\mathbf{w}^T]\mathbf{H}^T = \mathbf{H}\mathbf{H}^T$, which is not a polynomial on $\mathbf{S}$, but on both $\mathbf{S}$ and $\mathbf{S}^T$. As a result, it is no longer true that $\mathbf{C}_\mathbf{x}$ is diagonalized by the GFT associated with $\mathbf{S}$. Nonetheless, the model in (12) is still extremely useful because it 1) provides an intuitive explanation of the notion of graph stationarity; 2) can be used to establish connections with higher-order, autoregressive directed structural equation models (SEMs) in statistics: and 3) gives rise to efficient estimators that, rather than targeting the estimation of the full covariance, try to estimate the filter coefficients $\mathbf{h}$. Indeed, this latter point is also relevant for undirected graphs.

Although approaches that focus on the spectral definition of stationary processes require estimating the N eigenvalues of $\mathbf{C}_\mathbf{x}$ (i.e., the power spectral density of the process $\mathbf{x}$), the generative approaches based on (12) open the door to imposing additional structure on the generative filter. For example, one can consider a finite-impulse/infinite-impulse response filter with a number of coefficients L much smaller than the number of nodes N, which results in considerable gains in terms of either the sampling complexity or the estimation error.

The generative model in (12) can be generalized or constrained to fit a range of suitable scenarios. Focusing first on the input signal, cases of practical interest include 1) considering nonwhite input processes $\mathbf{w}$ with known covariance, 2) requiring $\mathbf{w}$ to not only be white but also independent, and 3)

particularizing the distribution of $\mathbf{w}$ to tractable and practically meaningful cases. Two examples that fall into the final category are modeling $\mathbf{w}$ as either a Gaussian or a (signed) Bernoulli vector, which is particularly relevant in the context of the diffusion of sparse signals. Alternatively, the model in (12) can be enlarged by considering other linear graph signal operators as generators, including the node- and edge-variant graph filters discussed in (5). Recent works have also proposed nonlinear generative models that exploit results in the deep learning literature to generate random signals over directed and undirected graphs. For example, one can take the architecture in (6)–(8), replace $\mathcal{T}_{\boldsymbol{\theta}^{(\ell)}}\{\cdot|\mathcal{G}\}$ with $\mathbf{H}^{(\ell)}=\Sigma_{l=0}^{L_\ell-1}h_l^{(\ell)}\mathbf{S}^l$, use a random realization of the white signal $\mathbf{w}$ as input, and then view the output of the GNN architecture as the random process to be modeled. Although characterizing how the coefficients $\{\mathbf{h}^{(\ell)}\}_{\ell=1}^{L_N}$ affect the statistical properties of the output is certainly relevant, equally interesting problems arise when the goal is to use a set of realizations of the output $\mathbf{x}$ to learn the parameters of the nonlinear generative model (i.e., the filter coefficients $\{\mathbf{h}^{(\ell)}\}_{\ell=1}^{L_N}$) that best fit the available observations.

Discovering directional influence among variables is at the heart of causal inference.

The statistical models briefly reviewed in this section accounted for nonsymmetric interactions among variables and can be leveraged, for example, to enhance covariance estimation schemes, to denoise a set of observed graph signals, or to interpolate (predict) all the values of a graph signal using as input observations collected at a subset of nodes. Perhaps less obvious but arguably equally important, the postulated models can also be used to infer the graph itself. Indeed, if one has access to a set $\mathcal{X}:=\{\mathbf{x}_r\}_{r=1}^R$ of R realizations of $\mathbf{x}$ and the graph is sufficiently sparse (so that the number of edges $|\mathcal{E}|$ is much smaller than N^2), one could identify the $L+|\mathcal{E}|$ degrees of freedom in (12) from the RN values in $\mathcal{X}$, provided that R is sufficiently large. This is partially the subject of the next section, which deals with the problem of inferring the topology of digraphs from a set of nodal observations.

Digraph topology inference

Capitalizing on the GSP advances surveyed thus far requires a specification of the underlying digraph. However, $\mathcal{G}$ is often unobservable and, accordingly, network topology inference from a set of (graph signal) measurements is a prominent yet challenging problem, even more so when the graph at hand is directed. Early foundational contributions can be traced back several decades to the statistical literature of graphical model selection (see, e.g., [4, Ch. 7] and the opening of the "Statistical Digraph SP" section). Discovering directional influence among variables is at the heart of causal inference, and identifying the cause-and-effect digraphs (so-termed structural causal models) from observational data are a notoriously difficult problem [6, Ch. 7 and 10].

Recently, the fresh modeling and signal representation perspectives offered by GSP have sparked renewed interest in the field. Initial efforts have focused mostly on learning undirected graphs, which naturally give rise to more tractable (and often uniquely identifiable) formulations [12]. Therefore, next, we outline a few noteworthy digraph topology-identification approaches that are relevant to (or are informed by) the GSP theme of this article. In accordance with our narrative's leitmotif, we emphasize the key differences with the undirected case and review the main challenges associated with the new formulations.

We initiate our exposition with structural equation modeling, which broadly encapsulates a family of statistical methods that describe causal relationships between interacting variables in a complex system. This is pursued through the estimation of linear relationships among endogenous as well as exogenous traits. SEMs have been extensively adopted in economics, psychometrics, social sciences, and genetics, among other disciplines (see, e.g., [28]). SEMs postulate a linear time-invariant network model of the following form, where the GSO is specified as the adjacency matrix $\mathbf{S}=\mathbf{A}$:

$$x_{it}=\sum_{j=1,j\neq i}^{N}S_{ij}x_{jt}+\omega_{ii}u_{it}+\epsilon_{it},\ i\in\mathcal{N}\Rightarrow\mathbf{x}_t=\mathbf{S}\mathbf{x}_t+\boldsymbol{\Omega}\mathbf{u}_t+\boldsymbol{\epsilon}_t, \tag{13}$$

where $\mathbf{x}_t=[x_{1t},\ldots,x_{Nt}]^T$ represents a graph signal of endogenous variables at discrete-time t and $\mathbf{u}_t=[u_{1t},\ldots,u_{Nt}]^T$ is a vector of exogenous influences. The term $\mathbf{S}\mathbf{x}_t$ in (13) models network effects, implying that x_{it} is a linear combination of the instantaneous values x_{jt} of node i's in-neighbors $j\in\mathcal{N}_i$. The signal x_{it} also depends on u_{it}, where weight ω_{ii} captures the level of influence of external sources, and we defined $\boldsymbol{\Omega}:=\text{diag}(\omega_{11},\ldots,\omega_{NN})$. Vector $\boldsymbol{\epsilon}_t$ represents measurement errors and unmodeled dynamics. Depending on the context, $\mathbf{x}_t$ can be thought of as an output signal, while $\mathbf{u}_t$ corresponds to the excitation or control input. In the absence of noise and letting $\boldsymbol{\Omega}=\mathbf{I}$ for simplicity, (13) becomes $\mathbf{x}_t=\mathbf{H}\mathbf{u}_t$, where $\mathbf{H}:=(\mathbf{I}-\mathbf{S})^{-1}$ is a polynomial graph filter, as in (4).

Given snapshot observations $\mathcal{X}:=\{\mathbf{x}_t,\mathbf{u}_t\}_{t=1}^T$, SEM parameters $\mathbf{S}$ and $\boldsymbol{\omega}:=[\omega_{11},\ldots,\omega_{NN}]^T$ are typically estimated via penalized least squares, for instance, by solving

$$\begin{aligned}&\hat{\mathbf{S}}=\underset{\mathbf{S},\boldsymbol{\omega}}{\text{argmin}}\ \sum_{t=1}^{T}\|\mathbf{x}_t-\mathbf{S}\mathbf{x}_t+\boldsymbol{\Omega}\mathbf{u}_t\|_2^2+\alpha\|\mathbf{S}\|_1,\\&\text{s.t. }\boldsymbol{\Omega}=\text{diag}(\boldsymbol{\omega}),\quad S_{ii}=0,\ i=1,\ldots,N,\end{aligned} \tag{14}$$

where the ℓ_1-norm penalty promotes sparsity in the adjacency matrix. Both edge sparsity and endogenous inputs play a critical role in guaranteeing that the SEM parameters (13) are uniquely identifiable (see also [28]). Acknowledging the limitations of linear models, [29] leverages kernels within the SEM framework to model nonlinear pairwise dependencies among network nodes (see the "Applications" section for results on the identification of gene-regulatory networks).

Although SEMs capture only contemporaneous relationships among the nodal variables (i.e., SEMs are memoryless),

sparse vector autoregressive models (SVARMs) account for linear time-lagged (causal) influences instead (see, e.g., [30]). Specifically, for a given model order L and unknown sparse-evolution matrices $\{\mathbf{S}^{(l)}\}_{l=1}^{L}$, SVARMs postulate a multivariate linear dynamical model of the form $\mathbf{x}_t = \Sigma_{l=1}^{L}\mathbf{S}^{(l)}\mathbf{x}_{t-l} + \boldsymbol{\epsilon}_t$. Here, a directed edge from vertex j to i is typically said to be present in $\mathcal{G}$ if $S_{ij}^{(l)} \neq 0$ for all $l = 1, \ldots, L$. The aforementioned AND rule is often explicitly imposed as a constraint during the estimation of SVARM parameters through the requirement that all matrices $\mathbf{S}^{(l)}$ have a common support. This can be achieved, for instance, via a group lasso penalty, which promotes sparsity over edgewise coefficients $\mathbf{s}_{ij} := [S_{ij}^{(1)}, \ldots, S_{ij}^{(L)}]^T$ jointly [30]. The sparsity assumption is often well justified due to physical considerations or for the sake of interpretability, but here (as well as with SEMs), it is also critical to reliably estimate $\mathcal{G}$ from limited and noisy time-series data $\mathcal{X} := \{\mathbf{x}_t\}_{t=1}^{T}$.

SVARMs are also central to popular digraph topology-identification approaches based on the principle of Granger causality (see, e.g., [6, Ch. 10]). Said principle is based on the concept of precedence and predictability, where node j's time series is said to "Granger-cause" the time series at node i if the knowledge of $\{x_{j,t-l}\}_{l=1}^{L}$ improves the prediction of x_{it} compared to using only $\{x_{i,t-l}\}_{l=1}^{L}$. Such a form of causal dependence defines the status of a candidate edge from j to i, and it can be assessed via judicious hypothesis testing [28]. Recently, a notion different from Granger's was advocated to associate a graph with causal network effects among vertex time series, effectively blending VARMs with graph filter-based dynamical models. The so-termed causal graph process (CGP) introduced in [31] also considers $\mathbf{S} = \mathbf{A}$ and has the form

$$\mathbf{x}_t = \sum_{l=1}^{L}\sum_{i=0}^{l} h_{li}\mathbf{S}^i\mathbf{x}_{t-l} + \boldsymbol{\epsilon}_t = (h_{10}\mathbf{I} + h_{11}\mathbf{S})\mathbf{x}_{t-1} + \cdots + (h_{L0}\mathbf{I} + \cdots + h_{LL}\mathbf{S}^L)\mathbf{x}_{t-L} + \boldsymbol{\epsilon}_t, \tag{15}$$

where $\mathbf{S}$ is the (possibly asymmetric) adjacency matrix encoding the unknown graph topology. The CGP model corresponds to a generalized VARM with coefficients given by $\mathbf{H}_l(\mathbf{S}, \bar{\mathbf{h}}) := \Sigma_{i=0}^{l} h_{li}\mathbf{S}^i$, where $\bar{\mathbf{h}} := [h_{10}, h_{11}, \ldots, h_{li}, \ldots h_{LL}]^T$. This way, the model can possibly account for multihop nodal influences per time step. Unlike SVARMs, matrices $\mathbf{H}_l(\mathbf{S}, \bar{\mathbf{h}})$ need not be sparse for larger values of l, even if $\mathbf{S}$ is itself sparse. Given data $\mathcal{X} := \{\mathbf{x}_t\}_{t=1}^{T}$ and a prescribed value of L, to estimate $\mathbf{S}$, one solves the nonconvex optimization problem

$$\hat{\mathbf{S}} = \underset{\mathbf{S},\bar{\mathbf{h}}}{\operatorname{argmin}} \sum_{t=L+1}^{T}\left\|\mathbf{x}_t - \sum_{l=1}^{L}\mathbf{H}_l(\mathbf{S}, \bar{\mathbf{h}})\mathbf{x}_{t-l}\right\|^2 + \alpha\|\mathbf{S}\|_1 + \beta\|\bar{\mathbf{h}}\|_1. \tag{16}$$

Similar to the sparse SEMs in (14) and SVARMs, the estimator encourages sparse graph topologies. Moreover, the ℓ_1-norm regularization on the filter coefficients $\bar{\mathbf{h}}$ effectively implements a form of model-order selection. A divide-and-conquer heuristic is advocated in [31] to tackle the challenging problem (16), whereby one 1) identifies the filters $\mathbf{H}_l := \mathbf{H}_l(\mathbf{S}, \bar{\mathbf{h}})$ so that $\mathbf{x}_t \approx \Sigma_{l=1}^{L}\Sigma_{i=0}^{l}\mathbf{H}_l\mathbf{x}_{t-l}$, exploiting that $\mathbf{H}_l$ and $\mathbf{H}_{l'}$ commute for all l, l'; 2) recovers a sparse $\mathbf{S}$ using the estimates $\{\hat{\mathbf{H}}_l\}$ and leverages the shift-invariant property of graph filters; and 3) estimates $\bar{\mathbf{h}}$ given $\{\hat{\mathbf{H}}_l, \hat{\mathbf{S}}\}$ via the lasso. For full algorithmic details and their accompanying convergence analysis, see [31].

In [32], observations from M network processes are modeled as the outputs of a polynomial graph filter [i.e, $\mathbf{x}_m = \mathbf{H}\mathbf{w}_m$, as in (4)] excited by (unobservable) zero-mean independent graph signals $\mathbf{w}_m$ with arbitrarily correlated nodal components. Observations of the output signals along with prior statistical information on the inputs are first utilized to identify the nonsymmetric diffusion filter $\mathbf{H}$. Such a problem entails solving a system of quadratic matrix equations, which can be recast as a smooth quadratic minimization subject to Stiefel manifold constraints (see [32] for details). Given an estimate $\hat{\mathbf{H}}$, the approach used in [32] to infer the digraph topology is to find a generic GSO $\mathbf{S}$ that satisfies certain desirable topological properties and commutes with $\mathbf{H}$. For instance, by focusing on the recovery of sparse graphs, one solves

$$\hat{\mathbf{S}} = \underset{\mathbf{S}}{\operatorname{argmin}}\ \|\mathbf{S}\|_1, \quad \text{s.t. } \mathbf{S} \in \mathcal{S}, \quad \|\hat{\mathbf{H}}\mathbf{S} - \mathbf{S}\hat{\mathbf{H}}\|_F \leq \epsilon, \tag{17}$$

where $\mathcal{S}$ is a convex set specifying the type of GSO sought (say, the adjacency matrix of a digraph), and the constraint $\|\hat{\mathbf{H}}\mathbf{S} - \mathbf{S}\hat{\mathbf{H}}\|_F \leq \epsilon$ encourages the filter $\mathbf{H}$ to be a polynomial in $\mathbf{S}$ while accounting for estimation errors (see [2], [12], and [31]). Imposing this last constraint offers an important departure from the related (undirected) graph-learning algorithms in [12], [33], and [2, Ch. 13], which identify the structure of network-diffusion processes from observations of stationary signals [cf. (12) but with symmetrical $\mathbf{H}$]. These approaches first estimate the eigenvectors of $\mathbf{H}$ and then constrain $\mathbf{S}$ to be diagonalized by those eigenvectors in a convex problem to recover the unknown eigenvalues. Although this naturally entails a search over the lower-dimensional space of GSO eigenvalues, the formulation (17) avoids computing an eigendecomposition and, more importantly, solving a problem over complex-valued variables. This was not an issue in [33] because the focus therein was on undirected graphs with real-valued spectrums. In closing, note that the graph-filtering model advocated in [32] is a special case of (15) provided that $\mathbf{S} = \mathbf{A}$ and, instead of multivariate time-series data, one relies on independent replicates from multiple network processes (obtained, e.g., via interventions, as in causal inference [6]).

Applications

We highlight four real-world applications of the methods surveyed in this article. The experiments were chosen to demonstrate the practical value of SP schemes applied to digraphs, with the diversity of the data sets considered (climate records, text excerpts, handwritten characters, and gene-expression levels) underscoring the versatility of the tools.

Frequency analysis for temperature-signal denoising

We consider a digraph $\mathcal{G}$ of the $N = 48$ contiguous United States (Alaska and Hawaii are excluded). A directed edge joins two states if they share a border, and the edge direction is set

so that the state whose barycenter is more to the south points to the one more to the north. As the graph signal $\mathbf{x} \in \mathbb{R}^{48}$, we consider the average annual temperature of each state [see Figure 3(a)]. The temperature map confirms that the latitude affects the average temperatures of the states, justifying the proposed latitude-based graph-construction scheme.

We determine a GFT basis $\mathbf{U}$ for this digraph via spectral-dispersion minimization as in (2) and test its utility in a denoising task. More specifically, our goal is to recover the temperature signal from noisy measurements $\mathbf{y} = \mathbf{x} + \mathbf{w}_y$, where the additive noise $\mathbf{w}_y$ is a zero-mean, Gaussian random vector with covariance matrix $10\mathbf{I}_N$. To achieve this, we implement a low-pass graph filter that retains the first K components of the signal's DGFT and eliminates the rest, i.e., $\tilde{\mathbf{h}} = [\tilde{h}_1, \ldots, \tilde{h}_N]^T$, where $\tilde{h}_k = \mathbb{I}\{k \leq K\}$, and K is a prescribed spectral window size. Hence, we estimate the true temperature signal as $\hat{\mathbf{x}} = \mathbf{U}\text{diag}(\tilde{\mathbf{h}})\tilde{\mathbf{y}} = \mathbf{U}\text{diag}(\tilde{\mathbf{h}})\mathbf{U}^T\mathbf{y}$.

The original signal $\mathbf{x}$ is bandlimited compared to the noisy signal $\mathbf{y}$, which spans a broader range of frequencies [see Figure 3(b)]. To better observe the low-pass property of $\mathbf{x}$, we also plot the cumulative energy of both $\mathbf{x}$ and $\mathbf{y}$, defined by the percentage of the total energy present in the first k frequency components for $k = 1, \ldots, N$. Setting the spectral window at $K = 3$, the average recovery error $e_f = \|\hat{\mathbf{x}} - \mathbf{x}\| / \|\mathbf{x}\|$ determined by 1,000 Monte Carlo simulations of independent noise was approximately 12%. Figure 3(c) shows a realization of the noisy graph signal $\mathbf{y}$ superimposed with the denoised temperature profile $\hat{\mathbf{x}}$, and it can be seen that, indeed, $\hat{\mathbf{x}}$ closely approximates $\mathbf{x}$. The recovery error increases when the edge directions are ignored (i.e., $\mathcal{G}$ is treated as undirected) and when they are selected randomly (i.e., every edge is directed but the specific orientation is chosen uniformly at random between the two possibilities), as opposed to following the south-to-north orientation that captures the temperate flow (see [14] for additional details and experiments).

GNNs for authorship attribution

We illustrate the performance of GNNs for classification in digraphs through an authorship attribution problem based on real data. The goal is, using a short text excerpt as input, to decide whether the text was written by a particular author. To capture the style of an author, we consider author-specific word adjacency networks (WANs), which are digraphs whose nodes are function words (i.e., prepositions, pronouns, conjunctions, and other words with syntactic importance but little semantic meaning [34]) and whose edges represent probabilities of directed coappearance of two function words within texts written by the author [see Figure 4(a)].

We select $N = 211$ functions words as nodes and build the WAN for Emily Brontë. More specifically, we count the number of times each pair of function words coappear in 10-word windows while also recording their relative order. We then

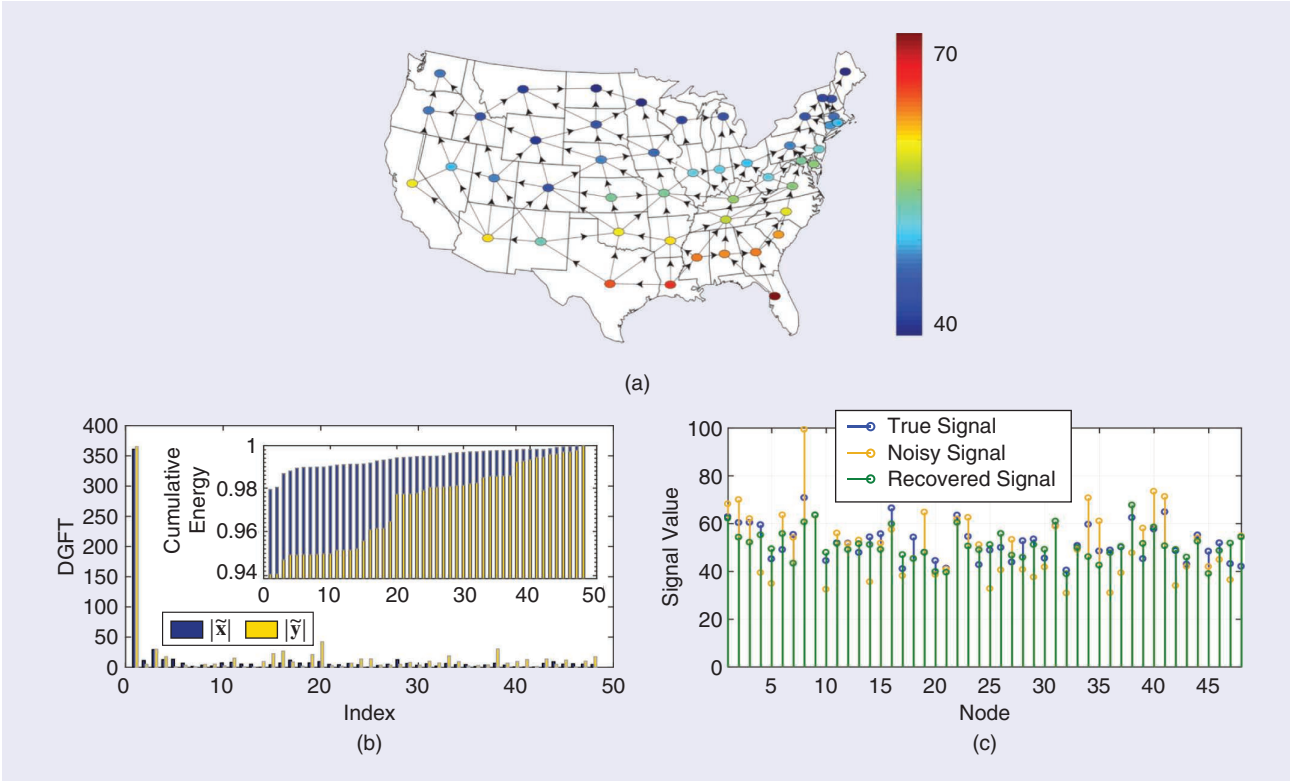

FIGURE 3. Temperature denoising using the DGFT [14]. (a) The graph signal of average annual temperature in Fahrenheit for the contiguous United States. In the depicted digraph, a directed edge joins two states if they share a border, and the edge directions go from south to north. (b) The DGFT of the original signal ($\tilde{\mathbf{x}}$) and the noisy signal ($\tilde{\mathbf{y}}$) along with their cumulative energy distribution across frequencies. (c) A sample realization of the true, noisy, and recovered temperature signals for a filter bandwidth $K = 3$.

normalize the counts out of each node to sum up to one, thus obtaining a weighted digraph whose weights are between 0 and 1. As for the graph signals, they are defined as each function word's count among 1,000 words. Splitting Emily Brontë's texts between training and test sets on an 80:20 ratio, her WAN is generated from function word coappearance counts in the training set only. The graph signals in the training set correspond to 1,000-word excerpts by Brontë and by a pool of 21 other contemporary authors. Each set of graph signals has an associated binary label where 1 indicates that the text has been written by Brontë and excerpts by the rest of the authors are labeled as 0. Test samples are defined analogously. The training and test sets consisted of 1,092 and 272 excerpts, respectively, both with equally balanced classes, and cross entropy was chosen as the loss function.

Several specific GNN architectures were compared in this experiment, all of which followed the general structure in (6)–(8) but for different choices of the nonlinearity $\sigma_{\mathcal{G}}^{(\ell)}$. Indeed, the popular pointwise rectified linear unit (ReLU) was contrasted with more sophisticated graph-localized (but not necessarily pointwise) median and maximum activation functions (see [34] for details). Figure 4(b) presents the authorship attribution accuracy results after conducting 10 rounds of simulations by varying the training and test splits. We can see that the median and maximum GNNs performed consistently better than the ReLU GNNs on discerning between texts written by Brontë and other authors in the pool. Localized activation functions outperformed the pointwise ReLU, with smaller average test errors as well as smaller deviations around this average. Equally important, the simulations also show that their associated error was 1–2% lower than that achieved by NN architectures that symmetrized the WAN. This superior performance underscores the importance of leveraging the digraph structure in the architecture of GNNs, not only in the linear operators via the incorporation of graph filters but also in the determination of nonlinearities.

Graph sampling for handwritten digit recognition

Our goal here is to employ the sampling theory introduced in the "Sampling and Reconstruction" section to classify handwritten digits with minimal labels, as developed in [21]. More precisely, we consider a digraph whose $N = 10,000$ nodes correspond to gray-scaled images in the Modified National Institute of Standards and Technology (MNIST) data set equally distributed among the 10 classes (0–9-digit characters). The edges are obtained from a 12-nearest-neighbor construction computed from the Euclidean distance between vector representations of the images. The graph is directed by construction because one node being in the 12-nearest neighborhood of another node does not guarantee that the relation in the opposite direction holds. This directionality can be especially relevant in the treatment of outliers in the embedded space, where every outlier still has an incoming neighborhood of size 12 but does not belong to the incoming neighborhood of other nodes, thus having a minimal effect in the label propagation. The edges are then weighted using a normalized Gaussian kernel so that, within each neighborhood of size 12, the closer connections have a larger weight. Intuitively, images representing the same digit tend to have similar pixel values and, hence, are more likely to belong to the neighborhood of each other. Thus, if we consider the value of the signal at a given node to be the digit represented by the image associated with that node, the whole graph signal will be piecewise constant in the graph and thus amenable to being reconstructed from observations

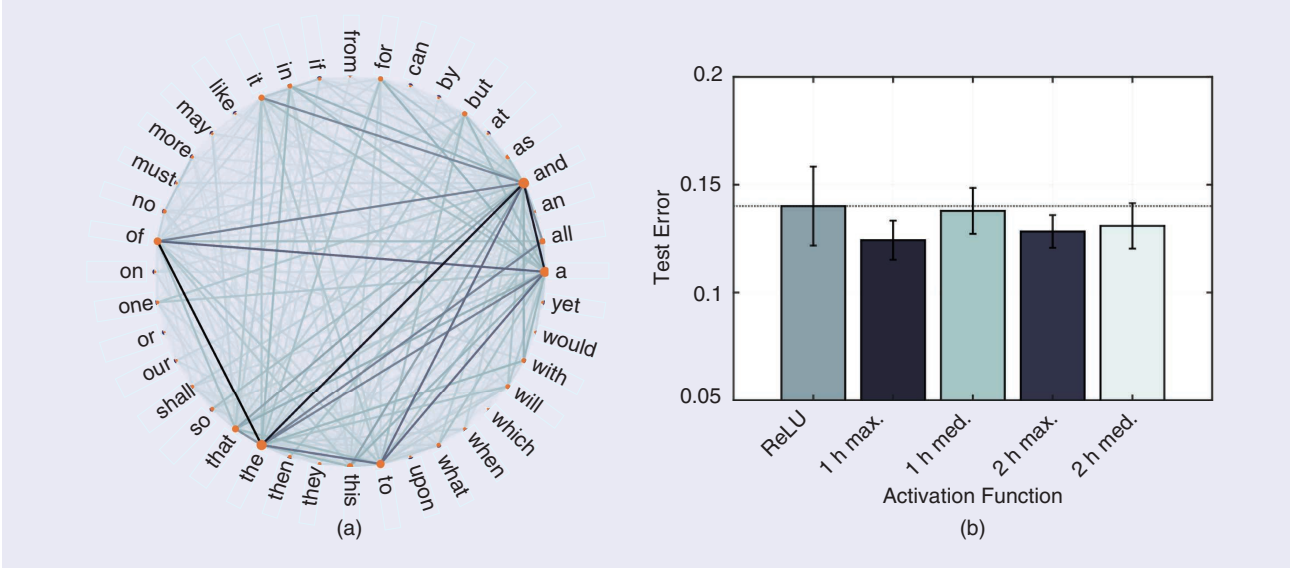

FIGURE 4. Identifying the author of a text using GNNs [34]. (a) An example of a WAN with 40 function words as nodes built from the play "The Humorous Lieutenant" by John Fletcher. The radius of the nodes is proportional to the word count, and the darker the edge color the higher the edge weight. Directionality has been ignored for ease of representation, but the GNNs are defined on the directed WAN. (b) The authorship attribution test error in GNN architectures with localized activation functions for the classification of Emily Brontë versus her contemporaries. Max.: maximum; med.: median.

at a few nodes. Furthermore, to account for the fact that the signal values are categorical, instead of considering a graph signal of dimension $\mathbf{x} \in \mathbb{R}^N$, we consider the alternative binary matrix representation $\mathbf{X} \in \mathbb{R}^{N \times 10}$, where $X_{ij} = +1$ if the ith image is a picture of the digit j and $X_{ij} = -1$ otherwise. Each column of $\mathbf{X}$ is modeled as a band-limited signal that can be written as the linear combination of the K leading columns of $\mathbf{V}$, that is, the eigenvectors of the nonsymmetric adjacency matrix $\mathbf{A}$.

The graph representation of the MNIST digits is shown in Figure 5(a), where the edges were removed for clarity, and the coordinates of each node are given by the corresponding rows of the first three columns of the iGFT $\mathbf{V}$. The enlarged black nodes indicate the optimal choice for 10 samples. Optimality, in this case, refers to the design of $\mathbf{C}_{\mathcal{M}}$ to maximize the minimum singular value of $\mathbf{C}_{\mathcal{M}}\mathbf{V}_{\mathcal{K}}$ (see the "Sampling and Reconstruction" section). Given that we have to (pseudo-)invert this matrix for reconstruction, a good condition number entails a robust behavior in the presence of noise. It is apparent that the images representing the same digit form clusters and that the optimal samples boil down to choosing representative samples from each cluster. The same procedure can be repeated for the United States Postal Service's (USPS') handwritten digits data set consisting of $N = 11{,}000$ images to obtain Figure 5(b). For both cases, one can compute the classification accuracy obtained from the reconstructed graph signals for a different number of optimal samples [see Figure 5(c)].

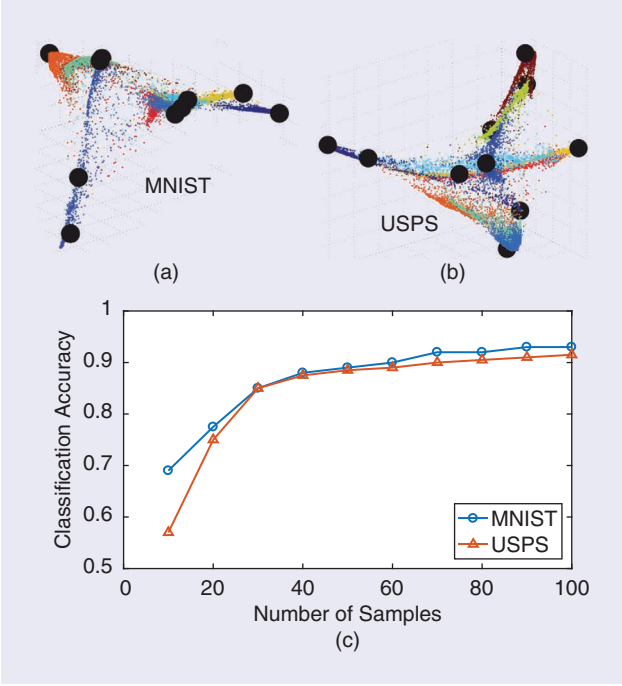

FIGURE 5. Semisupervised learning for handwritten digit classification via the sampling of graph signals [21]. (a) A 3D representation of the MNIST images colored by true class (digits 0–9). The 10 enlarged nodes correspond to the identified optimal samples. (b) Analogous to (a), but for the USPS data set. (c) Classification accuracy as a function of the number of samples used for interpolation for both data sets.

As expected, the accuracy increases with the number of samples. Furthermore, note that even when observing only 50 samples (0.5 and 0.45% of the MNIST and USPS data sets, respectively), the reconstruction accuracy is nearly 0.9, highlighting the importance of incorporating the graph structure via optimal samplers that can accommodate digraphs. This method was shown to outperform other graph-based, active semisupervised learning techniques (see [21] for additional details and experiments).

Kernel-based topology inference for gene-expression data

Consider now the problem of identifying gene-regulatory topologies, where nodes represent individual genes and directed edges encode causal regulatory relationships between gene pairs. Due to the inherent directional nature of regulatory interactions [4, Ch. 7.3], we must recover a digraph as opposed to an undirected relational structure. In this context, we compare the inferred digraphs recovered when implementing different kernels for SEM inference. The experiments were performed on gene-regulatory data collected from $T = 69$ unrelated Nigerian individuals under the International HapMap project (see [29] and the references therein for additional details). From the 929 identified genes, expression levels and genotypes of the expression quantitative trait loci (eQTLs) of $N = 39$ immune-related genes were selected and normalized. The genotypes of eQTLs were considered as exogenous inputs $\mathbf{u}_t$, whereas the gene-expression levels were treated as the endogenous variables $\mathbf{x}_t$ [compare (13)].

Figure 6 depicts the identified topologies, where the different graphs correspond to different choices for the kernel, and the visualizations only include nodes that have at least a single incoming or outgoing edge. More precisely, Figure 6(a) portrays the resulting network based on a linear SEM, while Figure 6(b) and (c) illustrate the results from nonlinear SEMs based on a polynomial kernel of second order and a Gaussian kernel with unit variance, respectively. In the three cases, the identified networks are very sparse, and the nonlinear approaches unveil all of the edges identified by the linear SEMs, along with a number of additional edges. Clearly, considering the possibility that interactions among genes may be driven by nonlinear dynamics, nonlinear frameworks encompass linear approaches and facilitate the discovery of causal (directed) patterns not captured by linear SEMs. The newly unveiled gene-regulatory interactions could potentially be the subject of further studies and direct experimental corroboration by geneticists to improve our understanding of causal influences among immune-related genes across humans.

Emerging-topic areas and conclusions

Contending that signals defined on digraphs are of paramount practical importance, this article outlined recent approaches to model, process, and learn from these graph signals. Accordingly, this article stretched in a comprehensive and unifying manner all the way from the definition of GFTs and graph signal operators especially designed for digraphs to the problem of inferring the digraph itself from the observed signals. A wide

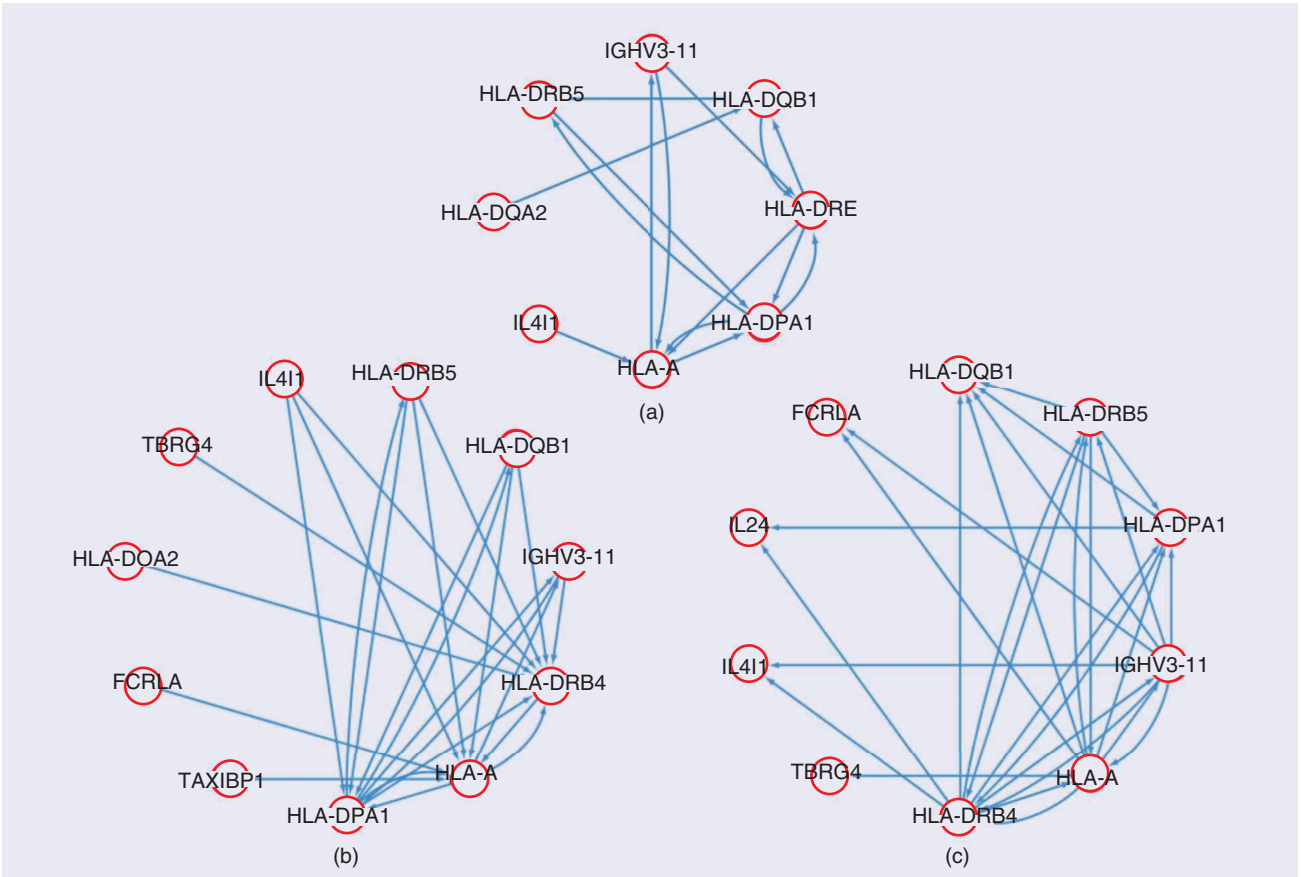

FIGURE 6. Inferring (directed) gene-regulatory networks from expression data [29]. The networks are inferred following the SEM formulation in (14) for (a) a linear kernel, (b) a polynomial kernel of order two, and (c) a Gaussian kernel with unit variance.

range of signal-recovery problems was selectively covered, focusing on inverse problems in digraphs, including sampling, deconvolution, and system identification. A statistical viewpoint for signal modeling was also discussed by extending the definition of weak stationarity of random graph processes to the directed domain. The last stop was to review recent results that applied the tools surveyed in this article to the problem of learning the topology of a digraph from nodal observations, an approach that can lead to meaningful connections between GSP and the field of causal inference in statistics. A common theme in the extension of established GSP concepts to the less-explored realm of digraphs is that the definitions and notions that rely heavily on spectral properties are challenging to generalize, whereas those that can be explicitly postulated in the vertex domain are more amenable to be extended to digraphs.

A diverse gamut of potential research avenues naturally follows from the developments presented. Efficient approaches for the computation of the multiple GFTs for digraphs (akin to the fast Fourier transform in classical SP) would facilitate the adoption of this methodology in large-scale settings. The incorporation of nonlinear (median, Volterra, and NNs) graph signal operators as generative models for the solution of inverse problems is another broad area of promising research. Deep generative models for signals defined in regular domains (such as images) have shown remarkable success over the last several years, and part of that success can be extended to our more challenging domain. Equally interesting is that the use of deep learning to generate the graphs themselves (as opposed to the graph signals) has recently gained traction so that, as discussed in this article, one can conceive NN architectures that learn (and even generate) digraphs from training graph signals while encoding desirable topological features. One final direction of future research is the extension of the concepts discussed in this article to the case of higher-order directed relational structures. The generalization of GSP to hypergraphs through tensor models and simplicial complexes has been explored in recent years, but their analysis in directed scenarios is almost uncharted research territory.

Acknowledgments

The work presented in this article was supported by the Spanish Federal Grants Klinilycs and SPGraph (TEC2016-75361-R and PID2019-105032GB-I00) and NSF Awards CCF-1750428 and ECCS-1809356. Antonio G. Marques is the corresponding author.

Authors

Antonio G. Marques (antonio.garcia.marques@urjc.es) received his five-year diploma and doctorate degree in telecommunications engineering (both with highest honors) from

Carlos III University of Madrid, Spain, in 2002 and 2007, respectively. Currently, he is a full professor with the Department of Signal Theory and Communications, King Juan Carlos University, Madrid, Spain. He is a member of EURASIP and the ELLIS society. He was also the recipient of the 2020 EURASIP Early Career Award. His research interests lie in the areas of signal processing, machine learning, and network science. He is a Member of IEEE.

Santiago Segarra (segarra@rice.edu) received his B.Sc. degree in industrial engineering with highest honors (valedictorian) from the Buenos Aires Institute of Technology, Argentina, in 2011, his M.Sc. degree in electrical engineering from the University of Pennsylvania (Penn), Philadelphia, in 2014, and his Ph.D. degree in electrical and systems engineering from Penn in 2016. Currently, he is the W.M. Rice Trustee Assistant Professor in the Department of Electrical and Computer Engineering at Rice University, with a courtesy appointment in computer science. His research interests include network theory, data analysis, machine learning, and graph signal processing.

Gonzalo Mateos (gmateosb@ece.rochester.edu) received his B.Sc. degree in electrical engineering from the University of the Republic, Montevideo, Uruguay, in 2005 and his M.Sc. and Ph.D. degrees in electrical engineering from the University of Minnesota, Minneapolis, in 2009 and 2012, respectively. Currently, he is an associate professor with the Department of Electrical and Computer Engineering, the University of Rochester, New York, where he is also the Asaro Biggar Family Fellow in Data Science. His research interests lie in the areas of statistical learning from graph data; network science; decentralized optimization; and graph signal processing with applications to social, communication, power grid, and brain network analytics.

References

[1] A. Ortega, P. Frossard, J. Kovačević, J. M. F. Moura, and P. Vandergheynst, "Graph signal processing: Overview, challenges, and applications," *Proc. IEEE*, vol. 106, no. 5, pp. 808–828, 2018. doi: 10.1109/JPROC.2018.2820126.

[2] P. Djuric and C. Richard, *Cooperative and Graph Signal Processing: Principles and Applications*. San Diego, CA: Academic Press, 2018.

[3] D. Shuman, S. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 83–98, 2013. doi: 10.1109/MSP.2012.2235192.

[4] E. D. Kolaczyk, *Statistical Analysis of Network Data: Methods and Models*. New York: Springer-Verlag, 2009.

[5] S. Segarra, G. Mateos, A. G. Marques, and A. Ribeiro, "Blind identification of graph filters," *IEEE Trans. Signal Process.*, vol. 65, no. 5, pp. 1146–1159, 2017. doi: 10.1109/TSP.2016.2628343.

[6] J. Peters, D. Janzing, and B. Schölkopf, *Elements of Causal Inference: Foundations and Learning Algorithms*. Cambridge, MA: MIT Press, 2017.

[7] A. Sandryhaila and J. M. F. Moura, "Discrete signal processing on graphs," *IEEE Trans. Signal Process.*, vol. 61, no. 7, pp. 1644–1656, 2013. doi: 10.1109/TSP.2013.2238935.

[8] A. Sandryhaila and J. M. F. Moura, "Discrete signal processing on graphs: Frequency analysis," *IEEE Trans. Signal Process.*, vol. 62, no. 12, pp. 3042–3054, 2014. doi: 10.1109/TSP.2014.2321121.

[9] H. Sevi, G. Rilling, and P. Borgnat, "Harmonic analysis on directed graphs and applications: From Fourier analysis to wavelets," 2018, arXiv:1811.11636v2.

[10] F. Chung, "Laplacians and the Cheeger inequality for directed graphs," *Ann. Comb.*, vol. 9, no. 1, pp. 1–19, 2005. doi: 10.1007/s00026-005-0237-z.

[11] J. A. Deri and J. M. F. Moura, "Spectral projector-based graph Fourier transforms," *IEEE J. Sel. Topics Signal Process.*, vol. 11, no. 6, pp. 785–795, 2017. doi: 10.1109/JSTSP.2017.2731599.

[12] G. Mateos, S. Segarra, A. G. Marques, and A. Ribeiro, "Connecting the dots: Identifying network structure via graph signal processing," *IEEE Signal Process. Mag.*, vol. 36, no. 3, pp. 16–43, 2019. doi: 10.1109/MSP.2018.2890143.

[13] S. Sardellitti, S. Barbarossa, and P. Di Lorenzo, "On the graph Fourier transform for directed graphs," *IEEE J. Sel. Topics Signal Process.*, vol. 11, no. 6, pp. 796–811, 2017. doi: 10.1109/JSTSP.2017.2726979.

[14] R. Shafipour, A. Khodabakhsh, G. Mateos, and E. Nikolova, "A directed graph Fourier transform with spread frequency components," *IEEE Trans. Signal Process.*, vol. 67, no. 4, pp. 946–960, 2019. doi: 10.1109/TSP.2018.2886151.

[15] S. Segarra, A. G. Marques, and A. Ribeiro, "Optimal graph-filter design and applications to distributed linear network operators," *IEEE Trans. Signal Process.*, vol. 65, no. 15, pp. 4117–4131, 2017. doi: 10.1109/TSP.2017.2703660.

[16] M. Coutino, E. Isufi, and G. Leus, "Advances in distributed graph filtering," *IEEE Trans. Signal Process.*, vol. 67, no. 9, pp. 2320–2333, 2019. doi: 10.1109/TSP.2019.2904925.

[17] S. Segarra, A. G. Marques, G. R. Arce, and A. Ribeiro, "Center-weighted median graph filters," in *Proc. IEEE Global Conf. Signal and Information Processing (GlobalSIP)*, Dec. 2016, pp. 336–340. doi: 10.1109/GlobalSIP.2016.7905859.

[18] F. Gama, A. G. Marques, G. Leus, and A. Ribeiro, "Convolutional neural network architectures for signals supported on graphs," *IEEE Trans. Signal Process.*, vol. 67, no. 4, pp. 1034–1049, 2019. doi: 10.1109/TSP.2018.2887403.

[19] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric deep learning: Going beyond Euclidean data," *IEEE Signal Process. Mag.*, vol. 34, no. 4, pp. 18–42, 2017. doi: 10.1109/MSP.2017.2693418.

[20] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *IEEE Trans. Neural Netw. Learn. Syst.*, early access, Mar. 24, 2020. doi: 10.1109/TNNLS.2020.2978386.

[21] S. Chen, R. Varma, A. Sandryhaila, and J. Kovačević, "Discrete signal processing on graphs: Sampling theory," *IEEE Trans. Signal Process.*, vol. 63, no. 24, pp. 6510–6523, 2015. doi: 10.1109/TSP.2015.2469645.

[22] S. Segarra, A. G. Marques, G. Leus, and A. Ribeiro, "Reconstruction of graph signals through percolation from seeding nodes," *IEEE Trans. Signal Process.*, vol. 64, no. 16, pp. 4363–4378, 2016. doi: 10.1109/TSP.2016.2552510.

[23] L. Le Magoarou, R. Gribonval, and N. Tremblay, "Approximate fast graph Fourier transforms via multilayer sparse approximations," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 4, no. 2, pp. 407–420, 2017. doi: 10.1109/TSIPN.2017.2710619.

[24] A. G. Marques, S. Segarra, G. Leus, and A. Ribeiro, "Sampling of graph signals with successive local aggregations," *IEEE Trans. Signal Process.*, vol. 64, no. 7, pp. 1832–1843, 2016. doi: 10.1109/TSP.2015.2507546.

[25] S. Chen, A. Sandryhaila, J. M. F. Moura, and J. Kovačević, "Signal recovery on graphs: Variation minimization," *IEEE Trans. Signal Process.*, vol. 63, no. 17, pp. 4609–4624, 2015. doi: 10.1109/TSP.2015.2441042.

[26] F. J. Iglesias, S. Segarra, S. Rey-Escudero, A. G. Marques, and D. Ramirez, "Demixing and blind deconvolution of graph-diffused sparse signals," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2018, pp. 4189–4193. doi: 10.1109/ICASSP.2018.8462154.

[27] Y. Zhu, F. J. Iglesias, A. G. Marques, and S. Segarra, "Estimation of network processes via blind graph multi-filter identification," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 5451–5455. doi: 10.1109/ICASSP.2019.8683844.

[28] G. B. Giannakis, Y. Shen, and G. V. Karanikolas, "Topology identification and learning over graphs: Accounting for nonlinearities and dynamics," *Proc. IEEE*, vol. 106, no. 5, pp. 787–807, 2018. doi: 10.1109/JPROC.2018.2804318.

[29] Y. Shen, B. Baingana, and G. B. Giannakis, "Kernel-based structural equation models for topology identification of directed networks," *IEEE Trans. Signal Process.*, vol. 65, no. 10, pp. 2503–2516, 2017. doi: 10.1109/TSP.2017.2664039.

[30] A. Bolstad, B. D. V. Veen, and R. Nowak, "Causal network inference via group sparse regularization," *IEEE Trans. Signal Process.*, vol. 59, no. 6, pp. 2628–2641, 2011. doi: 10.1109/TSP.2011.2129515.

[31] J. Mei and J. M. F. Moura, "Signal processing on graphs: Causal modeling of unstructured data," *IEEE Trans. Signal Process.*, vol. 65, no. 8, pp. 2077–2092, 2017. doi: 10.1109/TSP.2016.2634543.

[32] R. Shafipour, S. Segarra, A. G. Marques, and G. Mateos, "Directed network topology inference via graph filter identification," in *Proc. IEEE Data Science Workshop (DSW)*, June 2018, pp. 210–214. doi: 10.1109/DSW.2018.8439888.

[33] S. Segarra, A. G. Marques, G. Mateos, and A. Ribeiro, "Network topology inference from spectral templates," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 3, no. 3, pp. 467–483, 2017. doi: 10.1109/TSIPN.2017.2731051.

[34] L. Ruiz, F. Gama, A. G. Marques, and A. Ribeiro, "Invariance-preserving localized activation functions for graph neural networks," *IEEE Trans. Signal Process.*, vol. 68, pp. 127–141, Nov. 2019. doi: 10.1109/TSP.2019.2955832.

SP