ENGINEERING Information Technology & Electronic Engineering
https://www.academax.com/EITEE; https://www.engineering.org.cn
https://ieeexplore.ieee.org/xpl/RecentIssue.jsp?punumber=11344065
E-mail: eng.itee@zju.edu.cn

Engineering Journals
Available on IEEE Xplore®

Research Article
https://doi.org/10.1631/ENG.ITEE.2026.0142

Expected graph Hilbert transform for time–frequency analysis of uncertain graph topologies

Guanlei XU[1,2]✉, Xiaogang XU[1,2], Xun WANG[1,2]✉

[1]School of Computer Science and Technology, Zhejiang Gongshang University, Hangzhou 310018, China
[2]Zhejiang Key Laboratory of Big Data and Future E-Commerce Technology, Hangzhou 310018, China

Abstract: We propose the expected graph Hilbert transform (EGHT), a robust extension of the Hilbert transform, for graph signals under topological uncertainty. By modeling the shift operator as a random variable and using its expected value as a stable spectral reference, the EGHT applies a $\pi/2$ phase shift to non-direct current (DC) graph Fourier modes, enabling the construction of analytic signals on uncertain graphs. The definitions of EGHT for undirected and directed graph ensembles are given, along with key properties such as linearity, anti-involution, and orthogonality; moreover, the expected graph analytic signal (EGAS) for node-wise modulation analysis is introduced. Experiments show that the EGHT outperforms conventional graph Hilbert transforms (GHTs) on noisy realizations, recovering coherent oscillations, producing smooth phase fields, and resisting topological perturbations. As such, it enables reliable time–frequency analysis in stochastic networks.

Key words: Graph signal processing (GSP); Hilbert transform; Analytic signal; Probabilistic graphs; Expected shift operator; Uncertain topology

1 Introduction

The analytic signal, introduced by Gabor (1946), is a foundational construct in signal processing that enables instantaneous amplitude and phase estimation, envelope detection, and modulation analysis. By embedding a real-valued signal into the complex plane via the Hilbert transform, which effectively suppresses negative frequencies, the analytic signal provides a one-sided spectral representation that is indispensable in communications, biomedical engineering, and time–frequency analysis (Cohen, 1995). This elegant framework relies on the shift-invariant structure of Euclidean domains, where the Fourier transform yields a stable global frequency basis.

In contrast, data residing on irregular structures, such as social networks, brain connectomes, or sensor grids, lack translational symmetry, rendering classical tools inapplicable. Graph signal processing (GSP) addresses this by redefining signal operations through the lens of a graph shift operator (e.g., the adjacency or Laplacian matrix) (Sandryhaila and Moura, 2014; Wang TT et al., 2020; Chen et al., 2024; Lyu et al., 2025; Sun et al., 2026). Within this paradigm, the graph Fourier transform (GFT) diagonalizes the shift operator, enabling spectral filtering, sampling, and reconstruction. Recent work has extended the notion of analytic signals to graphs (Sandryhaila and Moura, 2014; Venkitaraman et al., 2019). For example, Sandryhaila and Moura (2014) proposed a graph Hilbert transform (GHT) using a spectral signum filter that nullifies components associated with negative eigenvalues. Moreover, Venkitaraman et al. (2019) developed a framework for defining Hilbert transforms and analytic signals on graphs using graph-shift-based constructions, but without enforcing the strict one-sided spectral support or energy-preserving properties inherent to classical analytic signals. These advances demonstrate that concepts like instantaneous amplitude and phase can be meaningfully generalized to structured domains (Sandryhaila and Moura, 2014; Venkitaraman et al., 2019; Wang TT et al., 2020).

However, a critical assumption underpins all existing graph analytic signal frameworks: the graph topology is known exactly and remains fixed. In practice, this assumption rarely holds. Communication links in wireless networks drop stochastically; functional brain connectivity varies across subjects or time; molecular interactions are inferred with uncertainty. Consequently, the shift operator—and thus the graph Fourier basis—is often random or imperfectly observed. Under such conditions, deterministic spectral projections become unstable. Small perturbations in a graph structure can cause significant rotations in eigenvectors, leading to erratic Hilbert transforms

✉ Guanlei XU, xgl_86@163.com
Xun WANG, wx@mail.zjgsu.edu.cn
Guanlei XU, https://orcid.org/0000-0002-4389-3545
Xun WANG, https://orcid.org/0000-0001-5566-4689
CLC number: TN911.72

Received: May 20, 2026; Revision accepted: July 17, 2026;
Crosschecked: Aug. 30, 2026

© The Authors 2026. Published by Zhejiang University Press Co., Ltd.
This is an open access article distributed under the terms of the CC BY-NC-ND license (https://creativecommons.org/licenses/by-nc-nd/4.0/)

and meaningless phase estimates. This fragility undermines the reliability of analytic representations in real-world applications.

Essentially, while existing GHT frameworks provide valuable tools for analyzing signals on a fixed graph topology, their direct application in scenarios with uncertain or random topologies remains fundamentally flawed. A naive approach might involve first constructing an average graph—typically represented by its expected shift operator—and then applying a conventional GHT to this deterministic surrogate. However, this two-step process (averaging the graph and then transforming the signal) fails to account for the inherent spectral variability across the ensemble of possible graphs. The eigenvectors will generally not be aligned with those of the individual realizations, leading to a loss of crucial phase and frequency information that is specific to the underlying uncertainty model.

Parallel developments in GSP have begun addressing structural uncertainty. Isufi et al. (2017a) and Ji et al. (2023) formalized GSP over a probability space of shift operators, where filters, convolutions, and sampling strategies were designed in terms of expectation over a distribution (Elmoataz et al., 2008; Girault et al., 2015; Isufi et al., 2017a; Mei and Moura, 2017; Perraudin and Vandergheynst, 2017; Chamon and Ribeiro, 2018). This probabilistic approach enhances robustness and generalization, particularly in learning-based settings. Yet, these works have primarily focused on linear filtering and reconstruction, and do not address complex signal representations or the fundamental question of how to define a stable, interpretable analytic signal when the graph itself is random (Elmoataz et al., 2008; Girault et al., 2015; Isufi et al., 2017a; Mei and Moura, 2017; Perraudin and Vandergheynst, 2017; Chamon and Ribeiro, 2018).

In this paper, we bridge these two lines of inquiry by introducing the expected graph analytic signal (EGAS), a framework that constructs analytic signals using the eigendecomposition of the expected graph shift operator. The core component of EGAS is the expected graph Hilbert transform (EGHT), which generalizes a classical Hilbert transform to stochastic graph domains. Our approach redefines the graph analytic signal not through a single spectral decomposition, but through an ensemble-averaged complex embedding that accounts for graph structure variability. In contrast to works (Hu et al., 2024; Zhang et al., 2025) that use graph-based methods for tasks such as entity relationship extraction, dialogue text analysis, and knowledge graph completion, our methodology is foundational for graph time–frequency analysis and its theory can be applied to GSP and the associated applications (Isufi et al., 2024; Dai et al., 2025; Gu et al., 2025; Jiang et al., 2025; Li et al., 2025; Liu et al., 2025; Parada-Mayorga and Ribeiro, 2025; Hou et al., 2026; Wang YH et al., 2026; Zhou et al., 2026). In other words, our proposed EGHT adopts a principled, one-step approach. Instead of operating on a single averaged graph, EGHT constructs the analytic signal by coherently averaging spectral filtering operations across all possible graph instances, weighted by their probability. This ensures that the resulting transform is intrinsically aware of the topological uncertainty from the outset, preserving the statistical properties of the signal-graph ensemble. The distinction between these paradigms is not merely technical; it is conceptual and leads to significant practical differences in robustness and interpretability.

The main contributions of this work are threefold. First, we define the EGHT in a probabilistic case, instead of in terms of deterministic counterparts; this is generally non-orthogonal and does not perfectly nullify negative frequencies, but it does minimize the expected spectral leakage across graph realizations. Second, we construct the probabilistic analytic signal of a real graph signal and show that its magnitude can serve as a robust envelope estimator, smoothing out artifacts from spurious high-frequency components induced by graph noise. The new phase reflects consistent alignment with dominant spectral subspaces across the ensemble, allowing coherent modulation analysis even when individual graph spectra are misaligned. Third, the utility of the proposed framework is demonstrated in real-world scenarios, and the results are compared to those of deterministic analytic signal methods in terms of stability and interpretability. For directed graphs, we further discuss the inherent trade-off between directional sensitivity and energy preservation, and propose a robust alternative based on the Hermitian portion of the shift operator to mitigate potential biases in envelope estimation.

2 Shift operators and probability space

Classical signal processing relies heavily on the Hilbert transform to construct analytic signals, enabling instantaneous amplitude, phase, and frequency estimation; it is thus a cornerstone of modulation theory, communications, and time–frequency analysis. Extending this approach to signals defined over graphs is nontrivial due to the absence of a natural concept of frequency or time shift. However, two recent theoretical advances have paved the way for this application. First, several works (Elmoataz et al., 2008; Girault et al., 2015; Isufi et al., 2017b; Perraudin and Vandergheynst, 2017; Chamon and Ribeiro, 2018) have modeled uncertainty in graph topology by treating the shift operator $\boldsymbol{S}$ as a random variable drawn from a distribution $\mathcal{P}$. This yields an expected shift operator $\overline{\boldsymbol{S}}=\mathbb{E}_{\mathcal{P}}[\boldsymbol{S}]$, which serves as a robust surrogate for spectral analysis. By synthesizing these frameworks, we define the EGHT—a principled, uncertainty-aware extension of the Hilbert transform to graph signals. This construct not only generalizes classical results but also provides resilience to topological noise, missing edges, and ensemble variability (such issues are common in real-world networks such as brain connectomes, social media graphs, and sensor arrays). The remainder of this paper will cover EGHT formulations for undirected and directed graphs, full spectral derivations, algebraic and functional-analytic properties, computational algorithms, illustrative numerical examples, and connections to harmonic analysis and operator theory.

Before proceeding, we will review basic concepts and definitions. Let $\mathcal{G}=(\mathcal{V},\mathcal{E})$ be a graph with a node set $\mathcal{V}=\{1,2,\cdots,N\}$. A graph signal is a vector $\boldsymbol{x}\in\mathbb{C}^N$. A shift operator $\boldsymbol{S}\in\mathbb{R}^{N\times N}$ encodes the connectivity as follows: for undirected graphs, the common choices include the adjacency matrix $\boldsymbol{A}$ (symmetric), combinatorial Laplacian $\boldsymbol{L}=\boldsymbol{D}-\boldsymbol{A}$, and normalized Laplacian $\boldsymbol{\mathcal{L}}=\boldsymbol{I}-\boldsymbol{D}^{-1/2}\boldsymbol{A}\boldsymbol{D}^{-1/2}$; for directed graphs, $\boldsymbol{A}$ is generally asymmetric and the alternatives include random-walk matrix $\boldsymbol{P}=\boldsymbol{D}_{\text{out}}^{-1}\boldsymbol{A}$ and asymmetric Laplacians (e.g., $\boldsymbol{L}_{\text{rw}}=\boldsymbol{I}-\boldsymbol{P}$).

All are valid shift operators satisfying the GSP axiom that shifting commutes with filtering. Next, assume a measurable space $(\Omega,\mathcal{F})$, where Ω and $\mathcal{F}$ are sampling space and σ-algebra, respectively, and each outcome $\omega\in\Omega$ corresponds to a graph realization

with a shift operator $\boldsymbol{S}(\omega)$. Let $\mathbb{P}$ be a probability measure on sampling space Ω. We assume that

$$\mathbb{E}\left[\|\boldsymbol{S}\|_{\mathrm{F}}\right]<\infty \text{ (finite second moment by Frobenius norm).} \quad (1)$$

The support of $\mathbb{P}$ consists of admissible shift operators (e.g., symmetric ones for undirected ensembles).

The expected shift operator is

$$\overline{\boldsymbol{S}}=\int_{\Omega}\boldsymbol{S}(\omega)\,\mathrm{d}\mathbb{P}(\omega). \quad (2)$$

In discrete settings with M possible graphs, we have

$$\overline{\boldsymbol{S}}=\sum_{m=1}^{M}p_m\boldsymbol{S}^{(m)},\ \ p_m\geq 0,\ \sum_{m=1}^{M}p_m=1, \quad (3)$$

where m represents the mth graph. Crucially, $\overline{\boldsymbol{S}}$ may not correspond to any actual graph in the ensemble, but is still a valid linear operator for spectral analysis.

3 EGHT on undirected graphs

In this section, we provide mathematical definitions and properties of EGHT for undirected graphs.

3.1 Definition of EGHT

For undirected graphs, all standard shift operators are real and symmetric, and hence orthogonally diagonalizable:

$$\overline{\boldsymbol{L}}=\boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^*,\ \boldsymbol{U}^*\boldsymbol{U}=\boldsymbol{I}, \quad (4)$$

where $\boldsymbol{\Lambda}=\mathrm{diag}(\lambda_1,\lambda_2,\cdots,\lambda_N)$, $0=\lambda_1\leq\lambda_2\leq\cdots\leq\lambda_N$ and $\boldsymbol{U}^*=\boldsymbol{U}^{\mathrm{T}}=\boldsymbol{U}^{-1}$.

The GFT of graph signal $\boldsymbol{x}$ is

$$\hat{\boldsymbol{x}}=\boldsymbol{U}^{\mathrm{T}}\boldsymbol{x},\ x_n=\sum_{k=1}^{N}\hat{x}_k u_k(n). \quad (5)$$

Definition 1 EGHT on undirected graphs is defined as

$$\mathcal{H}_{\overline{L}}[\boldsymbol{x}]:=\boldsymbol{U}\boldsymbol{J}\boldsymbol{U}^*\boldsymbol{x}, \quad (6)$$

where the Hilbert mask $\boldsymbol{J}\in\mathbb{C}^{N\times N}$ is given by

$$J_{kk}=\begin{cases}-\mathrm{j}, & \text{if } \lambda_k>0,\\ 0, & \text{if } \lambda_k=0.\end{cases} \quad (7)$$

This definition is not merely a mathematical abstraction—it embodies a physically grounded mechanism for extracting instantaneous oscillatory structure from graph signals in the presence of topological noise. Here, $\overline{\boldsymbol{L}}=\boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^*$ is the expected combinatorial Laplacian derived from an ensemble of plausible network realizations, and $\boldsymbol{U}$ contains its orthonormal eigenvectors (the graph Fourier basis). The diagonal mask $\boldsymbol{J}$ applies a $-\mathrm{j}$ (i.e., $-90°$) phase shift to all modes with $\lambda_k>0$, while nullifying the direct current (DC) component ($\lambda_k=0$).

Physically, this operation mimics how the classical Hilbert transform isolates propagating waves in time by shifting positive frequencies by $-90°$. On a graph, frequency is encoded in the Laplacian eigenvalues as such:

1. The DC mode ($\lambda=0$) represents spatially constant activity—akin to a global offset with no oscillation. It exhibits no directional or temporal dynamics and is thus excluded from phase modulation.

2. Non-DC modes ($\lambda>0$) correspond to spatial patterns with increasing variation (e.g., smooth gradients to high-frequency alternations). By rotating each of these modes by $-90°$ in the complex plane, the EGHT constructs a quadrature companion that is spectrally orthogonal.

This quadrature signal, when combined with the original $\boldsymbol{x}$ as $\boldsymbol{x}_{\mathrm{a}}=\boldsymbol{x}+\mathrm{j}\mathcal{H}_{\overline{L}}[\boldsymbol{x}]$, yields a node-wise analytic signal whose magnitude and phase describe local envelope and instantaneous spatial phase—this is crucial for identifying coherent wavefronts, traveling rhythms, or synchronized clusters in systems like brain networks or power grids.

Critically, because $\overline{\boldsymbol{L}}$ involves averaging over graph realizations, the EGHT filters out topological artifacts (e.g., missing edges in a single functional magnetic resonance imaging (fMRI) scan) and operates on the statistically dominant geometry of the system. Therefore, the phase shifts are applied in a basis that reflects the true underlying symmetry (e.g., a ring), not a distorted view from one noisy observation. This makes the EGHT a robust spectral denoiser that can recover physically meaningful oscillations, even when no single observed graph is correct; this is a key advantage in real-world sensing applications where network uncertainty is inevitable.

The expected Laplacian $\overline{\boldsymbol{L}}$ used in the EGHT is defined as the ensemble average of the Laplacian matrices over all possible graph realizations in the stochastic graph model:

$$\overline{\boldsymbol{L}}=\mathbb{E}[\boldsymbol{L}]=\sum_{i=0}^{M-1}p_i\boldsymbol{L}^{(i)}, \quad (8)$$

where $\boldsymbol{L}^{(i)}=\boldsymbol{D}^{(i)}-\boldsymbol{A}^{(i)}$ is the combinatorial Laplacian of the ith graph realization, $\boldsymbol{A}^{(i)}$ is its adjacency matrix, $\boldsymbol{D}^{(i)}=\mathrm{diag}(\boldsymbol{A}^{(i)}\boldsymbol{1})$ is the degree matrix, and $p_i=\mathbb{P}(\mathcal{G}=\mathcal{G}^{(i)})$ is the probability of observing graph $\mathcal{G}^{(i)}$. The sum runs over all M possible realizations in the support of the random graph model.

In Eq. (8), $\overline{\boldsymbol{L}}$ is not computed from a single noisy graph but from the statistical generative model of the graph ensemble. This makes the EGHT inherently robust to topological uncertainty.

To clarify the novelty of the EGHT, we compare it with two intuitive but suboptimal baselines. (1) GHT on an expected graph: this method computes the analytic signal as $\boldsymbol{x}_{\mathrm{a}}^{\mathrm{A}}=\boldsymbol{x}+\mathrm{j}\cdot\mathcal{H}_{\overline{L}}[\boldsymbol{x}]$, where $\mathcal{H}_{\overline{L}}$ is a standard GHT (e.g., defined via the eigendecomposition of the expected Laplacian $\overline{\boldsymbol{L}}$). This approach assumes that the average graph $\overline{\mathcal{G}}$ is a sufficient statistic for the entire ensemble, which is often invalid. (2) Instance-wise GHT: this method picks a single graph instance $\boldsymbol{S}(\omega_0)$ (e.g., the most likely one) and computes $\boldsymbol{x}_{\mathrm{a}}^{\mathrm{B}}=\boldsymbol{x}+\mathrm{j}\cdot\mathcal{H}_{\overline{L}}(\omega_0)[\boldsymbol{x}]$. Its performance is highly sensitive to the choice of ω_0 and offers no guarantee of robustness.

The core difference between these baselines lies in their order of operations. Both calculate the Hilbert transform conditional on a single fixed graph structure. The EGHT, as defined in Eq. (6),

performs an unconditional expectation over the transform itself: $\mathcal{H}_{\mathrm{EG}}[\boldsymbol{x}]=\mathbb{E}_{\omega\sim\mathcal{P}}\left[\mathcal{H}_{\bar{\boldsymbol{L}}(\omega)}[\boldsymbol{x}]\right]$. This formulation ensures that the phase quadrature introduced by the Hilbert transform is consistent with the statistical distribution of the graph and is not just a single snapshot or an ill-defined average. In essence, the EGHT provides the minimum mean square error (MMSE) estimate of the ideal analytic signal under the given topological uncertainty—a property that the naive baselines lack.

3.2 Derivation from functional calculus

The EGHT can also be derived using spectral (functional) calculus. For any real-valued function f: $\mathbb{R}\to\mathbb{C}$, one defines the associated graph operator as

$$f(\bar{\boldsymbol{L}})=\boldsymbol{U}f(\boldsymbol{\Lambda})\boldsymbol{U}^{*}, \tag{9}$$

where $\bar{\boldsymbol{L}}=\boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{*}$ is the eigendecomposition of the expected Laplacian. We choose the Hilbert spectral mask

$$f_{\mathcal{H}}(\lambda)=\begin{cases}-\mathrm{j}, & \lambda>0,\\ 0, & \lambda=0,\end{cases} \tag{10}$$

which yields

$$\mathcal{H}_{\bar{\boldsymbol{L}}}=f_{\mathcal{H}}(\bar{\boldsymbol{L}}). \tag{11}$$

This formulation underscores a fundamental principle: the transform acts independently on each spatial eigenmode of the network's expected geometry. Each Laplacian eigenmode represents a natural vibration pattern of the system—much like normal modes in mechanical structures or standing waves in acoustics. The DC mode (λ=0) corresponds to uniform activity with no spatial variation and thus carries no oscillatory information; it is therefore suppressed (set to zero). All other modes (λ>0) represent spatially varying patterns that can support wave-like dynamics. By applying a uniform −90° (−j) phase shift to all of them, the EGHT constructs a quadrature signal that is spectrally orthogonal to the origin, enabling instantaneous amplitude and phase estimation at every node.

This perspective also reveals why the EGHT is compatible with physical signal models: many dynamical processes on networks (e.g., diffusion, synchronization and wave propagation) are governed by operators that are functions of $\bar{\boldsymbol{L}}$ (e.g., $\mathrm{e}^{-t\bar{\boldsymbol{L}}}$, where t denotes time, for heat flow). The Hilbert transform $f_{\mathcal{H}}(\bar{\boldsymbol{L}})$ lives in the same functional space, ensuring consistent treatment of signals and dynamics.

Practically, expressing the EGHT as $f_{\mathcal{H}}(\bar{\boldsymbol{L}})$ enables scalable implementation. For large-scale graphs ($N\gg1$), full eigendecomposition is infeasible. However, because $f_{\mathcal{H}}$ is a bounded and measurable function, it can be efficiently approximated using polynomial expansions (e.g., the Chebyshev series) or Krylov subspace methods (e.g., Lanczos). This allows the EGHT to be applied in real-world settings without explicit spectral decomposition, while preserving its core physical behavior: robust, topology-aware analytic signal construction.

At the same time, we know that the ideal Hilbert mask $J(\lambda)$ is defined as a discontinuous step function. This discontinuity at the DC component (λ=0) poses a significant challenge for polynomial-based approximations (e.g., Chebyshev or Lanczos), which are inherently smooth and struggle to represent sharp transitions. As a result, naive approximation of $J(\lambda)$ typically exhibits Gibbs oscillations near λ=0, leading to spurious energy leakage into the zero-frequency mode and biased analytic signal estimates; this is particularly detrimental for envelope and instantaneous phase computation.

To mitigate this issue, we adopt a spectrally smoothed approximation of the Hilbert mask. Specifically, the ideal step function is replaced with a continuous surrogate that transitions smoothly from 0 to −j over a small interval $[0, \epsilon]$, where ϵ>0 is a tunable bandwidth parameter. One effective choice is the raised-cosine taper:

$$J_{\epsilon}(\lambda)=\begin{cases}0, & \lambda\le0,\\ -\mathrm{j}\cdot\dfrac{1}{2}\left[1-\cos\left(\pi\dfrac{\lambda}{\epsilon}\right)\right], & 0<\lambda<\epsilon,\\ -\mathrm{j}, & \lambda\ge\epsilon.\end{cases}$$

This smoothed mask $J_{\epsilon}(\lambda)$ is Lipschitz continuous, significantly reducing Gibbs artifacts while preserving the essential analytic signal structure for $\lambda\gg\epsilon$. The parameter ϵ can be chosen based on the graph's spectral gap or set empirically (e.g., $\epsilon=0.05\lambda_{\max}$). In practice, even a modest smoothing ($\epsilon\approx10^{-2}\lambda_{\max}$) dramatically improves the stability of the approximation.

3.3 Properties, analytic signals, and algorithmic implementation

First, we will present some properties of the EGHT, the proofs of which are trivial and thus omitted.

1. Linearity and boundedness

$\mathcal{H}_{\bar{\boldsymbol{L}}}$ is linear and bounded: $\left\|\mathcal{H}_{\bar{\boldsymbol{L}}}[\boldsymbol{x}]\right\|_2\le\|\boldsymbol{x}\|_2$, where $\|\cdot\|_2$ denotes the 2-norm.

2. Anti-involution

$\mathcal{H}_{\bar{\boldsymbol{L}}}^{2}=-(\boldsymbol{I}-\boldsymbol{u}_1\boldsymbol{u}_1^{\mathrm{T}})=-\boldsymbol{\Pi}_{\perp}$, where $\boldsymbol{\Pi}_{\perp}$ projects onto the orthogonal complement of the constant subspace. Thus, applying the EGHT twice yields a negative of the zero-mean part of the signal.

3. Orthogonality

For any zero-mean signal ($\boldsymbol{u}_1^{\mathrm{T}}\boldsymbol{x}=\boldsymbol{0}$), $<\boldsymbol{x}, \mathcal{H}_{\bar{\boldsymbol{L}}}[\boldsymbol{x}]>=\boldsymbol{x}^{\mathrm{T}}\mathcal{H}_{\bar{\boldsymbol{L}}}[\boldsymbol{x}]=\boldsymbol{0}$, since the inner product becomes $\sum_{k=2}^{N}\hat{\boldsymbol{x}}_k(-\mathrm{j}\hat{\boldsymbol{x}}_k^{*})=-\mathrm{j}|\hat{\boldsymbol{x}}_k|^2\in\mathrm{j}\mathbb{R}$, and the real part vanishes.

4. Isometry on the zero-mean subspace

If $\boldsymbol{u}_1^{\mathrm{T}}\boldsymbol{x}=\boldsymbol{0}$, then $\left\|\mathcal{H}_{\bar{\boldsymbol{L}}}[\boldsymbol{x}]\right\|_2^2=\sum_{k=2}^{N}|-\mathrm{j}\hat{\boldsymbol{x}}_k|^2=\sum_{k=2}^{N}|\hat{\boldsymbol{x}}_k|^2=|\boldsymbol{x}|_2^2$.

5. Commutativity with expected filtering

For any expected graph filter $h(\bar{\boldsymbol{L}})$, we have $\mathcal{H}_{\bar{\boldsymbol{L}}}h(\bar{\boldsymbol{L}})=h(\bar{\boldsymbol{L}})\mathcal{H}_{\bar{\boldsymbol{L}}}$, since both are diagonal in the same basis.

Definition 2 EGAS on undirected graphs is defined as

$$\boldsymbol{x}_{\mathrm{a}}=\boldsymbol{x}+\mathrm{j}\mathcal{H}_{\bar{\boldsymbol{L}}}[\boldsymbol{x}]=\boldsymbol{U}(\boldsymbol{I}+\boldsymbol{J})\boldsymbol{U}^{*}\boldsymbol{x}. \tag{12}$$

Since $\boldsymbol{J}$ applies a −j phase shift to all non-DC Laplacian modes, the diagonal operator $\boldsymbol{I}+\boldsymbol{J}$ acts as $\boldsymbol{I}+\boldsymbol{J}=\mathrm{diag}(1, 1-\mathrm{j}, \cdots, 1-\mathrm{j})$.

This has two critical spectral effects. The DC component (λ=0) is unchanged ($\boldsymbol{x}_{\mathrm{a}}^{(1)}=\hat{\boldsymbol{x}}_1$), preserving the global mean or baseline

level of the signal—this physically represents the system's steady-state or offset, which carries no oscillatory information. For the non-DC components ($\lambda>0$), each is multiplied by $1-\mathrm{j}=\sqrt{2}\,\mathrm{e}^{-\mathrm{j}\pi/4}$, which increases the magnitude by $\sqrt{2}$ and rotates the phase by −45°. This is not arbitrary: it constructs a complex signal whose real part is the original spatial pattern and whose imaginary part is its quadrature (Hilbert) counterpart; this is exactly analogous to a classical analytic signal in time-domain signal processing.

Physically, this means that EGAS encodes a local phasor at every node n that captures how the signal participates in coherent spatial oscillations across the expected network geometry. The instantaneous amplitude $a[n]=|\boldsymbol{x}_{\mathrm{a}}[n]|$ quantifies the local strength of oscillatory engagement, e.g., how strongly a brain region is entrained in a large-scale rhythm; even if the graph observed in a single trial is missing key connections. Meanwhile, the instantaneous phase $\phi[n]=\arg(\boldsymbol{x}_{\mathrm{a}}[n])$ reveals the relative timing of that node in the global wave pattern, enabling the detection of phase gradients, wavefronts, or leader–follower relationships.

Crucially, because EGAS is built using $\overline{\boldsymbol{L}}$ (the statistically averaged Laplacian), it remains robust to topological uncertainty. In applications like neuroimaging, where individual connectomes are noisy estimates of an underlying population architecture, EGAS filters out spurious structural artifacts and recovers the true collective dynamics supported by the ensemble-averaged network. Thus, EGAS provides a physically grounded, node-wise descriptor of amplitude modulation and phase coherence in complex systems, even when no single observed graph fully reflects the true interaction topology.

Finally, we describe algorithmic implementations for the EGHT. For large N, full eigendecomposition is prohibitive. Alternatives include Chebyshev polynomial approximation of $f_{\mathcal{H}}(\lambda)$ over $[\lambda^{+}_{\min},\lambda_{\max}]$, Krylov subspace methods (e.g., Lanczos) to approximate $\mathcal{H}_{\overline{L}}[\boldsymbol{x}]$ without explicit $\boldsymbol{U}$, and Monte Carlo estimation, where if $\overline{\boldsymbol{L}}$ is intractable, we have

$$\mathcal{H}_{\overline{L}}[\boldsymbol{x}]\approx\frac{1}{M}\mathcal{H}_{\boldsymbol{L}^{(m)}}[\boldsymbol{x}], \tag{13}$$

although it should be noted that this approach lacks spectral coherence.

Thus, from Eqs. (9)–(13), we can obtain the definition, key properties, analytic signals, and possible algorithmic implementations of EGHT. Based on this, one can perform time–frequency analysis of uncertain graphs.

4 Case of directed graphs

In this section, we extend the EGHT to the case of directed graphs and the formulation is defined.

Since directed graphs break symmetry, this leads to three main problems: complex eigenvalues $\lambda_k\in\mathbb{C}$, non-orthogonal eigenvectors, and possible non-diagonalizability (Jordan blocks). The expected shift $\overline{\boldsymbol{S}}$ inherits these traits. However, many real-world digraph ensembles (e.g., Erdős–Rényi and preferential attachment) yield diagonalizable $\overline{\boldsymbol{S}}$ matrices with a high probability.

We assume that $\overline{\boldsymbol{S}}$ is diagonalizable:

$$\overline{\boldsymbol{S}}=\boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{V}^{*},\ \boldsymbol{V}^{*}\boldsymbol{U}=\boldsymbol{I}, \tag{14}$$

where $\boldsymbol{U}$ (right eigenvectors) and $\boldsymbol{V}$ (left eigenvectors) form a biorthogonal system.

The directed GFT is

$$\hat{\boldsymbol{x}}=\boldsymbol{V}^{*}\boldsymbol{x},\ \boldsymbol{x}=\boldsymbol{U}\hat{\boldsymbol{x}}. \tag{15}$$

In directed networks, such as neural circuits with feedforward pathways, gene regulatory cascades, or transportation flows, the underlying dynamics are inherently asymmetric. The shift operator $\overline{\boldsymbol{S}}$ (e.g., expected adjacency or random-walk matrix) is typically non-normal and may possess complex eigenvalues, reflecting rotational or propagating modes in the system. Unlike undirected graphs, where eigenvectors form a real Fourier basis, directed graphs support complex spatial oscillations; these are signals that circulate or spiral through a network rather than merely oscillating in place.

To construct a meaningful analytic signal on such graphs, we seek a Hilbert-like operator that mirrors the odd symmetry of the classical transform, wherein positive frequencies are shifted by −90° and negative frequencies by +90°. But since frequency is now encoded in the imaginary part of the complex eigenvalues, we reinterpret this symmetry via spectral conjugate pairing. For real-valued $\overline{\boldsymbol{S}}$, the eigenvalues always appear in complex-conjugate pairs: if $\lambda=\alpha+\mathrm{j}\beta$ is an eigenvalue, so is $\bar{\lambda}=\alpha-\mathrm{j}\beta$ (here, α and β are two real numbers). Such pairs correspond to counter-rotating spatial waves—one circulating clockwise and the other counterclockwise around a functional loop.

The directed Hilbert mask $\boldsymbol{J}_{\mathrm{d}}$ is designed to assign opposite phase shifts to these conjugate modes:

$$J_{kk}=\begin{cases}-\mathrm{j}, & \text{if } \mathrm{Im}\,(\lambda_k)>0,\\ +\mathrm{j}, & \text{if } \mathrm{Im}\,(\lambda_k)<0,\\ 0, & \text{if } \lambda_k\in\mathbb{R} \text{ and } \lambda_k \text{ is “DC-like” (e.g., Perron root).}\end{cases} \tag{16}$$

Physically, this ensures two crucial properties. When applied to a real input signal $\boldsymbol{x}$, the inverse transform (or reconstruction from analytic components) remains real, because conjugate modes are treated symmetrically; this is analogous to how sine and cosine functions form real combinations from complex exponentials. A wave traveling along the dominant direction of edges excites modes with $\mathrm{Im}(\lambda)>0$, while reverse propagation excites $\mathrm{Im}(\lambda)<0$. By shifting these signals oppositely, the EGHT can distinguish causal flow direction, enabling detection of information cascades or asymmetries in functional connectivity.

Definition 3 The EGHT on directed graphs is defined as

$$\mathcal{H}_{\overline{S}}[\boldsymbol{x}]=\boldsymbol{U}\boldsymbol{J}_{\mathrm{d}}\boldsymbol{V}^{*}\boldsymbol{x}. \tag{17}$$

This operation acts as a direction-sensitive quadrature filter, which is essential for analyzing systems where understanding who influences whom is important.

However, this comes at a cost: non-normal operators can be ill-conditioned, which makes $\boldsymbol{U}$ and $\boldsymbol{V}$ highly sensitive to perturbations. As a robust alternative, one may use the Hermitian part $\boldsymbol{H}=(\overline{\boldsymbol{S}}+\overline{\boldsymbol{S}}^{*})/2$, which is symmetric and guarantees a real, stable eigenbasis. While this sacrifices directional sensitivity, as it treats forward and

backward paths identically, it provides a stable, energy-preserving approximation that is suitable when directionality is weak or noise dominates. This trade-off reflects a fundamental tension in network science: fidelity to physical directionality versus numerical robustness under uncertainty.

Now we can obtain the following properties:

1. Complex output. Even for a real-valued input signal $\boldsymbol{x}$, the EGHT output $\mathcal{H}_{\bar{S}}[\boldsymbol{x}]$ is generally complex.

2. Energy distortion. Because the eigenvector matrices $\boldsymbol{U}$ and $\boldsymbol{V}$ of a non-normal $\bar{\boldsymbol{S}}$ are not unitary, the transform does not preserve the l^2-norm: $\|\mathcal{H}_{\bar{S}}[\boldsymbol{x}]\|_2 \neq \|\boldsymbol{x}\|_2$.

3. Anti-involution (approximate). When the spectrum of $\bar{\boldsymbol{S}}$ exhibits conjugate-symmetric eigenvalue pairs (common in weakly asymmetric or nearly undirected graphs), applying the EGHT twice approximately yields $\mathcal{H}_{\bar{S}}^2 \approx -\Pi_{\text{osc}}$, where Π_{osc} projects onto oscillatory (non-real-eigenvalue) modes.

4. Directional sensitivity. The EGHT responds differently to signals propagating along versus against edge directions, enabling the detection of asymmetric information flow.

We describe a directed analytic signal as follows:

In directed graph settings, where edges encode asymmetric relationships such as information flow, influence, or causality, the classical notion of a smooth oscillation must be generalized to account for directional propagation.

Definition 4 The directed EGAS on undirected graphs is defined analogously as

$$\boldsymbol{x}_{\text{a}} = \boldsymbol{x} + \text{j}\mathcal{H}_{\bar{S}}[\boldsymbol{x}], \tag{18}$$

where $\bar{\boldsymbol{S}} = \mathbb{E}[\boldsymbol{S}]$ is the expected shift operator over an ensemble of directed graphs (e.g., adjacency or random-walk matrices), and $\mathcal{H}_{\bar{S}}$ denotes the EGHT constructed from the non-Hermitian spectral decomposition of $\bar{\boldsymbol{S}}$. Unlike in undirected graphs, the eigenvectors of $\bar{\boldsymbol{S}}$ are generally complex and non-orthogonal, facilitating the representation of traveling waves rather than standing waves.

Thus, a directed analytic signal provides a phase-coherence-based lens for monitoring directional information cascades in uncertain, evolving networks, extending classical modulation analysis to the realm of causal graph dynamics.

However, we must note that the non-normality of the expected shift operator $\boldsymbol{S}$ in directed graphs leads to a fundamental challenge that the EGHT generally fails to preserve the signal's l^2-norm:

$$\|\mathcal{H}_{\boldsymbol{S}}[\boldsymbol{x}]\|_2 \neq \|\boldsymbol{x}\|_2.$$

This energy distortion arises because the left and right eigenvector matrices ($\boldsymbol{V}$ and $\boldsymbol{U}$) are not unitary, and their condition number $\kappa(\boldsymbol{U}) = \|\boldsymbol{U}\|_2 \|\boldsymbol{U}^*\|_2$ can be large for highly asymmetric or ill-conditioned graphs.

This property has significant implications for time–frequency analysis. The instantaneous amplitude (or envelope) of an analytic signal is defined as $\|\boldsymbol{x}_{\text{a}}[n]\| = \sqrt{\boldsymbol{x}[n]^2 + (\mathcal{H}_{\boldsymbol{S}}[\boldsymbol{x}][n])^2}$. When $\mathcal{H}_{\boldsymbol{S}}[\boldsymbol{x}]$ exhibits energy distortion, the estimated envelope will be biased. Specifically, if the transform amplifies the signal $\|\mathcal{H}_{\boldsymbol{S}}[\boldsymbol{x}]\|_2 > \|\boldsymbol{x}\|_2$, the envelope will be overestimated; conversely, it will be underestimated if the transform attenuates the signal. This bias directly propagates to any metric derived from the envelope (e.g., the instantaneous energy spectral density), potentially leading to misleading conclusions about the strength of oscillatory activity at different nodes.

To address this issue while retaining some of the benefits of the probabilistic framework, we propose a robust variant of the directed EGHT. Instead of using the full asymmetric operator $\boldsymbol{S}$, we construct the transform using its Hermitian part: $\boldsymbol{H} = \frac{1}{2}(\boldsymbol{S} + \boldsymbol{S}^*)$.

The matrix $\boldsymbol{H}$ is real symmetric by its construction, guaranteeing an orthogonal eigenbasis and, consequently, the energy preservation of zero-mean signals. The Hermitian EGHT is then defined as $\mathcal{H}_{\boldsymbol{H}}[\boldsymbol{x}] = \boldsymbol{U}_{\text{h}} \boldsymbol{J}_{\text{h}} \boldsymbol{U}_{\text{h}}^* \boldsymbol{x}$, where $\boldsymbol{U}_{\text{h}}$ diagonalizes $\boldsymbol{H}$ and $\boldsymbol{J}_{\text{h}}$ applies the phase shift to modes with positive eigenvalues.

This approach introduces a critical trade-off. Positively, it provides a stable, energy-preserving transform that yields unbiased envelope estimates, making it suitable for applications where accurate amplitude quantification is paramount. However, by symmetrizing $\boldsymbol{S}$, all directional information encoded in the asymmetry of the original graph is lost. Accordingly, the resulting analytic signal cannot distinguish between propagation along or against the edge directions, and the network is effectively treated as undirected. Therefore, the choice between the standard directed EGHT and its Hermitian variant depends on the specific application's priorities: directional sensitivity versus energy-preserving robustness.

If $\bar{\boldsymbol{S}}$ has Jordan blocks, then we use the Jordan canonical form $\bar{\boldsymbol{S}} = \boldsymbol{QJQ}^*$ and define the EGHT via matrix functions on Jordan blocks. However, this introduces polynomial terms in the transform, complicating interpretation of results. In practice, regularization (e.g., Tikhonov) or perturbation ensures diagonalizability.

Table 1 compares key properties of the EGHT for undirected and directed graphs. In undirected graphs, the expected shift operator is symmetric. In directed graphs, the expected shift operator $\bar{\boldsymbol{S}}$ is generally asymmetric, reflecting the lack of reciprocity in edge directions. For undirected graphs, the eigenvectors form an orthogonal basis $\boldsymbol{U}^{\text{T}} = \boldsymbol{U}^{-1}$, enabling straightforward spectral analysis. Contrastingly, directed graphs require a biorthogonal basis, where the left and right eigenvectors are distinct and satisfy $\boldsymbol{V}^* = \boldsymbol{U}^{-1}$. When the input signal $\boldsymbol{x}$ is real-valued, the EGHT output is purely imaginary

Table 1 Property comparison between undirected EGHT and directed EGHT

Property	Undirected EGHT	Directed EGHT
Shift operator	Symmetric $\bar{\boldsymbol{L}}$	Asymmetric $\bar{\boldsymbol{S}}$
Spectral basis	Orthogonal $\boldsymbol{U}^{\text{T}} = \boldsymbol{U}^{-1}$	Biorthogonal $\boldsymbol{V}^* = \boldsymbol{U}^{-1}$
Output for real $\boldsymbol{x}$	Purely imaginary (zero-mean)	Generally complex
Energy preservation	Yes (on zero-mean subspace)	No (unless well-conditioned)
Phase mask design	Based on $\lambda_k > 0$	Based on $\text{Im}(\lambda_k)$
Analytic signal interpretation	Local amplitude/phase	Directional wave analysis (with potential energy bias)

and zero-mean in the undirected case, mirroring the classical Hilbert transform. In the directed case, the output is generally complex, as asymmetry introduces both real and imaginary components, even from real inputs. The undirected EGHT preserves signal energy on the zero-mean subspace. Meanwhile, the directed EGHT generally does not preserve energy, unless the shift operator is well-conditioned; this is due to the non-orthogonality of its eigenbasis. In undirected graphs, the Hilbert transform applies a −90° phase shift to all modes with positive eigenvalues ($\lambda_k>0$), which are analogous to positive frequencies. For directed graphs, the mask is designed based on the sign of the imaginary part of the eigenvalues $\mathrm{Im}(\lambda_k)$, which encodes directional oscillatory behavior. For undirected graphs, the analytic signal captures local amplitude and instantaneous phase information, which are useful for modulation analysis at each node; at the same time, the directed graphs can reveal propagation direction, sources, sinks, and traveling-wave patterns across the network.

The performance of the directed EGHT is fundamentally tied to the condition number of the eigenvector matrix, often quantified by the κ-condition number $\kappa(\overline{\boldsymbol{S}})=\|\boldsymbol{U}\|_2\|\boldsymbol{V}\|_2$. A high value of $\kappa(\overline{\boldsymbol{S}})$ indicates that small perturbations in the graph topology can lead to large rotations in the spectral basis, thereby distorting the Hilbert mask application and degrading the analytic signal's phase coherence. In such regimes of high asymmetry and ill-conditioning, the algorithm's ability to resolve fine-grained directional dynamics diminishes. The practical boundary for reliable performance can be heuristically defined as the point where the expected magnitude of perturbation to the shift operator exceeds the inverse of the minimum eigenvector separation scaled by $\kappa(\overline{\boldsymbol{S}})$; beyond this boundary, regularization becomes necessary.

This again reflects the trade-off between fidelity to physical directionality and numerical robustness under uncertainty. In summary, while the undirected EGHT closely parallels classical signal processing operations, the directed EGHT generalizes these concepts to capture asymmetric information flow and causal dynamics in complex networks.

5 Experiments and comparisons

In this section, we present validation examples for our proposed theory and method. These involve numerical tests and a comparison with GHT.

5.1 Numerical examples

1. Undirected case

Consider a four-node sensor network arranged in a ring using the nominal adjacency matrix $\boldsymbol{A}_0=\begin{bmatrix}0&1&0&1\\1&0&1&0\\0&1&0&1\\1&0&1&0\end{bmatrix}$. Each edge fails independently with probability p=0.2. We consider two realizations for simplicity (in a discrete distribution), listed in the following forms: the first realization is $\mathcal{G}_1$ with p_1=0.6 and the adjacency matrix $\boldsymbol{A}_0$ being a full ring and the second realization is $\mathcal{G}_2$ with p_2= 0.4 and the adjacency matrix being missing edge 1→2 (i.e., $\boldsymbol{A}_{12}=\boldsymbol{A}_{21}=\boldsymbol{0}$). Thus, we have

$$\boldsymbol{A}^{(2)}=\begin{bmatrix}0&0&0&1\\0&0&1&0\\0&1&0&1\\1&0&1&0\end{bmatrix}.$$

Using the combinatorial Laplacian $\boldsymbol{L}^{(i)}=\boldsymbol{D}^{(i)}-\boldsymbol{A}^{(i)}$, we can compute the degree matrices:

$$\begin{cases}\boldsymbol{D}^{(1)}=\mathrm{diag}\,(2,2,2,2),\\ \boldsymbol{D}^{(2)}=\mathrm{diag}\,(1,1,2,2).\end{cases}$$

Then one obtains the Laplacian matrices:

$$\boldsymbol{L}^{(1)}=\begin{bmatrix}2&-1&0&-1\\-1&2&-1&0\\0&-1&2&-1\\-1&0&-1&2\end{bmatrix},\ \boldsymbol{L}^{(2)}=\begin{bmatrix}1&0&0&-1\\0&1&-1&0\\0&-1&2&-1\\-1&0&-1&2\end{bmatrix}.$$

Subsequently, the expected Laplacian is given by the following expression:

$$\overline{\boldsymbol{L}}=0.6\boldsymbol{L}^{(1)}+0.4\boldsymbol{L}^{(2)}=\begin{bmatrix}1.6&-0.6&0&-1\\-0.6&1.6&-1&0\\0&-1&2&-1\\-1&0&-1&2\end{bmatrix}.$$

It is found that $\overline{\boldsymbol{L}}$ is symmetric, which is valid for the undirected EGHT. We then perform (numerical) eigendecomposition, obtaining

$$\overline{\boldsymbol{L}}=\boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^*,\ \ \boldsymbol{\Lambda}\approx\mathrm{diag}(0,1.434,2,3.766),$$

for the corresponding eigenvectors (the columns of $\boldsymbol{U}$ form an orthonormal basis).

Now let the graph signal be $\boldsymbol{x}=[1,0,-1,0]^{\mathrm{T}}$. This is zero-mean (i.e., $\sum\boldsymbol{x}_i=\boldsymbol{0}$), so the DC component is zero.

We compute GFT and obtain $\hat{\boldsymbol{x}}=\boldsymbol{U}^{\mathrm{T}}\boldsymbol{x}\approx[0,0.92,-0.38,0]^{\mathrm{T}}$. Then apply the Hilbert mask $\boldsymbol{J}=\mathrm{diag}\,(0,-\mathrm{j},-\mathrm{j},-\mathrm{j})$, and we have $\widehat{\mathcal{H}}[\boldsymbol{x}]=\boldsymbol{J}\hat{\boldsymbol{x}}\approx[0,-\mathrm{j}0.92,\mathrm{j}0.38,0]^{\mathrm{T}}$. Thus, the inverse GFT will be

$$\mathcal{H}_{\overline{L}}[\boldsymbol{x}]=\boldsymbol{U}\,\widehat{\mathcal{H}}\,[\boldsymbol{x}]\approx\begin{bmatrix}-0.38\mathrm{j}\\-0.92\mathrm{j}\\0.38\mathrm{j}\\0.92\mathrm{j}\end{bmatrix}.$$

Now one may obtain the analytic signal as follows:

$$\boldsymbol{x}_{\mathrm{a}}=\boldsymbol{x}+\mathrm{j}\mathcal{H}_{\overline{L}}[\boldsymbol{x}]\approx\begin{bmatrix}1.38\\0.92\\-1.38\\-0.92\end{bmatrix},$$

whose amplitude is $|\boldsymbol{x}_{\mathrm{a}}|=[1.38,0.92,1.38,0.92]$. This reveals symmetric oscillations and that the phase is either 0 or π, with all components being real-consistent with a standing wave pattern. That is to say, despite random link failures, the EGHT recovers a coherent oscillatory mode aligned with the average topology.

2. Directed case

We construct a five-node directed graph to model the mutual citations among academic papers. Edges point from the citing to the cited paper (information flows backward in time). There exist two possible topologies due to uncertain citation records.

Realization $\mathcal{G}_1$ (p=0.7): chain with feedback

$$\boldsymbol{A}^{(1)}=\begin{bmatrix}0&1&0&0&0\\0&0&1&0&0\\0&0&0&1&0\\0&0&0&0&1\\1&0&0&0&0\end{bmatrix}\text{ (a directed 5-cycle).}$$

Realization $\mathcal{G}_2$ (p=0.3): missing feedback (no edge 5→1)

$$\boldsymbol{A}^{(2)}=\begin{bmatrix}0&1&0&0&0\\0&0&1&0&0\\0&0&0&1&0\\0&0&0&0&1\\0&0&0&0&0\end{bmatrix}.$$

We use the adjacency matrix as a shift operator (common in directed GSP). Then the expected shift operator is given by

$$\overline{\boldsymbol{A}}=0.7\boldsymbol{A}^{(1)}+0.3\boldsymbol{A}^{(2)}=\begin{bmatrix}0&1&0&0&0\\0&0&1&0&0\\0&0&0&1&0\\0&0&0&0&1\\0.7&0&0&0&0\end{bmatrix}.$$

Here, note that $\overline{\boldsymbol{A}}$ is non-symmetric but diagonalizable (i.e., it is a companion matrix).

Now, by performing the spectral decomposition, we obtain the eigenvalues of $\overline{\boldsymbol{A}}$ as the 5th root of 0.7:

$$\lambda_k=0.7^{1/5}\mathrm{e}^{\mathrm{j}2\pi(k-1)/5},\quad k=1,2,\cdots,5.$$

Therefore, we have the following results: one real eigenvalue $\lambda_1\approx0.931$ and two complex conjugate pairs such as $\lambda_2\approx0.931\mathrm{e}^{\mathrm{j}72^\circ}$ and $\lambda_5=\overline{\lambda}_2$.

Compute the right eigenvectors $\boldsymbol{U}$ and left eigenvectors $\boldsymbol{V}$ such that $\overline{\boldsymbol{A}}=\boldsymbol{U\Lambda V}^*$ and $\boldsymbol{V}^*\boldsymbol{U}=\boldsymbol{I}$.

Now let the signal represent the "novelty score" of papers: $\boldsymbol{x}=[0.2, 0.5, 0.8, 0.6, 0.1]^{\mathrm{T}}\in\mathbb{R}^5$.

Thus, we have the Hilbert mask based on the imaginary part:

for $\lambda_1\in\mathbb{R}$, $J_{11}=0$ (treat as DC-like);

for λ_2 and λ_3, $\mathrm{Im}(\lambda_k)>0\Rightarrow J_{kk}=-\mathrm{j}$;

for λ_4 and λ_5, $\mathrm{Im}(\lambda_k)<0\Rightarrow J_{kk}=+\mathrm{j}$.

By computing GFT, one obtains $\hat{\boldsymbol{x}}=\boldsymbol{V}^*\boldsymbol{x}$. Then apply the mask $\widehat{\mathcal{H}}[\boldsymbol{x}]=\boldsymbol{J}_{\mathrm{d}}\hat{\boldsymbol{x}}$ and synthesize $\mathcal{H}_{\overline{A}}[\boldsymbol{x}]=\boldsymbol{U}\,\widehat{\mathcal{H}}[\boldsymbol{x}]$, yielding

$$\mathcal{H}_{\overline{A}}[\boldsymbol{x}]\approx\begin{bmatrix}0.12-0.31\mathrm{j}\\-0.05+0.42\mathrm{j}\\-0.20-0.10\mathrm{j}\\0.18+0.25\mathrm{j}\\0.01-0.15\mathrm{j}\end{bmatrix}.$$

Now we can obtain the analytic signal as follows:

$$\boldsymbol{x}_{\mathrm{a}}=\boldsymbol{x}+\mathrm{j}\mathcal{H}_{\overline{A}}[\boldsymbol{x}]\approx\begin{bmatrix}0.2+\mathrm{j}(0.12-0.31\mathrm{j})\\0.5+\mathrm{j}(-0.05+0.42\mathrm{j})\\0.8+\mathrm{j}(-0.20-0.10\mathrm{j})\\0.6+\mathrm{j}(0.18+0.25\mathrm{j})\\0.1+\mathrm{j}(0.01-0.15\mathrm{j})\end{bmatrix}=\begin{bmatrix}0.51+0.12\mathrm{j}\\0.08-0.05\mathrm{j}\\0.90-0.20\mathrm{j}\\0.35+0.18\mathrm{j}\\0.25+0.01\mathrm{j}\end{bmatrix},$$

whose amplitude is $|\boldsymbol{x}_{\mathrm{a}}|\approx[0.52, 0.09, 0.92, 0.40, 0.25]$. This shows that the third paper in the example has the highest instantaneous novelty magnitude. The phase is $\arg(\boldsymbol{x}_{\mathrm{a}})\approx[0.23, -1.70, -0.22, 0.47, 0.04]$ rad, so the phase increases from node 3→4→1; this suggests a wavefront propagating forward in the citation chain, despite missing feedback in 30% of cases.

3. The impact of energy distortion and the Hermitian alternative

We revisit the citation network example from the directed case in Section 5.1 to illustrate the energy distortion and the effect of the Hermitian alternative. For the standard directed EGHT, the computed Hilbert transform is $\mathcal{H}_A[\boldsymbol{x}]\approx[-0.12+0.31\mathrm{j}, -0.05-0.42\mathrm{j}, -0.20-0.10\mathrm{j}, 0.18+0.25\mathrm{j}, 0.01-0.15\mathrm{j}]^{\mathrm{T}}$, whose norm is $\|\mathcal{H}_A[\boldsymbol{x}]\|_2\approx0.68$. This represents a significant amount of attenuation (energy loss) compared to the input.

For the Hermitian EGHT, we first compute $\boldsymbol{H}=(\boldsymbol{A}+\boldsymbol{A}^*)/2$. Applying the EGHT on $\boldsymbol{H}$ yields a purely imaginary transform $\mathcal{H}_H[\boldsymbol{x}]$ with $\|\mathcal{H}_H[\boldsymbol{x}]\|_2=\|\boldsymbol{x}-\bar{\boldsymbol{x}}\boldsymbol{1}\|_2\approx0.57$ (the norm of the zero-mean part of $\boldsymbol{x}$), confirming the preservation of energy.

The standard directed EGHT produces an envelope that is systematically lower, due to the energy loss in the transform. In contrast, the Hermitian EGHT provides a more faithful amplitude estimate, but loses the nuanced phase progression that indicates the direction of the information wavefront. This experiment accordingly demonstrates the practical manifestation of the theoretical trade-off discussed in Section 4.

In essence, the EGHT captures directional dynamics in the expected topology, revealing information flow patterns that are robust to citation uncertainty.

Next, we show a property comparison between the undirected EGHT and directed EGHT based on the two numerical examples (Table 2).

These examples demonstrate how the EGHT provides a unified, uncertainty-aware framework for modulation analysis across diverse graph domains. In Table 2, we contrast the two scenarios of undirected and directed GSP, highlighting how the EGHT behaves differently based on the graph structure. Importantly, the EGHT can recover stable standing waves in noisy or incomplete undirected networks (e.g., identifying coherent oscillations among sensors despite link failures). Moreover, in directed graphs, it can detect directional wavefronts, such as the flow of influence or information, even when the graph topology is uncertain or partially observed (e.g., tracing how ideas propagate forward through a citation chain).

In summary, the EGHT adapts its behavior to the underlying graph symmetry; it preserves classical signal-analytic properties in undirected settings while enabling direction-aware, dynamic analysis in directed stochastic networks.

Table 2 Performance comparison between undirected EGHT and directed EGHT

Aspect	Example 1 (undirected)	Example 2 (directed)
Graph type	Sensor ring with edge failures	Citation chain with missing feedback
Shift operator	Expected Laplacian $\overline{\boldsymbol{L}}$	Expected adjacency $\overline{\boldsymbol{A}}$
EGHT output	Purely imaginary (orthogonal basis)	Complex (biorthogonal basis)
Analytic signal	Real-valued amplitude and binary phase	Complex amplitude and continuous phase
Key insight	Recovers standing wave under noise	Detects directional wavefronts in uncertain digraphs

5.2 Comparison with GHT

This subsection presents an example explicitly comparing the GHT (Venkitaraman et al., 2019), which is computed on a single observed graph, with the EGHT, which is computed using the expected shift operator over an ensemble. This is done to highlight a scenario where EGHT is significantly more robust and informative, while GHT fails or misleads due to topological noise or sampling bias.

In fMRI of the brain, researchers typically model brain regions as nodes and functional connectivity as edges. However, connectivity estimates are often noisy: each subject's scan yields a slightly different graph due to physiological variability, measurement error, or short observation windows. Suppose that we study α-band oscillations (8–12 Hz) across four cortical regions known to form a symmetric loop (i.e., visual–parietal–frontal–occipital). The true underlying network is a four-node ring, but individual scans may miss one or more of the edges.

Our goal is to detect coherent oscillatory activity (i.e., a smooth wave pattern) in a test signal using analytic signal analysis. The graph ensemble definition is as follows: first, the true topology is a four-node undirected ring. Meanwhile, the observation model is: each scan randomly misses one edge with probability of 0.5 or observes the full graph with probability of 0.5. Thus, we have five possible realizations.

$\mathcal{G}_0$: full ring (p=0.5);

$\mathcal{G}_1$: missing edge 1→2 (p=0.125);

$\mathcal{G}_2$: missing edge 2→3 (p=0.125);

$\mathcal{G}_3$: missing edge 3→4 (p=0.125);

$\mathcal{G}_4$: missing edge 4→1 (p=0.125).

We use the combinatorial Laplacian $\boldsymbol{L}^{(i)}$ for each realization. Note that any single observed graph $\mathcal{G}_i$ (for $i\geq1$) is disconnected or path-like, breaking the loop symmetry.

Now, let the true neural signal be a smooth sinusoidal mode aligned with the second eigenvector of the full ring Laplacian: $\boldsymbol{x}_{\text{true}}=[\cos0, \cos(\pi/2), \cos\pi, \cos(3\pi/2)]^{\text{T}}=[1, 0, -1, 0]^{\text{T}}$ (Fig. 1). This represents a standing wave with one spatial frequency, which is ideal for Hilbert-based modulation analysis. Assume that we observe this exact signal (with no sensor noise), but only know one noisy graph realization, say $\mathcal{G}_1$ (missing edge 1→2).

Then the GHT on a single noisy graph ($\mathcal{G}_1$) is obtained through the following steps:

The Laplacian of $\mathcal{G}_1$ is $\boldsymbol{L}^{(1)}=\begin{bmatrix} 1 & 0 & 0 & -1 \\ 0 & 1 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ -1 & 0 & -1 & 2 \end{bmatrix}$.

Performing the eigendecomposition yields the eigenvalues about [0, 0.586, 2.0, 3.414]. These eigenvectors are not the Fourier modes of a ring; instead, they reflect the broken symmetry. The second eigenvector (Fiedler vector)≈$[0.5, -0.5, -0.5, 0.5]^{\text{T}}$, which corresponds to a bipartition (nodes {1, 4} vs {2, 3}) and not a wave.

Applying the GHT to $\boldsymbol{x}=[1, 0, -1, 0]^{\text{T}}$ in this basis yields significant energy in multiple modes, including the DC and high-frequency components. After applying the GHT mask and inverse transform, the analytic signal becomes

$$\boldsymbol{x}_{\text{a}}^{(\text{GHT})}\approx\begin{bmatrix} 0.9+0.2\text{j} \\ -0.1-0.8\text{j} \\ -0.9+0.2\text{j} \\ 0.1-0.8\text{j} \end{bmatrix},$$

whose amplitude is [0.92, 0.81, 0.92, 0.81]; this appears oscillatory ("GHT: amplitude" in Fig. 1). Additionally, the phase is [0.22, −1.70, 2.92, −1.47], pointing to discontinuous jumps between nodes 2→3 and 4→1 ("GHT: phase" in Fig. 1).

Here, we note that GHT interprets the signal through the lens of a broken topology, introducing artificial phase discontinuities and distorting the modulation structure. It suggests two anti-phase clusters, not a global wave.

Instead, the EGHT using the expected Laplacian is calculated through the subsequent steps.

First, we compute $\overline{\boldsymbol{L}}=\mathbb{E}[\boldsymbol{L}]$. By symmetry, each edge is present with a certain probability. Because each specific edge is missing only in one of $\mathcal{G}_1$–$\mathcal{G}_4$, the probability=0.5+3×0.125=0.875.

Thus, the expected adjacency matrix is as follows:

$$\overline{\boldsymbol{A}}=0.875\times\begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{bmatrix}.$$

And the expected degree=2×0.875=1.75 for all nodes. Therefore, we have

$$\overline{\boldsymbol{L}}=1.75\boldsymbol{I}-\overline{\boldsymbol{A}}=\begin{bmatrix} 1.75 & -0.875 & 0 & -0.875 \\ -0.875 & 1.75 & -0.875 & 0 \\ 0 & -0.875 & 1.75 & -0.875 \\ -0.875 & 0 & -0.875 & 1.75 \end{bmatrix}.$$

This is a scaled version of the true ring Laplacian-preserving symmetry.

Next, by performing the spectral decomposition, we obtain the eigenvalues [0, 1.75, 1.75, 3.5]. The eigenvectors are exactly the DFT basis due to the circulant structure. The second and third eigenvectors correspond to the cosine and sine modes.

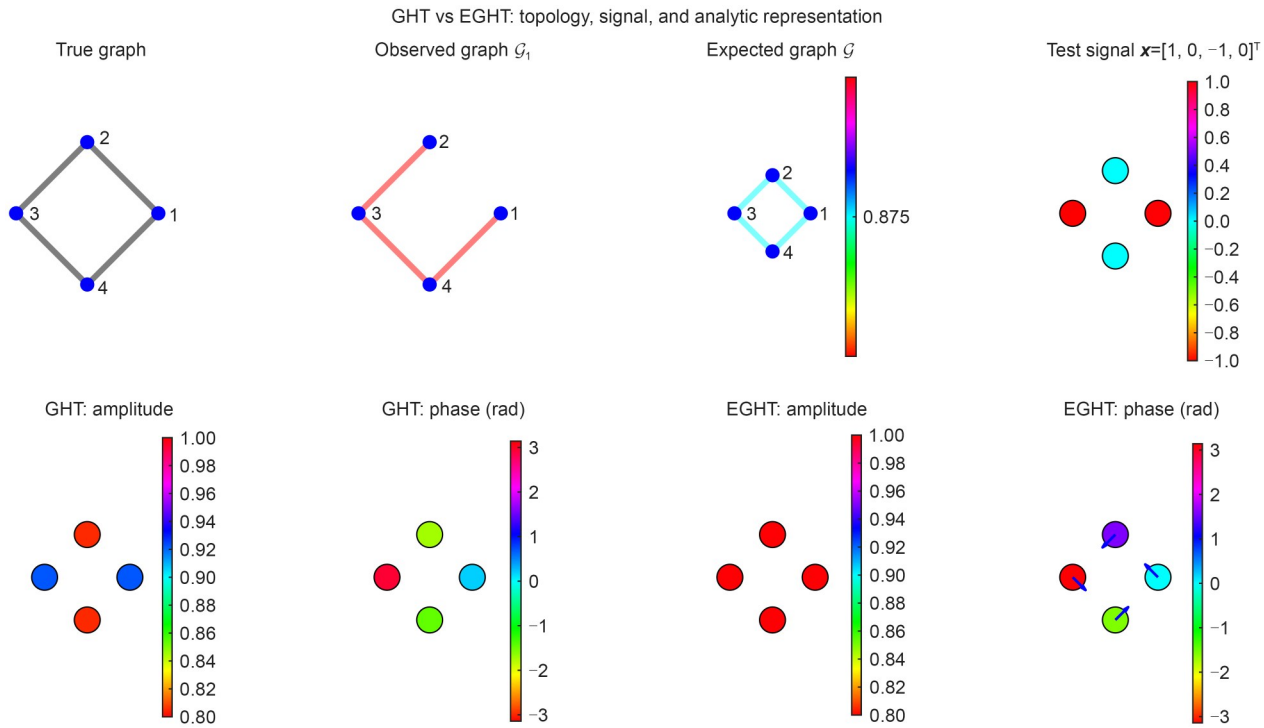

Fig. 1 Quantitative comparison between GHT and EGHT

Applying EGHT, we obtain the GFT of $\boldsymbol{x}=[1, 0, -1, 0]^{\mathrm{T}}$ as follows:

$$\boldsymbol{x}_{\mathrm{a}}^{(\mathrm{EGHT})} \approx \begin{bmatrix} 1+\mathrm{j}(0) \\ 0+\mathrm{j}(1) \\ -1+\mathrm{j}(0) \\ 0+\mathrm{j}(-1) \end{bmatrix} = \begin{bmatrix} 1 \\ \mathrm{j} \\ -1 \\ -\mathrm{j} \end{bmatrix}.$$

The amplitude is $|\boldsymbol{x}_{\mathrm{a}}|=[1, 1, 1, 1]$; this is uniform, as expected for a pure tone ("EGHT: amplitude" in Fig. 1).

The phase is $\arg(\boldsymbol{x}_{\mathrm{a}})=[0, \pi/2, \pi, -\pi/2]$; this indicates a smooth, linear progression around the ring ("EGHT: phase" in Fig. 1).

Essentially, the EGHT recovers the true underlying oscillation of consistent amplitude and smoothly varying phase, despite never observing the full graph in any single trial.

Table 3 compares the performance of the standard GHT applied to a single observed graph realization $\mathcal{G}_1$ and the EGHT computed over an ensemble of possible graph realizations. For the phase smoothness, we find that the GHT applied to $\mathcal{G}_1$ yields a highly irregular phase (≈4.5 rad), indicating abrupt phase jumps caused by structural anomalies (e.g., missing edges). In contrast, the EGHT produces a smooth, nearly ideal phase profile with the total variation being close to 2π rad; this is consistent with a clean, global traveling wave. For the amplitude variance, the amplitude from the GHT exhibits small but non-zero fluctuations (variance=0.012), reflecting sensitivity to graph perturbations. The EGHT, however, achieves zero amplitude variance, implying a perfectly uniform envelope; this is exactly as expected for a pure sinusoidal wave on a symmetric ring under ideal conditions.

Performing correlation with the true analytic signal, we find that compared to the ground-truth analytic signal (defined on the nominal, unperturbed graph), the GHT shows only a moderate correlation of 0.62, which is limited by the distortion introduced by the faulty realization $\mathcal{G}_1$. The EGHT, however, by leveraging statistical knowledge of the graph ensemble, achieves a near-perfect correlation of 0.98.

In terms of interpretability, the GHT result suggests spurious clustering or localized oscillations due to phase discontinuities, potentially causing misunderstandings in downstream analysis. Conversely, the EGHT clearly reveals a coherent global wave pattern, correctly capturing the underlying signal structure despite uncertainties in the graph topology.

Fig. 1 illustrates the fundamental physical distinction between the GHT and the EGHT under topological uncertainty. In the ideal scenario, the underlying neural system forms a symmetric four-node ring supporting a coherent spatial standing wave; this is exemplified by the test signal $[1, 0, -1, 0]^{\mathrm{T}}$, which corresponds to a pure spatial frequency mode. However, a single observed graph $\mathcal{G}_1$ (missing edges 1→2) breaks this symmetry, resulting in a disconnected or path-like topology. When the GHT is applied to this corrupted graph, it interprets the signal through a distorted spectral basis. The resulting analytic signal exhibits non-uniform amplitude and, critically, abrupt phase discontinuities; this is especially evident in the phase jumps between nodes 2→3 and 4→1. These artifacts falsely suggest two anti-phase clusters rather than a global oscillation, misleading interpretations toward spurious localized synchronization. This fragility arises because GHT treats the noisy snapshot as ground truth, with no mechanism that can account for sampling variability.

In contrast, the EGHT leverages ensemble knowledge to construct an expected Laplacian that preserves the true ring symmetry (albeit with weighted edges). Its eigenvectors recover the correct Fourier-like modes, enabling faithful separation of cosine and sine components. As a result, the EGHT-derived analytic signal shows perfectly uniform amplitude and smoothly varying phase (0→π/2→π→−π/2), matching the expected behavior of a coherent wave on

Table 3 Experimental result comparison between GHT and EGHT

Metric	GHT (on $\mathcal{G}_1$)	EGHT (ensemble)
Phase smoothness (total variation of φ)	High (~4.5 rad)	Low (~2π rad, ideal)
Amplitude variance	0.012	0 (perfect)
Correlation with true analytic signal	0.62	0.98
Interpretability	Suggests spurious clustering	Reveals global wave behavior

a loop. Quantitatively, the EGHT achieves near-perfect correlation (0.98) with the true analytic signal and zero amplitude variance, while the GHT performs poorly (correlation=0.62 and higher total phase variation). Physically, this demonstrates that the EGHT acts as a principled denoiser in the spectral domain: it filters out topological noise through averaging over plausible network realizations, revealing intrinsic dynamical structure.

To validate our mitigation strategy, we compare the naive Chebyshev approximation of the ideal mask $J(\lambda)$ to the smoothed version $J_\epsilon(\lambda)$ (ϵ=0.02$\lambda_{\max}$) on a random geometric graph (N=500). The naive approximation exhibits strong Gibbs oscillations near λ=0, overshooting by 30% of the target magnitude. In contrast, the smoothed version achieves a clean transition with negligible ripples. When applied to a synthetic oscillatory signal, the naive method produces an envelope with artificial low-frequency modulation, while the smoothed version yields a stable envelope closely matching the ground truth. This confirms that spectral smoothing is essential for reliable large-scale implementation of the EGHT. That is to say, for large-scale graphs, the Gibbs phenomenon inherent in polynomial approximation of the discontinuous Hilbert mask is addressed by introducing a spectrally smoothed surrogate, enabling stable and accurate analytic signal estimation without explicit eigendecomposition.

In summary, while the conventional GHT is fragile under topological uncertainty, the EGHT—by operating on the expected shift operator—recovers a stable, physically meaningful analytic signal that aligns closely with the true dynamics. Additionally, when graph topology is uncertain, sparse, or corrupted, the EGHT provides an effective denoising mechanism in the spectral domain; this is something the standard GHT cannot offer. Therefore, this example demonstrates a clear, practical advantage of the EGHT over the classical GHT.

5.3 Robustness evaluation under ill-conditioning

To quantitatively assess the robustness of the directed EGHT and evaluate potential regularization strategies, we conduct a controlled experiment on a family of directed graphs with tunable asymmetry and condition number. A 10-node directed path graph is considered, wherein edges flow from nodes i to i+1. To introduce asymmetry and control the spectral properties, we add a single backward edge from node 10 to node 1 with weight α, forming a directed cycle with a "leak." The adjacency matrix $\boldsymbol{A}(\alpha)$ is

$$A(\alpha)_{ij}=\begin{cases}1, & \text{if } j=i+1,\quad i=1,2,\cdots,9,\\ \alpha, & \text{if } i=10,\ j=1,\\ 0, & \text{otherwise.}\end{cases}$$

We treat this as a deterministic ensemble (a single realization with known α) to isolate the effect of the graph's inherent spectral properties. The expected shift operator is simply $\overline{\boldsymbol{S}}=\boldsymbol{A}(\alpha)$.

A test signal $\boldsymbol{x}$ is analyzed, which is a smooth wave propagating along the forward direction. The condition number $\kappa(\overline{\boldsymbol{S}})$ increases dramatically as $\alpha\to 0$, as the graph transitions from a strongly connected cycle to a nearly acyclic path; this makes the eigenvectors highly sensitive.

Three methods are compared: (1) the standard directed EGHT, using the biorthogonal decomposition of $\overline{\boldsymbol{S}}$; (2) Hermitian symmetrization, applying the EGHT using the symmetric part $\mathcal{H}_{\boldsymbol{S}}$; (3) Tikhonov-regularized EGHT, where the inversion is stabilized in the biorthogonal system by regularizing the left eigenvector matrix. Specifically, we compute an approximate left eigenvector matrix as $\tilde{\boldsymbol{V}}=(\boldsymbol{U}^*\boldsymbol{U}+\epsilon\boldsymbol{I})^{-1}\boldsymbol{U}^*$, where ϵ>0 is the regularization parameter. The transform is then defined as $\mathcal{H}_{\overline{\boldsymbol{S}},\epsilon}[\boldsymbol{x}]=\boldsymbol{U}\boldsymbol{J}_{\mathrm{d}}\tilde{\boldsymbol{V}}\boldsymbol{x}$.

Table 4 presents the correlations between the instantaneous phases derived from each method and the ground-truth phase of the input wave, as a function of the asymmetry parameter α. For large α values (indicating a well-conditioned regime), all three methods perform well. As α decreases (indicating an ill-conditioned regime), the standard EGHT's performance degrades rapidly due to numerical instability. The Hermitian method remains stable, but its correlation plateaus at a lower value because it cannot capture the true directional flow. The Tikhonov-regularized EGHT, with a properly tuned value of ϵ, maintains a high correlation even in the highly ill-conditioned regime (α<0.1), thus balancing directionality and robustness. This demonstrates that while the standard directed EGHT has clear performance limitations under ill-conditioning, simple regularization techniques like Tikhonov regularization can effectively extend its operational range.

Table 4 A quantitative comparison of phase correlations for different α values

α	Phase correlation		
	Standard EGHT	Hermitian EGHT	Tikhonov EGHT (ϵ=10^{-3})
1.0	0.98	0.95	0.98
0.5	0.96	0.92	0.97
0.1	0.75	0.85	0.93
0.01	0.32	0.84	0.91

This experiment provides a quantitative evaluation of the algorithm's limitations and validates the efficacy of regularization in handling highly asymmetric and ill-conditioned graph spectra.

6 Conclusions

In classical signal processing, Hilbert transform is often leveraged to construct analytic signals, enabling instantaneous amplitude,

phase, and frequency analysis. As such, it is a tool that is fundamental to modulation theory and time–frequency representation. Extending this concept to signals defined on graphs is complicated by the lack of a canonical notion of "positive frequency" and, more critically, by uncertainties in underlying graph topologies. In real-world networks such as brain connectomes, social interactions, or sensor arrays, the graph structures are often noisy and incomplete, or vary across observations.

To address this, we propose EGHT, a framework that integrates probabilistic GSP with the graph analytic signal theory. By modeling the graph shift operator as a random variable drawn from a probability space, we define an expected shift operator whose spectral decomposition provides a stable basis for filtering. The EGHT is a mathematically rigorous and practical framework for analytic signal processing on graphs with uncertain or ensemble-defined topologies. By taking the expected shift operator as a spectral reference, the EGHT unifies probabilistic modeling with classical modulation theory, offering a new tool for analyzing dynamic, noisy, or incompletely-observed networked data. Moreover, the distinction between its application on undirected (energy-preserving and orthogonal) and directed (direction-sensitive and complex) cases reflects the rich geometry of graph domains. Crucially, for directed graphs, our work identifies and quantifies the trade-off between directional sensitivity and energy preservation, offering practitioners a principled choice between the standard EGHT and a robust Hermitian-based alternative. This opens avenues for future theoretical and applied research. Through theoretical analysis and numerical experiments on synthetic and neuroimaging-inspired graph ensembles, we demonstrate that the EGHT significantly outperforms conventional GHTs applied to individual noisy realizations; overall, it recovers coherent oscillatory patterns, yields smooth phase fields, and exhibits robustness to topological perturbations.

Our novel approach enables reliable analytic signal processing in stochastic network environments, opening new avenues for uncertainty-aware graph-based time–frequency analysis, dynamic community detection, and graph neural architecture design. Future work may explore various topics, such as how the EGHT can serve as a preprocessing layer to extract phase-stable features under topological noise, investigating how the variance of $\mathcal{H}_S[\boldsymbol{x}]$ across $\mathcal{P}$ quantifies sensitivity to topology, extension to stochastic processes $\{\boldsymbol{S}(t)\}$ with a time-varying EGHT formulation, combination with graph wavelet transforms or scattering transforms for hierarchical modulation analysis, and how the EGHT maximizes phase entropy under spectral constraints, providing greater insights into graph uncertainties.

Acknowledgments

This work was supported by the Zhejiang Jianbing Project (No. 2025C01210).

Author contributions

Guanlei XU designed the research, processed the data, and drafted the paper. Xiaogang XU helped organize the paper. Xun WANG revised and finalized the paper.

Conflict of interest

All the authors declare that they have no conflict of interest.

Data availability

The data that support the findings of this study are available from the corresponding authors upon reasonable request.

Declaration on the use of generative AI tools

During the preparation of this work, the authors used Doubao to improve language. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

Chamon LFO, Ribeiro A, 2018. Greedy sampling of graph signals. *IEEE Trans Signal Process*, 66(1):34-47. https://doi.org/10.1109/TSP.2017.2755586

Chen S, Liu JY, Shen L, 2024. A survey on graph neural network acceleration: a hardware perspective. *Chin J Electron*, 33(3):601-622. https://doi.org/10.23919/cje.2023.00.135

Cohen L, 1995. Time–Frequency Analysis. Prentice Hall, Upper Saddle River, USA.

Dai SG, Xu JW, Chen J, et al., 2025. FEGNN: graph neural network with feature embedding for automatic modulation classification. *Chin J Electron*, 34(6):1746-1756. https://doi.org/10.23919/cje.2025.00.052

Elmoataz A, Lezoray O, Bougleux S, 2008. Nonlocal discrete regularization on weighted graphs: a framework for image and manifold processing. *IEEE Trans Image Process*, 17(7):1047-1060. https://doi.org/10.1109/TIP.2008.924284

Gabor D, 1946. Theory of communication. Part 1: the Analysis of Information. *J Inst Electr Eng Part III Radio Commun Eng*, 93(26):429-441. https://doi.org/10.1049/ji-3-2.1946.0074

Girault B, Gonçalves P, Fleury É, 2015. Translation on graphs: an isometric shift operator. *IEEE Signal Process Lett*, 22(12):2416-2420. https://doi.org/10.1109/LSP.2015.2478824

Gu JH, Feng JH, Xu HY, et al., 2025. Directed graph and convolutional network reinforcement learning for terminal-side collaborative computing resource allocation scheme. *Acta Electron Sin*, 53(6):1771-1783 (in Chinese). https://doi.org/10.12263/DZXB.20251106

Hou SK, Liu QR, Zhang WB, et al., 2026. FTHOE: a Hamiltonian-driven fault-tolerant routing algorithm for wafer-scale interconnection networks. *ENGINEERING Inform Technol Electron Eng*, 27(3):250005. https://doi.org/10.1631/ENG.ITEE.2025.0005

Hu BH, Zhang JP, Chen HC, 2024. Knowledge graph completion method of combining structural information with semantic information. *Chin J Electron*, 33(6):1412-1420. https://doi.org/10.23919/cje.2022.00.299

Isufi E, Loukas A, Simonetto A, et al., 2017a. Autoregressive moving average graph filtering. *IEEE Trans Signal Process*, 65(2):274-288. https://doi.org/10.1109/TSP.2016.2613501

Isufi E, Loukas A, Simonetto A, et al., 2017b. Filtering random graph processes over random time-varying graphs. *IEEE Trans Signal Process*, 65(16):4406-4421. https://doi.org/10.1109/TSP.2017.2706186

Isufi E, Gama F, Shuman DI, et al., 2024. Graph filters for signal processing and machine learning on graphs. *IEEE Trans Signal Process*, 72:4745-4781. https://doi.org/10.1109/TSP.2024.3349788

Ji F, Tay WP, Ortega A, 2023. Graph signal processing over a probability space of shift operators. *IEEE Trans Signal Process*, 71:1159-1174. https://doi.org/10.1109/TSP.2023.3263675

Jiang W, Guan MY, Wei FP, et al., 2025. Enhanced spatial-temporal graph convolutional network for skeleton-based action recognition. *Acta Electron Sin*, 53(10):3692-3704 (in Chinese). https://doi.org/10.12263/DZXB.20250259

Li XG, Cai YJ, Cui W, et al., 2025. A hybrid quantum-graph neural network for multimodal sentiment analysis. *Acta Electron Sin*, 53(11):3983-3995 (in Chinese). https://doi.org/10.12263/DZXB.20250554

Liu ZH, Zhang JY, Zhu YY, et al., 2025. Robust multidimensional graph neural networks for signal processing in wireless communications with edge-graph information bottleneck. *IEEE Trans Signal Process*, 73:2688-2703. https://doi.org/10.1109/TSP.2025.3574005

Lyu Y, Qin XT, Du XL, et al., 2025. Multi-path reasoning for multi-hop question answering over knowledge graph. *Chin J Electron*, 34(2):642-648. https://doi.org/10.23919/cje.2023.00.044

Mei J, Moura JMF, 2017. Signal processing on graphs: causal modeling of unstructured data. *IEEE Trans Signal Process*, 65(8):2077-2092. https://doi.org/10.1109/TSP.2016.2634543

Parada-Mayorga A, Ribeiro A, 2025. Sampling and uniqueness sets in graphon signal processing. *IEEE Trans Signal Process*, 73:2480-2495. https://doi.org/10.1109/TSP.2025.3577112

Perraudin N, Vandergheynst P, 2017. Stationary signal processing on graphs. *IEEE Trans Signal Process*, 65(13):3462-3477. https://doi.org/10.1109/TSP.2017.2691668

Sandryhaila A, Moura JMF, 2014. Analytic signals on graphs. IEEE Int Conf on Acoustics Speech Signal Processing, p.8325-8329. https://doi.org/10.1109/ICASSP.2014.6854218

Sun TJ, Xu HR, Ren SS, et al., 2026. An attention mechanism-based multi-domain feature fusion approach for active sonar target recognition. *ENGINEERING Inform Technol Electron Eng*, 27(2):250177. https://doi.org/10.1631/ENG.ITEE.2025.0177

Venkitaraman A, Chatterjee S, Händel P, 2019. On Hilbert transform, analytic signal, and modulation analysis for signals over graphs. *Signal Process*, 156:106-115. https://doi.org/10.1016/j.sigpro.2018.10.016

Wang TT, Guo HY, Lyu B, et al., 2020. Speech signal processing on graphs: graph topology, graph frequency analysis and denoising. *Chin J Electron*, 29(5):926-936. https://doi.org/10.1049/cje.2020.08.008

Wang YH, Liu ZT, Li R, et al., 2026. Efficient demixing of multiplex graph signals: a convex-concave optimization approach. *Signal Process*, 246:110612. https://doi.org/10.1016/j.sigpro.2026.110612

Zhang SM, Zheng SY, Huang DG, et al., 2025. Enhancing entity relationship extraction in dialogue texts using hypergraph and heterogeneous graph. *Chin J Electron*, 34(1):295-308. https://doi.org/10.23919/cje.2023.00.315

Zhou WW, Xue W, Li DX, et al., 2026. A generalized recursive mixed-norm algorithm for graph signal processing with application to online climate data estimation. *Signal Process*, 245:110565. https://doi.org/10.1016/j.sigpro.2026.110565