Raman signal extraction from CARS spectra using a learned-matrix representation of the discrete Hilbert transform

CHARLES H. CAMP JR.

Biosystems and Biomaterials Division National Institute of Standards and Technology 100 Bureau Dr., Gaithersburg, MD 20899 charles.camp@nist.gov

Abstract: Removing distortions in coherent anti-Stokes Raman scattering (CARS) spectra due to interference with the nonresonant background (NRB) is vital for quantitative analysis. Popular computational approaches, the Kramers-Kronig relation and the maximum entropy method, have demonstrated success but may generate significant errors due to peaks that extend in any part beyond the recording window. In this work, we present a learned matrix approach to the discrete Hilbert transform that is easy to implement, fast, and dramatically improves accuracy of Raman retrieval using the Kramers-Kronig approach.

1. Introduction

Coherent anti-Stokes Raman scattering (CARS) microscopy and spectroscopy have long been proposed as a method for label-free chemical analysis and imaging at orders of magnitude higher speeds than traditional Raman micro/spectroscopy and at higher spatial resolution than afforded by infrared absorption [1-3]. CARS uses coherent stimulation of molecular vibrations that enables both more efficient scattering of the probe beam and non-isotropic radiation of the output signal, which dramatically improves collection efficiency [4]. Additionally, the cogeneration of a so-called nonresonant background (NRB) acts as a stable homodyne amplifier, greatly amplifying the signal [5,6] — without which CARS would often have limited advantage over spontaneous Raman [7]. Though beneficial, the NRB distorts spectral peak shapes and contributes to the CARS signal having a quadratic polynomial dependence on analyte concentration, potentially spoiling quantitative analysis capabilities. To combat such challenges, optical methods to reduce the NRB were originally proposed; though, they reduced the overall CARS signal [5,8–10]. Later research revealed that computational methods, based on electromagnetic [11] or information theory [12], could remove the distortion from the NRB without a reduction in signal strength [13]. Additionally, these computational methods result in a signal that is linear with analyte concentration, enabling quantitative analysis of biological, chemical, and polymer samples [14-18]. Ideally, these methods require independent measurements of NRB spectra, which is not currently possible, leading to spectroscopic errors and distortions [19]; though, experimental [16] and computational [19] mitigations have been developed. Even with a known NRB, these methods suffer from artifacts when spectroscopic peaks (and "wings"/"tails") extend beyond the recorded spectroscopic window. These artifacts distort spectral lineshapes and can even generate new features with similar intensities to the analyte signals. In this work, we present a new method of performing the Hilbert transform for the Kramers-Kronig (KK) relation method of extracting Raman spectral features from CARS spectra. The Learned Discrete Hilbert Transform (LeDHT) uses a synthetically generated set of single-peak training data (with known Hilbert transforms) to learn a matrix representation of the Hilbert transform. After learning ("training" in machine learning parlance), arbitrary new input spectra can be transformed via matrix multiplication, which is computationally fast and efficient. Furthermore, even with synthetic training data, the LeDHT matrix is usable on experimental data with superior performance as will be demonstrated with CARS spectra of glycerol.

2. Theory

2.1. CARS spectroscopy and Raman signal extraction using the Kramers-Kronig relation

CARS is a third-order nonlinear optical mechanism in which the difference frequency between a "pump" photon and a "Stokes" photon is capable of stimulating a Raman (vibrational or rotational) transition from which a "probe" photon can inelastically scatter [2,5,6,20–22]. From a macroscopic and scalar perspective, the generated CARS spectral intensity, $I_{CARS}(\omega)$, may be described as [19]:

$$I_{\text{CARS}}(\omega) \propto \left| \left\{ \left[E_{\text{S}}(\omega) \star E_{\text{p}}(\omega) \right] \chi^{(3)}(\omega) \right\} \ast E_{\text{pr}}(\omega) \right|^{2} \approx \left| \tilde{C}_{\text{st}}(\omega) \right|^{2} \left| \tilde{\chi}^{(3)}(\omega) \right|^{2}$$
(1)

$$\chi^{(3)}(\omega) = \chi_{\rm r}(\omega) + \chi_{\rm nr}(\omega) \tag{2}$$

$$\tilde{C}_{st}(\omega) \equiv \frac{[E_{s}(\omega) \star E_{p}(\omega)] * E_{pr}(\omega)}{\int E_{pr}(\omega) d\omega}$$
(3)

$$\tilde{\chi}^{(3)}(\omega) \equiv \chi^{(3)}(\omega) * E_{\rm pr}(\omega), \tag{4}$$

where E_p , E_s , and E_{pr} are the pump, Stokes, and probe spectral fields (in frequency ω), respectively; $\chi^{(3)}$ is the third-order nonlinear susceptibility; and \star and * are cross-correlation and convolution operators, respectively. The third-order nonlinear susceptibility can be separated into a summation of vibrationally (or rotationally) resonant (χ_r) and nonresonant (χ_{nr}) components, the latter of which is responsible for the aforementioned NRB. For simplicity, we have defined \tilde{C}_{st} that describes the effective stimulation intensity and bandwidth, and $\tilde{\chi}^{(3)}$ is the nonlinear susceptibility with intensity and spectral resolution as set by the probe source. Though neglected here, one could also account for detector and filter spectral response characteristics through modification of this term as well.

Spontaneous Raman spectroscopy records peaked spectra that are proportional to $\Im\{\chi_r(\omega)\}$ [23], but CARS spectra are proportional to $|\chi^{(3)}(\omega)|^2$; thus, computational or experimental approaches need be taken to ascertain the phase of $\chi^{(3)}$. The two primary computational approaches (also known as "phase-retrieval" methods) are the maximum entropy method (MEM) [12] and the Kramers-Kronig relation (KK) [11], which have been shown to be functionally equivalent [13]. Mathematically, using the KK relation as derived in [19], the retrieved nonlinear susceptibility (normalized to the NRB) is:

$$\frac{\chi_{\rm r}(\omega)}{\chi_{\rm nr}(\omega)} = \underbrace{\sqrt{\frac{I_{\rm CARS}(\omega)}{I_{\rm NRB}(\omega)}}}_{A(\omega)} \exp\left[i\underbrace{\mathcal{H}\left\{\frac{1}{2}\ln\frac{I_{\rm CARS}(\omega)}{I_{\rm NRB}(\omega)}\right\}}_{\phi(\omega)}\right],\tag{5}$$

where $I_{\text{NRB}}(\omega)$ is the spectrum of the NRB, \mathcal{H} is a Hilbert transform, ϕ is the phase, and A is the amplitude (unitless ratio). The retrieved χ_r is normalized by χ_{nr} in order to remove \tilde{C}_{st} (in both MEM and KK methods)— a quantity that has not been shown to be independently measurable but may be assumed to be constant between different CARS measurements [14]. Without loss of generality, we will treat the NRB as measurable in the der of this manuscript.

2.2. Hilbert transform

The Hilbert transform is a widely used linear operator implemented to extract analytic representations of signals from real-valued measurements. Simply put, causal systems may be represented by complex-valued functions in the frequency domain with a defined relationship between the real and complex components [24]. The Hilbert transform lets one calculate the imaginary (real) portion of this function when only the real (imaginary) portion are directly measurable. For example, a molecular system with inhomogeneous broadening may present itself as a Gaussian lineshape through a spectroscopy that measures the real (imaginary) part of the system response, and its imaginary (real) part – with analogous information – would be a Dawson lineshape. There is also a similar relationship for when one only measures the magnitude (phase) under certain conditions [25]. The Hilbert transform is used across a myriad of application spaces such as vibrational analysis [26–28], biomedical signal processing (e.g., cochlear implants [29], neural [30, 31] and cardiac monitoring [32]), and spectroscopies (e.g., optical [11, 19, 33, 34], terahertz [35], impedance [36], dielectric [37], magnetic resonance [38, 39]), where it is often referred to as a Kramers-Kronig "transform" or "relation". Further, the Hilbert transform is an integral part of the analysis tool known as the Hilbert-Huang transform: a NASA-developed method most frequently applied to geophysical analyses [40, 41]. A significant hurdle to the accurate use of the Hilbert transform, however, is that the discrete implementation may deviate significantly from the continuous form.

The Hilbert transform, \mathcal{H} , of a function f(x) on the real line is defined as

$$\mathcal{H}{f}(x) = \frac{\mathcal{P}}{\pi} \int_{-\infty}^{\infty} \frac{f(x')}{x - x'} dx' = f(x) * \frac{\mathcal{P}}{\pi x},\tag{6}$$

where * is a convolution and \mathcal{P} is the Cauchy principle value. Though this may be solved analytically for known functions [42], under more experimentally relevant conditions with discrete data there are a myriad of proposed approaches to the discrete Hilbert transform (DHT). Broadly speaking, these methods approximate the integral in Eq. 6 to a more tractable form [43, 44], model the input signal as a linear superposition of functions with analytically known Hilbert transforms [45–47], or solve the equivalent problems using the discrete Fourier transform [48, 49], which is a particular case of modeling the signal with known functions (sin and cos).

Using the Fourier transform is the most common approach in-so-much that it is broadly available in such computational software as GNU Octave, MATLAB, SciPy, and SAS. The aforementioned software packages use a fast-Fourier transform (FFT)-based method developed by Marple [48] that uses symmetry properties of the discrete Fourier transform (DFT) to calculate the DHT. Another Fourier transform approach by Henrici [49], which is implemented in LabView (and in previous work of the author [19]), uses the (continuous) Fourier transformed version of Eq. 6 as the starting point for the computation with a FFT. The mathematical forms and the equivalency of these two methods are derived in the Supplemental Document subsection 1.A.

2.3. Errors in the DHT

Significant challenges to accurate DHT computation arise in the case of signals that are nonbandlimited, non-compactly supported, and non-periodic — "general inputs" per [50] — such as is common in spectroscopies. The DHT on signals with these characteristics may have dramatic errors near the window edge — "end-"/"edge-effect" errors; though in fact, the error exists across the entire signal. These errors, as will be demonstrated, may cause significant errors in quantitative analysis of peak heights and relative ratios and may even generate new features.

Alternative DHT algorithms beyond Henrici and Marple primarily focus on computational efficiency and/or accuracy away from the window-edge; though, some methods show improvements for window-centered signals [51, 52]. To address window-edge errors, padding schema have been proposed [53] as well as analytical [54] and machine learning methods [55, 56] to artificially extend the input signal as to push the edge-effect errors further from the original data. These methods, however, do not stop the generation of the edge effect errors, but rather move it further away from the window center to an artificially larger window. This, in some

circumstances may be adequate, but as previously mentioned the edge effect errors permeate the entire window.

The aforementioned errors arise from the discretization and truncation of the input data. For example, the use of the DFT/FFT to effectively perform a circular convolution would see a difference in values from each edge of the signal window as a discontinuity. Similarly, a sudden change in derivative between the edges of the window would act as a form of discontinuity that would ultimately affect the correct calculation (related and bounded by the Paley-Wiener theorem [50]). Two properties of the DHT that are derived below (the second being previously unpublished to our knowledge), provide a means to calculate the total error of using the DHT (supporting derivations are found within the Section 1 of the Supplemental Document):

- 1. The mean of the DHT of a signal is zero
- 2. The variance (and standard deviation) of the input signal is the same as the DHT of the signal.

Using the mathematical definition of the discrete Fourier transform (DFT) from [50] (see also the Supplemental Document Section 1), the mean of the DHT of a length N signal, $f[n] \in \mathbb{R}^N$, may be shown to always be zero:

$$\langle \mathcal{H}\{f\}[n]\rangle = \frac{1}{N} \sum_{n} \mathcal{H}\{f\}[n] = \mathcal{F}\{H\{f\}\}[0] = -i \operatorname{sgn}[0]F[0] = 0,$$
 (7)

where \mathcal{F} is the DFT, \mathcal{H} is the DHT, F[k] is the Fourier-transformed signal at frequency $k \in \mathbb{Z}$, and 'sgn' is the sign/signum function. For this derivation we have used the Henrici version of the DHT (see the Supplemental Document section 1.A).

Next, we show that the variance of a signal is the same as that of the DHT of the signal, which is not necessarily true for the continuous Hilbert transform. First, using Parseval's theorem:

$$\frac{1}{N} \sum_{n} |\mathcal{H}\{f\}[n]|^{2} = \sum_{k} |\mathcal{F}\{\mathcal{H}\{f\}\}[k]|^{2} = \sum_{k} (\vec{1} - \delta[k]) |\mathcal{F}\{f\}[k]\}|^{2}$$
$$= \frac{1}{N} \sum_{n} |f[n]|^{2} - |F[0]|^{2}$$
$$= \frac{1}{N} \sum_{n} |f[n]|^{2} - \left|\frac{1}{N} \sum_{n} f[n]\right|^{2} = \sigma_{f}^{2}, \tag{8}$$

where $\hat{1}$ is a vector of 1's, δ is the unit impulse response (or unit sample function), and σ_f is the standard deviation (σ_f^2 is the variance) of the signal f[n]. We have used the derivation of the Fourier transform of the Hilbert transformed signal that is derived in the Supplemental Document section 1.B. We can use this finding to analyze the variance of the DHT-transformed signal:

$$\sigma_{\mathcal{H}\{f\}}^{2} = \frac{1}{N} \sum_{n} |\mathcal{H}\{f\}[n]|^{2} - \left| \underbrace{\frac{1}{N} \sum_{n} \mathcal{H}\{f\}[n]}_{= 0} \right|^{2} = \sigma_{f}^{2}, \tag{9}$$

where the second term is zero owing to first property of the DHT that was derived. This conservation of variance between the original signal and the DHT-transformed is not found when using analytically solved Hilbert transformed signals and is an artefact.

Though not provided in this manuscript, a detailed analysis of the DHT through analyzing bounds of decay relations in the DFT/FFT (e.g., see Chapter 6.7 in [50]) can provide insights and bounds onto the shape of the errors for known functions.

2.4. LeDHT

As described above, the DHT has properties that adversely affect its accuracy — especially in the common case of the spectrum extending to or beyond the recording window edge. Additionally, the DHT has fundamental properties that are aberrant when compared to the continuous Hilbert transform. Thus the aim of the LeDHT is to calculate the Hilbert transform on discrete data (spectra) with significantly fewer errors than the DHT. For CARS spectra, this means more accurate phase retrieval and Raman signal extraction than conventional methods. At its core, the LeDHT is a square matrix, $\mathbf{H} \in \mathbb{R}^{N \times N}$, which applies the transform via matrix multiplication: $G = F\mathbf{H}$, where F is a matrix of input spectra and G are the calculated Hilbert transforms. Both F and $G \in \mathbb{R}^{M \times N}$, where M are the number of spectra and N is the length of each spectrum. This matrix approach evolved from the fact that linear discrete transforms may be described by matrix formulations [57] (e.g., Hilbert [58, 59] and Fourier [50, 60] transforms). These matrix forms, however, are equivalent and generate the same errors and distortions. In the LeDHT, though, the matrix is not prescribed but rather learned through linear regression (least-squares) of synthetic training data with known [continuous] Hilbert transforms. This matrix form may be learned once for a given spectroscopic system and reused indefinitely – assuming one's choice of lineshape(s) and training width ranges and center ranges are still optimal for the sample and optical systems. The three steps to training and applying the LeDHT are described pictographically in Fig. 1 and will be elaborated on below.

The first step to implementing the LeDHT is the generation of synthetic training spectra. These M spectra will be used to learn the square LeDHT matrix, $\mathbf{H} \in \mathbb{R}^{N \times N}$ $(M \gg N)$, that describes the relationship between a spectrum of length N and its Hilbert transform (See step 1 of Fig. 1). The linear nature of the LeDHT offers significant benefits when generating synthetic data to span the available training space (i.e., possible lineshape parameter values). Of the most consequential benefits is that each training spectrum need only a single peak as more complex spectra can be composed of a linear superposition of single-peak spectra. This greatly reduces the number of possible unique spectra for training when compared to nonlinear models such as neural network-based approaches [61,62] that require full multi-peak spectra. As an example, consider the use of Gaussian lineshapes with corresponding scaled $(2\pi^{-1/2})$ Dawson functions for the Hilbert transforms. These lineshapes have 2 free parameters: center position and a width parameter. Assuming $n \in \mathbb{Z}^N$, center position is an integer value on this line, the width is a positive integer value that also enables at least 3 samples per full-width at half-maximum (FWHM), and the maximum FWHM is N, the number of unique peaks is $(N-1)(N/2-1) = N^2/2 - 3N/2 + 1$. For a 1000-point signal, 498 501 unique lineshapes exist under these conditions. Assuming double-point precision, the complete training dataset would be under 4 GB. This is dramatically lower than if full, multi-peak spectra were needed, such as with neural network approaches [61, 62], where there are $\approx 2.9 \times 10^{85}$ unique spectra for a 1000-point signal length with 1 to 15 possible peaks.

Generation of the synthetic training data has several considerations such as lineshape function and any constraints on training parameter space. From earlier experiments, we found that Dawson lineshapes, of which Gaussian lineshapes are the Hilbert transform, generated more optimal results than dispersive and Lorentzian lineshape pairs. These lineshape types have only two training space dimensions to consider: center position (frequency) and width. A potentially more general lineshape for future inquiry would be using the Faddeeva function in which real part is a Voigt function and the imaginary part is its Hilbert transform. From a practical perspective of this work, the three parameters (two related to width and one to spectral position) would have greatly expanded our training space, limited the fraction of which we could sample (limited by available memory of the computer). The second important consideration are any constraints on the training space parameters. In this work, the minimum peak width (standard deviation for a Gaussian) was 4 spectral pixels, and the maximum was 80 spectral pixels. Additionally, any



Fig. 1. Infographic describing the three steps of implementing the LeDHT: generation of synthetic training data, learning an optimal matrix form of the Hilbert transform, and application to new spectra via matrix multiplication. Note: Scatter plot points sub-sampled (5x) for visual clarity.

generated peak must be at least half a FWHM from the window edge; thus, as peaks are centered closer to the window edge, the maximum allowable width is narrower. This gives the training space parameter map a trapezoidal quality as seen in Fig. 1. Additionally, we limited spectral position and width values to integer values (unit of spectral pixels in *n*). A description of how the training data was algorithmically generated is provided within the Methods Section.

The second step in implementing the LeDHT is solving for the **H** matrix in $G = F\mathbf{H}$, where F are the input training spectra and G are the known Hilbert transforms. Generally speaking, there are numerous methods to solve for an optimal LeDHT matrix. In this work, we are using ordinary least squares (OLS) which minimizes the residual sum-of-squares: $\mathbf{H} = \operatorname{argmin}_{\mathbf{H}} ||G - F\mathbf{H}||^2$. Beneficially, OLS is the most common linear regression method; thus, there are numerous efficient algorithms for its calculation. We have also tested common "regularized" least squares methods such as ridge and LASSO regression with success though with inferior results for the data within this manuscript. For different lineshapes and different spectroscopic conditions, these other methods could potentially be superior.

Finally, the LeDHT can be applied to new data via matrix multiplication, which is computationally inexpensive: $G_{exp} = F_{exp}\mathbf{H}$, where F_{exp} is a newly measured input spectrum/spectra and G_{exp} is the LeDHT Hilbert-transformed spectrum to be calculated. Ideally, the solved-for LeDHT matrix for a given spectroscopy system does not need to be re-learned between experiments and can be reused indefinitely, again, as long as the training dataset still optimally reflects the expected future input spectral lineshape characteristics.

3. Methods and materials

3.1. Processing and software

All simulated and experimental data processing was performed on a Dell Precision 7730 laptop with a 6-core i7-8850H processor at 2.6 GHz with 64 GB of RAM. All processing of simulated and experimental data was performed using Python 3.8.5, NumPy, SciPy [63], scikit-learn [64], Matplotlib, and in-house developed Python packages. The traditional DHT used in this manuscript is the Marple version as available in SciPy. The LeDHT method, several other implementations of DHT methods (and padding schema), and Jupyter Notebook examples have been made available as an open-source Python package, *Hilbert*, at [65].

3.2. Generation of synthetic data

Synthetic data is generated in a stochastic manner in which the free parameters are randomly selected and accepted based on whether it is a unique combination and is within given bounds. For the simulations, *N* is 401 with *n* in the range [-200, 200]. The generated training data is composed of Dawson function lineshapes and their Hilbert transforms: Gaussian lineshapes. The minimum acceptable Gaussian standard deviation is 4 *n*-pixels and the maximum is 80. Furthermore, the center location of Gaussian lineshapes is bounded [-200,200] (see Fig. 1 "Training Space" subimage) but that for a given peak, the center position must be at least half a FWHM from the window edge. Under these conditions, 100 000 of these unique single-lineshape spectra are generated. For further stability and improved results, we augmented the data by repeating the dataset 3 times, adding white Gaussian noise (standard deviation of 1×10^{-12}) and adding a constant offset sampled from a normal distribution. Thus the total training data was 300 000 single-peak spectra.

An additional 100 000 single-peak spectra were generated as test set for validation. This validation was repeated 3 times, for a total of 300 000 spectra. These spectra were noiseless and did not have a DC offset as to represent the ideal, expected spectral case. Results with test data that did contain noise and DC offset provided similar results to those presented in the Results Section.

For the simulated multi-peak spectra experiments, the same training data was used as for the single-peak experiments, but the test data (i.e., multi-peak spectra) was generated from linear summations of Dawson function lineshapes (and their corresponding Hilbert transform Gaussian lineshapes). Each randomly generated spectrum was composed of 1 to 15 peaks (uniform sampling). Each peak had an amplitude that was uniformly sampled in the range of 1 to 2 (no units), a width that was uniformly sampled between 4 and 26 (integer valued), and a center position that needed to be at least half a FWHM from the window edge.

For working with the experimental BCARS spectra that are 250 to 460 n-pixels long, we generated 150 000 Dawson and Gaussian lineshapes limited to widths of 3 to 50 n-pixels. Next we replicated (3x) and augmented the data in a similar manner as previously described, adding white Gaussian noise (standard deviation 10^{-3}), and a DC offset sampled from a normal distribution. This provided a total of 450 000 training spectra.

3.3. CARS spectrometer

In this work, spectra of neat glycerol (BioUltra, Sigma) were collected on a previously developed broadband CARS (BCARS) microscope that is described in detail elsewhere [6]. The glycerol was mounted between a coverslip (VWR) and glass slide (Ward's) separated by an adhesive well (GraceBio). Each spectrum was collected with a 3.5 ms dwell time and with average power of the probe and supercontinuum of 20.1 mW and 11.2 mW, respectively. Spectra of the coverslip were collected and used for normalization — a process that effectively suppresses a co-generated, coherent nonresonant background [6, 19]. Additionally, dark spectra (250) were collected under the same conditions except with the laser sources not temporally overlapped (but still illuminating the sample). The glycerol spectrum shown in this work is an average of 97 individual spectra.

4. Results

4.1. Simulated

As an initial demonstration, we will compare the DHT, the DHT with extrema-value padding (DHT-Pad) – the most common padding scheme in the CARS community – and the LeDHT. Figure 2(a) is a representative Dawson function spectrum that is spectrally centered, and Figs. 2(b-d) show the respective outputs of the DHT, DHT-Pad, and LeDHT. The DHT generates the most severe errors, with new downward-facing features at the spectroscopic window's edges. The magnitude of these errors are such that the mean of the DHT spectrum is zero as described previously in the Theory section. The DHT-Pad, by visual inspection, agrees with the true Hilbert transform with an apparent slight constant offset. Finally, the LeDHT appears to agree with the true Hilbert transform. Using mean squared error (MSE) as the error metric, the LeDHT is approximately 6 orders of magnitude better than the DHT-Pad and 8 orders of magnitude improvement over the DHT.

Next, we examine the results of using the same input Dawson spectrum but shifted closer to a window edge, as shown in Fig. 2(e). Similar to before, the DHT produces erroneous downward facing features that are significant in amplitude as shown in Fig. 2(f). Additionally, the DHT peak maximum value is approximately 20 % less than the true value. With the spectral peak near the window edge, now the DHT-Pad is demonstrating significant inaccuracies. The peak maximum is significantly above the ideal case (> 20% at the peak maximum) and there is a clear distortion in the baseline. Figure 2(h) again demonstrates the superior performance of the LeDHT. The MSE of the LeDHT is approximately 550 times lower than the DHT-Pad, and nearly 1 200 times lower than for the DHT.

To further examine the peak position-error relationship, Fig. 2(i) shows the MSE (on a log scale) as a function of peak position from the 300 000 test data spectra. The LeDHT has the



Fig. 2. (a,e) Example Dawson function input spectra. (b,f) DHT results showing significant errors. (c,g) Results using the DHT with extrema-value padding (DHT-Pad). (d,h) Results using the LeDHT. (i) Mean MSE (\pm 1 standard deviation) calculated for the 300 000 stochastically generated test spectra as a function of peak center position. (j) Histogram of optimal method based on minimum MSE for test spectra dataset. Y-axis shows fraction, parenthetical values are number of spectra.

lowest mean MSE of all methods across all peak positions. The next best method is DHT-Pad, which outperforms DHT across the entire window save the edges. And finally, we calculated for each of the 300 000 test spectra, which method had the best performance as shown in Fig. 2(j). The LeDHT had the best performance in 299 974 of 300 000 spectra (\approx 99.99 %), with the DHT having better MSE 26 times (\approx 0.01 %). All 26 results were from narrow peaks centered at the window edges, for which improved training data with some added emphasis on these cases would likely improve LeDHT performance. The DHT-Pad was never the optimal choice, which indicates that the LeDHT always outperforms the DHT-Pad. Section 2.A of the Supplemental Document contains the converse scenario in which input spectra are single Gaussian peaks and the resultant Hilbert transforms (Dawson functions). Like the results presented here, the LeDHT is similarly superior. Section 2.C of the Supplemental Document also demonstrates using the Gaussian training set above on test data that is Lorentzian lineshapes. Surprisingly, the LeDHT still outperforms the DHT-Pad on 300 000 test spectra.

Next we analyzed the DHT, DHT-Pad, and LeDHT performance when applied to synthetically generated multi-peak spectra. 30000 spectra were generated stochastically, each with 1 to 15 peaks. Figure 3 (a) shows an example single spectrum composed of a linear combination of 13 Dawson function lineshapes. Figures 3(b-d) shows the corresponding transforms. Looking at the DHT, again one can see downward sloping features at the window edge that are larger than any real peak (though negative) with over 1000 % relative error at the window edges (over 5800 % at the right side). Additionally, one can see that the entire DHT spectrum is offset below ideal outcome leading to a minimum relative error of ≈ 29.5 %. Similarly, the DHT-Pad spectrum is entirely shifted above the ideal spectrum with a minimum relative error of ≈ 8.5 %. The DHT-Pad also shows amplitude errors that become worse at the window edges: greater than 300 % relative error at the left edge and 550 % at the right edge. In contrast, the LeDHT is significantly less biased with a minimum relative error of ≈ 0.04 %. The window edge extrema relative errors were 13 % (left) and 0.15 % (right). Figure 3(e) shows box plots of MSE generated by the three methods. The LeDHT median MSE is ≈ 4.5 orders of magnitude less than that of the DHT, and approximately 3 orders of magnitude less than the DHT-Pad. Figure 3(f) further affirms the superiority of LeDHT, identifying that the LeDHT outperformed, based on MSE, the other methods in all 30000 test spectra. Section 2.B of the Supplemental Document contains the converse scenario in which input multi-peak spectra are composed of Gaussian lineshapes and the resultant Hilbert transforms, displaying similarly superior results.

As a final note, we compared the computational time of the DHT, DHT-Pad, and LeDHT methods. We do not assume that any of our algorithms are optimal; thus, there is likely much room for improvement. For 100 000 single-peak spectra, the DHT and DHT-Pad took approximately 4 seconds and 15 seconds, respectively over 3 trials. The LeDHT, took 10 seconds to train on 300 000 spectra, and only 0.4 seconds to 0.5 seconds to calculate 100 000 test spectra. And as mentioned previously, the LeDHT can be trained once and reused in future experiments; thus, the LeDHT was approximately an order of magnitude faster than the DHT when already trained, and over 30 times faster than DHT-Pad.

4.2. Experimental CARS spectra

Finally, we applied the LeDHT to experimental CARS spectra of neat glycerol, varying the proximity of the spectrum to the window edge to demonstrate the superior repeatability and stability of the LeDHT method. Though this scenario may seem artificial, for broadband (>3500 cm⁻¹) or low-frequency (<500 cm⁻¹) applications, peaks extending to the window edge become probable (especially in chemical and biological samples). The Raman phase (ϕ , Eq. 5) and Raman:NRB ratio (χ_R/χ_{NR} , Eq. 5) were calculated as described in the Theory and Methods section. These quantities were calculated at different signal bandwidths: from ≈ 2080 cm⁻¹ down to ≈ 1020 cm⁻¹ in 23 increments. The training data was composed of 450000



Fig. 3. (a) Example synthetic spectrum composed of 13 Dawson functions. (b,c,d) The true Hilbert transform and the calculated Hilbert transforms using DHT, the DHT applied to extrema-value padded spectra (DHT-Pad), and the LeDHT, respectively. (e) Box plots of MSE for each method. Note the whiskers displays the extreme range of values. (f) Histogram of the optimal method (based on minimum MSE) for 30 000 test spectra. Y-axis shows fraction, parenthetical values are number of spectra.

Gaussian lineshapes augmented with noise and DC bias as described in the Methods section.

Figure 4(a) shows the input to the Hilbert transforms, $-0.5 \ln(I_{CARS}/I_{NRB})$, where the NRB was recorded from coverslip glass. Note the negative sign when compared to Eq. 5, due to our experimental data vectors being oriented with decreasing ω vectors (i.e., the ω vector is numerically high-valued to low-valued). Figures 4(a) and (g) show the Raman phase and Raman:NRB ratio, respectively, as extracted using the DHT. As expected, the DHT generates large negative features at the window extrema. Figures 4(b) and (h) show the DHT-Pad results that are by inspection significantly improved. One can see, though, that as the bandwidth of the spectrum is changed, the resulting Hilbert transform is modulated in amplitude, especially at the baseline near 2600 cm⁻¹ and at the OH-stretch maximum near 3400 cm⁻¹. Conversely, the peak at 2949 cm-1 is stable. A common approach in CARS spectroscopy is to set the baseline to 0 at the window edge. Doing so, in this case, at approximately 2600 cm⁻¹ would adversely cause the stable amplitude peak at 2949 cm⁻¹ to modulate. Finally, Figs. 4(c) and (i) show the LeDHT results that are significantly more stable with window width and only appear to fluctuate mildly within the OH-stretch region. To qualify the stability of the LeDHT algorithm, Figs. 4(d) and (j) present violin plots of the peak phase and Raman:NRB at 2887 cm⁻¹, 2949 cm⁻¹, and 3400 cm⁻¹ across the various window widths. In all six plots, the vertical extent of the LeDHT violin (indicating variation) is significantly smaller than for DHT and DHT-Pad. For example, for the 3400 cm^{-1} peak, the mean \pm standard deviation phase for the DHT is (0.50 ± 0.12) rad, (0.8 ± 0.04) rad for the DHT-Pad, and (0.83 ± 0.01) rad for the LeDHT. Thus in short, the LeDHT standard deviation is about 12 times less than the DHT and 4 times less than the DHT-Pad. For the Raman:NRB calculations, the LeDHT standard deviation is similarly 10 times less than the DHT and 3 times smaller than the DHT-Pad.



Fig. 4. (a) Black line associated with right y-axis: negative half natural log of experimental BCARS spectrum from glycerol normalized to coverslip, which is the input to the Hilbert transform for phase calculations. Green lines: DHT-calculated phases when the input is truncated at different spectral widths. (b) Phase calculations using extrema-valued padded spectra into the DHT (DHT-Pad). (c) Phase calculation using the LeDHT. (d-f) Violin plots of the phase magnitude for three prominent peaks. (g-i) Raman-to-NRB ratio calculations using DHT, DHT-Pad, and LeDHT methods, respectively. (j-1) Violin plots of the extracted Raman-to-NRB ratio for prominent peaks.

5. Discussion and conclusion

In this work, we have presented a learned matrix approach for calculating the Hilbert transform with superior performance compared to the DHT and a padded DHT schema. These demonstrations were performed on a variety of synthetic signals as well as experimental CARS spectra. And though the LeDHT demonstrated orders-of-magnitude improvement in MSE on the synthetic datasets and heightened stability with experimental data, there are still numerous opportunities for further gains and new application spaces to consider.

There are many extensions and improvements to LeDHT that could further enhance and expand the method. For example, the training data could be augmented to explicitly diminish noise. In this case, it may be advantageous to repeat the training dataset with random additions of noise solely on the input. Our initial attempts at this were successful though areas of low-signal and high-noise were affected in a similar manner to traditional [convolution] filtering methods (e.g., bandpass, Savitzky–Golay) with erroneous lineshape-like features being generated within the noise. Traditionally, we have had more success with separate processes of statistical denoising [19].

Another opportunity for improvement is methods to effectively train with smaller datasets, which in turn enables the generation of larger and larger training datasets with unique features. For examples, we have experimented with generating large training datasets and factorizing with singular value decomposition (SVD), which effectively generates a much smaller training dataset with each entry being more complex than a single lineshape. The 300 000 training spectra used for Figs. 2 and 3, for example, could be reduced down to 401 (matching the size of N) – a nearly 750x reduction. In the long term, this could enable a training dataset to be originally generated from numerous different lineshapes, potentially enabling a more universal LeDHT matrix that could be shared.

We are also exploring alternative regression/optimization methods for finding the LeDHT matrix. This matrix may be more optimal in terms of MSE, or it may be more optimal in the sense of having a symmetry (e.g., a Procrustes problem) that can be more easily human-interpreted or implemented in hardware. Thus far we have worked on the former problem: implementing Ridge and LASSO regression methods; though linear regression has performed most optimally.

And finally, we are looking to apply LeDHT to other signal types and spectroscopic modalities. Although CARS is a relatively niche technique, methods such as NMR [38, 39], dielectric [37], and terahertz [35] spectroscopies have also demonstrated use of the DHT.

In closing, we have developed LeDHT and demonstrated significantly improved MSE over the traditional DHT and with signal extrema padding. Our method is fast, stable, and improves opportunities for quantitative analysis of spectroscopic signals. Additionally, an open-source Python library for this method is freely available to enable community development and foster adoption.

Acknowledgements. The author would like to thank the peer-reviewers for their effort and insightful comments. Additionally, the author would like to thank Young Jong Lee, Nancy Lin, Sheng Lin-Gibson, Donald Atha, John Henry Scott, and the internal review board (WERB) at NIST for their helpful reviews and recommendations towards the preparation of this manuscript.

Disclaimer. Any mention of commercial products or services is for experimental clarity and does not signify an endorsement or recommendation by the National Institute of Standards and Technology. The authors declare no conflicts of interest.

Disclosures. The authors declare no conflicts of interest.

Data availability. The scripts, data, and library used in this manuscript are publicly available at [65].

Supplemental document. See Supplement 1 for supporting content.

References

- A. Zumbusch, G. R. Holtom, and X. S. Xie, "Three-dimensional vibrational imaging by coherent anti-stokes raman scattering," Phys. Rev. Lett. 82, 4142–4145 (1999).
- J.-X. Cheng, A. Volkmer, L. D. Book, and X. S. Xie, "Multiplex coherent anti-Stokes Raman scattering microspectroscopy and study of lipid vesicles," The J. Phys. Chem. B 106, 8493–8498 (2002).
- G. W. Wurpel, J. M. Schins, and M. Muller, "Chemical specificity in three-dimensional imaging with multiplex coherent anti-stokes raman scattering microscopy," Opt Lett 27, 1093–1095 (2002).

- G. I. Petrov, R. Arora, V. V. Yakovlev, X. Wang, A. V. Sokolov, and M. O. Scully, "Comparison of coherent and spontaneous raman microspectroscopies for noninvasive detection of single bacterial endospores," Proc Natl Acad Sci U S A 104, 7776–7779 (2007).
- M. Muller and A. Zumbusch, "Coherent anti-stokes raman scattering microscopy," Chemphyschem 8, 2156–2170 (2007).
- C. H. Camp Jr, Y. J. Lee, J. M. Heddleston, C. M. Hartshorn, A. R. Hight Walker, J. N. Rich, J. D. Lathia, and M. T. Cicerone, "High-speed coherent Raman fingerprint imaging of biological tissues," Nat. Photonics 8, 627–634 (2014).
- M. Cui, B. R. Bachler, and J. P. Ogilvie, "Comparing coherent and spontaneous raman scattering under biological imaging conditions," Opt Lett 34, 773–775 (2009).
- J. X. Cheng, L. D. Book, and X. S. Xie, "Polarization coherent anti-stokes raman scattering microscopy," Opt Lett 26, 1341–1343 (2001).
- F. Ganikhanov, C. L. Evans, B. G. Saar, and X. S. Xie, "High-sensitivity vibrational imaging with frequency modulation coherent anti-stokes raman scattering (fm cars) microscopy," Opt Lett 31, 1872–1874 (2006).
- D. Pestov, R. K. Murawski, G. O. Ariunbold, X. Wang, M. Zhi, A. V. Sokolov, V. A. Sautenkov, Y. V. Rostovtsev, A. Dogariu, Y. Huang, and M. O. Scully, "Optimizing the laser-pulse configuration for coherent Raman spectroscopy," Science 316, 265–268 (2007).
- Y. X. Liu, Y. J. Lee, and M. T. Cicerone, "Broadband CARS spectral phase retrieval using a time-domain Kramers-Kronig transform," Opt. Lett. 34, 1363–1365 (2009).
- E. M. Vartiainen, "Phase retrieval approach for coherent anti-Stokes Raman scattering spectrum analysis," J. Opt. Soc. Am. B 9, 1209–1214 (1992).
- M. T. Cicerone, K. A. Aamer, Y. J. Lee, and E. Vartiainen, "Maximum entropy and time-domain Kramers-Kronig phase retrieval approaches are functionally equivalent for CARS microspectroscopy," J. Raman Spectrosc. 43, 637–643 (2012).
- F. Masia, A. Glen, P. Stephens, P. Borri, and W. Langbein, "Quantitative chemical imaging and unsupervised analysis using hyperspectral coherent anti-stokes raman scattering microscopy," Anal Chem 85, 10820–10828 (2013).
- H. A. Rinia, M. Bonn, M. Müller, and E. M. Vartiainen, "Quantitative CARS spectroscopy using the maximum entropy method: The main lipid phase transition," ChemPhysChem 8, 279–287 (2007).
- A. Karuna, F. Masia, P. Borri, and W. Langbein, "Hyperspectral volumetric coherent anti-Stokes Raman scattering microscopy: quantitative volume determination and NaCl as non-resonant standard," J. Raman Spectrosc. 47, 1167–1173 (2016).
- D. Schafer, J. A. Squier, J. Van Maarseveen, D. Bonn, M. Bonn, and M. Müller, "In situ quantitative measurement of concentration profiles in a microreactor with submicron resolution using multiplex CARS microscopy," J. Am. Chem. Soc. 130, 11592–11593 (2008).
- Y. J. Lee, D. Moon, K. B. Migler, and M. T. Cicerone, "Quantitative image analysis of broadband CARS hyperspectral images of polymer blends," Anal. Chem. 83, 2733–2739 (2011).
- C. H. Camp Jr, Y. J. Lee, and M. T. Cicerone, "Quantitative, comparable coherent anti-Stokes Raman scattering (CARS) spectroscopy: correcting errors in phase retrieval," J. Raman Spectrosc. 47, 408–415 (2016).
- J. S. Gomez, "Coherent Raman spectroscopy," in *Modern Techniques in Raman Spectroscopy*, J. J. Laserna, ed. (John Wiley and Sons, Chichester, UK, 1996), pp. 305–342.
- C. H. Camp Jr and M. T. Cicerone, "Chemically sensitive bioimaging with coherent Raman scattering," Nat. Photonics 9, 295–305 (2015).
- E. O. Potma and S. Mukamel, "Theory of coherent Raman scattering," in *Coherent Raman Scattering Microscopy*, J.-X. Cheng and X. S. Xie, eds. (CRC, Boca Raton, FL, 2013), pp. 3–42.
- W. M. Tolles, J. W. Nibler, J. R. McDonald, and A. B. Harvey, "A Review of the theory and application of coherent anti-Stokes Raman spectroscopy (CARS)," Appl. Spectrosc. 31, 253–271 (1977).
- 24. J. S. Toll, "Causality and the dispersion relation: logical foundations," Phys. Rev. 104, 1760–1770 (1956).
- 25. D. Y. Smith, "Dispersion relations for complex reflectivities*†," J. Opt. Soc. Am. 67, 570 (1977).
- 26. X. Fang and H. Luo, "An improved Hilbert transform for nonlinear vibration signal analysis," in 50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, (American Institute of Aeronautics and Astronautics, Reston, Virigina, 2009), May.
- 27. M. Feldman, "Hilbert transform in vibration analysis," Mech. Syst. Signal Process. 25, 735-802 (2011).
- 28. H. Luo, X. Fang, and B. Ertas, "Hilbert transform and its engineering applications," AIAA J. 47, 923–932 (2009).
- 29. K. Nie, A. Barco, and F.-G. Zeng, "Spectral and temporal cues in cochlear implant speech perception," Ear Hear. 27, 208–217 (2006).
- P. Tass, M. G. Rosenblum, J. Weule, J. Kurths, A. Pikovsky, J. Volkmann, A. Schnitzler, and H.-J. Freund, "Detection of n:m phase locking from noisy data: application to magnetoencephalography," Phys. Rev. Lett. 81, 3291–3294 (1998).
- M. Le Van Quyen, J. Foucher, J. P. Lachaux, E. Rodriguez, A. Lutz, J. Martinerie, and F. J. Varela, "Comparison of Hilbert transform and wavelet methods for the analysis of neuronal synchrony," J. Neurosci. Methods 111, 83–98 (2001).
- D. Benitez, P. A. Gaydecki, A. Zaidi, and A. P. Fitzpatrick, "The use of the Hilbert transform in ECG signal analysis," Comput. Biol. Medicine 31, 399–406 (2001).

- 33. T. S. Robinson, "Optical constants by reflection," Proc. Phys. Soc. Sect. B 65, 910–911 (1952).
- 34. D. M. Roessler, "Kramers-Kronig analysis of reflection data," Br. J. Appl. Phys. 16, 1119–1123 (1965).
- Y. Divin, M. Lyatti, A. Snezhko, U. Poppe, and V. Pavlovskiy, "Thz hilbert-transform spectrum analyzer based on high-t_c josephson junction in stirling cryocooler," IEEE Transactions on Appl. Supercond. 23, 1800204–1800204 (2013).
- 36. P. Agarwal, M. E. Orazem, and L. H. Garcia-Rubio, "Application of measurement models to impedance spectroscopy: III. Evaluation of consistency with the Kramers-Kronig relations," J. The Electrochem. Soc. 142, 4159–4168 (2019).
- N. Axelrod, E. Axelrod, A. Gutina, A. Puzenko, P. B. Ishai, and Y. Feldman, "Dielectric spectroscopy data treatment: I. frequency domain," Meas. Sci. Technol. 15, 755–764 (2004).
- M. Mayzel, K. Kazimierczuk, and V. Y. Orekhov, "The causality principle in the reconstruction of sparse nmr spectra," Chem. Commun. 50, 8947–8950 (2014).
- R. R. Ernst, "Numerical hilbert transform and automatic phase correction in magnetic resonance spectroscopy," J. Magn. Reson. (1969) 1, 7 – 26 (1969).
- 40. N. E. Huang, Z. Shen, S. R. Long, M. C. Wu, H. H. Snin, Q. Zheng, N. C. Yen, C. C. Tung, and H. H. Liu, "The empirical mode decomposition and the Hubert spectrum for nonlinear and non-stationary time series analysis," Proc. Royal Soc. A: Math. Phys. Eng. Sci. 454, 903–995 (1998).
- N. E. Huang and Z. Wu, "A review on hilbert-huang transform: Method and its applications to geophysical studies," Rev. Geophys. 46, RG2006 (2008).
- 42. A. D. Poularikas, "Hilbert transform," in *The Handbook of Forumulas and Tables for Signal Processing*, A. D. Poularikas, ed. (CRC, Boca Raton, FL, 1999), Chap. 15.
- R. Kress and E. Martensen, "Anwendung der Rechteckregel auf die reelle Hilberttransformation mit unendlichem Intervall," ZAMM - J. Appl. Math. Mech. / Zeitschrift für Angewandte Math. und Mech. 50, 61–64 (1970).
- 44. J. A. C. Weideman, "Computing the hilbert transform on the real line," Math. Comput. 64, 745–762 (1995).
- 45. C. A. Micchelli, Y. Xu, and B. Yu, "On computing with the Hilbert spline transform," Adv. Comput. Math. 38, 623–646 (2013).
- 46. F. Stenger, "Approximations via Whittaker's cardinal function," J. Approx. Theory 17, 222-240 (1976).
- C. Zhou, L. Yang, Y. Liu, and Z. Yang, "A novel method for computing the Hilbert transform with Haar multiresolution approximation," J. Comput. Appl. Math. 223, 585–597 (2009).
- 48. L. Marple, "Computing the discrete-time "analytic" signal via FFT," IEEE Transactions on Signal Process. 47, 2600–2603 (1999).
- 49. P. Henrici, Applied and computational complex analysis Vol III (Wiley-Interscience, New York, 1986).
- 50. W. L. Briggs and V. E. Henson, *The DFT: An Owner's Manual for the Discrete Fourier Transform* (Society for Industrial and Applied Mathematics, 1995).
- R. Bilato, O. Maj, and M. Brambilla, "An algorithm for fast Hilbert transform of real functions," Adv. Comput. Math. 40, 1159–1168 (2014).
- M. B. Abd-el Malek and S. S. Hanna, "The Hilbert transform of cubic splines," Commun. Nonlinear Sci. Numer. Simul. 80, 104983 (2020).
- D. Huang, J. Zhao, and J. Su, "Practical implementation of Hilbert-Huang Transform algorithm," Acta Oceanol. Sinica 22, 1–14 (2003).
- M. Meissner, "Accuracy issues of discrete Hilbert transform in identification of instantaneous parameters of vibration signals," Acta Phys. Polonica A 121, A–164–A–167 (2012).
- J. Cheng, D. Yu, and Y. Yang, "Application of support vector regression machines to the processing of end effects of Hilbert-Huang transform," Mech. Syst. Signal Process. 21, 1197–1211 (2007).
- Y. Deng, W. Wang, C. Qian, Z. Wang, and D. Dai, "Boundary-processing-technique in EMD method and Hilbert transform," Chin. Sci. Bull. 46, 954–960 (2001).
- K. Riley, M. Hobson, and S. Bence, *Mathematical Methods for Physics and Engineering* (Cambridge, Cambridge, UK, 2002), 2nd ed.
- F. Burris, "Matrix formulation of the discrete hilbert transform," IEEE Transactions on Circuits Syst. 22, 836–838 (1975).
- 59. S. C. Dutta Roy, "Alternative matrix formulation of the discrete hilbert transform," Proc. IEEE 64, 1435–1435 (1976).
- D. Kahaner, "Matrix description of the fast fourier transform," IEEE Transactions on Audio Electroacoustics 18, 442–450 (1970).
- C. M. Valensise, A. Giuseppi, F. Vernuccio, A. De La Cadena, G. Cerullo, and D. Polli, "Removing non-resonant background from CARS spectra via deep learning," APL Photonics 5, 061305 (2020).
- R. Houhou, P. Barman, M. Schmitt, T. Meyer, J. Popp, and T. Bocklitz, "Deep learning as phase retrieval tool for cars spectra," Opt. Express 28, 21002–21024 (2020).
- 63. P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," Nat. Methods 17, 261–272 (2020).

- 64. F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," J. Mach. Learn. Res. 12, 2825–2830 (2011).
- 65. C. H. C. Jr., "Hilbert," GitHub (2021) [retrieved 21 May 2022], https://github.com/usnistgov/Hilbert.

Raman signal extraction from CARS spectra using a learned-matrix approach to the discrete Hilbert transform: supplemental document

1. MATHEMATICAL DERIVATIONS

In this work, we are considering an odd-length signal f[n] where $-(N-1)/2 \le n \le (N-1)/2$ and are using the following definitions for the DFT and iDFT per [1]:

$$F[k] = \mathcal{F}\{f\}[k] := \frac{1}{N} \sum_{n=-\frac{N-1}{2}}^{\frac{N-1}{2}} f[n]e^{-i\frac{2\pi}{N}kn}$$
(S1)

$$f[n] = \mathcal{F}^{-1}\{F\}[n] := \sum_{k=-\frac{N-1}{2}}^{\frac{N-1}{2}} F[k]e^{i\frac{2\pi}{N}kn},$$
(S2)

A. DHT of Henrici and Marple are equivalent

For an odd-length signal the DHT implementations of Marple [2] and Henrici [3] are:

$$\text{Marple}: \mathcal{H}\{f\}[n] = \text{Im} \left\{ \mathcal{F}^{-1} \left\{ \begin{matrix} 0 & k < 0 \\ \mathcal{F}\{f\}[k] & k = 0 \\ 2\mathcal{F}\{f\}[k] & 1 \le k \le \frac{N-1}{2} \end{matrix} \right\} \right\}$$
(S3)

Henrici :
$$\mathcal{H}{f}[n] = \mathcal{F}^{-1}{-i\operatorname{sgn}[k]\mathcal{F}{f}[n]}[n],$$
 (S4)

where one should note that if $g \in \mathbb{R}^N$, than Henrici's output is purely real-valued (see Section 1.C). Also note that in Marple's original work, only even-length signals were considered but it provided a conceptual layout to apply to odd-length signals. The cases structure of the Marple implementation may be re-written as:

$$\mathcal{H}{f}[n] = \operatorname{Im}{\mathcal{F}^{-1}{(\operatorname{sgn}[k]+1)\mathcal{F}{f}[k]}} = \operatorname{Im}{\mathcal{F}^{-1}{\operatorname{sgn}[k]\mathcal{F}{f}[k]}} + \underbrace{\operatorname{Im}{\mathcal{F}^{-1}{\mathcal{F}{f}[k]}}}_{=\operatorname{Im}{f}=0, \text{ if } g \in \mathbb{R}^{N}} = \operatorname{Re}{\mathcal{F}^{-1}{-i \operatorname{sgn}[k]\mathcal{F}{f}[k]}[n]},$$
(S5)

where we have used the fact that $Im\{y\} = Im\{y' + i y''\} = Re\{-i y\} = y''$, where *y* is a complex-valued entity (with *y*' and *y''* as its real and imaginary parts, respectively).

For an even-valued N, the original form of the Marple implementation is utilized:

Marple:
$$\mathcal{H}{f}[n] = \operatorname{Im} \left\{ \mathcal{F}^{-1} \left\{ \begin{matrix} 0 & k < 0 \\ \mathcal{F}{f}[k] & k = 0 \\ 2\mathcal{F}{f}[k] & 1 \le k < \frac{N}{2} \\ \mathcal{F}{f}[k] & k = \frac{N}{2} \end{matrix} \right\} \right\}.$$
 (S6)

Similar to as in Eq. S5, we can re-write this implementation as:

$$\mathcal{H}\{f\}[n] = \operatorname{Im}\{\mathcal{F}^{-1}\{\left(\operatorname{sgn}[k]+1-\delta\left\lfloor\frac{N}{2}\right\rfloor\right)\mathcal{F}\{f\}[k]\}\}$$

$$= \operatorname{Im}\{\mathcal{F}^{-1}\{\operatorname{sgn}[k]\mathcal{F}\{f\}[k]\}\} + \underbrace{\operatorname{Im}\{\mathcal{F}^{-1}\{\mathcal{F}\{f\}[k]\}\}}_{=\operatorname{Im}\{f\}=0, \text{ if } g \in \mathbb{R}^{N}} - \operatorname{Im}\{\underbrace{\mathcal{F}^{-1}\{\delta\left\lfloor\frac{N}{2}\right]\mathcal{F}\{f\}\}}_{\in\mathbb{R}^{N}, \text{ if } g \in \mathbb{R}^{N}}$$

$$= \operatorname{Re}\{\mathcal{F}^{-1}\{-i \operatorname{sgn}[k]\mathcal{F}\{f\}[k]\}[n]\};$$
(S7)

thus, for even signals, the Marple and Henrici implementations are equivalent. It should be noted that in our numerical work, we find the difference between the results on the order of machine epsilon for double floating point arithmetic or exactly 0.0 in some cases.

B. The Fourier transform of the Hilbert transformed signal

$$|\mathcal{F}{\mathcal{H}{f}}[k]|^{2} = |-j \operatorname{sgn}[k] \mathcal{F}{f}[k] \}|^{2} = (\vec{1} - \delta[k]) |\mathcal{F}{f}[k] \}|^{2},$$
(S8)

where $\vec{1}$ is a 1-vector of length N. Note that this finding disagrees with that of [4] in which the square of the sign-function was implied to be a constant unity.

C. For a real-valued input, the DHT output is real-valued

$$\mathcal{H}\{f\}[n] = \sum_{k} -i \operatorname{sgn}[k] e^{i\frac{2\pi}{N}kn} \frac{1}{N} \sum_{n'} f[n] e^{-i\frac{2\pi}{N}kn'} = \sum_{k} -i \operatorname{sgn}[k] e^{i\frac{2\pi}{N}kn} [F_{\text{even}}[k] + i F_{\text{odd}}[k]],$$
(S9)

where we use F_{even} and F_{odd} to represent the DFT of the even and odd components of the input signal *g*. Importantly, the DFT has symmetry properties such that F_{even} and F_{odd} have even and odd symmetry, respectively [1].

Assuming that *N* is odd-valued such that summations span the same range negatively and positively, using symmetry of the signal and the DFT:

$$\mathcal{H}{f}[n] = \sum_{k} -i \operatorname{sgn}[k] \left[\cos\left(\frac{2\pi}{N}kn\right) + i \sin\left(\frac{2\pi}{N}kn\right) \right] [F_{\text{even}}[k] + i F_{\text{odd}}[k]]$$

$$= \sum_{k} \left\{ -i \operatorname{sgn}[k] \cos\left(\frac{2\pi}{N}kn\right) F_{\text{even}}[k] + \operatorname{sgn}[k] \sin\left(\frac{2\pi}{N}kn\right) F_{\text{even}}[k] + \operatorname{sgn}[k] \cos\left(\frac{2\pi}{N}kn\right) F_{\text{odd}}[k] + i \operatorname{sgn}[k] \sin\left(\frac{2\pi}{N}kn\right) F_{\text{odd}}[k] + i \operatorname{sgn}[k] \sin\left(\frac{2\pi}{N}kn\right) F_{\text{odd}}[k] \right\}$$

$$= \sum_{k} \operatorname{sgn}[k] \sin\left(\frac{2\pi}{N}kn\right) F_{\text{even}}[k] + \operatorname{sgn}[k] \cos\left(\frac{2\pi}{N}kn\right) F_{\text{odd}}, \quad (S10)$$

which is real-valued.

D. An even (odd) signal input to the DHT produces an odd (even) symmetry output

Using Eq. S10 and comparing $\mathcal{H}{f}[n]$ and $\mathcal{H}{f}[-n]$ as applied to a signal with even-symmetry f_{even} :

$$\mathcal{H}\{f_{\text{even}}\}[n] = \sum_{k} \text{sgn}[k] \sin\left(\frac{2\pi}{N}kn\right) F_{\text{even}}[k]$$
(S11)

$$\mathcal{H}{f_{\text{even}}}[-n] = \sum_{k} -\text{sgn}[k] \sin\left(\frac{2\pi}{N}kn\right) F_{\text{even}}[k]$$
(S12)

$$\therefore \mathcal{H}{f_{\text{even}}}[n] = -\mathcal{H}{f_{\text{even}}}[-n];$$
(S13)

thus an even input generates an odd-symmetry DHT output.

Similarly for an odd-symmetry signal f_{odd} :

$$\mathcal{H}\{f_{\text{odd}}\}[n] = \sum_{k} \operatorname{sgn}[k] \cos\left(\frac{2\pi}{N}kn\right) F_{\text{odd}}[k]$$
(S14)

$$\mathcal{H}{f_{\text{odd}}}[-n] = \sum_{k} \operatorname{sgn}[k] \cos\left(\frac{2\pi}{N}kn\right) F_{\text{even}}[k]$$
(S15)

$$\therefore \mathcal{H}{f_{\text{odd}}}[n] = \mathcal{H}{f_{\text{odd}}}[-n];$$
(S16)

thus an odd input generates an even-symmetry DHT output.

2. SIMULATED SPECTRA

A. Gaussian input spectra

In the main text of the manuscript, input spectra were of a Dawson form with a peaked lineshape (Gaussian) Hilbert transform. In this section, we consider the converse scenario, when input data is assumed to be Gaussian. The generation and constraints of the synthetic training and test data was identical to that of the main text. Similar to the main text results, the LeDHT is superior 98.74 % and shown in Fig. S1.

B. Multi-peak spectra

Similar to above, we demonstrate the scenario with multi-peak input spectra composed of Gaussian lineshapes. The generation and constraints of the synthetic training and test data was identical to that of the main text. In the main text the LeDHT was always superior for 30 000 spectra, and under this new scenario the LeDHT is superior to either the DHT or DHT-Pad 99.57 % of the 30 000 test spectra (see Fig. S2).

C. Training on Gaussian-Dawson lineshape pairs and testing on Lorentzian-dispersive lineshape pairs

In the main text, the lineshape types for training and testing were the same. In this section, we consider training with one type and testing with another. Ideally, the LeDHT matrix describes a relationship – a transform – and is not strictly a database for fitting input spectra. Thus, one would expect the LeDHT to still work, though possibly sub-optimally. Surprisingly, the LeDHT is superior to the DHT and DHT-Pad for all 300 000 test spectra (see Fig. S3).



Fig. S1. (a,e) Example Gaussian input spectra. (b,f) DHT results showing significant errors. (c,g) Results using the DHT with extrema-value padding (DHT-Pad). (d,h) Results using the LeDHT. (i) Mean MSE (\pm 1 standard deviation) calculated for the 300 000 stochastically generated test spectra as a function of peak center position. (j) Histogram of optimal method based on minimum MSE for test spectra dataset. Y-axis shows fraction, parenthetical values are number of spectra.



Fig. S2. (a) Example synthetic spectrum composed of 13 Gaussian functions. (b,c,d) The true Hilbert transform and the calculated Hilbert transforms using DHT, the DHT applied to extrema-value padded spectra (DHT-Pad), and the LeDHT, respectively. (e) Box plots of MSE for each method. Note the whiskers displays the extreme range of values. (f) Histogram of the optimal method (based on minimum MSE) for 30 000 test spectra. Y-axis shows fraction, parenthetical values are number of spectra.



Fig. S3. (a,e) Example dispersive lineshape input spectra. (b,f) DHT results showing significant errors. (c,g) Results using the DHT with extrema-value padding (DHT-Pad). (d,h) Results using the LeDHT. The LeDHT was trained with Gaussian and Dawson lineshapes though the new input data is dispersive with a Lorentzian Hilbert transform. (i) Mean MSE (\pm 1 standard deviation) calculated for the 300 000 stochastically generated test spectra as a function of peak center position. (j) Histogram of optimal method based on minimum MSE for test spectra dataset. Y-axis shows fraction, parenthetical values are number of spectra.

REFERENCES

- 1. W. L. Briggs and V. E. Henson, *The DFT: An Owner's Manual for the Discrete Fourier Transform* (Society for Industrial and Applied Mathematics, 1995).
- 2. L. Marple, "Computing the discrete-time "analytic" signal via FFT," IEEE Transactions on Signal Process. 47, 2600–2603 (1999).
- 3. P. Henrici, *Applied and computational complex analysis Vol III* (Wiley-Interscience, New York, 1986).
- 4. A. D. Poularikas, "Hilbert Transform," in *The Handbook of Forumulas and Tables for Signal Processing*, A. D. Poularikas, ed. (CRC, Boca Raton, FL, 1999), chap. 15.