

Analysis of complex multidimensional optical spectra by linear prediction

E. SWAGEL,^{1,2} J. PAUL,^{1,3} A. D. BRISTOW,^{1,3} AND J. K. WAHLSTRAND^{4,*}

¹Associate, Nanoscale Device Characterization Division, National Institute of Standards and Technology, Gaithersburg, MD 20899, USA

²Department of Physics, Brown University, Providence, RI 02912, USA

³Department of Physics and Astronomy, West Virginia University, Morgantown, WV 26506, USA

⁴Nanoscale Device Characterization Division, National Institute of Standards and Technology, Gaithersburg, MD 20899, USA

*jared.wahlstrand@nist.gov

Abstract: We apply Linear Prediction from Singular Value Decomposition (LPSVD) to two-dimensional complex optical data in the time-domain to generate spectra with advantages over discrete Fourier transformation (DFT). LPSVD is a non-iterative procedure that fits time-domain complex data to the sum of damped sinusoids, or Lorentzian peaks in the spectral domain. Because the fitting is linear, it is not necessary to give initial guess parameters as in nonlinear fits. Although LPSVD is a one-dimensional algorithm, it can be performed column-wise on two-dimensional data. The method has been extensively used in 2D NMR spectroscopy, where spectral peaks are typically nearly ideal Lorentzians, but to our knowledge has not been applied in the analogous optical technique, where peaks can be far from Lorentzian. We apply LPSVD to the analysis of zero, one, and two quantum electronic two-dimensional spectra from a semiconductor microcavity. The spectra consist of non-ideal, often overlapping peaks. We find that LPSVD achieves a very good fit even on non-ideal data. It reduces noise and eliminates discrete distortions inherent in the DFT. We also use it to isolate and analyze weak features of interest.

© 2022 Optical Society of America under the terms of the [OSA Open Access Publishing Agreement](#)

1. Introduction

Optical multidimensional coherent spectroscopy (MDCS) is a versatile technique for exploring light-matter interaction [1, 2]. Using multiple phase-stable femtosecond optical pulses to excite a sample, the nonlinear response is measured as a function of time delays between pulses. Spectra are spread onto two or more frequency axes, which allows effects that are subtle or invisible in linear spectroscopy to become readily apparent. Separation of homogeneous and inhomogeneous broadening is easily achieved, and contributions from various quantum pathways can be isolated. Many schemes for MDCS exist (for example, [3–7]), but all involve acquiring data in the time domain and transforming one or more axes using a discrete Fourier transform (DFT). DFT methods introduce spectral distortions inherent to the discrete nature of the method. For example, unless the data fall to zero during the scan, the DFT will present convolution with the step function, necessitating long scan time or artificial windowing [8, 9]. These distortions especially impact the lineshape of peaks, which is particularly problematic because lineshape contains information about broadening mechanisms and many-body interactions.

In nuclear magnetic resonance (NMR), an algorithm to improve on the DFT was developed that involves analyzing time domain data and fitting it to a linear combination of damped sinusoidal oscillators, or “poles.” This technique, linear prediction from singular value decomposition (LPSVD), has been used for decades in NMR spectroscopy [8–11]. It has also been applied to ultrafast optical measurements of coherent phonons [12]. LPSVD is a noniterative, autoregressive procedure which performs two linear least squares fits on a complex time-domain signal to

determine the frequency, damping rate, amplitude, and phase of the signal's constituent damped sinusoids. Because the fit is linear, it is efficient and requires no initial parameters except the number of poles. Additionally, LPSVD operates on complex time-domain data and effectively fits complex spectra, unlike many nonlinear fitting procedures, which work with spectral magnitude data. The employment of singular value decomposition, a generalization of eigenvalue decomposition to non-square matrices, enhances the stability and noise-filtering capabilities of linear prediction [10]. We note that LPSVD is related to compressive sensing [13] and orthogonal matching techniques [14, 15] but in contrast to those techniques, it is used on dense data sets.

Here, we apply linear prediction to zero-, one-, and two-quantum optical MDCS data from a semiconductor nanostructure sample. The technique has to the best of our knowledge not been previously applied to optical MDCS, perhaps because samples typically produce highly non-Lorentzian lineshapes, in contrast to NMR. We demonstrate that LPSVD successfully fits despite this high degree of non-ideality and is able to construct spectra without the discrete distortions of a DFT, while also rejecting noise. No modifications to the algorithm are needed for non-ideal data; peaks that are non-Lorentzian are fit by multiple poles. Because the spectrum is decomposed into poles, the spectral contributions of peaks can be separated, improving analysis of overlapping spectral peaks by allowing each to be analyzed in isolation. This technique is especially useful for isolation of weak spectral features. We show that LPSVD even isolates a weak feature not visible to the naked eye.

In the next section, we give a brief background on LPSVD, explain how it is used to construct 2D spectra, and discuss how it handles non-Lorentzian features. Then we give three examples using data taken on a semiconductor microcavity. First, we use LPSVD to calculate a zero-quantum spectrum, showing that LPSVD is able to remove distortions due to the DFT. Second, we use LPSVD to calculate one-quantum rephasing spectra, showing that it can be used to isolate weak features of interest. Finally, we use LPSVD to calculate two-quantum spectra.

2. Theory

Many previous papers have discussed linear prediction on complex data [8–11]. Multiple authors [16, 17] applied 1D linear prediction row-by-row to 2D sets of NMR data. Here we apply this approach to optical 2D spectra. Optical MDCS is virtually the same as multidimensional NMR mathematically, but the measurement scheme differs because of the relative difficulty of measuring an optical frequency field. In the experimental scheme we use, the complex field is retrieved using spectral interferometry, so the emission frequency axis is already in the spectral domain. We apply LPSVD to each column, once for each discrete emission frequency. Linear prediction can however be used on any one-dimensional time domain slice.

Linear prediction assumes that the signal is composed of the sum of damped sinusoids with white noise. We assume that the one-dimensional time domain data X_n is sampled with a uniform step Δt ,

$$X_n = \sum_{j=1}^m A_j e^{n\Delta t(-b_j+i\omega_j)+i\phi_j} + w_n, \quad (1)$$

where m is the number of sinusoids; A_j , b_j , ω_j , and ϕ_j are the amplitude, damping coefficient, frequency, and phase of the j th sinusoid, respectively; and w_n is a noise term. In the spectral domain, these sinusoids manifest as Lorentzian peaks.

The first step in the LPSVD algorithm is to perform singular value decomposition (SVD) on a data matrix, a Hankel matrix constructed as described in [18]. It can be proved that there is one nonzero singular value for each sinusoid given an ideal signal [10]. With noisy data, all singular values will be nonzero, and noise-related singular values may be distinguished by their smaller magnitudes and set to zero. Barkhuijsen *et al.* also proposed subtracting the mean of the noise-related singular values from the signal-related singular values to improve the signal-to-noise

ratio (SNR) [10]. The Hankel matrix is then reconstructed from the manipulated SVD. Use of SVD improves the numerical stability of the linear fitting process and increases the SNR [18]. In fact, the improvement from SVD makes linear prediction approach the upper Cramer-Rao fitting bound [19].

In the second step, an autoregressive model is employed, meaning every discrete data point is assumed to be a linear combination of the previous k points, a parameter chosen when fitting. A linear least squares fit gives the linear coefficients from the Hankel data matrix [18]. For a noiseless signal, k should be m , the number of sinusoids of which the signal is comprised. For noisy data, many coefficients should be used, with 3/4 the length of the signal as a general suggestion [19]. The frequency and damping coefficient may be extracted from these coefficients by finding the roots of the polynomial $1 - c_1z^1 - c_2z^2 - \dots - c_kz^k$, where c_1, \dots, c_k are the autoregression coefficients. Each of these roots will be of the form $e^{-b_j+i\omega_j}$, from which the natural logarithm gives the damping and angular frequency of the j th sinusoid.

The third step is a simple linear least squares fitting to find the quantities $A_j e^{i\phi_j}$ using the known values of X_n and of $-b_j + i\omega_j$ in Eq. (1). Not all roots found in the second step are used in the linear fit. Typically, Kumaresan and Tufts' condition of decaying sinusoids is enforced, i.e. all roots have absolute value less than or equal to one [18]. This condition makes sense physically and eliminates some noise-related sinusoids. However, when the signal decays to a constant rather than to zero, as in zero-quantum scans, we have found that keeping roots with slightly larger magnitudes results in a better fit. This will be discussed in more detail when we apply LPSVD to optical MDCS data.

A Python package implementing the algorithm, along with the data used to generate the figures presented in Section 3, can be found in Code 1 ([20]) [21].

2.1. Constructing spectra

Linear prediction is typically used in combination with DFT methods to enhance the resolution of the discrete spectrum, which is proportional to the total number of sampled points. The time-domain signal is extended forward with linear prediction and then a DFT is applied [8, 10, 22] to produce a spectrum. With more points in the time domain, the DFT has higher resolution. There are, however, discrete distortions inherent to DFT methods [8]. Using linear prediction to extend the signal forward enhances the spectral resolution and filters some of the noise, but the discrete distortions and some noisy data will inherently affect the DFT.

Synthesizing the spectrum directly from the linear prediction results, on the other hand, eliminates these discrete effects. The one-dimensional spectral equation can be found through an analytical Fourier transformation of Eq. (1):

$$\hat{S}(\omega) = \frac{1}{2\pi} \sum_{j=1}^m A_j \frac{e^{i\phi_j}}{i(\omega - \omega_j) - b_j}. \quad (2)$$

A two-dimensional spectrum may be constructed column-wise.

In a review of NMR applications, Koehl found that direct linear prediction spectrum construction was rarely used [8], although Led and Gesmar mentioned the technique [9]. This is likely because linear prediction has previously been applied mostly to NMR, in which the data are highly ideal, meaning each pole fits one peak (see Sec. 2.2), and further analysis may be performed directly from the fitting parameters of each pole found using LPSVD.

2.2. Handling non-ideality

Optical MDCS spectra in condensed matter systems are often highly distorted from the ideal line-shape calculated from simple level schemes using the Bloch equations because of inhomogeneous broadening [23, 24], many-body effects [25, 26], and other factors such as propagation effects [27].

So it might seem odd to effectively fit slices of the spectra to Eq. (2). Although our data are highly non-ideal, we generally find that linear prediction closely matches the time-domain signal and the DFT-generated one-dimensional spectral slices and full two-dimensional spectrum (in data sets where discrete distortions are minimal).

This works because each peak in the spectrum can be fit to more than one pole. In the context of NMR, Balcou and Balcou showed that a variety of non-Lorentzian peaks may be approximated with the sum of two Lorentzians, but there are many stable pairs of Lorentzian peaks which match the non-Lorentzian peak with identical accuracy [28]. In practice, this means that the amplitudes, damping, and phases of individual poles do not necessarily carry physical significance, especially two nearby poles. Instead, linewidths and other spectral data may be extracted numerically from the synthesized spectrum, as would be done for a DFT-generated spectrum. The LPSVD generated spectrum is less noisy and rejects artifacts due to early truncation of scans.

3. Applications

We demonstrate LPSVD on experimental data from a semiconductor microcavity. The experimental apparatus uses a multidimensional optical nonlinear spectrometer (MONSTR [5]), a system that produces four phase-stable pulses through active feedback. Details of our experimental implementation are given in [29]. The sample is a single 8-nm $\text{In}_{0.04}\text{Ga}_{0.96}\text{As}$ quantum well embedded in a λ microcavity that is slightly wedged so that the detuning Δ between the cavity mode and the exciton energy can be adjusted continuously. The cavity mirrors are distributed Bragg reflectors composed of GaAs/AlAs bilayers (14.5 bilayers on one side and 12 on the other). Details of the sample and typical detuning dependent multidimensional spectra can be found in [30].

We acquire complex four-wave mixing spectra through spectral interferometry with a reference pulse. Spectra are captured as a function of a time delay between two of the three excitation pulses. It is most natural to use LPSVD along this time axis, once for each discrete emission frequency. Note that for other techniques, which operate entirely in the time domain [6], it might be optimal to perform LPSVD along another time axis.

3.1. Illustrative example: a zero quantum spectrum

Zero quantum scans [1], where the time delay T between the second and third pulses is scanned, are generally used to explore population dynamics and Raman coherence. Figures 1(ab) show the full DFT and LPSVD spectra for cross-linear (VHHV) polarization at near zero detuning. The peaks at lower (higher) emission energy are associated with the lower (upper) polariton branch. Peaks near zero absorption energy usually correspond to population dynamics, while peaks at nonzero absorption energy correspond to Raman coherence. To show the LPSVD time domain fit, we plot in Fig. 1(c) a time-domain vertical slice (shown as a dashed green line in Fig. 1(b)) through R^+ , split into real and imaginary parts. Figures 1(def) show the same vertical slice through the spectrum, with the contribution of each pole depicted using a dotted black line.

This figure illustrates much of the discussion in previous sections on linear prediction's advantages over DFT methods. First, as is clear, linear prediction fits well the signal in the time domain. To achieve this quality of fit for zero quantum scans, we had to keep roots with magnitude less than 1.1 rather than the usual threshold of 1.0. This is equivalent to keeping some poles with damping rates b that are slightly negative. Note that the terms with $b < 0$ also have extremely low frequency, so low that they oscillate over less than half a period over the scan time. We attribute this to population dynamics that cause the non-oscillating component of the signal to shrink and/or grow before coming to steady state at time delays outside our measurement range. Because the real part decays to a constant, the DFT reflects convolution with the step function, as can be seen in the "wiggles" in the DFT spectrum at low absorption energy in Fig. 1(d) [8]. Note that we employed an inverse tangent window to reduce this distortion, but it still exists and

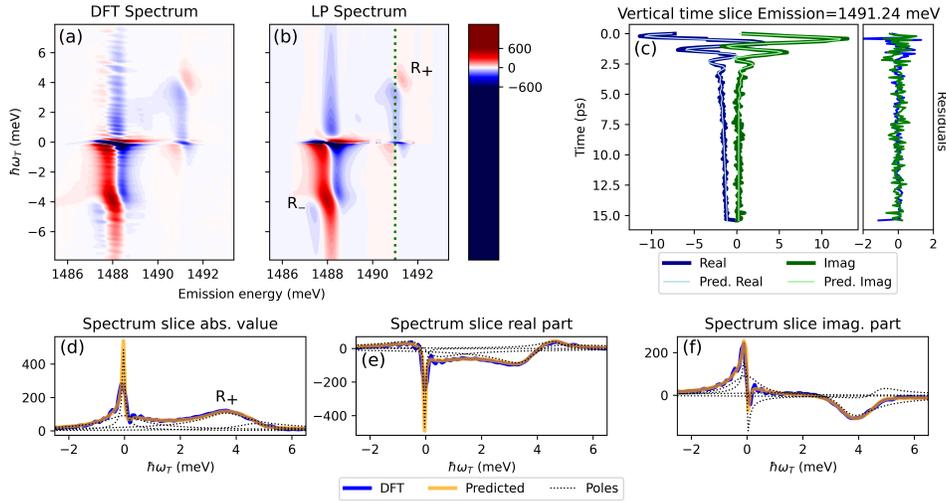


Fig. 1. Experimental spectra (magnitude) generated from a zero quantum scan by (a) DFT and (b) LPSVD. Four poles were used in the LPSVD. The real part of the complex 2D spectrum is plotted in (a) and (b). A vertical slice at a particular emission energy (green dashed line in part (b)) is shown for (c) the time domain signal and (def) the spectrum. In (d), the amplitude contribution of each pole is shown by the black dotted line (phase is obscured), while the magnitude of the sum of these contributions (including phase offsets) is shown in orange. In (ef), the real and imaginary parts are plotted for the poles and for the spectrum, highlighting the fit of multiple poles to a single non-ideal peak. All units not given are arbitrary units (e.g. for the colorbar in (ab)).

resolution is hurt by this windowing. The spectrum generated by linear prediction does not have this issue. Also note that removing the constant offset in the real part is insufficient because the rapid exponential decay, corresponding to the wide zero frequency pole in Fig. 1(d), remains.

Lastly, as noted previously [28], multiple LP poles can be used to fit non-Lorentzian peaks without any modification to the algorithm. This can be seen in the black dotted lines in Fig. 1(def), as there are two low frequency LP poles used to fit the single physical zero frequency signal, and two poles near the Raman peak. Because LPSVD fits so well to the time-domain data, it is clear that the zero frequency peak constructed is reliable even though it disagrees with the DFT. As mentioned above, the zero frequency peaks arise in part from population dynamics, which are typically best studied in the time domain.

3.2. Isolating weak features: one quantum spectra

As is clear from the previous section, decomposition of the spectrum into poles allows straightforward isolation of spectral contributions of weak peaks. Here we demonstrate that application of LPSVD. In one quantum scans of the microcavity sample, four peaks appear in the 2D spectra, corresponding to lower and upper polariton branches (LP and UP along the diagonal plus corresponding cross peaks LP-UP and UP-LP), along with a much weaker biexciton-related feature [30]. The polariton peak positions change sensitively with detuning Δ , while the biexciton remains relatively fixed at an absorption energy around 1491 meV. The detuning Δ is adjusted by moving the position of the laser focus on the sample (recalling that the microcavity is wedged). We plot the spectral locations of the polariton peaks as a function of spot position. The position

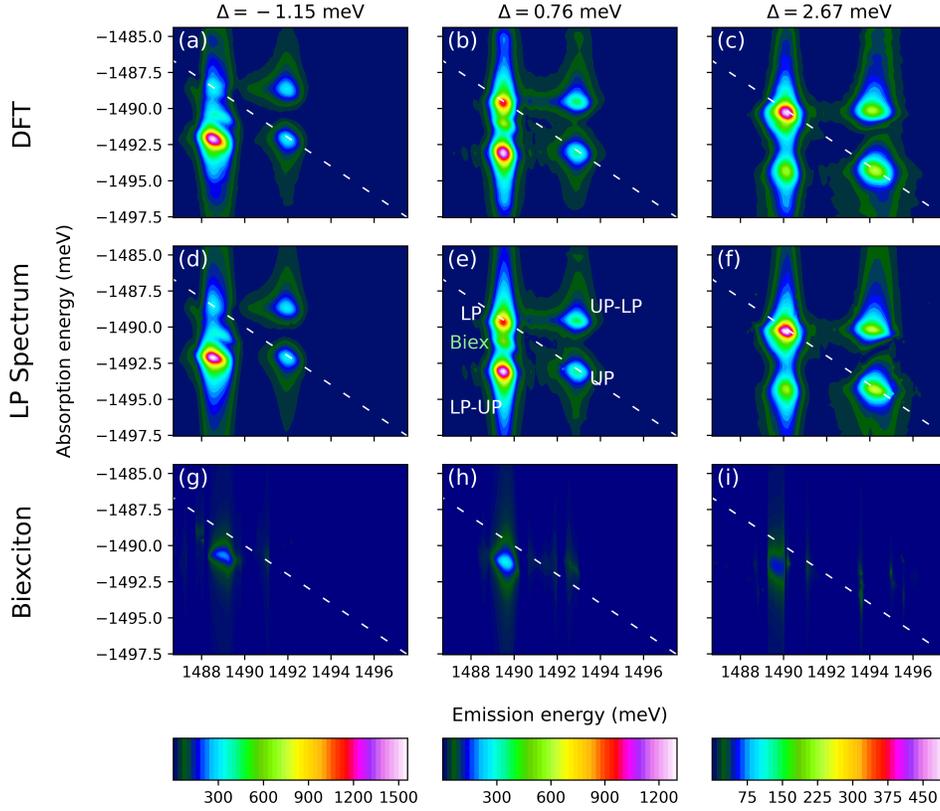


Fig. 2. Experimental spectra generated from one quantum scans by (a-c) DFT and (d-f) LPSVD. Six poles were used for LPSVD. (g-i) Spectra generated from only the biexciton poles, selected by frequency as described in the main text. Subfigures g-i are displayed at twice the brightness for clarity. The color scale is in arbitrary units.

at which the peak separation is minimized corresponds to $\Delta = 0$, and the peak separation is the vacuum Rabi splitting E_{VRS} . The detuning can then be found from the peak positions by solving $E_{\text{UP/LP}} = 1/2[2E_X + \Delta \pm (\Delta^2 + E_{\text{VRS}})^{1/2}]$ [30].

At small positive values of Δ , the biexciton peak is difficult to distinguish from the wide tail of nearby stronger polariton peaks. This is illustrated in Figs. 2(abc), which show DFT-generated rephasing spectra for cross-circular polarization ($\sigma^- \sigma^+ \sigma^+ \sigma^-$) at three values of Δ . Note that at $\Delta = 2.67$ meV [Fig. 2(c)], the biexciton peak is not visible to the naked eye. Corresponding spectra generated by LPSVD are shown in Fig. 2(def). The predicted spectra reproduce all the features of the spectrum.

At low values of Δ , the biexciton feature is more easily distinguished from the polariton features, and we can guess its position for the scan with $\Delta = 2.67$ meV. Figure 2(ghi) show spectra where only the posited biexciton-related poles are used, distinguished by their absorption energy. The technique performs remarkably well near zero detuning [Fig. 2(gh)], where the signal is relatively strong. The isolated biexciton peaks visually match the shape of the partially visible biexciton peaks in Fig. 2(de). At $\Delta = 2.67$ meV [Fig. 2(i)], where it is impossible to discern by eye, LPSVD is able to produce a signal exactly where we expect the biexciton to be. Because of this, we may trust the amplitude of the biexciton in the synthesized spectrum. However, there are sharp vertical cutoffs visible which hinder lineshape analysis. This results from the failure of

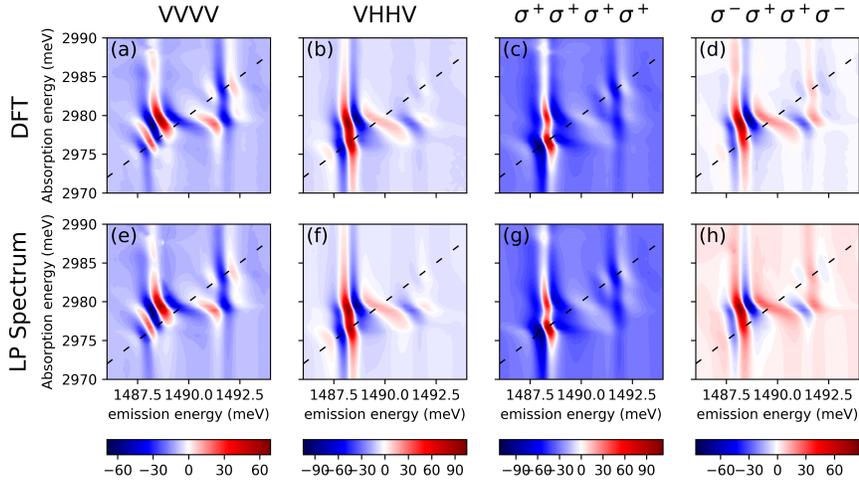


Fig. 3. Experimental spectra (real part) generated from two quantum scans with various light polarization configuration by (a-d) DFT and (e-h) LPSVD. Four poles were used for fitting all polarizations except for VVVV, which required 14 poles. Colorbar values are in arbitrary units.

LPSVD to find a pole above the threshold of noise when the emission frequency is too far from the peak, and is a symptom of the column-wise application of the algorithm. Nevertheless, the fit is sufficient to isolate the biexciton and find its amplitude and phase in a spectrum in which it would not otherwise be possible.

3.3. Highly non-ideal spectra: two quantum spectra

Finally, we show that linear prediction can fit to highly non-ideal and noisy data. For completeness, we draw this example from two quantum scans, where the time delay between the second and third pulses is scanned. In two quantum scans of the microcavity, a strong polarization dependence is evident. This dependence is illustrated in Fig. 3(a-d). In Fig. 3(e-h), corresponding spectra generated by LPSVD are shown. As is evident, the spectra are in close agreement. Although some differences can be seen, especially at the tails of peaks, these are mostly due to noise in the DFT spectrum that is not present in the linear prediction spectrum. Even for VVVV polarization (Fig. 3(e)), where the spectrum is highly non-ideal, the LPSVD-generated spectrum closely resembles the DFT spectrum. The caveat, however, is that Fig. 3(f-h) were fitted using four poles, whereas Fig. 3(e) required 14 poles to get the best fitting. The use of more poles is necessary with low SNR and high non-ideality, as the largest singular values are not necessarily well-fitted to signal-related Lorentzians (See Sec. 3.4).

3.4. Performance considerations

LPSVD efficiently iterates across 1024 columns of length 125 in seconds on a modern laptop on a single thread. Because it operates column-by-column, the procedure would be trivially parallelizable. The DFT, in contrast, takes milliseconds. Substantial time savings can be had by only using linear prediction to fit columns with emission energy near the spectral features of interest. In our scans, this cuts runtime by up to a factor of four.

One of the advantages of LPSVD is the fact that one does not need to guess initial parameters of the fits. However, hand tuning of the number of poles is typically required, depending on the spectrum's ideality and signal-to-noise ratio. The number of poles does not necessarily directly

correspond to the number of spectral peaks. Therefore, linear prediction must be initially run with multiple numbers of poles to see which number of poles produces the best fit by analysis of residuals in the time domain. This typically must only be done once for a type of scan (e.g. zero quantum scans for $\sigma^+\sigma^+\sigma^+\sigma^+$ polarization). Generally, two poles for each peak is a good initial guess. We emphasize that spectral information such as linewidths and peak positions should be extracted numerically from the synthesized spectrum rather than from pole parameters.

4. Conclusion

In summary, we have applied LPSVD to optical MDCS and demonstrated the technique's ease of use and applicability to the analysis of non-ideal optical electronic multidimensional spectra. We showed linear prediction's advantages over DFT methods in filtering noise and eliminating discrete distortions. Additionally, linear prediction proves extremely useful in isolating weak spectral features. All of the data shown here came from a microcavity sample, but we have also successfully used LPSVD on data from a multiple quantum well sample and observed similar performance. While LPSVD is more computationally intensive than the DFT, the time it takes to perform the analysis is much smaller than the acquisition time, and it can be used to perform real time analysis. With enough signal-to-noise ratio (and sufficiently ideal lineshapes), LPSVD can allow one to generate accurate spectra from truncated scans. For our non-ideal spectra we found that we could remove up to 25% of the data points at the end of our scans without significantly changing the spectra constructed by LPSVD. This implies that application of LPSVD to spectroscopy of this system could shorten data acquisition time by up to 25%.

We applied one-dimensional LPSVD column-wise to 2D spectra, which although practically effective ignores the relationship between columns when fitting. Some authors have attempted fully two-dimensional linear prediction [16, 17], but generally found that it is unstable and less accurate than one-dimensional methods. Additionally, it is computationally expensive. Nevertheless the development of simple 2D LPSVD techniques deserves further exploration.

Acknowledgments

E. S. was supported by the NIST Summer Undergraduate Research Fellowship program. This research was conducted through the support of NIST grant 70NANB18H238.

Disclosures

The authors declare no conflict of interest.

Data availability

Data underlying the results presented in this paper are available in Code 1, Ref. [20].

References

1. S. Mukamel, "Multidimensional femtosecond correlation spectroscopies of electronic and vibrational excitations," *Annu. Rev. Phys. Chem.* **51**, 691–729 (2000).
2. G. Moody and S. T. Cundiff, "Advances in multi-dimensional coherent spectroscopy of semiconductor nanostructures," *Adv. Physics: X* **2**, 641–674 (2017).
3. J. Hybl, A. Ferro, and D. Jonas, "Two-dimensional Fourier transform electronic spectroscopy," *J. Chem. Phys.* **115**, 6606–6622 (2001).
4. J. Vaughan, T. Hornung, K. Stone, and K. Nelson, "Coherently controlled ultrafast four-wave mixing spectroscopy," *J. Phys. Chem. A* **111**, 4873–4883 (2007).
5. A. D. Bristow, D. Karaiskaj, X. Dai, T. Zhang, C. Carlsson, K. R. Hagen, R. Jimenez, and S. T. Cundiff, "A versatile ultrastable platform for optical multidimensional Fourier-transform spectroscopy," *Rev. Sci. Instrum.* **80**, 073108 (2009).
6. G. Nardin, T. M. Autry, K. L. Silverman, and S. T. Cundiff, "Multidimensional Coherent Photocurrent Spectroscopy of a Semiconductor Nanostructure," *Opt. Express* **21**, 28617–28627 (2013).

7. J. O. Tollerud and J. A. Davis, "Coherent multi-dimensional spectroscopy: Experimental considerations, direct comparisons and new capabilities," *Prog. Quantum Electron.* **55**, 1–34 (2017).
8. P. Koehl, "Linear prediction spectral analysis of NMR data," *Prog. Nucl. Magn. Reson. Spectrosc.* **34**, 257–299 (1999).
9. J. J. Led and H. Gesmar, "Application of the linear prediction method to NMR spectroscopy," *Chem. Rev.* **91**, 1413–1426 (1991).
10. H. Barkhuijsen, R. de Beer, W. M. M. J. Bovée, and D. van Ormondt, "Retrieval of frequencies, amplitudes, damping factors, and phases from time-domain signals using a linear least-squares procedure," *J. Magn. Reson.* (1969) **61**, 465–481 (1985).
11. H. Barkhuijsen, R. de Beer, and D. van Ormondt, "Improved algorithm for noniterative time-domain model fitting to exponentially damped magnetic resonance signals," *J. Magn. Reson.* (1969) **73**, 553–557 (1987).
12. F. Wise, M. Rosker, G. Millhauser, and C. Tang, "Application of linear prediction least-squares fitting to time-resolved optical spectroscopy," *IEEE J. Quantum Electron.* **23**, 1116–1121 (1987).
13. S. Adhikari, C. L. Cortes, X. Wen, S. Panuganti, D. J. Gosztola, R. D. Schaller, G. P. Wiederrecht, and S. K. Gray, "Accelerating Ultrafast Spectroscopy with Compressive Sensing," *Phys. Rev. Appl.* **15**, 024032 (2021).
14. Y. C. Eldar, P. Kuppinger, and H. Bolcskei, "Block-Sparse Signals: Uncertainty Relations and Efficient Recovery," *IEEE Transactions on Signal Process.* **58**, 3042–3054 (2010).
15. H. Frostig, T. Bayer, Y. C. Eldar, and Y. Silberberg, "Revealing true coupling strengths in two-dimensional spectroscopy with sparsity-based signal recovery," *Light. Sci. & Appl.* **6**, e17115 (2017).
16. J. Gorcester and J. H. Freed, "Linear prediction and projection of pure absorption lineshapes in two-dimensional FTESR correlation spectroscopy," *J. Magn. Reson.* (1969) **78**, 292–301 (1988).
17. G. Zhu and A. Bax, "Two-dimensional linear prediction for signals truncated in both dimensions," *J. Magn. Reson.* (1969) **98**, 192–199 (1992).
18. R. Kumaresan and D. Tufts, "Estimating the parameters of exponentially damped sinusoids and pole-zero modeling in noise," *IEEE Transactions on Acoust. Speech, Signal Process.* **30**, 833–840 (1982).
19. D. Tufts and R. Kumaresan, "Estimation of frequencies of multiple sinusoids: Making linear prediction perform like maximum likelihood," *Proc. IEEE* **70**, 975–989 (1982).
20. E. Swagel, J. Paul, A. D. Bristow, and J. K. Wahlstrand, "Supplementary materials for Analysis of complex multidimensional optical spectra by linear prediction," figshare (2021), <https://osapublishing.figshare.com/s/d069a73003c57b3ce32a>
21. We adopted a python implementation of LPSVD by David P. Hoffman that can be found at <https://github.com/david-hoffman/dphtools>.
22. H. Feng, S. Cai, Z. Chen, M. Lin, and J. Feng, "Application of the forward linear prediction on high-resolution NMR spectra in inhomogeneous fields," *Spectrochimica Acta Part A: Mol. Biomol. Spectrosc.* **71**, 1027–1031 (2008).
23. M. E. Siemens, G. Moody, H. Li, A. D. Bristow, and S. T. Cundiff, "Resonance lineshapes in two-dimensional Fourier transform spectroscopy," *Opt. Express* **18**, 17699–17708 (2010).
24. J. D. Bell, R. Conrad, and M. E. Siemens, "Analytical calculation of two-dimensional spectra," *Opt. Lett.* **40**, 1157–1160 (2015).
25. D. Keusters and W. Warren, "Propagation effects on the peak profile in two-dimensional optical photon echo spectroscopy," *Chem. Phys. Lett.* **383**, 21–24 (2004).
26. H. Li, A. P. Spencer, A. Kortyna, G. Moody, D. M. Jonas, and S. T. Cundiff, "Pulse Propagation Effects in Optical 2D Fourier-Transform Spectroscopy: Experiment," *J. Phys. Chem. A* **117**, 6279–6287 (2013).
27. A. P. Spencer, H. Li, S. T. Cundiff, and D. M. Jonas, "Pulse Propagation Effects in Optical 2D Fourier-Transform Spectroscopy: Theory," *J. Phys. Chem. A* **119**, 3936–3960 (2015).
28. Y. Balcou and P. Balcou, "Study of ambiguities inherent to the spectral analysis of Voigt profiles—a modified Simplex approach," *Chemom. Intell. Lab. Syst.* **63**, 41–55 (2002).
29. J. K. Wahlstrand, G. M. Wernsing, J. Paul, and A. D. Bristow, "Automated polarization-dependent multidimensional coherent spectroscopy phased using transient absorption," *Opt. Express* **27**, 31790–31799 (2019).
30. B. L. Wilmer, F. Passmann, M. Gehl, G. Khitrova, and A. D. Bristow, "Multidimensional coherent spectroscopy of a semiconductor microcavity," *Phys. Rev. B* **91**, 201304(R) (2015).