Applying wavelet-based denoising techniques to electron spin resonance spectra

Easton K. Crockett ’28, J. Anthony Cortez ’28, Herbert J. Sipe, Jr., and Dylan Domel-White

Department of Chemistry, Department of Mathematics and Computer Science, Hampden-Sydney College, Hampden-Sydney, VA 23943

Abstract

In the field of electron spin resonance spectroscopy, collected data is often difficult to interpret due to noise. Several methods exist for removing noise from collected data, and recent research suggests a denoising method using wavelets may be more effective than traditional methods using the Fourier transform. The ability to denoise ESR data would be incredibly helpful, as it would eliminate the need for repetitive testing, and save time and material in the experimental process. Dr Sipe had previously collected ESR data, but was unable to interpret it due to noise, but this new method may provide a way to improve the legibility of the data. To apply this new method, we first needed to study the mathematical basis behind wavelets and the Fourier transform. We then applied the method to example noisy data. Additionally, we investigated a method for recreating the data without the need for denoising by extracting coupling constants from the noisy data. We found that the wavelet method is effective at denoising the ESR data but becomes less effective the more noise is present in the collected data. Coupling constant extraction proved difficult, but some progress was made in developing an effective method. However, we were able to develop a strong skill set in both mathematics and programming that will allow us to pursue these topics further in the future.

Electron Spin Resonance Spectroscopy

ESR spectroscopy is a method for analyzing free radicals, or molecules with unpaired electrons. By placing these molecules in a linearly increasing magnetic field, they will transition between their magnetic energy levels, and these transitions can be detected by measuring the energy absorption when they occur (Aryal, 2022). This method has a variety of applications, with its main use being in material science. The biggest issue faced in applying ESR spectroscopy is the presence of noise in the collected data. Noise can be thought of as random numbers of a certain magnitude that are added to each data point. A variety of environmental factors and sample imperfections can contribute large quantities of noise to collected readings, rendering them uninterpretable. To combat this, denoising methods can be applied, although these methods vary in effectiveness.

A secondary objective when analyzing collected ESR spectroscopy data is the extraction of coupling constants from the readings. These constants are unique to a specific molecule and allow for the creation of simulated spectra for that molecule (Roy & Srivastava, 2022). Figure 1 demonstrates how these constants can be visually discerned from plotted data. In the case where a given data set is too noisy to be effectively denoised, a method for extracting coupling constants without first denoising would be incredibly valuable.

Fourier Transform

The Fourier Transform of a function takes a function and transforms it into another function that describes the various frequency contributions present in the original function (Bogges & Narcowich, 2009). For our research, we were mostly working with the discrete Fourier transform (DFT), because ESR data is recorded as a list of discrete points. The DFT functions similarly to the Fourier transform but instead outputs a discrete function with the same number of points as the original to represent the frequency information (Bogges & Narcowich, 2009). Figure 2 displays an example ESR spectra and its Fourier transform.

The discrete Fourier transform can be used to denoise data by removing certain frequency contributions from the transform before reconstructing back to the original function. If the frequency contributions removed represent noise the result will be a less noisy spectra upon reconstruction. Figure 3 displays the effects of removing high frequency contributions from Figure 2.

However, this method is not effective for denoising ESR spectra, because it assumes that frequency contributions do not change throughout the original function, which is obviously not for true ESR spectra. This issue can be somewhat remedied by using the short time Fourier transform, which breaks a signal up into several different transforms of the same length, but wavelets provide a much simpler solution (Sahoo & Srivastava, 2025).

Wavelet Decomposition

A wavelet is simply a single oscillation of a wave. These wavelets are used to decompose a signal into its high and low frequency components. This is done using two different wavelets, one that detects high frequency signal information, and one that detects low frequencies. This breaks a signal into a high frequency detail component and a low frequency approximation component (Sahoo & Srivastava, 2025). Figure 4 depicts two of the wavelets that were used for this research. This process is done several times, decreasing the resolution of these wavelets each time to effectively capture lower frequency signal features. These resolution levels are called the decomposition levels. The process of decomposing a signal in this way is called the discrete wavelet transform.

Denoising a signal using the discrete wavelet transform involves thresholding the detail component of each decomposition level, because noise is often high frequency, and it will be largely represented in the detail component, while lower frequency signal features will not be (Srivastava et al., 2021). The signal can then be reconstructed using thresholded detail components to remove noise from the original signal.

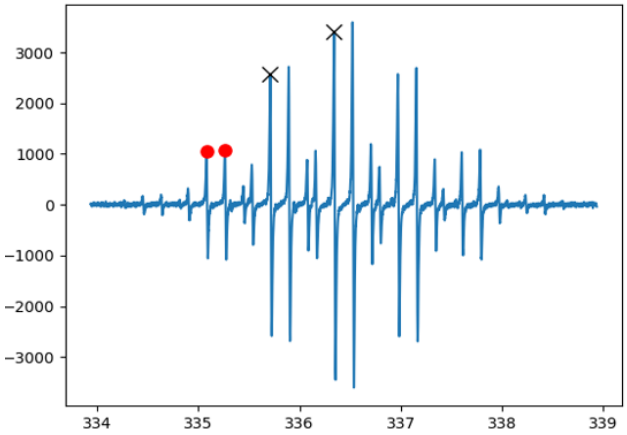

Figure 1. Phenalenyl with Coupling Constants Marked. Note the Coupling constants are measured as the distance between the x-axis crossings, but they are marked from peak-to-peak to improve readability.

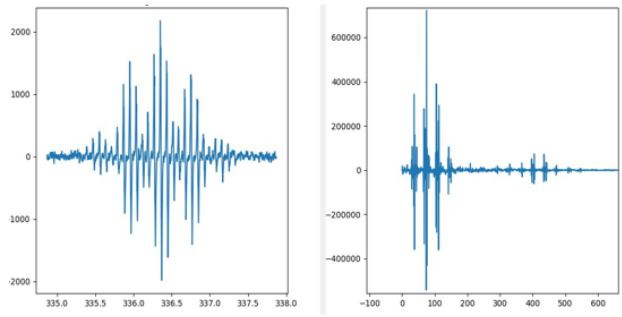

Figure 2. Input Function Next to the Discrete Fourier Transform of the Input Function. Note the DFT graph is zoomed in to better show frequency information.

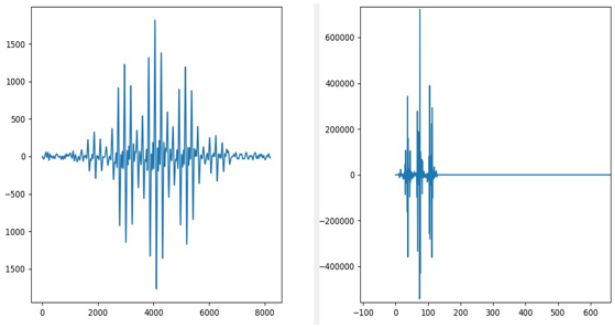

Figure 3. The Effects of Removing High Frequency Contributions from Figure 2

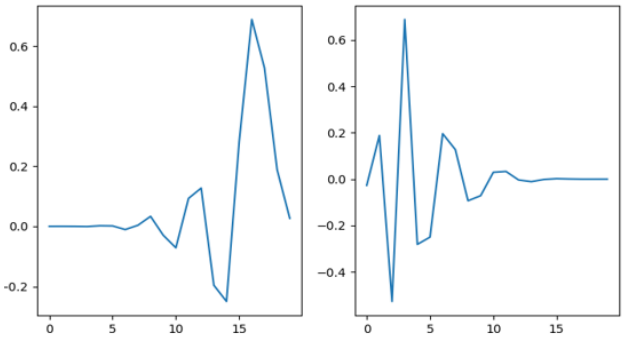

Figure 4. Low Frequency Detecting and High Frequency Detecting Daubechies Wavelets

Methods

Fourier denoising served as a precursor to wavelet denoising for this research. It involves transforming a discrete function from the time domain to the frequency domain, where certain frequency components can be filtered out, and then reconstructing the signal without those frequencies (Bogges & Narcowich, 2009). The component selection is done through either vertical or horizontal thresholding (horizontal thresholding may be referred to as windowing). Components selected this way are then reduced by some proportional amount or zeroed out entirely. This method is not very effective at denoising ESR spectra for the reasons mentioned above.

Wavelet-based Denoising

As stated above, the discrete wavelet transform decomposes an original data set into several detail and

approximation components that can be individually denoised and then reconstructed. There are several parameters that can be altered to achieve a better result.

- Wavelet Selection – A variety of wavelets can be used to decompose the original data set, and these vary in effectiveness. However, this research mainly used the Daubechies and Coiflet wavelets, as they have been empirically proven to be most effective for ESR spectra (Srivastava et al., 2021)
- Decomposition Level Selection – This refers to the number of levels that thresholding is performed on. As a data set is decomposed more and more, prominent signal features begin to be less visible, and denoising is no longer viable, because there is a risk of thresholding actual signal. The optimal decomposition level can be determined visually or through a calculation based on the sparsity of the data set (Srivastava et al., 2021)

Once these parameters are chosen, the detail components of each decomposition level can be denoised using vertical and horizontal thresholding. This research did not arrive at an objective method for determining these thresholds, rather they were chosen subjectively for each decomposition level. These denoising operations were performed in NERD, a web-based software developed by ACERT designed specifically for ESR data denoising (ACERT's website can be found here https://www.acert.cornell.edu/). Some charts were also produced using EZDenoise, a software developed over the course of this research that functions very similarly to NERD.

Coupling Constant Extraction

Developing a method for extracting coupling constants from collected ESR data was a secondary objective of this research. The literature included a prospective method, but not much work was completed in testing this method. A promising method involving the continuous wavelet transformation was also investigated.

The first method involved using specific levels of the discrete wavelet transform to extract the coupling constants. This method was developed by Dr. Srivastava and is detailed in the 2022 paper Hyperfine Decoupling of ESR Spectra Using Wavelet Transform.

The second method involved the continuous wavelet transform to visually observe coupling constants even from very noisy data. This method is most easily demonstrated visually, and an example is included in the results section below.

Results

The methods described above were applied to numerous data sets over the course of this research. Only two data sets are shown here for brevity. The first includes a level-by-level breakdown of the denoising process, but a similar process was used for every data set that was denoised.

Biphenylene

The coif3 wavelet was used for this denoising. The optimal decomposition The db6 wavelet was used for this denoising, because the features in the signal were present for more decomposition levels than for the coif3 wavelet. The sparsity-based decomposition level selection calculated an optimal decomposition level of ten, but level eight was sufficient, because the features completely disappeared from the approximation component at that level. All windows were chosen subjectively. This resulted in a 97.5% reduction in noise on edges of the signal, and a clear increase in legibility to the center. The use of signal windowing allows for targeted denoising on the edges of the signal, which accounts for the large noise reduction.

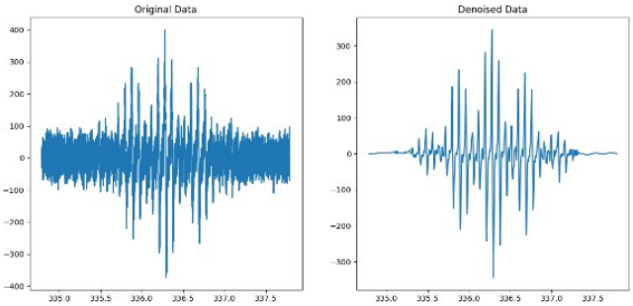

Figure 5. Noisy Biphenylene Data Compared to Denoised Data

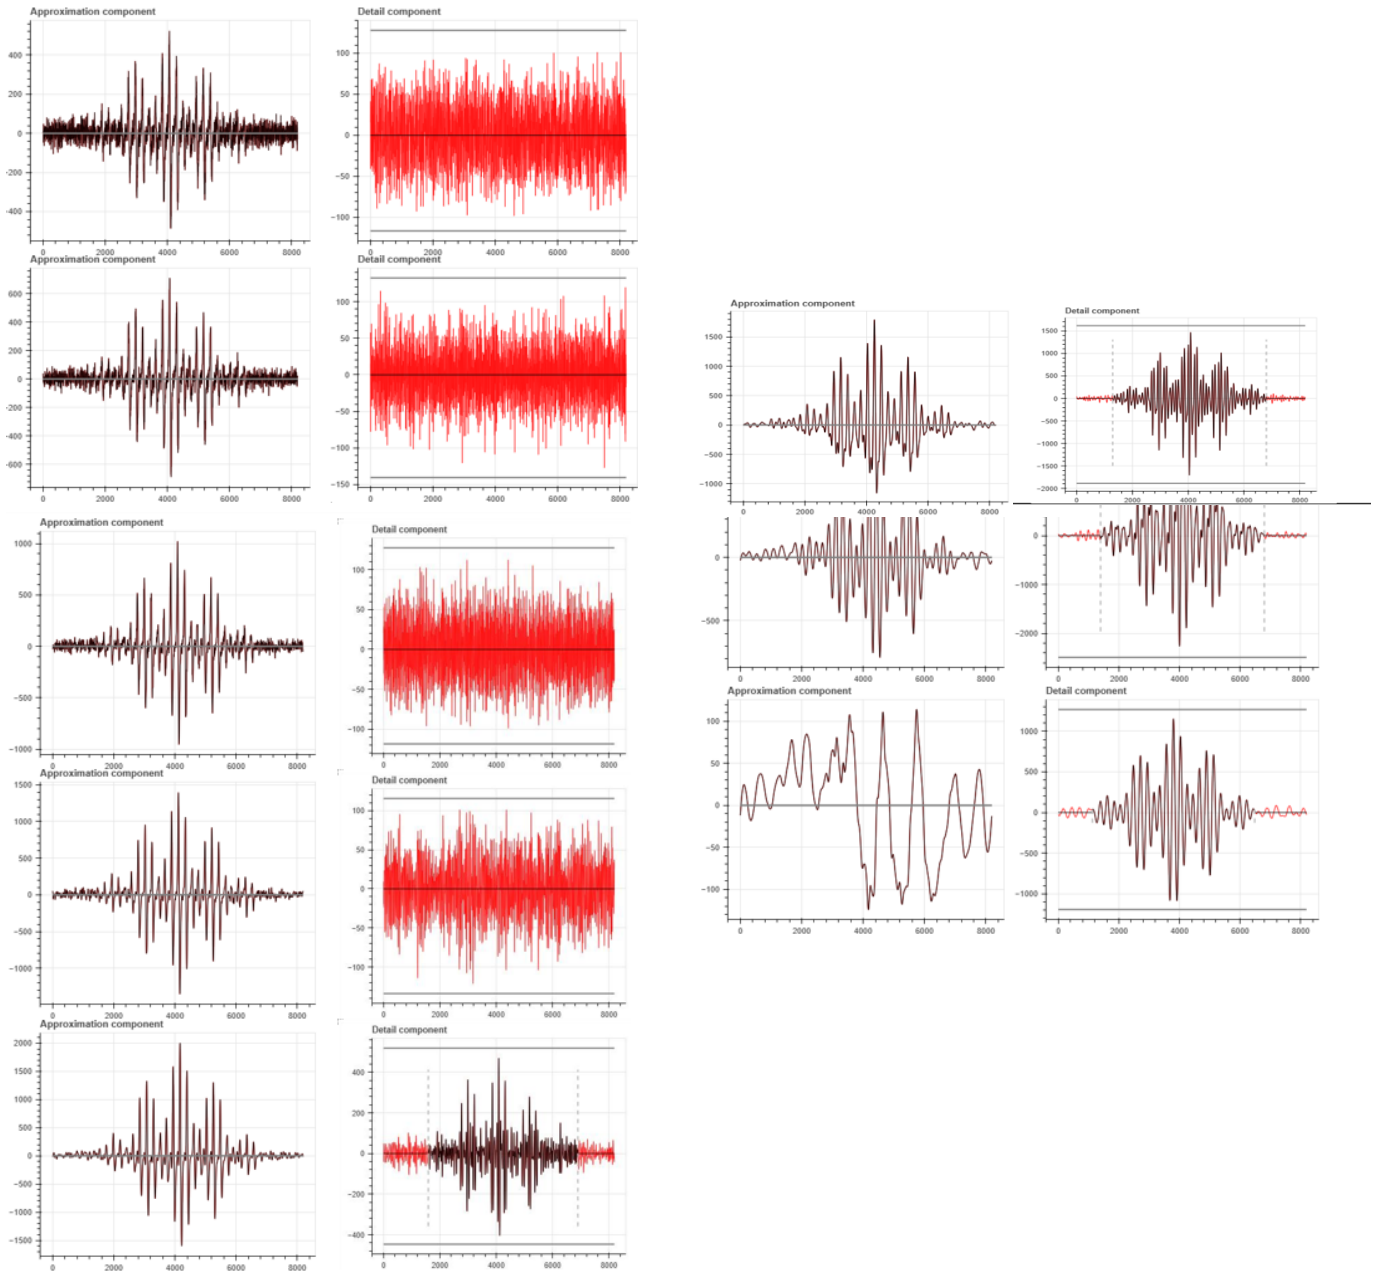

Figure 6. Level-by-Level Detail and Approximation Components. Levels 1-8 represented by rows. Note the Red sections of detail components indicate that those components were zeroed out.

Levels 1-5: Left Column.
Levels 6-8: Right Column.

Srivastava's Decoupling Method

Decoupling appears to occur at decomposition level 7. This produces a detail component and a smoothed signal component whose peak-to-peak distances were approximately equal to our smallest known constant. The second constant is clearer in this graph, but the measured value is around 4, which isn't very close to the known value.

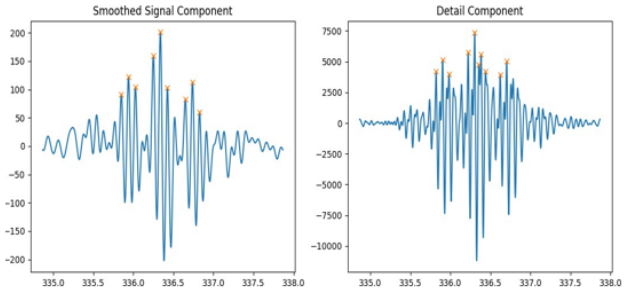

Figure 7. Hyperfine and Super Hyperfine Components of Biphenylene. Peak-to-peak distances are listed below.

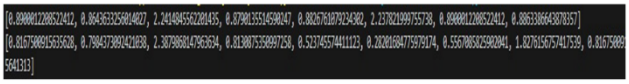

Figure 8. X-axis Distances Between Peak Markings in Figure X

Continuous Wavelet Transform Method

When we took the CWT of the original biphenylene data, the coupling constants were relatively easy to pick out of the heat map and see on the plotted 3d surface. The measured values were still significantly higher than the known value. This is likely due to error on our part, as it is difficult to determine the correct peaks to measure between, and that presented a major roadblock in our research which was never fully resolved.

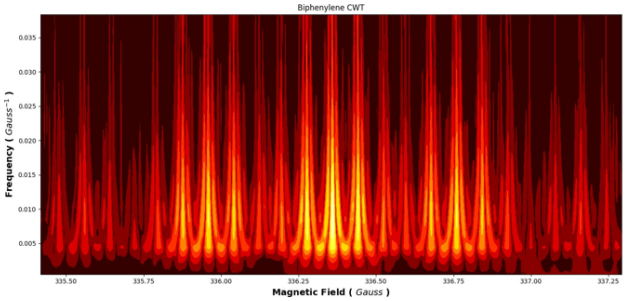

Figure 9. Continuous Wavelet Transform of Biphenylene Data. Note the coupling constants are measured as the distance between the brightest peaks on the CWT Measured Coupling Constants are 0.82 and 4. These measured coupling constants deviate slightly from the known values of 0.71 and 2.86, but the consistency across the two methods is encouraging, because it means that there are measurable trends in the data that we are detecting.

Anthracene

The coif3 wavelet was used for this denoising. The optimal decomposition level was chosen subjectively. Compared to the biphenylene data, this signal saw a similar reduction in noise on the edges, but the middle remained noisy. This can likely be attributed to the pervasiveness of the noise in the original data.

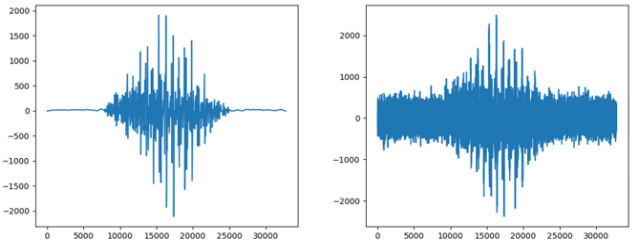

Figure 10. Noisy Anthracene Data Compared to Denoised Data

Srivastava's Decoupling Method

Decoupling appears to occur at decomposition level 7. This produces a detail component and a smoothed signal component whose peak-to-peak distances were approximately equal to our smallest known constant. However, we could not locate the other coupling constants in the graph.

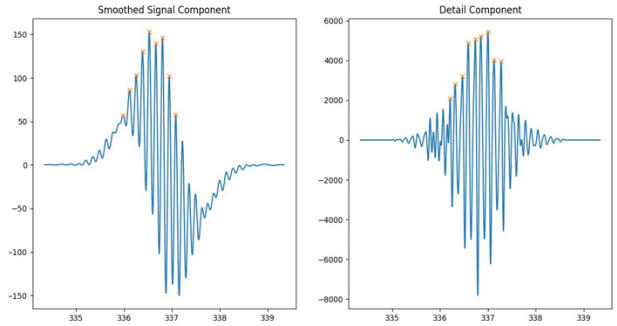

Figure 11. Hyperfine and Super Hyperfine Components of Anthracene with Peaks Marked. The orange marks on the graph come from marking the peaks above a certain magnitude. The numbers below correspond to the distances between those peaks.

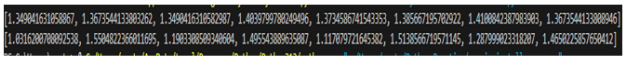

Figure 12. X-axis Distances Between Peak Markings in Figure X

Continuous Wavelet Transform Method

When we took the CWT of the original anthracene data, the coupling constants were relatively easy to pick out of the heat map and see on the plotted 3d surface.

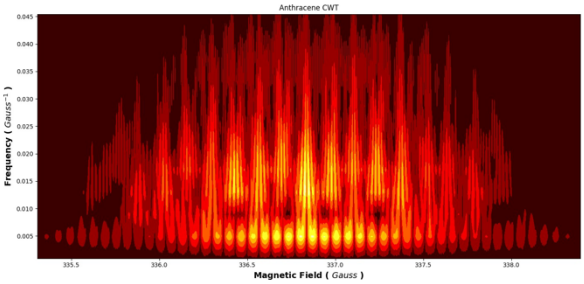

Figure 13. *Continuous Wavelet Transform of Anthracene Data.* Measured Coupling Constants are 1.42, 2.73, and 5.55. Although we were unable to extract all three coupling constants using Srivastava's method, the CWT method was very effective with each constant being very close to the known values of 1.51, 2.74, and 5.34.

Discussion

This was a very broad research topic with several core and complementary goals. In this section, these goals are assessed individually to determine the overall success of our research.

–Gain an understanding of the mathematical basis for Fourier transformations and wavelets – This goal was largely accomplished. While we lacked the understanding of mathematical proofs needed to fully grasp these concepts, we are now able to apply these concepts to practice questions and effectively summarize them for an uninformed audience.

–Denoise Dr. Sipe's collected data – This goal was partially accomplished. We were able to apply the denoising techniques discussed in Dr. Srivastava's papers and provide Dr. Sipe with denoised plots of data. However, the denoising method was most effective on simpler molecules and had disappointing results when applied to more complex molecules.

–Develop the Python skills necessary to work with provided data – This goal was accomplished. We began this project with little to no Python expertise, but by the end of our research we were able to effectively work with data, produce graphs, and develop a denoising program in Python.

Future Work

There are several topics related to this research that warrant further investigation. One being optimal decomposition level selection. The current method for optimal decomposition level selection often appears incorrect and relies on some unproven assumptions to determine the level. We discussed this issue during our research, but we didn't devote any significant time to addressing it, instead relying on visual analysis to determine the optimal level. Developing a more robust method could be a topic of future research.

Another topic of interest would be threshold selection. Detail component thresholds were often chosen subjectively during this research. An objective method for determining vertical thresholds is discussed in the paper A New Wavelet Denoising Method for Selecting Decomposition Levels and Noise Thresholds, but this method produces disappointing results.

Finally, coupling constant extraction would also be a viable option to further research. Several of the data sets we worked with during this research were too noisy to effectively denoise, but they also contained too many coupling constants to extract them easily. There is currently no good way to reproduce these readings, so a new extraction method is needed to salvage the data.

REFERENCES

Aryal, S. (2022, August 22). Electron Spin Resonance (ESR)- principle, instrumentation, applications. Microbe Notes. Retrieved February 24, 2025, from https://microbenotes.com/electron-spin-resonance-esr-principle-instrumentation-applications/

Bogges, A., & Narcowich, F. J. (2009). A First Course in Wavelets with Fourier Analysis (2nd Edition). John Wiley & Sons.

Roy, A. R., Srivastava, M. (2022). Hyperfine Decoupling of ESR Spectra Using Wavelet Transform. Magnetochemistry, 8, 32. https://doi.org/10.3390/magnetochemistry8030032

Sahoo, G. R., Srivastava, M. (2025) Signal Analysis and Denoising in Electron Spin Resonance Spectroscopy: Limitations of Fourier Transform and Introduction to Wavelet Transform.

Srivastava, M., Dzikovski, B., & Freed, J. H. (2021). Extraction of Weak Spectroscopic Signals with High Fidelity: Examples from ESR. The Journal of Physical Chemistry A, 125(20), 4480–4487. https://doi.org/10.1021/acs.jpca.1c02241

Srivastava, M., Georgieva, E., & Freed, J. (2017). A New Wavelet Denoising Method for Experimental Time-Domain Signals: Pulsed Dipolar Electron Spin Resonance. The Journal of Physical Chemistry A, 2452–2465. https://doi.org/10.1021/acs.jpca.7b00183