

Extracting the Nonlinear Response of Buried Interfaces from Electronic Sum-Frequency Generation Data

Benny Renard, Department of Chemistry, The University of Texas at Austin

Abstract

Organic semiconductors (OSCs) are a set of solution-processable materials that have applications in the design of flexible electronics, light emitting displays, and light-harvesting systems. However, in each of these applications, energy or charge must be transferred between the OSC and another material, such as an inorganic semiconductor. The electronic structure of the OSC at this junction can have a strong impact on the effectiveness of these processes, and as such there is a need to understand the electronic states present at these buried interfaces. Information about these states can be extracted by exciting transitions between them – for example, by probing the buried interfaces with light. Unfortunately, there are many molecules between the top of a material and its buried interface; even in the case of a thin film, the signal from a single beam of photons sent through the film would be skewed to reflect the transitions in the bulk much more heavily than those occurring at the interface between the materials. To get around this issue, an interface-specific technique, Electronic Sum-Frequency Generation (ESFG) is used.

In this technique, the signal being measured is the result of two beams of photons interacting at the interface; a third beam is generated, with a frequency that is the sum of the frequency of the two interacting beams. This process must occur within a material; specifically, this process is enhanced if a material has real electronic (or vibrational) transitions at the energies of either the incident photons, or at their summed energy. In addition, for centrosymmetric media – those materials with an inversion center, such as the bulk of air or glass – the signal of an even-ordered process (such as SFG with 2 incident beams) is necessarily 0^4 . That is to say, the bulk of the thin film or substrate measured will not contribute to the SFG signal. However, the interfaces between media are inherently non-centrosymmetric; thus, this technique allows the interfacial states to be selectively probed.

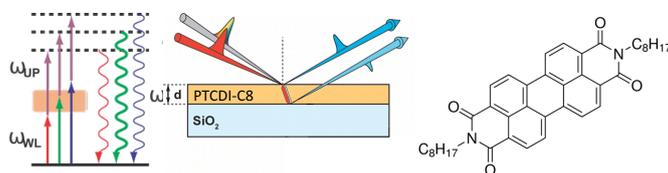


Figure 1: ESFG between a white light spectrum and an upconversion beam (left) occurring at the top and buried interfaces (right) of PTCDI-C8 and SiO₂ (structure shown at right).

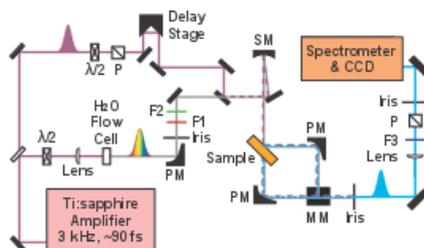


Figure 2: Experimental setup.

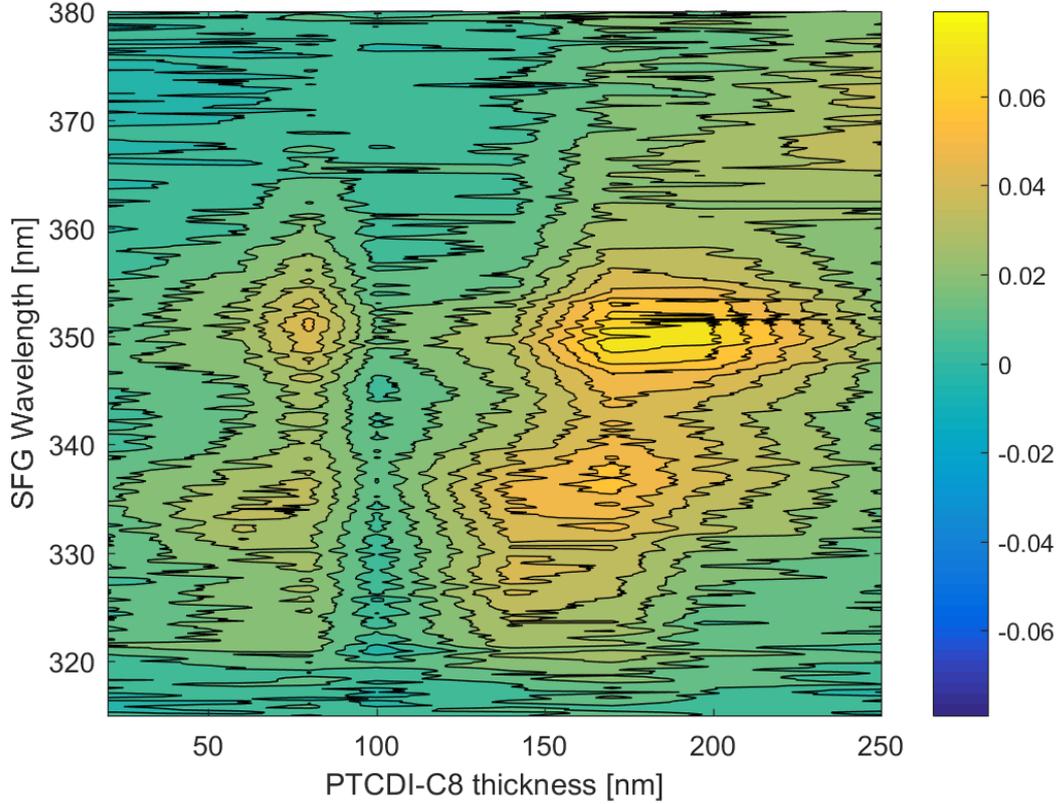


Figure 3: Measured experimental data. The thickness of the PTCDI-C8 layer is plotted along the x-axis, while the wavelength of the SFG (white light + 800nm upconversion) signal is plotted along the y-axis. Signals were measured from films with approximately 20, 60, 80, 100, 140, 170, 197, and 250 nm thick PTCDI-C8 layers. The colored contours represent the strength of the signal generated at a certain wavelength for a given thickness. Data taken by Aaron Moon and Ravindra Pandey.

Unfortunately, the information obtained in this fashion does not directly lead to an understanding of the electronic structure at the buried interface. We need to know the strength of the SFG response generated at the buried interface as a function of wavelength; from this information, we can determine information about the electronic states at this buried interface. However, the generated signal is complicated by thin film effects, causing data such as that seen in Figure 3. The precise location, relative strengths, and widths of the peaks that should be recovered might be obscured by the absorption of the generated signal as it travels back through the bulk, as well as the thicknesses used to measure the signals, which can lead to enhanced or diminished signals for different wavelengths of light. . It is therefore necessary to model these effects to extract the second-order nonlinear response ($\chi^{(2)}$) of a buried interface. Since this response could be arbitrarily complex, it is necessary to also make a decision regarding the complexity of the form of this response. Less complex $\chi^{(2)}$ functions will be easier to interpret and solve for, but may not be correct. More complex functions may allow for the correct result, but may also allow for many other uninteresting solutions.

In this work, several models are used to fit the second-order nonlinear response ($\chi^{(2)}$) of a PTCDI-C8/glass interface. The goal was to make these models powerful, interpretable, and extensible, to ensure that other researchers interested in fitting similar data could do so without too much trouble. From the models used, we can see that simple models, such as lorentzians, can be fairly easy to fit; however, they may leave behind interesting information that could otherwise be extracted from the

data.

Background

When light passes through a material, it has the opportunity to interact with that material through linear and nonlinear processes. The particular processes, and the extent to which they occur will depend on the properties of the material and the distribution of light incident on the material. Carefully controlling the light passing through a material can therefore be used to extract information about its properties. For example, the energy difference between the ground and excited electronic states in that material can be probed.

Under the electric dipole approximation, and representing the light as a sum of monochromatic plane waves, the following formula can be used to describe the polarization generated after interacting with a material⁴:

$$P(k, w) = \chi^{(1)}(k, w) \cdot E(k, w) + \chi^{(2)}(k = k_1 + k_2, w = w_1 + w_2) : E(k_1, w_1)E(k_2, w_2) + \dots$$

In this formula, k is a wavevector, w is a frequency, and the χ terms are tensors describing the physical media with which the light is interacting (such as air or glass). The first term computes the polarization generated by linear processes, such as the absorption of a single photon or elastic scattering; the second term describes the polarization generated by second-order nonlinear processes, such as the simultaneous absorption of two photons (as in sum-frequency generation).

In Sum-Frequency Generation (SFG), two photons of light are absorbed by a molecule at the same time, exciting it to a virtual excited state; the molecule cannot persist in this state, and so it falls down to its ground state, releasing a single photon with the same energy as the sum of the two initially absorbed photons. If the energy of one of these photons, or their sum, is consistent with a real transition in the molecule, whether an electronic or vibrational transition, this process will be enhanced; such a transition is known as a resonant transition¹. In Electronic SFG, the lower-energy photon (often infrared) is chosen to be off-resonance, so that it does not excite a vibrational transition in the molecule(s) of interest; then, the frequency of the higher-energy photon is allowed to vary across a wide range. The resulting spectrum will have peaks where the medium has an electronic transition, providing information about the electronic states that contribute to this signal. Since this is a second-order nonlinear process, centrosymmetric materials (those with a center of inversion, such as bulk air or glass) will not generate a signal. This happens because the material parameters ($\chi^{(2)}$) must stay the same when an inversion through the inversion center is carried out; however, the polarization generated, and electric fields incident, will flip signs in response to this inversion: $-P^{(2)} = \chi^{(2)} : (-E_1)(-E_2)$. Since $P^{(2)} = \chi^{(2)} : (E_1)(E_2)$, this means that $P^{(2)} = -P^{(2)} = 0$ for such materials; since E_1 and E_2 are not (in general) 0, the $\chi^{(2)}$ s must be 0 for the bulk of these centrosymmetric materials, which means that this process (SFG) will not occur in the bulk of these materials.

However, at the interfaces between materials, this inversion symmetry is necessarily broken; therefore, the molecules at the interfaces can generate a second-order polarization, and the associated $\chi^{(2)}$ s are not necessarily zero. Therefore, the signal generated by SFG within these materials necessarily comes from the interface.

However, thin film effects can make it difficult to interpret these spectra without additional processing, since even in the absence of transitions, the signal for certain wavelengths of light will be enhanced due to interference effects; in addition, there are two different interfaces generating signal – one at the surface, and one buried. We wish to discover the properties of this buried interface; it is therefore important to separate out the signal from each interface. Previous work by Massari et al³ provides a model to compute the signal that would be measured at a detector as a result of these thin film effects, given the bulk $\chi^{(1)}$ and interfacial $\chi^{(2)}$ terms, which respectively specify the linear and nonlinear properties of the system being studied. This allows the problem of extracting the interfacial $\chi^{(2)}$ terms to be reformulated as an optimization problem, in which some loss function of the measured and computed signals is minimized over the possible interfacial $\chi^{(2)}$ values.

Results

One goal of this project was to make the modeling code used to fit these spectra (and extract the $\chi^{(2)}$ s, the interesting optical properties of the interfaces) simpler to understand and use, to reduce the burden on any researchers who might want to use this same (or a similar) method to fit their data, and to make it so a small change in the experimental setup or modeling could be accomplished with a small change in the code.

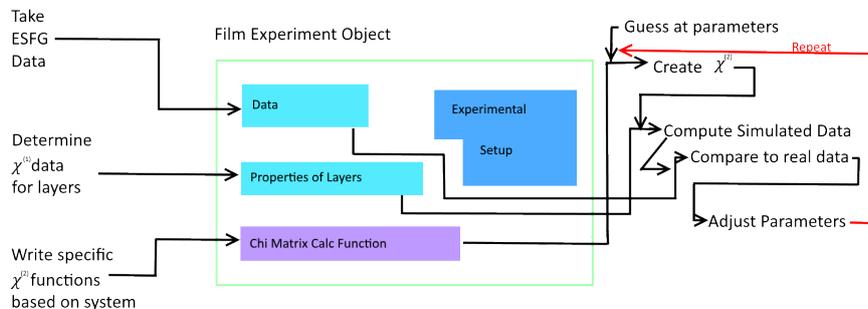


Figure 4: Basic workflow to use the fitting code. The steps on the left are performed by the researcher, and the steps on the right are performed by the code.

To accomplish this, most of the information that might otherwise be stored in globals or baked into the code was moved to a Film Experiment object, which would be passed around to all of the modeling code. The information in this object was split up into a few different types of parameters (seen in Figure 4): those that would likely be changed once based on an experimental setup (blue), those that would need to be edited for each new system examined (light blue), and those that might change more frequently as a researcher attempts to fit their data (purple). The intention was to separate the core modelling code from the components of the model that would likely need to be changed to accommodate a particular experimental setup. For example, changing the optical properties of the materials used, the loss function (for deciding how good or bad a given fit is), and the different classes of functions used to represent the $\chi^{(2)}$ s should not require one to understand how the code uses these parameters to arrive at a good fit.

In order to fit ESFG data and extract $\chi^{(2)}$ responses using this code, a researcher must do a few things. First, optical constants for the system being studied, as well as the order of the interfaces, must be determined. Then, a loss (error) function must be created. This uses the experimental data collected (in our case, SFG signal as a function of wavelength and thickness) to determine how good or bad a given fit is. This will tend to be some function of the difference between the data projected by a fit and the data actually measured. Depending on the quality of the data collected, some data points may be given more weight than others, or regularization terms may be added to penalize or encourage certain types of fits; by default, the RMS error between the collected and modelled data is used. Finally, a method for creating $\chi^{(2)}$ lineshapes from parameters must be selected or written. Initial guesses at these parameters are optimized using Matlab's `fmincon` and the earlier-specified loss function as the objective function; the optimized variables are then used to generate the final, fitted $\chi^{(2)}$ s.

The functions to produce $\chi^{(2)}$ lineshapes, called `chi_calc_functions` in the code, can be very important for getting good fits to the data. If one considers the emission of a molecule in an excited state with some finite lifetime, weakly coupled to a bath, one would expect a Lorentzian profile. Therefore, using a fairly simple model, the sum of two Lorentzians for each interface, gives the fit in Figure 5. The fit captures the basic features of the data, and the physical information we want (energies and strength of transitions) can be easily extracted from this fit; it assumes that there are

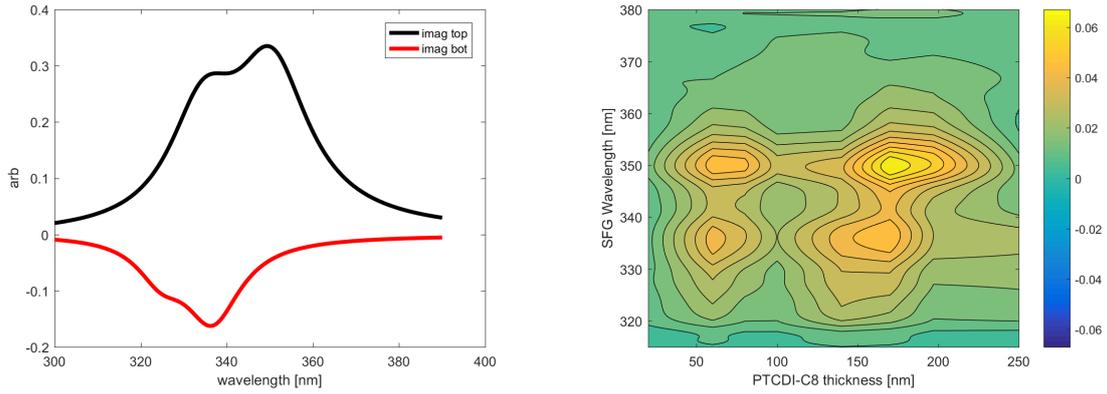


Figure 5: Left: A fit to two Lorentzians for each of the top (air-PTCDI) and bottom (PTCDI-glass) interfaces. Right: Data predicted from this fit.

a small number of transitions available at the interface, and that the top and bottom interfaces' $\chi^{(2)}$ responses do not change as a function of film thickness.

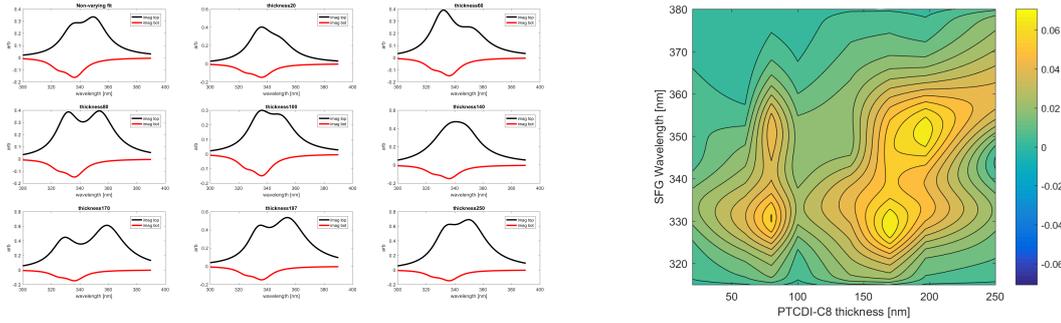


Figure 6: Left: A fit of two Lorentzians for the $\chi^{(2)}$ of each of the top (air-PTCDI) and bottom (PTCDI-glass) interfaces. The top interface is allowed to vary as a function of thickness, while the bottom function is kept constant. This would be assuming that the packing structure at the top interface changes as the films become thicker, while the structure at the bottom interface remains more or less constant. Right: Data predicted based on these $\chi^{(2)}$.

While this is likely a good assumption for the bottom interface, it may not be true for the top interface. As the film grows, the amount of disorder at the top interface may change, and the orientation of the molecules at this interface may change as well. Therefore, in Figure 6, the top interface's $\chi^{(2)}$ s are allowed to change as the film grows thicker, while the bottom $\chi^{(2)}$ s are not. While this fit is able to model the collected data more closely, the physical picture painted by the fit is not quite so reasonable. While one might expect the transitions present at the top interface to smoothly change from some characteristic set (say, with a particular amount of disorder or a particular characteristic structure) to some other set (with a different amount of disorder, or a different characteristic structure), the fits reported seem to involve a kind of oscillation of the energies of the transitions at that top interface as a function of thickness. This interpretation does not seem physically reasonable; it is much more likely that this actually represents overfitting.

Fitting with Lorentzians may make sense in some cases, but will not always be sufficient to describe a given system. As more detail is required in the $\chi^{(2)}$ functions, and as molecules with more complicated transitions are considered, it can become necessary to use rather involved calculations to

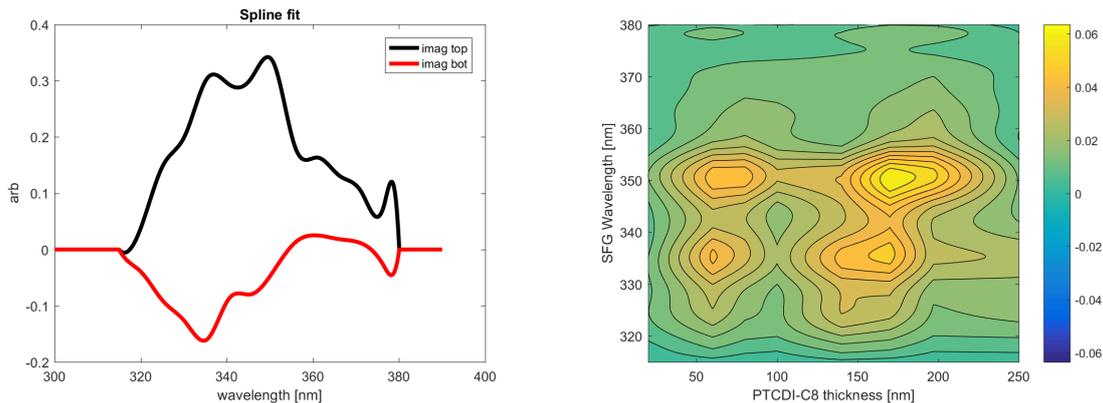


Figure 7: Left: A fit using cubic b-splines, with control points roughly every 4.5nm (14 in total), clamped at the ends. The Kramers-Kronig consistent real portion of the corresponding function is combined with these data to produce the chis used to fit the SPS data. Right: Data predicted based on these chis. Based on Figure 8, the rightmost peak should be disregarded.

determine a good lineshape for a particular system's $\chi^{(2)}$ s. Just as modelling a molecule's transitions more faithfully, with a high level of theory, can be used to allow for complicated lineshapes and an effective fit, so too can not explicitly forcing the $\chi^{(2)}$ s to have a particular lineshape. The fit displayed in Figure 7 attempts to depart from explicitly modelling the lineshape by using cubic b-splines to represent the $\chi^{(2)}$ s; such a model would allow an approximation of most any function, provided sufficiently many control points were used. In practice, allowing more parameters to vary will make the fits take longer, and make it more possible to end up with parameters that are not truly optimal. It can also become easier to model the noise in the data collected. However, it also allows for an extremely wide range of $\chi^{(2)}$ lineshapes, almost certainly including an approximation to the actual $\chi^{(2)}$'s shape; if the correct lineshape is actually arrived at in this manner, it can be used to extract interesting information about the electronic states at the buried interface, without requiring a high level of theory to determine these optimal parameters.

To help constrain the fits, it can be useful to add in regularization terms, which essentially punish 'bad' properties, or give a benefit for 'good' properties. As an example, in Figure 7, we can see a small spike on the right-hand side of the fit. Based on calculations from one of the members of our group, we do not expect transitions that far out coming from the coupling between PTCDI-C8 molecules, and so that peak seems aphysical. However, it is not possible to appropriately penalize this fit with the data, since some of the places where signal would be expected (if there really were a transition out there) is outside the data collection range of our experiment. To help test this, if the fit is allowed more leeway at the edges (as in Figure 8, it will prefer to add in very aphysical peaks in those locations. To produce Figure 9, an exponentially decaying penalty term was added for values at the edges of the fit; no penalty was applied for values between 320 and 370nm (to not penalize peaks where they might be reasonable). In this way, the spurious peak at the edge was suppressed, without impacting the rest of the fit drastically.

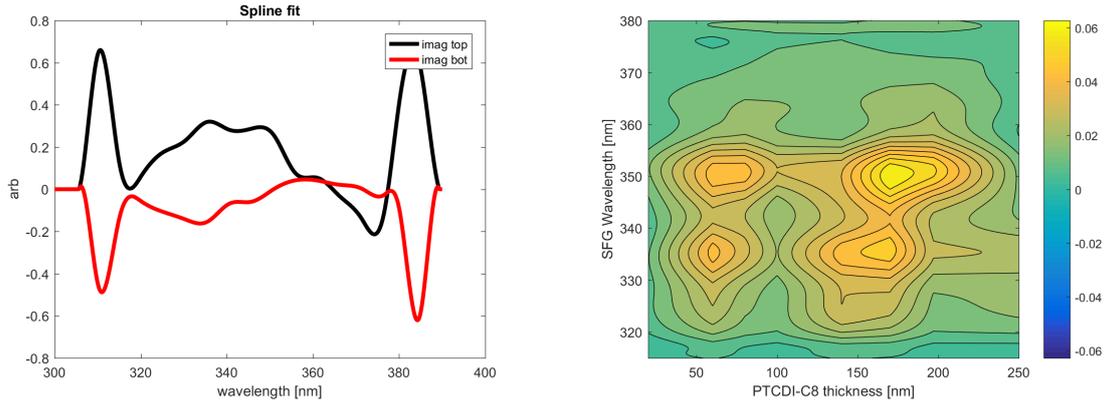


Figure 8: Left: A fit using cubic b-splines, with control points roughly every 4.5nm (18 in total), clamped at the ends (but further out than the previous fit). The Kramers-Kronig consistent real portion of the corresponding function is combined with these data to produce the chis used to fit the SPS data. Right: Data predicted based on these chis. We can see that these peaks at the edges do not cause too much of a change in this region for which data was taken; however, these peaks are highly aphysical.

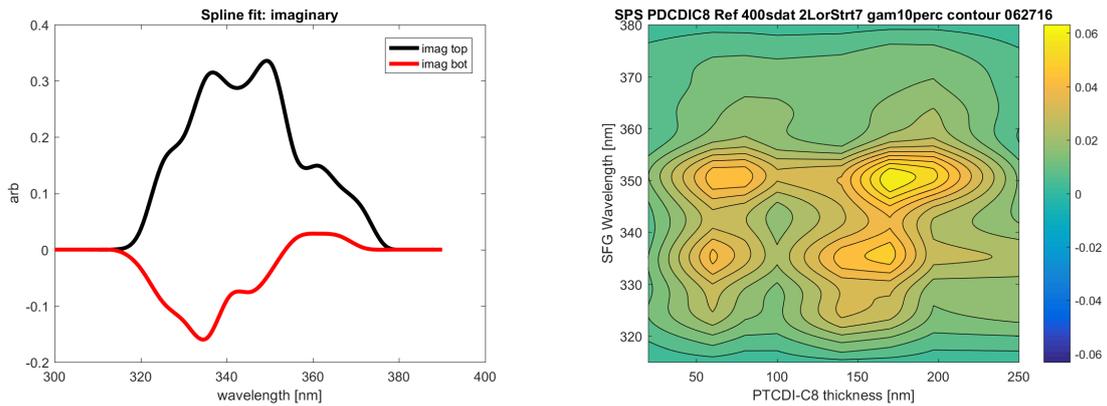


Figure 9: Left: A fit using cubic b-splines, with control points roughly every 4.5nm (18 in total), with an exponential regularization term pulling in the ends (instead of clamping). The Kramers-Kronig consistent real portion of the corresponding function is combined with these data to produce the chis used to fit the SPS data. Right: Data predicted based on these chis.

Discussion and Conclusions

Our primary goal is to use a tool – ESFG spectroscopy – to understand the electronic properties of buried interfaces. These properties stem from the interactions between molecules at these interfaces: based on the orientation of these molecules, the distribution of their electrons and unfilled orbitals and the overlap in orbitals between neighboring molecules, different electronic transitions (with different energies) become available. As the local environment of these molecules changes at different locations in the interface, the energy and strength of a given transition can be altered; to determine the $\chi^{(2)}$ s accurately from first principles would require modelling a large number of molecules, and understanding how they interact with each other – since these $\chi^{(2)}$ s are based on a large number of individual transitions. Such an analysis would be both computationally and theoretically demanding.

The spline fitting method used here essentially allows for a departure from a specific lineshape

model (thus reducing the level of theory and computation needed to consider complicated lineshapes), while still keeping the number of parameters relatively low and allowing for a reasonably fast fit. It also enforces a smoothness constraint; this is both generally physically realistic and can help prevent overfitting to noise in the data. Later, one could come back with a higher level of theory and attempt to fit these $\chi^{(2)}$ s, instead of fitting all of the experimental data measured (which might be more computationally demanding).

The $\chi^{(2)}$ s that need to be modelled are complex-valued functions. The imaginary and real components are related to each other by what are known as the Kramers-Kronig relations, which exist to preserve causality; a complete description of one is sufficient to define the other. It is not always possible to create simple relations that will account for this relationship in the case of nonlinear effects²; fortunately, sum frequency generation falls into a general class for which it is relatively simple to produce such relations¹, and for which the results are practically applicable. Essentially, the experimentally measurable quantities (signal generated due to the imaginary susceptibility) vary over a relatively small spectral range (small number of different “colors” of light), while the corresponding difficult-to-measure quantities (real susceptibilities) change over a wider spectral range. Therefore, if the imaginary susceptibilities can be determined, the corresponding real values can be determined as well. For the purposes of optimization, only one of these sets of variables needs to be allowed to vary; the other set can be computed from those variables.

Unfortunately, since both the real and imaginary components of the nonlinear susceptibility influence the end result of the calculation, but these parameters are not independent, it is possible for some features in the imaginary component (important for interpretation) to be present more to cause changes in the real component (and thereby fit the data better) than due to the presence of an actual electronic transition with the properties hinted at by that imaginary susceptibility. It can be difficult to separate these effects; a more thorough analysis may be needed to attempt to isolate or enhance the effect of the imaginary term in this situation. Looking at Figures 7 and 8, we can see that the peaks at the edges are large and aphysical; moreover, they seem to offset each other, with an opposite sign between the top and bottom interfaces. Since these peaks would cause a response outside the range for which ESFG signal was measured, this fit would not be penalized properly for suggesting non-existent transitions. Such peaks would, however, change the associated real components of the nonlinear susceptibility, and therefore allow for a better fit. Since these peaks would appear to be aphysical, they are suppressed in Figure 9 with a regularization term. This is a strategy that could be used to restrict the class of functions considered in optimization, by penalizing peaks in areas where other analyses have indicated the materials should not have transitions, or (in general) to incorporate information gleaned in other ways into the fits, to help arrive at a good $\chi^{(2)}$.

Finally, one of the goals of this work was to allow other researchers to use this general modelling approach to fit their own data. When writing code for research, it can be tempting to write something fairly specialized; then, when changes are needed to fit new experimental data or accommodate new ideas about the models being used, this core code can be changed. Unfortunately, this approach can lead to code that is either difficult to understand or difficult to modify for new tasks – the researcher might need to understand quite a bit of code to implement a small change. Some effort was made here to separate out the components of the code that represent the physical system, the $\chi^{(2)}$ modelling, and the thin film modelling. The first of these would need to be understood by any researcher who wanted to fit new data, while the second would need to be modified if a researcher wanted to fit complicated lineshapes and use their deeper understanding of their particular system to achieve a better fit. The final thin-film modelling component should not need to be modified unless a researcher wanted to use a more complicated model (for example, consider more than the dipole nature of light) or needed to accommodate some fundamentally different types of data. The hope is that this will allow a wider set of researchers to more easily fit their ESFG data.

Code

For a copy of the code used to generate these fits, please email Sean Roberts at roberts@cm.utexas.edu.

Acknowledgements

The author would like to thank the Roberts group for their help and mentorship throughout this thesis. Specifically, thanks to Sean Roberts, the Principle Investigator; to Aaron Moon, for providing the original version of the code that was modified for this work; to Aaron Moon and Ravindra Pandey, for measuring the ESFG data; to Jon Bender, for growing the films and determining the linear susceptibilities via ellipsometry; and to Daniel Cotton, for calculations that would help rule out peaks at the far edges of Figures 7 and 8.

The author would like to acknowledge the support from an Undergraduate Research Fellowship provided by the Office of the Vice President for Research at the University of Texas at Austin. The author would also like to acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing computational resources in support of this work.

Bibliography

1. Boyd, R. W. *Nonlinear Optics*, 3rd ed.; Academic Press; Massachusetts, 2008.
2. Hutchings, D. C.; Sheik-Bahae, M.; Hagan, D. J.; Van Stryland, E. W. Kramers-Kronig relations in nonlinear optics. *Optical and Quantum Electronics* 2002, 24 pp 1-30.
3. O'Brien, D. B. and Massari, A. M. Modeling multilayer thin film interference effects in interface-specific coherent nonlinear spectroscopies. *J. Opt. Soc. Am. B.* 2013, 30 (6) pp 1503-1512.
4. Shen, Y. R. *The Principles of Nonlinear Optics*; Wiley; New York, 1984.