

**HARDWARE / ALGORITHM INTEGRATION FOR PHARMACEUTICAL
ANALYSIS**

by

Casey Jake Smith

A Dissertation

Submitted to the Faculty of Purdue University

In Partial Fulfillment of the Requirements for the degree of

Doctor of Philosophy



Department of Chemistry

West Lafayette, Indiana

May 2020

**THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF COMMITTEE APPROVAL**

Dr. Garth J. Simpson, Chair

Department of Chemistry

Dr. Lynne Taylor

Department of Industrial and Physical Pharmacy

Dr. Julia Laskin

Department of Chemistry

Dr. Adam Wasserman

Department of Chemistry

Approved by:

Dr. Christine Hrycyna

Dedicated to my wonderful wife, Nalani and to my parents, Jake & Pam.

Thank you for believing in me.

ACKNOWLEDGMENTS

I have been incredibly blessed in my life with all the love and support that I have received. I could not have come this far in life without numerous great examples, leaders, and mentors.

Starting from those who influenced me chronologically, I would like to thank my high school Chemistry teacher, Mrs. Karen Lyman, who fostered a love of chemistry & exploration in me. She let me follow what interested me and helped me clean up a few messes in the aftermath. Ultimately, I pursued chemistry as a bachelor's degree because of her. I gained a few good friends during my time at the University of Utah. I am very grateful for the friendship and camaraderie of Nathan Odendahl, Shantel Leithead, and Emmett Barlow. I am thankful for the opportunity Dr. Joel Miller offered me to work in his lab and study organometallics. I especially appreciate the patience and instruction I received from Chris Kareis, the graduate student / post-doc, I was assigned to work with. What I learned from you was invaluable. You all encouraged me to chase my dreams, which led me to graduate school, and helped me raise my expectations that I could achieve them.

I am first and foremost grateful to Garth Simpson. He truly is the definition of a scholar and a gentleman in every sense of the word. Garth, you have been a constant source of optimism, encouragement, and enthusiasm for life & science. Being around you has made me want to be a better person and a better scientist. I am thankful to older graduate students such as Paul Schmitt and Justin Newman who helped me make the decision to join the Simpson Lab, and who subsequently helped me make the transition to graduate life. I am also thankful for the other group members and the good times we've had. Scott Griffin has been an especially good friend and a great co-worker. I am extremely thankful to the AMY facility, the best kept secret of the chemistry department. Without them, I would have taken at least twice as long to finish my graduate work. Their expertise and help cannot be understated. I especially want to thank Greg Eakins who was not only a huge help in building the tribo instrument (and various 3D printed parts), but who also became a close friend. I am also grateful to the support staff who have helped me and have become my friends over the years (in no particular order)– Mike Everly, Conrad Kliewer, Tom Hanna, Mark Carlsen, Ned Gangwer, Austin Johanning, Suzy Gustafson, Rob Reason, Randy Replogle, Darci DeCamp, Steve Scherer, Jane Johanning, Betty Hatfield, Debbie Packer, Lynn Rider, Betty Dexter, and Aaron Harkleroad. I would like to extend a special thank you to my collaborators, Paul

Stroud, Jeremy Hinds, Siwei Zhang, Julie Novak, and Liu Zhen. You all have helped me accomplish my research while providing me with valuable direction and relevance for my research. Working with you has been an absolute pleasure and I hope I can have co-workers like you in the future.

I would like to thank Tim Newby, Doug Dawson, and Ben Dunford for being personal mentors to me during my time at Purdue. You have taught me how to selflessly serve, how to be a good father, and how to be a righteous leader. During a time in my life where I was looking to become more than I was, you three came and showed me how. I also want to thank Dallin Morris for being an excellent friend and someone I could rely on.

My parents have been fantastic parents to me. I can never fully express how grateful I am to them. They laid the foundation of my life with love, support, ethics, and morals. They have shown me through their example how to be good people, how to achieve your goals, and how to be. Again, I can never thank you enough for raising me the way you have. I hope I can be half the parent you were. I would also like to thank my siblings who have always been some of my best friends. Thank you for listening to me explain my research and keeping me on my toes.

Nalani, you are the woman of my dreams and my best friend. During our time together here at Purdue, we have cried together, laughed together, waded through uncertainty and discouragement together, and figured out how to more closely be the people we want to be. Your example of constant improvement has driven me to push forward when I was tired or discouraged. Your support and belief in me buoyed me through times of discouragement and given me a calmness in my heart. Thank you for keeping me grounded to what matters most. I'm grateful I get to spend eternity with you. And I look forward to being that cute old couple together with you.

And finally, Penny – You joined our family towards the end of my graduate career. I knew life would change when you arrived; I never knew how much I would enjoy being your dad. Thank you for your morning kisses, waving goodbye when I go to work, and running up to me when I get home. You are a wonderful little girl. And I'm so thankful that you have joined our family.

TABLE OF CONTENTS

LIST OF TABLES	9
LIST OF FIGURES	10
LIST OF ABBREVIATIONS.....	16
ABSTRACT.....	18
CHAPTER 1. INTRODCUTION	19
1.1 Philosophy.....	19
1.2 Motivation for Research	19
1.3 Ph.D. Summary.....	22
CHAPTER 2. TRIBOLUMINESCENCE.....	25
2.1 Overview.....	25
2.2 Introduction.....	25
2.3 Motivation.....	26
2.4 Experimental Section	28
2.5 TL Instrument Fabrication & Characterization.....	29
2.6 Triboluminescence Theory	39
2.7 Applications of TL to Pharmaceutical Amorphous Solid Dispersions.....	45
2.8 Triboluminescence of Slurries	48
2.9 Other Applications of Triboluminescence	54
2.10 Acknowledgements	55
CHAPTER 3. ITERATIVE NON-NEGATIVE MATRIX FACTORIZATION.....	56
3.1 Introduction.....	56
3.2 Methodology.....	59
3.3 Algorithm.....	62
3.3.1 Baseline Correction	63
3.3.2 Matrix Factorization	63
3.3.3 Initialization.....	63
3.3.4 Solving for the $d_{\text{fit},1}$ element of the d_{fit} Vector.....	63
3.3.5 Solving for the IRF	64
3.3.6 Iterate to Convergence.....	64

3.3.7	Autocorrelation	64
3.4	Experimental	64
3.5	Application to Triboluminescence	66
3.6	Conclusions	75
3.7	Acknowledgements	75
CHAPTER 4.	CALIBRATION-FREE SHG IMAGE ANALYSIS	76
4.1	Introduction	76
4.2	Mathematical Theory	79
4.2.1	Particle Counting Regime	79
4.2.2	Integrated Intensity Regime	82
Regime I.	82
Regime II.	83
Regime III.	84
4.3	Experimental Details	85
4.4	Simulations	88
4.5	Application to Crystalline Volume Determination within a formulated ASD	93
4.6	Conclusions	98
4.7	Acknowledgements	99
CHAPTER 5.	STOCHASTIC DIFFERENTIAL SCANNING CALORIMETRY	100
5.1	Introduction	100
5.2	Experimental	102
5.2.1	SHG Microscopy	102
5.2.2	DSC measurements	103
5.2.3	Data analysis	104
5.2.4	Sample preparation	104
5.3	Application to Urea Crystals	104
5.4	Application to Trehalose	108
5.5	Conclusions	114
CHAPTER 6.	Generative ADVERSARIAL Linear Discriminant Analysis	116
6.1	Introduction	116
6.2	Mathematical Framework	117

6.2.1	Overview.....	117
6.2.2	Performing LDA.....	118
6.2.3	Generating Decoy Data.....	118
6.2.4	Launching an Attack.....	119
6.2.5	Iterative Algorithm	121
6.3	Experimental – Raman Data Set	121
6.4	Performing an attack on LDA.....	123
6.5	Iterative Generative Approach to Improving Loading Plots.....	131
6.6	Conclusions.....	142
APPENDIX A. CODE		143
APPENDIX B. BUILD LIST FOR TL INSTRUMENT		159
PUBLICATIONS.....		160
REFERENCES		161

LIST OF TABLES

Table 1. Unit Formula for Evacetrapib Tablets 87

LIST OF FIGURES

- Figure 2.1. Graphical depiction of the initial proof of concept instrument for TL analysis of pharmaceutically relevant materials. 30
- Figure 2.2: Image of the redesigned triboluminescence instrument. The instrument was housed in a sealable box to allow for measurements in a low light environment. The red plastic houses the solenoid, used to impact the sample. The green plastic provides support for a threaded rod and a support rod. The white plastic houses the stepper motor, used for sample translation. The black plastic piece holds the sample and is translated following each impact event. Underneath the black plastic piece, the PMT is aligned to collect any emitted light from TL..... 31
- Figure 2.3: Close up image of the hammer and sample stage (black). 32
- Figure 2.4. Example of the sample slides used in triboluminescence experiments. Top image is of the slide with sample loaded into the channel and sealed in with black electrical tape. The bottom image is the sample slide before loading with sample..... 34
- Figure 2.5: A) Representative trace from a 0.05% by weight griseofulvin sample. B) 11 impact events concatenated into a single graph are shown to illustrate the shot to shot variability. The green (lower), offset data trace is pure PEG, shown to demonstrate the baseline response in the absence of crystalline griseofulvin..... 35
- Figure 2.6: A histogram showing the time difference from the initial peak to the secondary peak. There are two groupings clearly seen at ~ 0.4 ms & ~ 1.9 ms..... 37
- Figure 2.7: Both TL (a) and SHG microscopy (b) were performed on powdered blends of griseofulvin in PEG to demonstrate the calibration curve. In each case, the inset shows 5 data points with concentrations at 1% or less..... 38
- Figure 2.8: Average integrated signal for 4 different particle sizes of sucrose crystals. The error bars represent the standard deviation of the mean slide value (e.g. 36 averaged impacts). 40
- Figure 2.9: A) Shows the SHG (blue, horizontal lines) and TL (red, vertical lines) signals from two polymorphic forms of clopidogrel bisulphate which are indistinguishable by bright field imaging. Both data sets were normalized to Form I. Both the Form II and the PEG data were rescaled by 100 to allow meaningful comparisons with the Form I result. B) Two representative images from SHG microscopy to demonstrate the difference in image contrast. The SHG results for Form II are rescaled by 5, showing only diffuse, low signal to noise, SHG signal. The bright punctum in the Form II image is attributed to an impurity and was excluded in the analysis. 41
- Figure 2.10: Triboluminescence signal from Griseofulvin. Used as an example to demonstrate deconvolution using the non-negative matrix factorization..... 42
- Figure 2.11: Showing the results of deconvolution. Blue lines are arrival times of photons. Zoom in on the right showing the density of photon arrival times at the first peak of the TL data. 43
- Figure 2.12: Results of autocorrelation on the raw TL data (red-top), the measured IRF (blue-middle), and the deconvolved photon events (green-bottom). The x-axis is on a log scale..... 44

Figure 2.13: Powder X-Ray Diffraction pattern of Griseofulvin HPMCAS-MF amorphous solid dispersion immediately following preparation. 46

Figure 2.14: A) Shows the crystallinity content of the griseofulvin and HPMCAS-MF ASD after being exposed to heat and constant humidity for 48 hours. B) The upper trace (blue) shows the TL signal from the stressed sample while the lower trace (green) shows the signal from the unstressed sample, offset for clarity. 47

Figure 2.15: A schematic cutaway of the SHG flow cell (left) and the TL flow cell (right). The holes in the center provide an optical window to take SHG images or to record emitted light from TL..... 50

Figure 2.16: A) SHG image of the flow cell, image taken while pump is off. B) SHG image showing blurring particles due to flow of liquid during imaging. 52

Figure 2.17: Number of impact events with signal above the noise threshold with respect to concentration of the solution. Solution contains griseofulvin in water. 54

Figure 3.1: Illustration of the NMF-F algorithm workflow..... 62

Figure 3.2: Example signal transient pulled from a larger set of TL measurements which illustrate the complexity and difficulty in distinguishing individual photon events. Transients on a microsecond time scale appear as impulses over the millisecond scale of the emission event. ... 67

Figure 3.3: Using simulated data zoomed in to microsecond time scales and normalized results for ease of comparison with the ground truth. A) shows the recovered transient (blue-top) compared to the original simulated transient (black-bottom). Where the black trace is displayed with negative amplitude for ease of comparison. In B), the ground truth (black-top, offset by 1 Norm.U.), recovered (blue dashed, offset by 0.5 Norm.U.), and initial guess (red-bottom) IRFs are compared and show good agreement. In C), the recovered photon events (blue-top) are compared with the ground truth events (black-bottom) and show nearly complete recovery with minimal error. The digital filter's accuracy and ability to recover photons in high density data is highlighted in D), by the ROC plots made with the same data and analyzed using NMF-F (blue), Wiener deconvolution (black), and BRL (red). Both the Wiener deconvolution and BRL were performed using the true IRF, while the NMF-F analysis was performed with a good guess from previous experiments. The AUC for NMF-F is 0.987 and 0.998 for the Wiener deconvolution. While the NMF-F performed better at lower false positive rates, it performs worse as the false positive rate increases. 69

Figure 3.4: An example of the algorithm's ability to recover an IRF similar to the ground truth. (a) Representative initial random guess used in the NMF-F and the BRL algorithm. (b) Recovered IRFs for the NMF-F (blue-middle) and the BRL (red-bottom) algorithms are compared with the ground truth (black-top), which are offset for ease of comparison. (c) Bar chart highlighting the difference in ability to recover an IRF similar to the ground truth of the NMF-F and BRL algorithms. NMF-F converged to a similar IRF to the ground truth 82.5% of 1000 attempts, while BRL failed in all attempts. 71

Figure 3.5: Same TL signal transient from Figure 3.2 with the recovered photon events (blue – lower). The inset illustrates how many photons are within the initial rise caused by a TL burst event..... 72

Figure 3.6: Amorphous solid dispersion of griseofulvin in HPMCAS-MF, stressed to induce crystallinity (black – upper) and the recovered photon events (blue - lower, inverted for clarity). 73

Figure 3.7: Results of autocorrelation on the raw TL data (red-top), the measured IRF (blue-middle), and the deconvolved photon events (green-bottom). The x-axis is on a log scale. 74

Figure 4.1: Red data points are the raw data measured from particle-counting simulations. Each point represents a simulated image. True particle counts range from 1 particle to 4000 particles per image. Blue data points are the estimated true particle count μ_n as estimated by Equation 4.4. Inset images are simulated images of 1000 and 3000 particles per image. 89

Figure 4.2: Each point represents the volume determination of a simulated image. Red x's are the raw measured volume. Blue diamonds are the recovered volumes using Equation 4.3. The solid black line is the ground truth value. 90

Figure 4.3: Comparison of the predicted signal to noise ratios for volume determination by particle-counting versus integrating the SHG intensity, normalized by the inherent variance in the number of particles from Poisson statistics (i.e., variance equal to the mean number of particles). All calculations were performed for a mean particle area relative to the field of view $\bar{A}_0 = 5.9 \times 10^{-5}$. The vertical bar indicates the upper limit for the particle-counting correction. 92

Figure 4.4: Shows representative images of the different physical mixtures of spiked crystalline evacetrapib in the ground ASD tablets. Some of the excipients are SHG-active as seen in the blank. All concentrations except the 100% were used for the analysis. The 100% image is included here for reference only. 94

Figure 4.5: Image Analysis of SHG micrographs allowed determination of crystalline fraction of evacetrapib in formulated tablets. The recovered volume obtained by particle counting is plotted in blue (0.05% mass fraction to 2.5%). The linear least squares fit is calculated using only the solid blue data set. The red diamond data set shows the recovered volume fraction from the intensity measurements obtained by autocalibration performed from the measurements indicated in the solid blue points, Equation 4.6. 96

Figure 4.6: The FT-IR absorbance of different concentrations of evacetrapib. The 1736.6 cm^{-1} peak was used for quantification (shown in the black box). Limit of detection using this technique was determined to be 0.8%. 97

Figure 4.7: The baseline corrected spectra from PXRD for various concentrations of crystalline evacetrapib in drug formulation. The area of the most intense peak, located at $12.05^\circ 2\theta$, was used for evaluation of the data. Limit of detection was determined to be 0.9%. 98

Figure 5.1: Schematic of the integrated SHG-DSC microscope. A pulsed laser beam (800 nm, 80 MHz) was scanned with a galvo-resonant scan pair, which was 4f coupled to the back of an objective and focused onto a DSC sample pan. The SHG signal was collected in the epi direction, isolated with a dichroic mirror and detected with a photomultiplier tube. 103

Figure 5.2: Distribution of crystalline urea volumes measured via image analysis on the SHG microscope. 104

Figure 5.3: SDSC analysis of urea by combined SHG microscopy and DSC. The solid black line is the DSC trace of the melting of urea with endothermic direction being down (left-axis). The thin colored lines show the normalized SHG intensity of individual urea particles (right-axis), acquired concurrently with DSC measurements. DSC and SHG measurements both indicate a structural transformation in the same temperature range. The inset shows a zoom-in of the temperature range 135.5-136.5°C. 106

Figure 5.4: SHG microscope images of urea crystals during SDSC analysis. SHG microscopy reveals the stochastic nature of the melting transition and enable analysis on a per-particle basis. 106

Figure 5.5: Comparison of the experimental DSC trace and the DSC trace generated from melting events. The best-fit double exponential IRF (yellow) is convolved with impulsive melting events (blue) of single-particles observed in SHG images to generate the purple trace, which is in good agreement with the experimental DSC data (red). The time constant and offset of double-exponential IRF were optimized to minimize difference between experimental and generated DSC traces. Best-fit values of 1.161 ± 0.019 s, 0.435 ± 0.013 s, and 3.088 ± 0.013 s were computed for the double-exponential time constants and the offset respectively. 108

Figure 5.6: The left image is the summed frames over the temperature range of 140-220°C. Each particle included in the analysis is enclosed in a box. The image on the right is the same image but on a log scale to better differentiate background and signal. 109

Figure 5.7: SDSC analysis of trehalose by SHG microscopy. The colored lines show the percentage of pixels in a trehalose particle above an SHG intensity threshold overlaid with the black line of the DSC trace with the endothermic direction being down. SDSC images were summed for the temperature range of 140-220°C and the logarithm of the summed image was cropped to the measured region of each particle (each cropped image points to the representative SHG trace). 111

Figure 5.8: The final DSC peak in the trehalose dihydrate experiment was analyzed using the IRF recovered from the urea experiment. Results of the convolution are shown as the solid blue and dashed orange lines. Neither separately nor together do they match the DSC peak. 114

Figure 6.1: Workflow of generative adversarial linear discriminant analysis algorithm. Initial data that falls into three classes is projected onto a reduced dimensional space. Random decoy data are then generated and projected onto this 2D space. The decoy data are then attacked to make them overlap with the initial data on the reduced dimensional space while minimizing the magnitude of the perturbation. LDA is performed again, but now with an additional class to separate out the decoy data. The process is iterated. 117

Figure 6.2: Conceptual illustration of adversarial spectroscopy. Major spectral peaks drive the position of spectra in lower-dimensional projections, demonstrated by the set of thin, long black arrows for Classes A and B, respectively. In panel C the addition of patterned perturbations in the vector δ optimally relocates the position from the initial Class A to Class B in this reduced dimensional space. 119

Figure 6.3: Pseudo code to describe the algorithm. 121

Figure 6.4: Projection of experimental Raman spectra in LDA-space (on the left), with the corresponding mean spectrum for each class (on the right), offset for clarity. The mean spectrum for each class is an average of 84 measurements. 123

Figure 6.5: The eigenvectors recovered from LDA to project Raman spectra onto the lower-dimensional LDA space. The underlying spectral features are present and support reliable classification in validation data sets but are largely obscured by noise from the use of finite training data. 124

Figure 6.6: Demonstration of an incremental attack from Class 1 (red) to Class 2 (blue). Shown by the x's progressively moving towards the mean of Class 2, indicated by hollow blue circle. The probability of the attack belonging to each class was calculated at each point, x. The probability of the purple x's in the box belonging to each class is shown in Figure 6.7. The purple x's denote the region of greatest uncertainty. 125

Figure 6.7: Zoom in of the probability of the perturbation belonging to each class in the region of greatest uncertainty, as shown by the boxed, purple x's in figure 6.6. $|\Delta|$ represents the degree of displacement toward the mean of the target class. 126

Figure 6.8: Comparison of the initial spectrum (red, top) and the applied perturbation (black, bottom). No offset applied. 126

Figure 6.9: Comparison of the initial spectrum (red, bottom) from Class 1 with the attacked spectrum (blue, middle) classified as Class 2, and the mean target spectrum (black, top) of Class 2. 127

Figure 6.10: Additional example of an adversarial attack, in which a Class 3 spectrum is perturbed such that it is misclassified as Class 2. A) The direction of the perturbation and the final position in LDA space are shown in the reduced dimensional space. B) The region of greatest uncertainty in classification is highlighted by the boxed, purple x's in panel A. $|\Delta| = 0$ is the initial unperturbed spectrum, and $|\Delta| = 1$ is the mean position of class 2 in LDA-space. C) The attacked spectrum (middle, blue) is classified as Class 2 with the initial spectrum (bottom, red) from Class 1 and the target mean spectrum from Class 2 (top, black). D) Optimized perturbation used to induce misclassification. 128

Figure 6.11: Demonstration of an attack on an initial spectrum from class 1 to induce a misclassification as Class 3. A) The direction of the perturbation and the final position in LDA space are shown in the reduced dimensional space. B) The region of greatest uncertainty in classification is highlighted by the boxed, purple x's in panel A. $|\Delta| = 0$ is the initial unperturbed spectrum, and $|\Delta| = 1$ is the mean position of class 3 in LDA space. At $|\Delta| = 0.38$ there is a >95% probability that the attacked spectrum belongs to Class 3 based on the position in LDA-space. C) Comparison of the initial spectrum (bottom, red) with the attacked spectrum at $|\Delta| = 0.38$, which bears clear visual similarity to the initial unperturbed spectrum. D) Optimized perturbation used to induce misclassification. 129

Figure 6.12: Attack strategies similar to those detailed previously were performed for an additional 10 initial spectra, selected at random, for randomly selected target classification and spectra. All results are quantitatively similar to those detailed previously. 130

Figure 6.13: Initial loading plots after performing LDA on the raw data set with additional spectra generated from knowledge of the mean and standard deviation.....	132
Figure 6.14: Initial placement in LDA space using the two eigenvectors shown in Figure 6.13	133
Figure 6.15: Decoy data projected onto the reduced LDA-space as purple squares. Note that the decoy data cluster around the background data which primarily consists of noise.	134
Figure 6.16: Example of a generated Decoy Spectrum. Purposefully noisy and drawn from a random uniform distribution.	134
Figure 6.17: The location of the decoy data following perturbation to the mean of each class in LDA-space. The optimal perturbation was calculated using Equation 7.4.....	135
Figure 6.18: LDA-space after the first iteration of GA-LDA. The 2 nd and 3 rd LDA axis are shown in this plot. Decoy data are shown in purple, with the original Raman spectra shown in blue, black, and red.....	136
Figure 6.19: Another view of the first iteration of GA-LDA. This time showing all three LDA axis to demonstrate the high degree of separation of decoy data from measured Raman spectra.	137
Figure 6.20: The loading plots following the first iteration of the GA-LDA. Note the improvement over the original loading plots when LDA was first performed.	138
Figure 6.21: 3D LDA-space after the 20th iteration of GA-LDA. Decoy data are shown in purple, with the original Raman spectra shown in blue, black, and red. The first eigenvector provides a large discrimination power to separate the decoy data from the Raman spectra.....	139
Figure 6.22: Another view of the decoy data in purple and the original Raman spectra in blue, black, and red. The 2 nd and 3 rd eigenvectors best show the resolution in the Raman spectra. ...	140
Figure 6.23: The loading plots (eigenvectors) after 20 iterations of the GA-LDA algorithm. Significant improvements have been made over the original LDA loading plots. The 2 nd and 3 rd eigenvectors now show well defined spectral features that correspond to the spectral features found in each of the 3 original classes of Raman spectra.	141

LIST OF ABBREVIATIONS

Active Pharmaceutical Ingredient	API
Amorphous Solid Dispersion	ASD
Area Under Cover	AUC
Barium Titanate	BaTiO ₃
Blind-Richardson Lucy	BRL
Charge Coupled Device	CCD
Deionized	DI
Differential Scanning Calorimetry	DSC
Direct Current	DC
Electroluminescence	EL
Field of View	FoV
Fourier Transform Infrared Spectroscopy	FTIR
Generative Adversarial Linear Discriminant Analysis	GA-LDA
Hydroxypropyl methylcellulose	HPMC
Instrument Response Function	IRF
Lead Zirconate Titanate	PZT
Limit of Detection	LoD
Linear Discriminant Analysis	LDA
Maximum Likelihood Estimate	MLE
Non-negative Matrix Factorization	NMF
Non-negative Matrix Factorization Filter	NMF-F
Photomultiplier Tube	PMT
Photon Counting	PC
Polyethylene Glycol	PEG
Polyvinylpyrrolidone	PVP
Powder X-Ray Diffraction	PXRD
Principle Component Analysis	PCA
Probability density function	PDF
Process Analytical Technologies	PAT
Relative Humidity	RH

Second Harmonic Generation	SHG
Second Order Nonlinear Imaging of Chiral Crystals	SONICC
Signal to Noise Ratio	SNR
Solid State Nuclear Magnetic Resonance	ssNMR
Stochastic Differential Scanning Calorimetry	SDSC
Thermal Analysis by Structure Characterization	TASC
Thermal Gravimetric Analysis	TGA
Transistor-Transistor Logic	TTL
Triboluminescence	TL
Ultraviolet Transparent	UVT

ABSTRACT

New experimental strategies and algorithmic approaches were devised and tested to improve the analysis of pharmaceutically relevant materials. These new methods were developed to address key bottlenecks in the design of amorphous solid dispersions for the delivery of low-solubility active pharmaceutical ingredients in the final dosage forms exhibiting high bioavailability. New hardware includes: i) triboluminescence based instrumentation for crystallinity quantification, ii) integration of differential scanning calorimetry with second harmonic generation. On the algorithm side, improvements include: i) iterative non-negative matrix factorization for blind deconvolution, ii) calibration-free quantification of crystallinity within second harmonic generation images of pharmaceutical materials, and iii) adversarial spectroscopy, an application of techniques from neural networks to common chemical analysis tools such as principle component analysis and linear discriminant analysis for improving these techniques as well as an interesting case to understand the underlying principles behind neural networks. The combination of hardware and algorithmic techniques have enabled crystallization detection and quantification within model pharmaceutical formulations such as ASDs.

CHAPTER 1. INTRODCUTION

1.1 Philosophy

This dissertation is submitted as a partial fulfillment of the requirements for the bestowal of a Doctorate in Philosophy degree. It is therefore relevant that philosophy is discussed within this document. There are two points of philosophy (or ways to live) that I think are critical, i) a belief that there is absolute truth and ii) leave things better than you found them.

Truth should always be sought after. The beautiful thing about being an analytical chemist is that it is a field primarily concerned with the determination of truth. This is especially true when reporting on a measured analyte. Each value is generally reported with its corresponding confidence interval or with an error bar on a plot. That is where truth is found – between the upper and lower bounds of the reported value (at least, 95% of the time). There is always truth in the world, however, shrinking the confidence interval usually requires great effort and great patience.

The second guiding principle in my life is that things should be left better than you found them. This applies to places, things, and most of all, people. There is no greater joy and satisfaction in my life knowing that someone / something / some place is better because of their interaction with me. This way of living is not always easy or straightforward, however, it is worthwhile. In that spirit, I hope this dissertation can help you in some way.

1.2 Motivation for Research

The solid state form of an active pharmaceutical ingredient (API) can profoundly influence bioavailability and efficacy (Babu & Nangia, 2011; Byrn et al., 1994). The desire for more potent drugs with fewer negative side effects has driven the pharmaceutical industry towards APIs with greater molecular complexity. However, these same highly effective and complex molecules often suffer from low aqueous solubility. Drugs with low solubility in their crystalline form generally exhibit slow dissolution kinetics and low bioavailability. Several strategies have been adopted to address this problem, including co-crystal formulations, the use of metastable crystal forms with faster dissolution kinetics, and preparing the API in an amorphous solid dispersion (ASD)(Baghel et al., 2016; Steed & Steed, 2015; Williams et al., 2013).

While co-crystals and meta-stable crystal forms have advantages in stability and processability, their use can still introduce complications. Meta-stable crystal forms suffer from the possibility of transitioning to more stable, but typically less soluble, crystal forms during manufacturing or storage. A well-known example of this problem is found in the manufacturing of ritonavir (Bauer et al., 2001). Co-crystals generally do not suffer from the same stability issues that meta-stable crystal forms do, because they are often in their most thermodynamically stable state. Co-crystals are solids containing two or more distinct molecular compounds that form a unique crystalline structure. Because co-crystals form a unique crystalline structure, they often exhibit unique thermodynamic and pharmacokinetic properties distinct from either of the single compound crystalline materials. However, the difficulty of finding appropriate molecules to form co-crystals with APIs has prevented widespread adoption of this technique (Aitipamula et al., 2012).

Amorphous solid dispersions are gaining rapid acceptance and widespread use as broadly applicable delivery platforms with relatively straightforward manufacturability. An ASD is prepared by casting the amorphous API in a compatible polymer matrix. The polymer matrix typically maintains the amorphous API in a meta-stable glassy state prior to dissolution (e.g., in the gastrointestinal tract for an oral formulation) (He & Ho, 2015). ASD formulations are generally meta-stable, in which crystallization is suppressed through kinetic rather than thermodynamic control. As such, even trace residual crystallinity in ASDs can negatively impact bioavailability and shelf-life by removing kinetic barriers for nucleation. The possibility of trace crystallinity impacting stability and dissociation leads to extensive stability testing of ASD dosage forms. Stability testing is often performed under accelerated conditions, with Arrhenius-type kinetic modeling enabling extrapolation to long-term predictions under ambient conditions. The reliability of the modeling hinges on precise quantification of the crystallinity over large dynamic ranges.

The most commonly used analytical techniques to quantify crystalline content within APIs and API formulations include the use of differential scanning calorimetry (DSC), powder X-Ray diffraction (PXRD), spectrochemical techniques including Raman, and solid state nuclear magnetic resonance (ssNMR). Unfortunately, these methods have detection limits around 1-2% crystallinity under routine conditions, which can be problematic when drug loadings are on the same relative level. The relatively high limit of detection reduces the dynamic range and complicates reliable stability predictions in accelerated stability testing efforts.

Recently, second harmonic generation (SHG) has been shown to rapidly quantify trace crystallinity within ASDs. The symmetry requirements of SHG, in which coherent signal is only generated from noncentrosymmetric assemblies, allows for a selective measurement of homochiral crystals with detection limits as low as parts per trillion under favorable conditions and routinely achieve parts per million detection limits in powders (Newman et al., 2015). SHG has been used to monitor changes in crystallinity in amorphous solid dispersions in spray-dried dispersions. In that work, the integrated SHG intensity was found to scale with the fraction of crystallinity arising during storage at high temperature and humidity, with limits of detection for crystallinity down to ~0.01% (Correa-Soto et al., 2017; Casey J. Smith et al., 2018). This capability greatly reduced the timescale required for early identification of conditions leading to crystal formation.

Despite successes in detection, precise quantification of residual crystallinity by SHG continuously spanning a wide dynamic range of crystalline content remains challenging. To date, quantification of crystallinity has relied almost exclusively on the integrated SHG intensity produced within a field of view (Chowdhury et al., 2016; Correa-Soto et al., 2017). This approach has significant limitations for developing predictive models for crystallization kinetics. Quantification of SHG by integrated SHG intensity also typically requires the need for calibration standards, in which measurements of the pure substance provide a reference point for connecting the measured SHG intensity back to the percent crystallinity. However, access to reliable standards is not always trivial. For many APIs, several different crystal forms are potentially kinetically accessible during crystal formation. Given the high sensitivity of SHG to crystal form, standards prepared based on the assumption of one polymorphic form would be inapplicable for measurements with a different crystal form. Algorithms have been developed to address bias from the particle size distribution (C. J. Smith et al., 2018), but represent yet another potential complication for routine quantification of crystallinity.

Trace crystallinity measurements that take advantage of particle counting techniques enable quantification of trace crystallinity content with high statistical certainty. In comparison to signal integration, particle counting approaches have the potential to lower the limits of detection by SHG by removing SHG intensity bias from particle size and orientation, as well as removing the noise floor through thresholding. Differences in particle size and orientation greatly influence the SHG intensity, but not the integrated area in the image. Despite the numerous advantages of particle counting in the limit of low crystallinity (<1%), the dynamic range is quite limited; as the overall

crystallinity increases, so too does the probability of observing spatially overlapping crystals. This effect ultimately sets the upper limit on the accessible dynamic range. Unfortunately, the bias from particle overlap introduces a measurement gap, in which neither technique (particle counting or signal integration) yields an optimal unbiased signal to noise ratio.

1.3 Ph.D. Summary

The research completed for my Ph.D. can be summarized as occurring in two general areas, i) analytical chemistry to inform the physical state of pharmaceutical formulations and ii) improvements to common chemometric techniques. Chapters are organized according to projects and chronological.

Chapter 2 covers the development of impact driven triboluminescence which enables rapid assessment of crystalline content within pharmaceutical powders. The expense and complexity of ultrafast optical methods complicates applications of SHG microscopy for routine and rapid on-line and at-line reaction monitoring. In addition, optical analysis of powders is often limited by a penetration depth on the order of the mean particle size, potentially introducing sample bias if the crystallinity is heterogeneously distributed throughout the powder. Alternative methods based on simpler and more robust measurement platforms would facilitate on-line reaction monitoring and process analytical technologies (PAT). As described in this chapter, triboluminescence was developed as a sensitive and selective complementary indicator for the presence of residual noncentrosymmetric crystals in model ASDs comprised of trace crystalline active pharmaceutical ingredients (APIs) cast within an amorphous polymer matrix. Additionally, the triboluminescence technique was adapted to support measurements of crystallinity in liquid slurries. Key outcomes include i) the development and characterization of the triboluminescence instrument and ii) demonstration of a limit of detection of 140 ppm.

Chapter 3 details the development of a digital filtering approach integrating non-negative matrix factorization (NMF) for blind deconvolution and the application of the method to deconvolve the time trace of triboluminescence events. The previous work on triboluminescence left some lingering ambiguities in the mechanism underpinning TL measurements. The most obvious difference between proposed TL mechanisms was the timescale at which they occur. The intentional impedance mismatch and the resulting pulse-overlap obscured the underlying time scales of TL events. This non-negative matrix factorization filter (NMF-F) approach was evaluated

on test data with known ground-truth results, then implemented to provide insights into the underlying TL measurement timescales of model pharmaceutical materials doxycycline hyclate, an antibiotic of the tetracycline class, and griseofulvin, an antifungal agent.

Chapter 4 showcases analytical formulas designed to correct for particle-particle overlap when performing image analysis to determine the volume of crystallinity within a sample. Wherein, signal integration and particle counting methods for crystallinity quantification are bridged through the derivation of a statistical model based on analytical expressions for the optimization of signal to noise. Specifically, a particle-counting algorithm capable of unbiased extension of the counting regime to overlap with the regime optimized for signal integration was developed. Based on the relative variance of the recovered volume, an analytical expression was derived for the percent crystallinity for which both bias-corrected particle counting and signal integration produced an identical average signal to noise ratio. Crystalline fractions higher than this cross-over level were optimally determined by signal integration and fractions below by bias-corrected particle counting. The average SHG-activity per unit volume measured in the particle counting regime served as an internal calibration standard for quantifying crystallinity in the signal integration regime, removing the need for external standards. Monte Carlo simulations were performed to assess the merits of the approach. Measurements were performed on evacetrapi, a cholesteryl ester transfer protein inhibitor, formulated tablet to experimentally demonstrate the bridging of the two regimes and autocalibration.

Chapter 5 demonstrates the integration of SHG microscopy with DSC analysis to monitor phase transitions on a per-particle basis. This multi-modal approach enabled simultaneous single-particle analysis for distinguishing the impact of crystal size and shape on phase transitions in a single experiment, complementing the ensemble-averaged information obtained from DSC. Furthermore, the strict symmetry requirements for SHG allowed the mapping of complex desolvation and crystal form transitions with sensitivity in the ppm regime. The merits of integrated SHG and DSC measurements were explored for a model system, then applied to map the complex phase transitions in trehalose dihydrate.

Chapter 6 details the development of an adversarial attack applied to Raman spectra of different polymorphic forms of an API (clopidogrel bisulfate) in a reduced dimensional space defined by linear discriminant analysis (LDA) which is directly analogous to the adversarial step

in a generative adversarial neural network. The extension of this attack was then applied iteratively to improve the stability and statistical reliability in data-limited training scenarios.

CHAPTER 2. TRIBOLUMINESCENCE

2.1 Overview

Triboluminescence is the emission of light upon mechanical disruption of a crystal. Triboluminescence only occurs in crystals that are piezoelectric, e.g. adopt space groups that are non-centrosymmetric and have a permanent dipole. In triboluminescence pressure to the point of fracturing is applied over a short time period. The rapid change in potential across the crystal due to fracturing results in a large transient current associated with the lattice displacement. For sufficiently large potential changes, plasma from dielectric breakdown can produce light emission in the form of TL. Following fabrication and characterization of a novel instrument to measure triboluminescence, TL is used as a sensitive and selective indicator for the presence of crystalline active pharmaceutical ingredients with an amorphous solid dispersion. An alternative instrument is developed and used for the measurement of dissolution of pharmaceutical formulations in biorelevant media.

2.2 Introduction

Triboluminescence was first recorded by Francis Bacon when he recorded the phenomenon, “It is well known that all sugar, whether candied or plain, if it be hard, will sparkle when broken or scraped in the dark.” (Bacon, 1620) Triboluminescence at first glance might appear to be a misnomer. The prefix “tribo” comes from Greek for “to rub”, however, the word triboluminescence can be applied to the emission of light upon mechanical action. The phenomenon was popularized by a party trick involving wintergreen Lifesavers and is often used as a common starting point on the discussion of TL. Additionally, triboluminescence has been recorded in native American (Ute) ceremonial rattles which contain quartz rocks, a piezoelectric material.(Fei et al.)

TL is the emission of light upon mechanical disruption of a crystal. In our studies, only crystalline material that is piezoelectric (e.g. non-centrosymmetric and has a permanent dipole) shows any TL activity. However, there are a couple of recorded exceptions to this in the literature. The piezoelectric effect is the underlying mechanism that allows for TL. Piezoelectric materials show a physical response to applied voltages, or vice versa, a voltage in response to a mechanical

stimulus. This physical response is due primarily to the permanent dipole in the crystalline lattice, upon an applied voltage, these dipoles will repulse one another, and the physical dimensions of the crystal will expand. In TL of piezoelectric materials, pressure to the point of fracturing is applied over a short time period. The rapid change in potential across the crystal results in a large transient current associated with the lattice displacement. For sufficiently large potential changes, plasma from dielectric breakdown can produce light emission in the form of TL. Evidence supporting dielectric breakdown as a mechanism for TL emission can be found in studies demonstrating plasma-generated electron impact excitation of molecular N₂ which is the most commonly seen line spectra associated with TL as well as the studies presented here.(Sweeting, 2001)

2.3 Motivation

The solid-state form of an active pharmaceutical ingredient (API) can profoundly influence bioavailability and efficacy. (Babu & Nangia, 2011; Byrn et al., 1994) The desire for more potent drugs with fewer negative side effects has driven the pharmaceutical industry towards APIs with greater molecular complexity, with a corresponding drop in aqueous solubility. A study published in 2011 estimated 80-90% of all new small molecules in the pharmaceutical pipeline suffer from poor aqueous solubility.(Babu & Nangia, 2011) Drugs with low solubility in their crystalline form generally exhibit slow dissolution kinetics, with correspondingly low bioavailability. Several strategies have been adopted to address this problem, including co-crystal formulations(Aitipamula et al., 2012; Steed & Steed, 2015), the use of metastable crystal forms with faster dissolution kinetics(Williams et al., 2013), and preparing the API in an amorphous solid dispersion (ASD)(Baghel et al., 2016).

While co-crystals and meta-stable crystal forms have advantages in stability and processability, their use can still introduce complications. Metastable crystal forms suffer from the possibility of transitioning to the more stable, but typically less soluble, crystal form during manufacturing or storage. Alternatively, co-crystals generally do not suffer from the same stability issues since they are often in their most thermodynamically stable state. However, the difficulty of finding appropriate molecules to form co-crystals with APIs has prevented the widespread adoption of the technique.

Amorphous solid dispersions are gaining rapid acceptance and widespread use. An ASD is prepared by casting the amorphous API in a compatible polymer matrix. The polymer matrix typically maintains the amorphous API in a metastable glassy state prior to dissolution (e.g., in the gastrointestinal tract for an oral formulation). (He & Ho, 2015) ASD formulations are metastable, in which crystallization is suppressed through kinetic rather than thermodynamic control. (He & Ho, 2015; Konno & Taylor, 2008) As such, even trace residual crystallinity in ASDs can negatively impact bioavailability and shelf-life by removing kinetic barriers for nucleation. The possibility of trace crystallinity impacting stability and dissociation leads to extensive stability testing of ASD dosage forms.

The most commonly used analytical techniques to quantify crystalline content within APIs and API formulations include the use of differential scanning calorimetry (DSC), (Johari et al., 1990) powder X-Ray diffractometry (PXRD), (Andronis et al., 1997) spectrochemical techniques including Raman, (Berry et al., 2008; Bugay, 2001; Ng et al., 2013; Taylor & Zografi, 1998) and solid state nuclear magnetic resonance (ssNMR). Unfortunately, these methods have detection limits around 1-2% crystallinity under routine conditions, which can be particularly problematic when drug loadings are low within a formulation.

Recently, second harmonic generation (SHG) has been shown to rapidly quantify trace crystallinity within ASDs. (Schmitt et al., 2015) The symmetry requirements of SHG, in which coherent signal is only generated from noncentrosymmetric crystals, allows for a selective measurement with detection limits as low as parts per trillion under favorable conditions and routinely achieves ppm limits in powders. (Wanapun et al., 2011) Trace crystallinity measurements that take advantage of particle counting techniques enable quantification of trace crystallinity content with high statistical certainty. (Chowdhury et al., 2016) SHG has been shown to exhibit limits of detection greater than 3 orders of magnitude lower than conventional benchtop methods. Advances in the advent of ultrafast fiber lasers have further improved access to affordable turn-key laser sources for SHG microscopy. However, the expense and complexity of ultrafast optical methods complicates applications of SHG microscopy for routine and rapid on-line and at-line reaction monitoring. In addition, optical analysis of powders is often limited by a penetration depth on the order of the mean particle size, (Chowdhury et al., 2016) potentially introducing sample bias if the crystallinity is heterogeneously distributed throughout the powder. Alternative methods

based on simpler and more robust measurement platforms would facilitate on-line reaction monitoring and process analytical technologies (PAT).

In this work, impact driven triboluminescence (TL) enables rapid assessment of crystalline content within pharmaceutical powders. TL is the emission of light upon mechanical disruption of a crystal. The piezoelectric effect greatly enhances TL activity in crystals that are symmetry-allowed for piezoelectricity.(Sweeting, 2001) In conventional TL of piezoelectric materials, pressure to the point of fracturing is applied over a short time period. The rapid change in potential across the crystal results in a large transient current associated with the lattice displacement. For sufficiently large potential changes, plasma from dielectric breakdown can produce light emission in the form of TL. Evidence supporting dielectric breakdown as a mechanism for TL emission can be found in studies demonstrating plasma-generated electron impact excitation of molecular N₂ which is the most commonly seen line spectra associated with TL.(Sweeting, 2001) In the present study, the presence of TL is used as a sensitive and selective indicator for the presence of residual noncentrosymmetric crystals in model ASDs comprised of trace crystalline active pharmaceutical ingredients (APIs) cast within an amorphous polymer matrix. Given the close similarities in the symmetry properties of both piezoelectricity and SHG, TL arising from the piezoelectric mechanism is anticipated to exhibit similar selectivity for chiral API crystals in a much simpler measurement platform.

2.4 Experimental Section

Sucrose purchased from Sigma-Aldrich was ground with a mortar and pestle and then sieved to obtain 63-125 μm , 125-250 μm , 250-500 μm , and >500 μm diameter particle fractions.

Polyethylene glycol (PEG), BioUltra 1,000, purchased from Fluka Analytical was ground with a mortar and pestle and then sieved to obtain 125-250 micron particles.

Griseofulvin purchased from Sigma-Aldrich was sieved to isolate particle sizes between 63 μm and 125 μm , and physically mixed by stirring with appropriate amounts of the sieved PEG to make mixtures with concentrations of 10, 5, 2, 1, 0.5, 0.1, and 0.05% by weight. Mixtures were prepared by serial dilution and stirred for 2 minutes with a spatula between additions.

Two different polymorphic forms of clopidogrel bisulphate were prepared in-house at Dr. Reddy's Laboratory (Hyderabad, India). Formulation protocols for clopidogrel bisulphate were followed to produce ~60 μm diameter spheroidal particles with a narrow size distribution. The

polymorphs of clopidogrel bisulphate were identical in their preparation except for a single step in the reaction, in which the Form II particles were controllably produced from the Form I particles through a polymorph transition.

To prepare the Griseofulvin HPMCAS-MF ASD samples, Griseofulvin was obtained from Sigma-Aldrich, and HPMCAS-MF was purchased as AQOAT from Shin Etsu. The spray dry formulation comprised of Griseofulvin (20% w/w) and HPMCAS-MF (80% w/w). The solid components were dissolved in Acetone (4% w/v). The solution was spray dried on a 4M8-TriX spray dryer (ProCepT NV, Zelzate, Belgium) with a two-fluid nozzle diameter of 0.6 mm. The dried particles were carried by air into a cyclone and then into a collection vessel. Three processing rates were controlled and monitored: 1) solution feed rate, 2) processing air flow rate, and 3) atomizing air flow rate. The processing air flow rate was set to be 0.40 m³/min. The atomization air and the solution feed flow rates were taken to be 3.0 L/min and 6 mL/min, respectively. The spray drying conditions were set to reach inlet temperature of 91 °C and outlet temperature of 50 °C. The cyclone differential pressure was set at 30 mbar and tangential air flow rate was chosen to be 0.1 m³/min while the cooling air was shut off. The spray dried products were further dried under vacuum at ambient temperature overnight.

2.5 TL Instrument Fabrication & Characterization

The ideal TL instrument would deliver a reproducible impact to the sample, then require little set up until the next sample could be run. Since the instrument deals primarily with powders of materials that could potentially be hazardous, the system needs to be able to contain the powder and prevent it from spreading during the application of force to the sample.

Initial proof of concept work was done on an instrument constructed by Scott Griffin & the author during the summer of 2015. The instrument consisted of a steel tube with holes at regular intervals. The holes were used to provide a consistent drop height for a brass ball which was used to impact the powdered sample. Underneath the tube was a clear plexiglass disc, whereupon the powder sample would be placed. Underneath the plexiglass disc, a Hamamatsu photomultiplier tube (PMT) was placed to collect light. The PMT was connected to an oscilloscope (ATS 9462, AlazarTech) which recorded the voltage transient following an impact event.

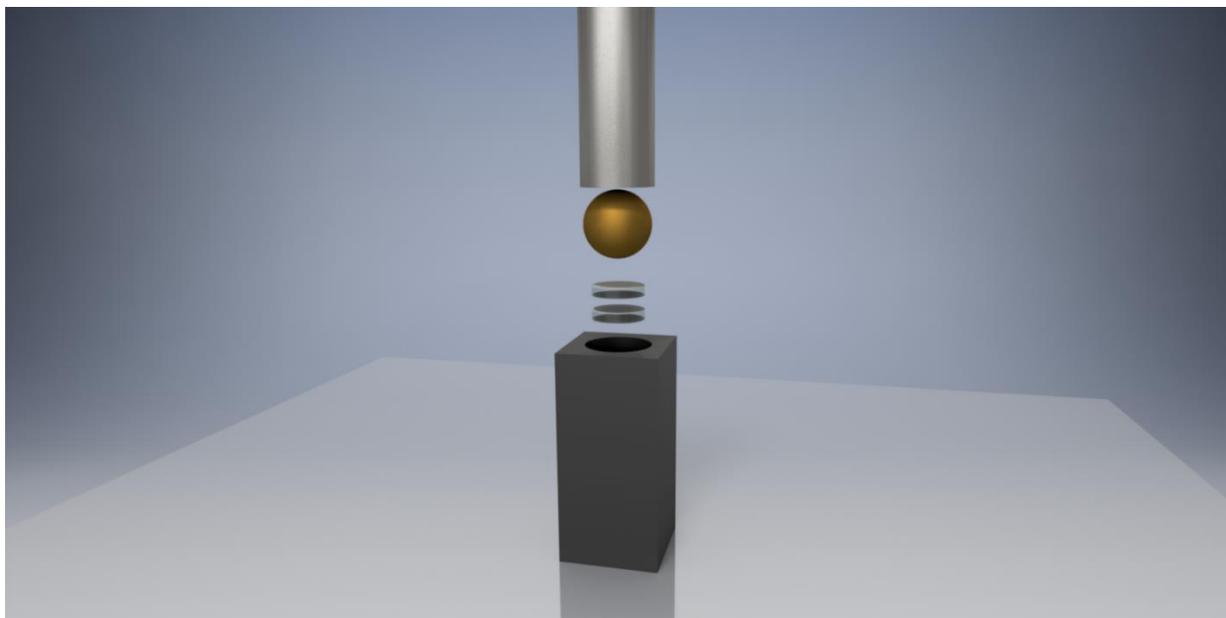


Figure 2.1. Graphical depiction of the initial proof of concept instrument for TL analysis of pharmaceutically relevant materials.

Figure 2.1 is a diagram of the proof-of-concept instrument. This initial instrument allowed for the proof of concept work to be completed. However, the instrument in **Figure 2.1** was not optimized for routine sample analysis. Only a small portion of sample had force applied to it, due to the curvature of the brass ball used to impart a force. More of the sample was probed per impact by the addition of a second plexiglass disc and sandwiching the sample between the two discs. However, the low duty cycle of the instrument remained as a larger, more pressing issue. A single impact event, $\sim 30 \mu\text{s}$ of recorded data, required ~ 10 minutes of preparation work. More concerning than the duty cycle was the lack of sample containment. Upon impact, the powdered sample would be spread over both plexiglass plates, the brass ball, and beyond. Cleaning required the disassembly of the instrument. The spread of powdered sample was not concerning when measuring sucrose, but the intended application to pharmaceutical formulations encouraged the design and fabrication of an improved instrument.

A list was compiled for the next generation of the instrument. The new instrument should have the following: i) a high duty cycle, ii) containment of the powder sample both for speed of the instrument and safety of the operator, iii) reduction of ambient light, iv) require less sample, v) more consistent force applied to each impact, and vi) digital data acquisition.

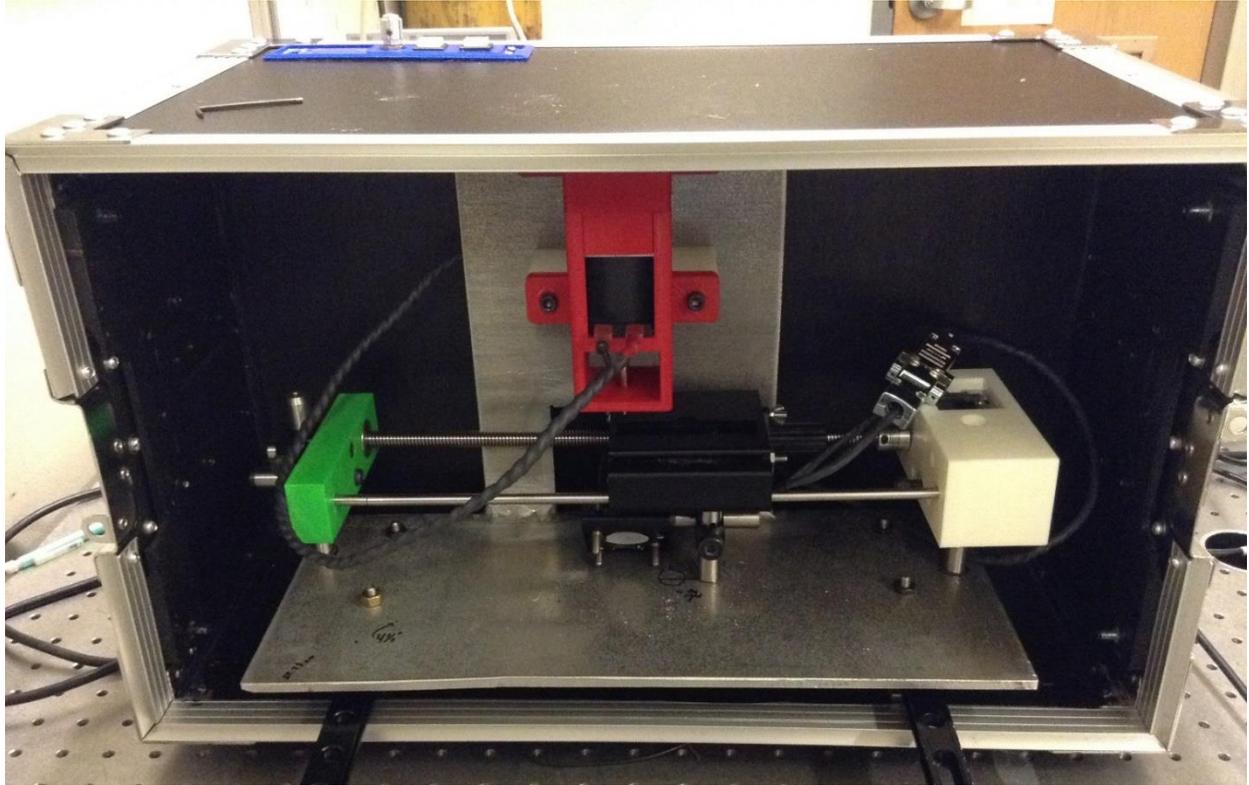


Figure 2.2: Image of the redesigned triboluminescence instrument. The instrument was housed in a sealable box to allow for measurements in a low light environment. The red plastic houses the solenoid, used to impact the sample. The green plastic provides support for a threaded rod and a support rod. The white plastic houses the stepper motor, used for sample translation. The black plastic piece holds the sample and is translated following each impact event. Underneath the black plastic piece, the PMT is aligned to collect any emitted light from TL.

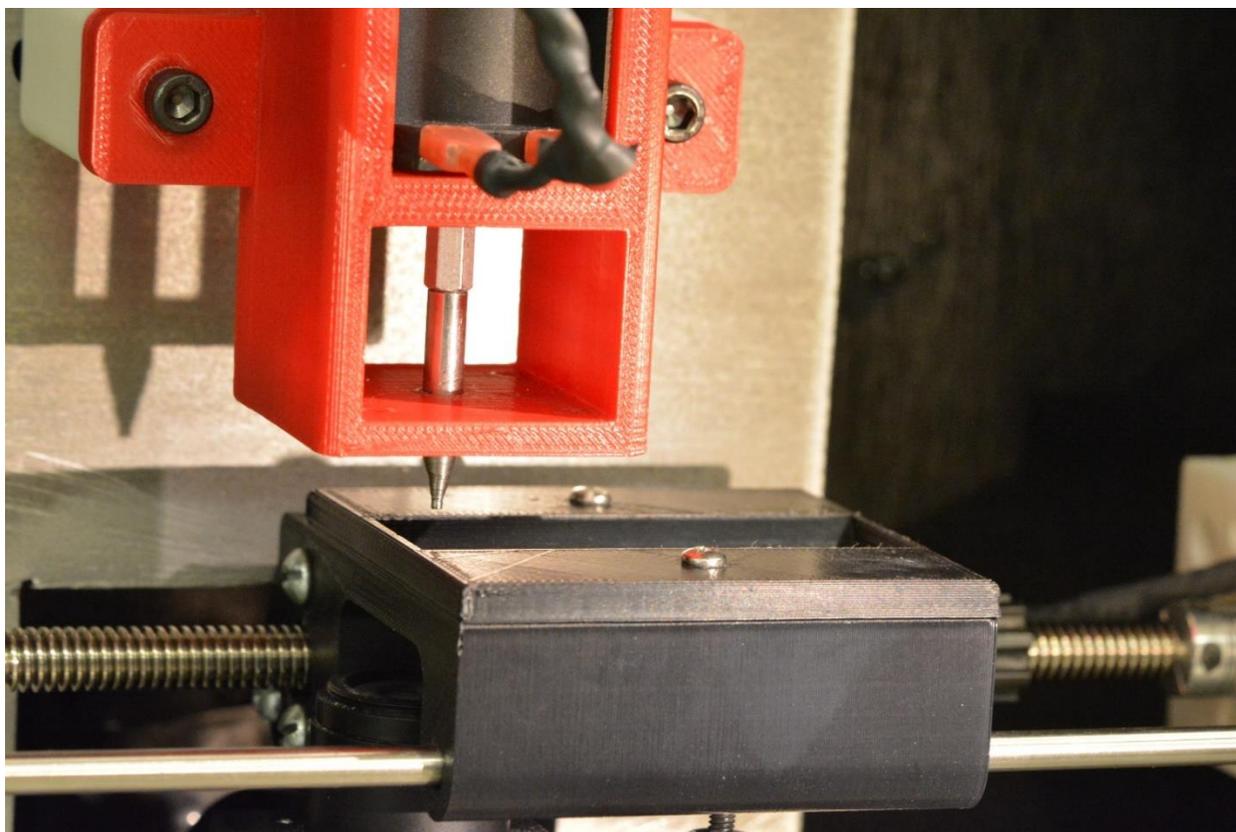


Figure 2.3: Close up image of the hammer and sample stage (black).

The final build sheet for the redesigned TL instrument can be found in Appendix B. Pictures of the redesigned instrument can be seen in **Figures 2.2 & 2.3**. Triboluminescence was induced in powdered samples using mechanical force delivered from a solenoid as depicted in **Figure 2.3**. The instrument consisted of: a solenoid (240 V DC, 20 mm stroke, pull type) housed in the red plastic frame with a custom power source (~200 V DC) designed and built by the Jonathan Amy Facility; a custom milled hammer attached to the solenoid; a custom built translation stage with a stepper motor to move the sample; and two lenses (focal length of 25mm, A coated) to collect and direct the generated light onto a photomultiplier tube (PMT, Hamamatsu H10721-210, Japan). A digital oscilloscope card (AlazarTech, Pointe-Claire, Quebec, Canada) recorded the voltage transients from the PMT at 100 mega-samples a second, timed to the impact events. An Arduino Uno R3 board (Arduino, Italy) with a coupled Adafruit motor shield (Adafruit, New York, New York) was used to control all timing. The Arduino board delivered a 5 volt TTL pulse to a mechanical relay, which closed the high voltage circuit to the solenoid, leading to an impact of the sample. A concurrent TTL pulse was sent to the digital oscilloscope card to initiate

the recording of data. A digitation window of 25 ms was used for each TL event. After each TL impact, the Arduino board controlled the stepper motor to translate the sample to provide an unperturbed area for the next TL event. Including sample translation time, each TL measurement took ~1s, while an entire sample took ~40 s for 37 impact events. The solenoid housing, stepper motor mounts, sample stage, and hammer guide were created by additive manufacturing (Lulzbot Mini, Colorado).

An impedance mismatch was intentionally introduced between the transmission line and the digitizing electronics. This mismatch provided a signal to noise advantages through two independent mechanisms: i) increasing the integrated voltage through the increased resistance of the digitizer according to Ohms Law, $V=IR$, and ii) simplifying digitization of the transients by increasing the temporal duration of each voltage transient according to the RC time constant associated with discharge of the transmission line capacitance through the register. The increase in voltage manifested as a series of reflections separated by the round-trip transit time through the transmission line of 100 ns, followed by eventual decoherence to a constant baseline offset. The resulting complex waveform produced by an impulsive photon absorption event still retains essentially the same high frequency information content of the original current transient from the detector, but with different frequencies shifted in phase to produce the observed impulse response function.

A sample holder was designed to improve the consistency in sample prep (e.g. density and packing), reduce ambient light noise, and to prevent the spread of API powder samples during TL measurements. UVT Acrylic slides (Ted Pella, Inc.) were used as a base and polystyrene microscope slides (Fisher Scientific, S67112A) were cut and solvent-welded to the base using acetone. See **Figure 2.4** for an example of the sample holder. To increase consistency in channel dimensions, a custom 3D-printed guide was created and used during sample slide preparation. Channel dimensions were 75 x 5 x 1 mm. Electrical tape was placed over the channel to enclose the slide while leaving a space below for the light to pass through. The electrical tape allowed for safe handling of powdered API samples and prevented the spread of API powder during TL measurements. Several types of tape were tested for their ability to contain the powdered sample. The testing consisted of striking a taped slide with and without powdered sample to assess the background signal (if any). All tested tapes were able to sufficiently contain the powdered sample. Among the types of tape tested were the following: black electrical tape, clear packing tape, clear

scotch tape, and duct tape. During testing it became apparent that an additional parameter was important to consider, the pliability of the tape. If the tape is too stiff, the energy from the hammer did not translate well to the sample and little to no TL signal was recorded. The clear packing tape and the clear scotch tape were found to have a background signal. It was not clear if the background signal was coming from the tape itself or from the translucent nature of the tape. Duct tape was found that it was not sufficiently pliable to transfer energy from the solenoid driven hammer to the sample. Electrical tape was chosen as the tape of choice because of its negligible background, ability to contain the powdered sample, and its pliability.

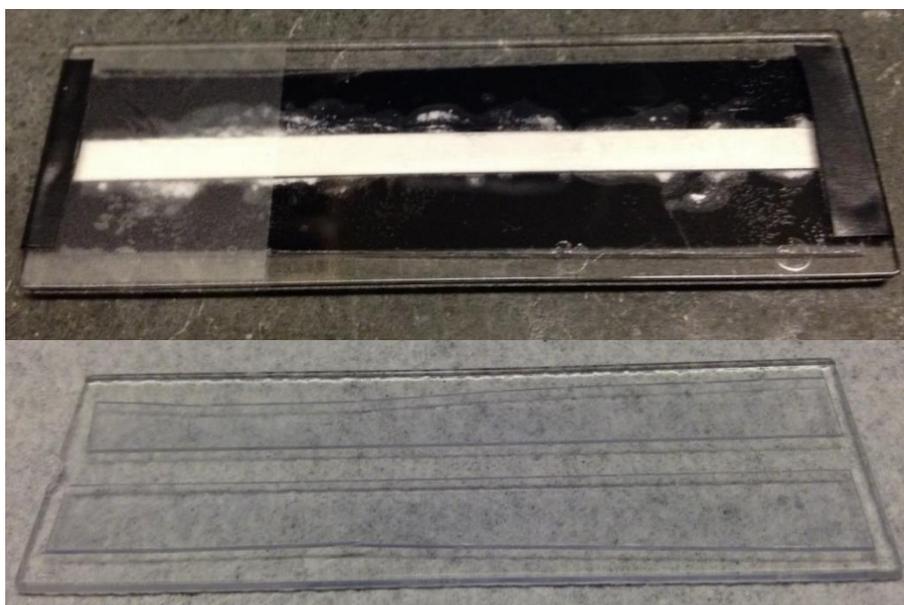


Figure 2.4. Example of the sample slides used in triboluminescence experiments. Top image is of the slide with sample loaded into the channel and sealed in with black electrical tape. The bottom image is the sample slide before loading with sample.

The energy imparted by the solenoid in a single impact was 0.5 Joules over a 7 mm² area, this was calibrated based on the response obtained from striking a similarly prepared control slide with a brass weight dropped from different known heights. Characterization of the solenoid found that maximum force was applied when samples were located 1-2 mm before its maximum stroke length.

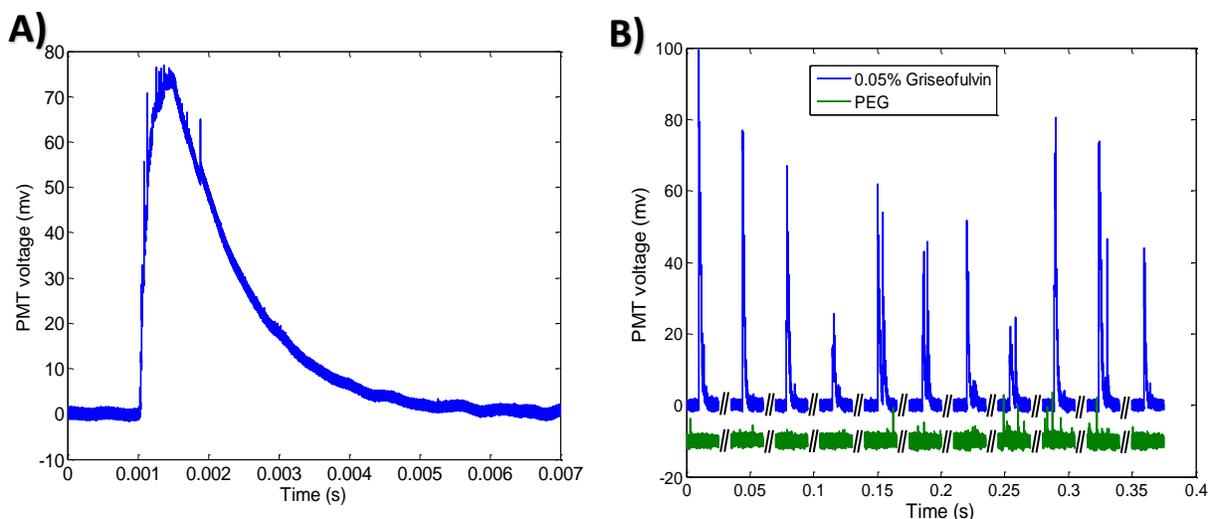


Figure 2.5: A) Representative trace from a 0.05% by weight griseofulvin sample. B) 11 impact events concatenated into a single graph are shown to illustrate the shot to shot variability. The green (lower), offset data trace is pure PEG, shown to demonstrate the baseline response in the absence of crystalline griseofulvin.

Representative time-traces of the PMT sensor outputs during impact events are shown in **Figure 2.5**. Several features are noteworthy. First, relatively long decay times arose from the intentional impedance mismatch of the PMT inputs to the oscilloscope cards to facilitate detection. Impedance matching resulted in transients from single-photon events only a few ns in duration and too short for digital capture at the sampling rates used. In addition, the time-traces do not correspond to a simple single-exponential decay impulsively initiated, but often present multiple peaks within a single impact event (**Figure 2.5B**). Multiple origins for the observations of distributions in emission times were considered: i) emission from multiple sources with slight differences in timing, ii) acoustic recurrences from reflection, and iii) multiple impacts from recoil of the impactor. Emission from multiple sources is consistent with the time for which the impactor translates through the powdered sample, 0.5 ms, determined by slow motion analysis of the impact event. Consideration of the speed of sound in the polyacrylic substrate of >2500 m/s (Selfridge, 1985) suggests a recurrence time of <400 ns, which was much shorter than the observations. The presence of multiple impacts from recoil and bouncing of the impacting piston were characterized by complementing the optical measurements with recordings of the acoustic response upon impact. These measurements revealed evidence for multiple collision events from recoil over the timeframe of the transients shown in **Figure 2.5**. Consequently, the temporal trends are tentatively

attributed to the presence of multiple sources stochastically producing light emission at different time points during the primary impact event and the occasional presence of a secondary recoil hit.

A significant shot-to-shot variability in the emission intensity was observed in **Figure 2.5B**. Two sources of variance were considered: i) variability in impact force, and ii) variability in the number, sizes, and orientations of impacted crystals. Measurements performed to assess the first of the two possible sources were obtained by recording the variance in acoustic amplitude on a lead-zirconate titanate (PZT) piezoelectric transducer (not shown). In brief, the relative standard deviation in shot-to-shot acoustic amplitude upon impaction was 6%, which was significantly lower than the relative standard deviation in detected intensity of 53%. Consequently, the major source of shot-to-shot variance was attributed to the stochastic nature of the sample coupled with the limited volume ($\sim 7 \text{ mm}^2$) of material sampled.

Many data traces contained multiple peaks. The origin of these multiple peaks were explored by characterizing the time frame of their occurrence. 319 impacts were visually analyzed for multiple peaks. The time difference between the first peak and the second peak was calculated and plotted on the histogram shown in **Figure 2.6**. There were two main groupings of when the secondary peak occurred, i) between 0 and 1 ms and ii) between 1.3 and 2.4 ms. Two mechanisms were explored to explain the presence of secondary peaks, i) multiple point sources emitting light at the piston completes its travel, and ii) a secondary recoil hit from the hammer. Using a high-speed camera, the time the hammer takes to travel through the sample was approximately 0.5 ms. Additionally, the high-speed camera also revealed that a secondary recoil hit would occur approximately 3 ms after the initial impact. These two mechanisms appear to explain the two time frames for the secondary peak well. The first cluster of peaks occurring between 0 and 1 ms being attributed to the hammer traveling through the powdered sample and leading to the emission of multiple point sources. The second cluster of peaks between 1.3 and 2.4 ms aligning with the presence of a secondary recoil hit.

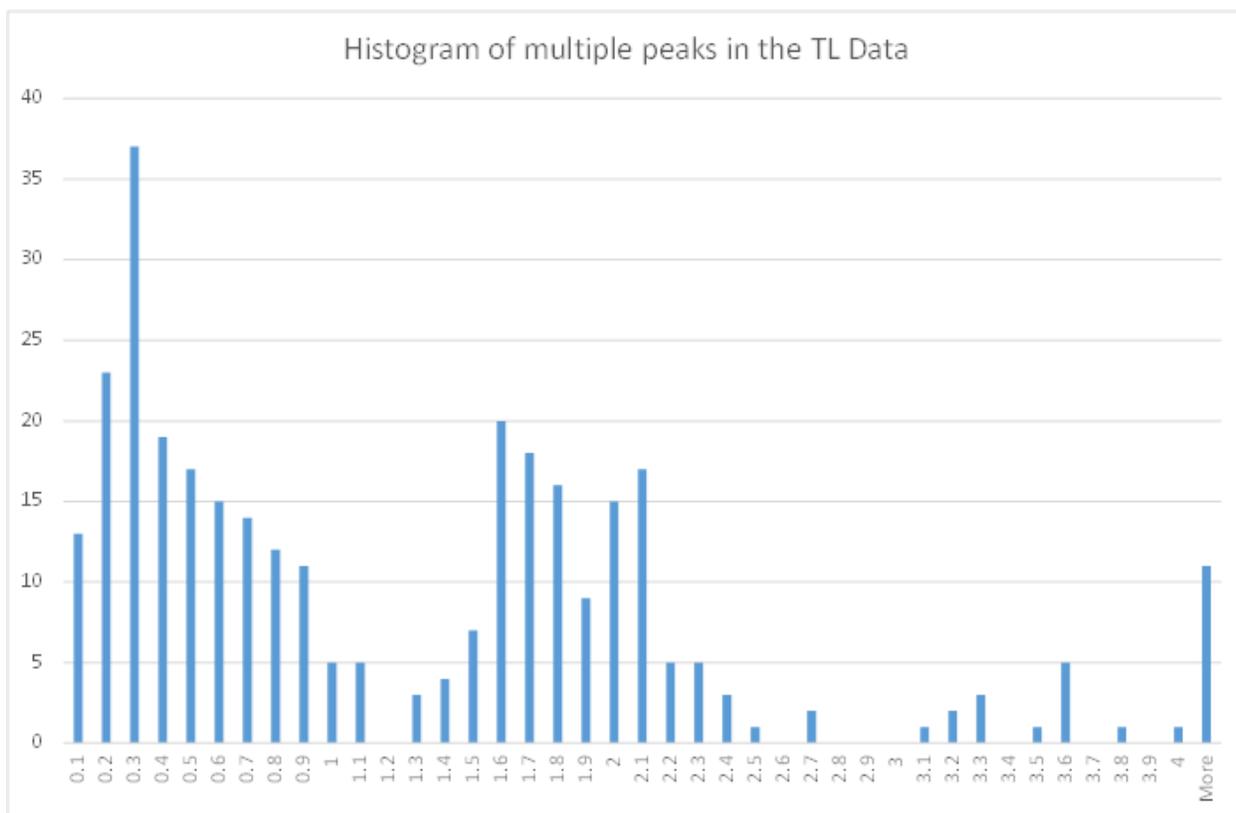
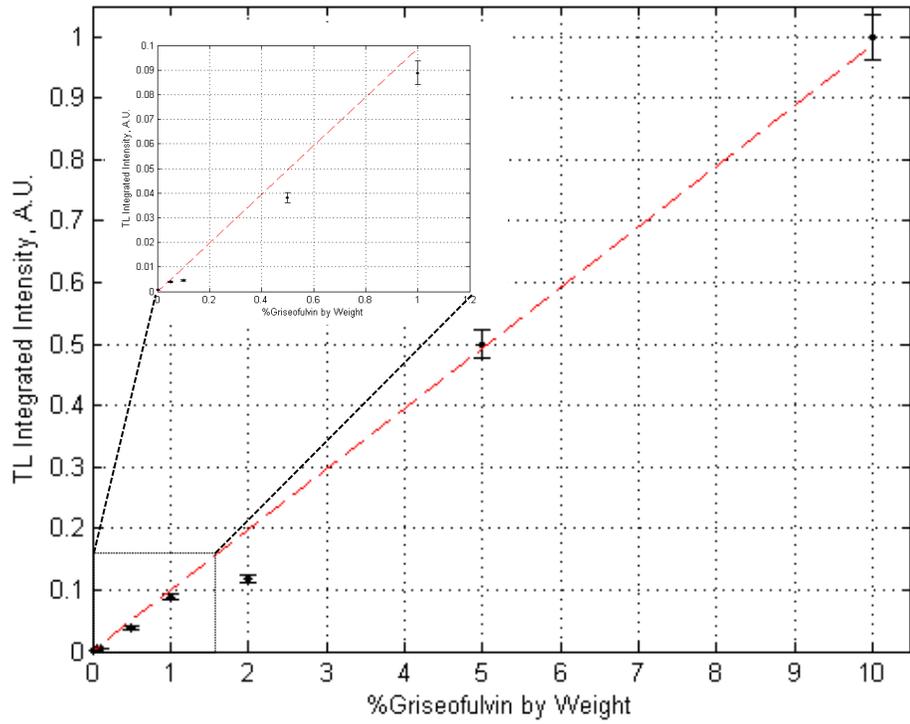


Figure 2.6: A histogram showing the time difference from the initial peak to the secondary peak. There are two groupings clearly seen at ~0.4 ms & ~1.9 ms.

Calibration of the instrument was performed to provide an estimate of the limits of detection, shown in the **Figure 2.7**. All calibration measurements shown were acquired for powdered blends of crystalline griseofulvin in polyethylene glycol (PEG), chosen as a simple model for an amorphous solid dispersion. PEG was chosen as an amorphous substrate due to the considerable interest in it for use in ASDs. (Guedes et al., 2011; Law et al., 2004) Despite its semi crystalline nature, PEG was found to yield little to no detectable background as shown in **Figure 2.5**. The absence of significant background TL activity is consistent with the symmetry expectations of the measurements; centrosymmetric crystalline content does not support the piezoelectric response that drives the observed TL. Lower limits of detection (LoD) were found to be 140 ppm, with the low values arising primarily from the low background TL observed within the API-free polymer blanks.

a)



b)

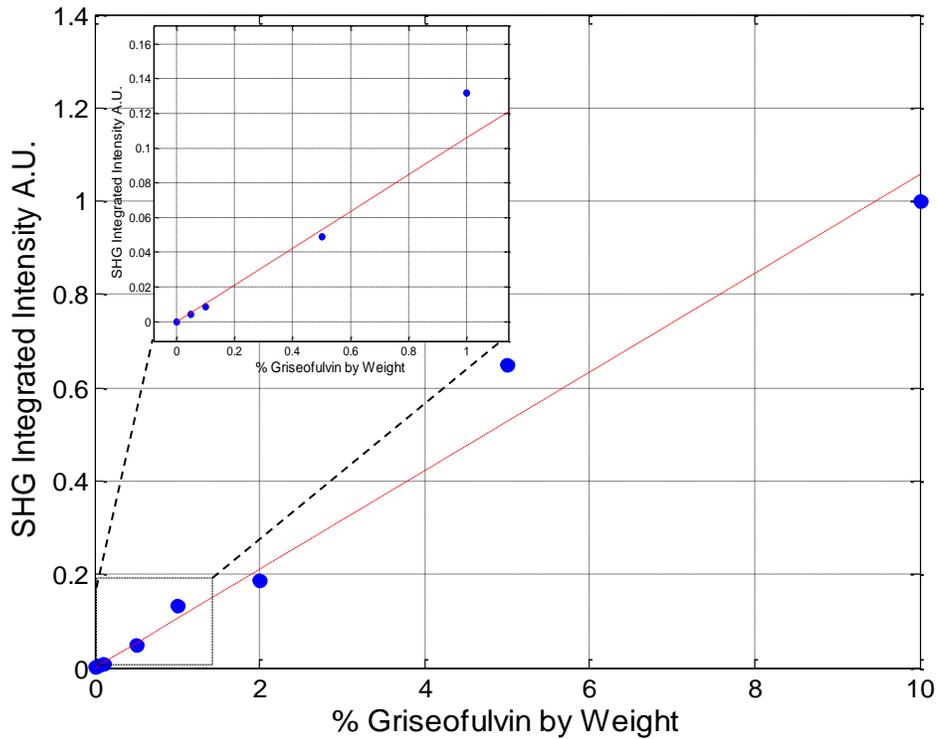


Figure 2.7: Both TL (a) and SHG microscopy (b) were performed on powdered blends of griseofulvin in PEG to demonstrate the calibration curve. In each case, the inset shows 5 data points with concentrations at 1% or less.

2.6 Triboluminescence Theory

Triboluminescence is the phenomenon of light being released upon application of force to a piezoelectric crystal. The force applied need to be enough to fracture the crystal in some way. However, the underlying mechanism for TL remains disputed in literature.

Based on previous reports suggesting several possible mechanisms underlying TL, a suite of studies was designed to assess the dominant contributions to the measurements of the model ASD. Two dominant mechanisms have emerged depending on the symmetry of the sample. The first mechanism considered is symmetry independent and is related to electroluminescence (EL). EL is understood to arise from the radiative recombination of electrons and holes, typically within a semiconductor or in localized trap states, caused by an applied current. In the case of TL, it has been conjectured that TL can also arise from electron-hole recombination, with the electrons and holes produced at defect sites upon lattice disruption. Evidence supporting this mechanism can be found from the large variety of materials capable of producing TL under ambient conditions and in vacuum, and from the similarities between the TL and EL spectra of some TL active crystals.⁴³ The second mechanism is selective to piezoelectric materials. In brief, the rapid mechanical action arising upon physical disruption of the crystal results in large and rapid lattice displacements. In piezoelectric materials, these displacements correspond to fast changes in capacitance, which if large enough, can produce electrical breakdown and plasma generation. While EL has been noted to contribute to the TL signal of some low band gap piezoelectric crystals (i.e., doped zinc sulfide), crystals with a high band gap in the UV are expected to have little to no contribution from EL. Two separate sets of experiments were performed to assess the dominant mechanism underlying the observations of TL in the ASD model systems investigated in this study, with the results from the first set shown in **Figure 2.8**.

In **Figure 2.8**, the TL intensity significantly increased with increasing crystal size in sieved, powdered samples. The total volume and chemical composition of the material was identical in all the samples investigated in **Figure 2.8**, differing only in the crystal size distribution. The trends in the figure are in good qualitative agreement with the piezoelectric mechanism, as increases in crystal size correspond to increases in both the standing potential across the crystal (for dipolar materials, consistent with the $P2_1$ space group of sucrose used in **Figure 2.8**) and the capacitance from the increase in area. As such, TL is expected to increase monotonically with increases in crystal size, as the entire crystal contributes to the potential change from the rapid displacements

associated with crystal collapse. In contrast, similar effects should not be expected in the strain-based mechanism that is independent of piezoelectricity. If the light-producing strain is highly localized relative to the size of the crystal, it is not clear that a change in crystal size would significantly impact the resulting light emission.

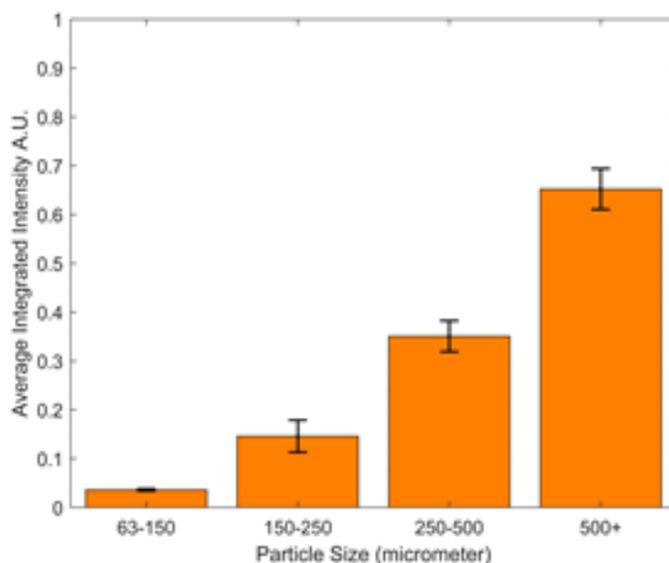


Figure 2.8: Average integrated signal for 4 different particle sizes of sucrose crystals. The error bars represent the standard deviation of the mean slide value (e.g. 36 averaged impacts).

The results from TL detection are compared with measurements probed by SHG microscopy, see **Figure 2.9**. The SHG results were produced by integrating the SHG intensity over a field of view of ~ 2 mm. Although significant improvements in the LoD by SHG have been reported by image analysis approaches (Chowdhury et al., 2016; Schmitt et al., 2015; Toth et al., 2012), the integrated intensity provided the most direct side-by-side comparison with the TL results. The LoD for similarly prepared samples was found to be 17 ppm by SHG microscopy (compare to 140 ppm via TL measurements), which is roughly comparable to the LoD reported in previous measurements of the integrated SHG intensity of powdered blends. (Chowdhury et al., 2016; Kestur et al., 2012; Kissick et al., 2011; Sweeting, 2001; Wanapun et al., 2010) As in the case of TL, the low LoD of SHG in powdered samples arose primarily by the absence of significant background in the API-free blank samples. The higher limits of detection from TL are attributed to the presence of some weak TL emission from the pure polymer samples.

In addition to the size-dependent studies, both TL and SHG microscopy measurements were performed on powders of two polymorphs of clopidogrel bisulphate with different anticipated piezoelectric properties to assess the role of crystal form. Form I and Form II of clopidogrel bisulphate have identical chemical composition, differing only in the packing arrangements within the lattice. The lower symmetry ($P2_1$) Form I particles produce significantly brighter SHG than the Form II ($P2_12_12_1$) particles, generally consistent with expectations from previous quantum chemical calculations, which suggest brighter SHG with lower symmetry. (Hardy et al., 1981) Because of the close symmetry similarities between SHG and piezoelectricity, it is reasonable to expect similar trends to arise in TL produced from the piezoelectric mechanism. In **Figure 2.9**, a comparison of the SHG and TL activities from the two powders differing only in crystal forms indicates close correlation between the SHG and TL, consistent with this expectation. The Form II and PEG signals were multiplied by 100 to show the difference in intensity between Form II and PEG. Overall, in combination with the size-dependent results, these measurements collectively suggest that the detected TL is dominated by piezoelectric contributions, which also explains the origins of the specificity to the chiral crystalline content.

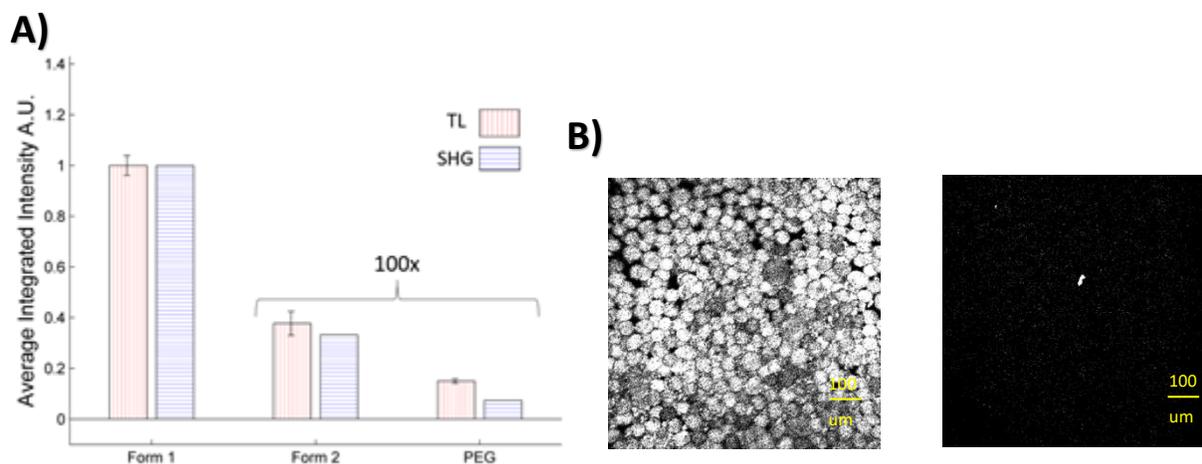


Figure 2.9: A) Shows the SHG (blue, horizontal lines) and TL (red, vertical lines) signals from two polymorphic forms of clopidogrel bisulphate which are indistinguishable by bright field imaging. Both data sets were normalized to Form I. Both the Form II and the PEG data were rescaled by 100 to allow meaningful comparisons with the Form I result. B) Two representative images from SHG microscopy to demonstrate the difference in image contrast. The SHG results for Form II are rescaled by 5, showing only diffuse, low signal to noise, SHG signal. The bright punctum in the Form II image is attributed to an impurity and was excluded in the analysis.

A third experiment was designed to test the inherent response time of the light. Using NMF-F (see Chapter 3) the characteristic time response of TL was found to be in the same range as a dielectric breakdown from lightning. A representative time-trace for a TL measurement is shown in **Figure 2.10**. Several key aspects of the measurement are noteworthy. First, at least three separate burst-events are detectable from the time-trace, the first of which at ~ 0.1 ms being the largest. Second, the intentional impedance mismatch introduces a relatively long exponential decay on the detected signals arising from burst events. Thirdly and less obviously at this scale, the time-trace includes high-frequency spikes arising from the absorption of individual photons. While the impedance mismatch was advantageous for detection of the individual transients, the resulting temporal convolution substantially complicated efforts to characterize the intrinsic burst profile of individual TL events and inform models for describing the mechanisms of TL production. Next, the NMF-F was applied to the measured triboluminescence signal shown in **Figure 2.10**.

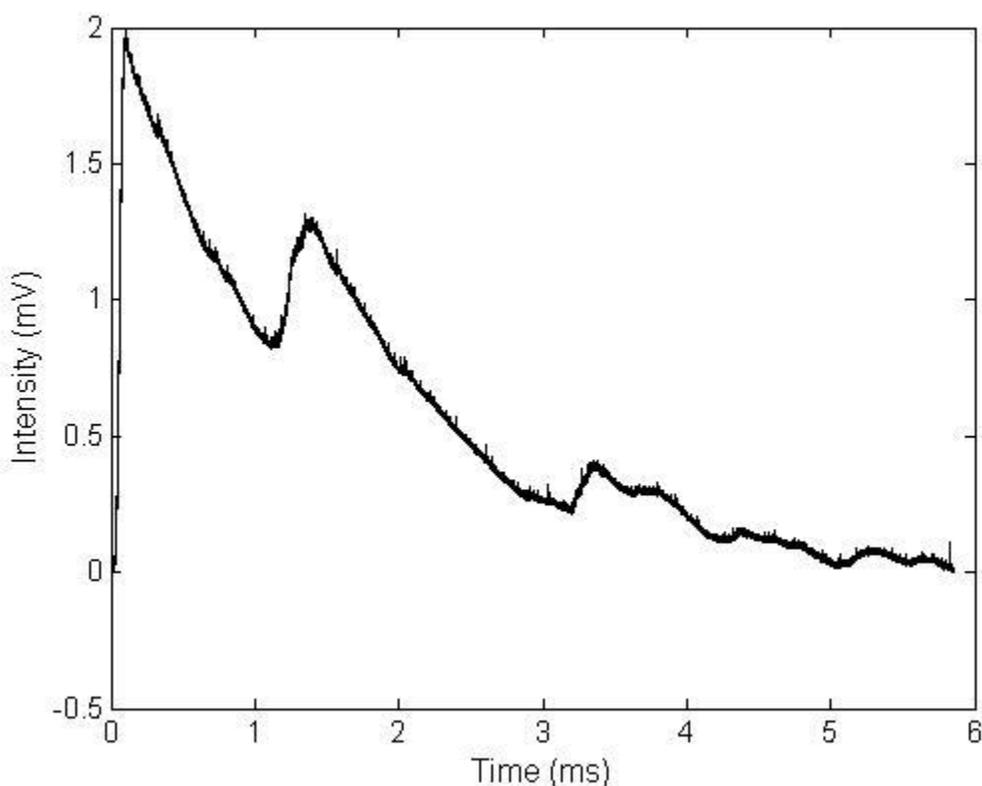


Figure 2.10: Triboluminescence signal from Griseofulvin. Used as an example to demonstrate deconvolution using the non-negative matrix factorization.

The results of the deconvolution are shown in **Figure 2.11**, in which the photon arrival times are compared with the original data. The expanded portion of **Figure 2.11** allows identification of individual photon-events at the beginning of a burst generated from a triboluminescence event. From inspection of the results, the deconvolution algorithm notably increases the localization of photon events and suppresses noise from convolution with the IRF, from the baseline offset, and from the white electronic noise. Following deconvolution, autocorrelation analysis of the measured TL time-traces was performed to yield information on the characteristic time-scale for the individual burst events, and in turn insights into the mechanisms of TL in pharmaceutical powders. The results of the autocorrelation analysis are shown in **Figure 2.12**, revealing a characteristic timescale for the burst events on the order of $\sim 50 \mu\text{s}$.

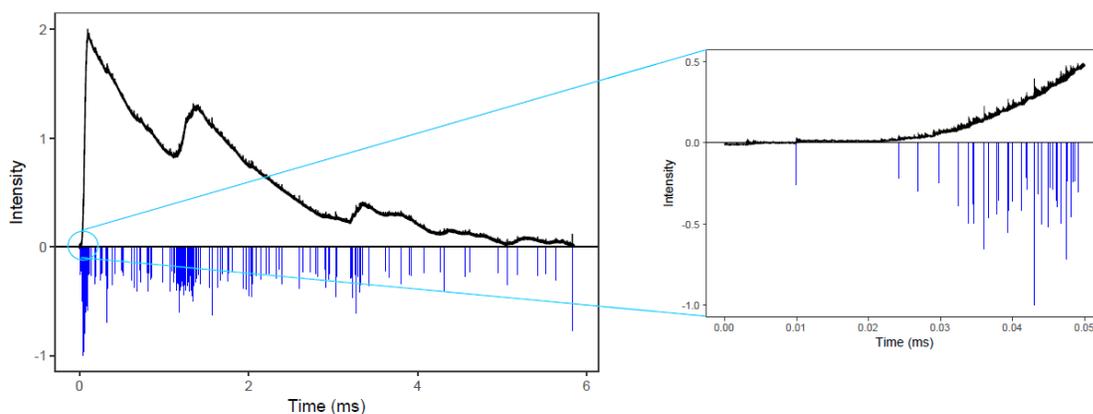


Figure 2.11: Showing the results of deconvolution. Blue lines are arrival times of photons. Zoom in on the right showing the density of photon arrival times at the first peak of the TL data.

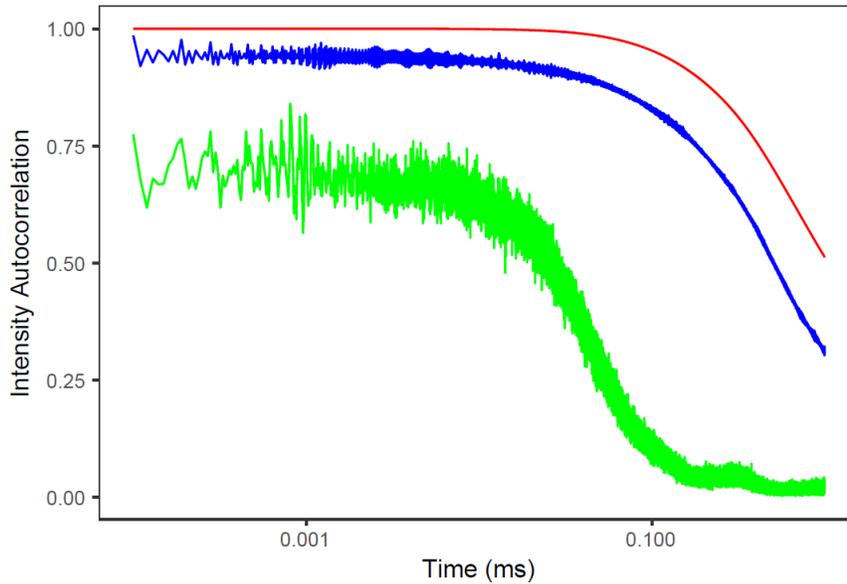


Figure 2.12: Results of autocorrelation on the raw TL data (red-top), the measured IRF (blue-middle), and the deconvolved photon events (green-bottom). The x-axis is on a log scale.

Inspection of the deconvolved autocorrelogram, see **Figure 2.12**, for TL bursts provides several insights into the possible mechanisms driving the observed signal production. Notably, longer timescale recurrences in the autocorrelogram that might arise from ringing events were not observed, consistent with TL arising stochastically in $\sim 50 \mu\text{s}$ bursts. High-speed imaging was performed to assess the piston speed at the point of impact, and from that the timescale for traversal through the sample. The maximum piston speed was $\sim 2 \text{ m/s}$ corresponding to the piston taking $\sim 0.5 \text{ ms}$ to travel through the 1 mm thick sample. This timescale is significantly longer than the individual burst events from autocorrelation analysis, but comparable to the entire envelope of bursts produced by a single impact event. The $50 \mu\text{s}$ time scale is in qualitative agreement with the $\sim 20 \mu\text{s}$ temporal envelopes reported in laser induced breakdown spectroscopy experiments, where emission of excited gas-phase AIO caused by a spark created by a laser pulse was measured over time in different bath gases. (Piehler et al.) Consistent with this collective set of results, we can propose the following mechanism for the observed TL. During the $\sim 500 \mu\text{s}$ traversal of the piston through the powdered sample, individual crystalline particles independently produced TL. The temporal profiles of these individual TL bursts were consistent with localized plasma production from dielectric breakdown, presumably from rapid changes in capacitance upon lattice collapse.

In some TL events, multiple bursts were observed, consistent with the collapse of several different independent particles during the ~ 500 μs impact event. The routine observation of the most pronounced TL burst amplitudes relatively early in the impact window is consistent with the larger particles collapsing at earlier times, subsequently followed by bursts associated with smaller crystals as shown by the secondary peaks in **Figure 2.10**.

If individual TL bursts arise from the stochastic collapse of individual particles as suggested by these measurements, a TL mechanism based on capacitive discharge should enable quantitative connections to be drawn between the crystal size and the amplitude of the individual TL burst events. Such connections would raise the intriguing possibility of probing the crystal size distribution from the histogram of intensities for the individual TL bursts. This histogram would be dependent on the size distribution of individual crystals, which is distinct and complementary to the conglomerate particle size distribution most easily accessible from alternative analysis approaches.

2.7 Applications of TL to Pharmaceutical Amorphous Solid Dispersions

In this work, impact driven triboluminescence (TL) enables rapid assessment of crystalline content. This shows application to pharmaceutical powders. Because TL has been shown to be selective to noncentrosymmetric crystalline materials, TL will have applications in the pharmaceutical field. Since most active pharmaceutical ingredients are relatively complex organic molecules, a high percentage of them will crystallize into space groups that are TL active. A survey of Merck APIs showed that roughly 60% should be TL active. In the present study, the presence of TL is used as a sensitive and selective indicator for the presence of residual noncentrosymmetric crystals in model ASDs comprised of trace crystalline active pharmaceutical ingredients (APIs) cast within an amorphous polymer matrix. Given the close similarities in the symmetry properties of both piezoelectricity and SHG, TL shows promise as an initial on-line measurement for the detection of crystallinity with pharmaceutical formulations, with additional studies being performed by SHG and other complementary methods.

An amorphous griseofulvin sample in HPMCAS-MF was prepared by collaborators at Merck and confirmed to be amorphous using PXRD following preparation. See **Figure 2.13** for PXRD data. The lack of sharp peaks and large rolling features is indicative of an amorphous sample.

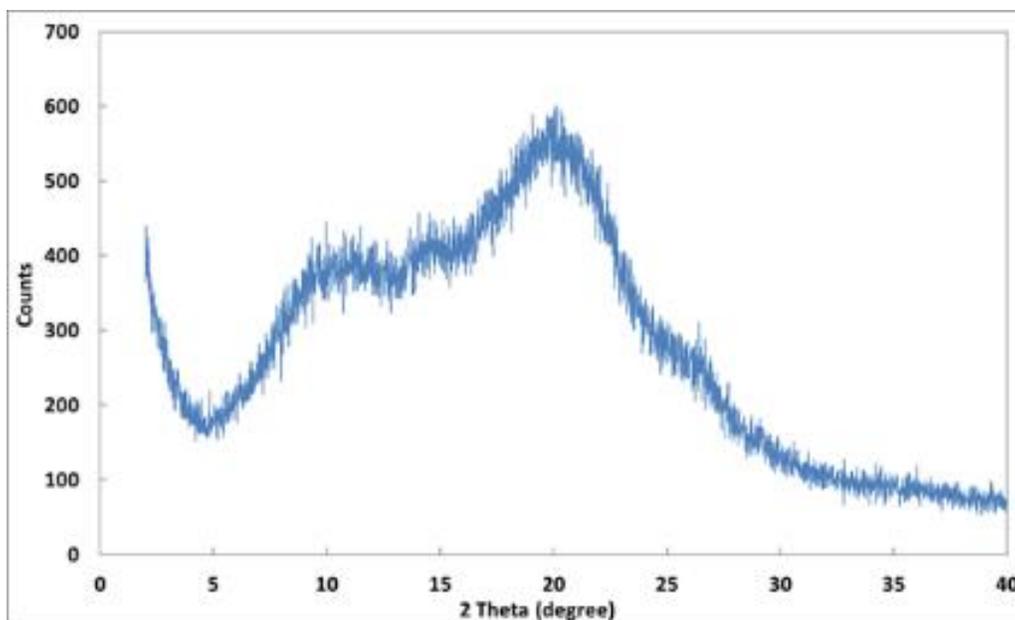


Figure 2.13: Powder X-Ray Diffraction pattern of Griseofulvin HPMCAS-MF amorphous solid dispersion immediately following preparation.

In addition to PXRD, the ASD was found to be amorphous via SHG and TL. To stress the sample and induce crystallinity, the spray dried ASD was heated at 65° C and constant humidity (via a saturated NaCl solution) for 48 hours at which point SHG and TL measurements were performed on the material, see **Figure 2.14** for TL & SHG results following the stressing. SHG image analysis allowed for quantitation of the crystalline content, which was determined to be 0.11%. TL analysis showed detectable signal with large shot-to-shot variability.

From these combined results, TL may be well-suited for rapid screening during on-line reaction monitoring of amorphous solid dispersions. In practice, ASD's are designed to contain only amorphous content, such that any residual trace crystallinity is generally considered disadvantageous. TL is able to detect low levels of crystallinity (140 ppm demonstrated in this work), making it appropriate for detecting crystallinity in nominally amorphous materials. While the large measurement variance from sample stochastics complicates quantitative powders analysis by TL as conceived herein, the high selectivity still suggests a role for outlier detection. In such cases, rapid measurements of TL signals greater than the background may serve as an initial indicator for subsequent more detailed analysis (e.g., by SHG microscopy or spectrochemical characterization).

Triboluminescence has been shown to enable detection of crystallinity in amorphous solid dispersions with detection limits of 140 ppm. The speed of the measurements and simplicity of the instrument make it applicable to process analytical technologies. The crystal size and polymorphic form dependence of TL parallel SHG measurements and are both consistent with a mechanism dominated by the piezoelectric properties of the crystallites within the powder. The symmetry demands of piezoelectricity also underpin the origin of the selectivity of TL to the crystalline API fraction. The capabilities of TL suggest applications for outlier detection during process monitoring of amorphous solid dispersions. Studies are currently underway to characterize a diverse portfolio of different API and excipient materials to better assess the breadth applications that are potentially compatible with TL analyses. Studies are currently underway to characterize a diverse portfolio of different API and excipient materials to better assess the breadth of applications that are potentially compatible with TL analyses.

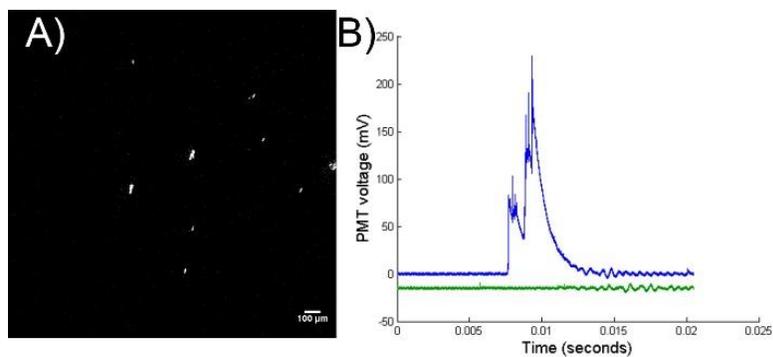


Figure 2.14: A) Shows the crystallinity content of the griseofulvin and HPMCAS-MF ASD after being exposed to heat and constant humidity for 48 hours. B) The upper trace (blue) shows the TL signal from the stressed sample while the lower trace (green) shows the signal from the unstressed sample, offset for clarity.

Complementary second harmonic generation (SHG) microscopy results were performed using a SONICC microscope from Formulatrix (Bedford MA), modified in-house for powders analysis. Laser power was set to 150 mW at the sample (170 fs, 50 MHz) and the signal was integrated for 500 ms over 512×512 pixels in a 2.0 mm² field of view. Modifications to the instrument as described previously allowed SHG to be collected back through the same objective used for delivery of the fundamental (i.e., in the epi direction) to facilitate measurements of powders. Recently, second harmonic generation (SHG) has been shown to rapidly quantify trace

crystallinity within ASDs.(Schmitt et al., 2015) The symmetry requirements of SHG, in which coherent signal is only generated from noncentrosymmetric crystals, allows for a selective measurement with detection limits as low as parts per trillion under favorable conditions and routinely achieves ppm limits in powders.(Wanapun et al., 2011) Trace crystallinity measurements that take advantage of particle counting techniques enable quantification of trace crystallinity content with high statistical certainty.(Savitzky & Golay, 1964) SHG has been shown to exhibit limits of detection greater than 3 orders of magnitude lower than conventional benchtop methods. Advances in the advent of ultrafast fiber lasers have further improved access to affordable turn-key laser sources for SHG microscopy.(Becker et al., 2004; Kissick et al., 2010; Piehler et al., 2005) However, the expense and complexity of ultrafast optical methods complicates applications of SHG microscopy for routine and rapid on-line and at-line reaction monitoring. In addition, optical analysis of powders is often limited by a penetration depth on the order of the mean particle size,(Savitzky & Golay, 1964) potentially introducing sample bias if the crystallinity is heterogeneously distributed throughout the powder. Alternative methods based on simpler and more robust measurement platforms would facilitate on-line reaction monitoring and process analytical technologies (PAT).

2.8 Triboluminescence of Slurries

ASDs are prepared by casting the amorphous API in a polymer matrix such as polyethylene glycol (PEG), polyvinylpyrrolidone (PVP), or hydroxypropylmethylcellulose (HPMC). ASDs offer stability to the amorphous powder by kinetically preventing recrystallization.³ ASDs are rapidly gaining acceptance and use within the pharmaceutical industry because of their accelerated dissolution rates and higher apparent solubilities. These benefits are achieved by maintaining a supersaturation state in the gut. However, ASD formulations are still not in their thermodynamically most stable state, and consequently, even trace residual crystallinity can negatively impact bioavailability and shelf-life by removing nucleation barriers. This may lead to a shorter than intended time of supersaturation within the GI tract, or completely prohibit the ASD from achieving supersaturation conditions. The analysis of the impacts of trace crystallinity of ASDs during dissolution and uptake in the body remains a large concern of pharmaceutical companies. Instruments with the ability to rapidly detect trace crystallinity during dissolution can help researchers make well informed decisions earlier in the formulation process.

Dissolution testing commonly uses a spectrochemical method to measure the concentration of the API in solution. This method is lacking in that it is limited to solutions that do not accurately represent the GI tract (e.g. the solution must be optically transparent) and that the crystallization kinetics (if any) are not readily observed (Zimper et al., 2010).

Second harmonic generation (SHG) can be a powerful tool for quantifying trace crystallinity within ASDs.¹² SHG is the frequency doubling of the incident light and it is a nonlinear process. The symmetry requirements of SHG, in which coherent signal is only generated from noncentrosymmetric crystals (that is crystals lacking an inversion center), allows for a selective measurement with detection limits as low as parts per trillion under favorable conditions.¹³ Additionally, it has been shown that SHG can measure crystallinity content even in a turbid media that mimics the GI tract. (Elkhabaz et al., 2018; Ulcickas & Simpson, 2020)

Even though SHG offers an attractive solution to many of the problems of the pharmaceutical industry, the relative delicacy of optical methods complicates applications of SHG microscopy for routine and rapid on-line and at-line reaction monitoring. Alternative methods based on simpler and more robust measurement platforms would facilitate on-line reaction monitoring and process analytical technologies. In addition to the SHG measurements, we present an instrument based on impact-driven triboluminescence (TL) to rapidly access the quantitative presence of crystalline content within biorelevant media. TL is the emission of light upon mechanical disruption. The piezoelectric effect greatly enhances TL activity in crystals that are symmetry-allowed for piezoelectricity.¹⁴ Piezoelectricity describes the generation of an electric field resulting from the physical distortion of a material, or conversely the presence of a material displacement when in the presence of an electric field. The presence of TL can be used as a sensitive and selective indicator for the presence of residual non-centrosymmetric crystals of model ASDs comprised of trace crystalline APIs cast within an amorphous polymer matrix. Given the close similarities in the symmetry properties of both piezoelectricity and SHG, TL is positioned to be a quick qualitative indicator of the presence of crystallinity within turbid media during dissolution testing.

As a comprehensive strategy to assist in the quantification of trace crystallinity within ASDs during dissolution testing. This two-pronged approach which uses SHG and TL to gain an understanding of the crystallization kinetics and effects of trace crystallinity on dissolution offers an attractive solution for researchers looking to optimize their pharmaceutical formulations.

Previous studies that have used SHG as a method for detecting crystallinity within biorelevant media simply sampled and measured crystallinity at set time points (Elkhabaz et al., 2018). In order to get a better sense of the complex kinetics during dissolution testing, a SHG & TL compatible flow cell was constructed. The SHG flow cell was made from stainless steel and contained an optical window in both the top and bottom made from clear plastic. The top and bottom windows were made with appropriate o-rings such that the cell is watertight. The flow cell contained grooves to fit an o-ring at the intersection of the top and bottom plates to ensure a liquid tight seal. The groove was designed to be the same area as the pipe used to deliver the solution to and from the flow cell in order to avoid any bottleneck in the design. See **Figure 2.15** for a diagram of the SHG flow cell.

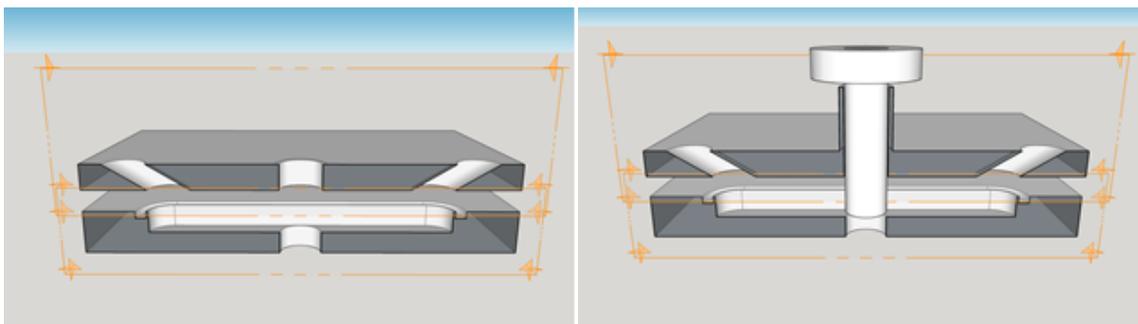


Figure 2.15: A schematic cutaway of the SHG flow cell (left) and the TL flow cell (right). The holes in the center provide an optical window to take SHG images or to record emitted light from TL.

A similar strategy was undertaken for construction of the TL flow cell the exception being the top plate has a piston that is able to impart a force to the particles within the impact area. Appropriate measures were taken to ensure that the flow cell remained watertight during TL impact events. The durability of the TL flow cell was accessed by doing repeated impact events until failure. After ~12k hits, the o-ring on the piston failed, showing that the instrument is significantly robust for routine measurements. The flow of liquid was driven using a peristaltic pump, which causes some SHG images to be recorded with blurred particles.

The electronics to drive the impactor was constructed with the help of the AMY facility. TL was induced in the slurry samples using mechanical force delivered from a solenoid as detailed previously. Complimentary SHG microscopy results were performed using a modified SHG

SONICC microscope from Formulatrix (Bedford, MA). Further modifications to the system were performed for this study. The original system is set up to take images of a $\sim 2\text{mm}^2$ area, for greater speed and simplicity to the measurement, the system was modified to just perform a line scan using the fast axis resonant mirror scanner (8 kHz). Data was collected using the line scanner and assembled (line by line) into 512 x 512 images for image analysis. Image analysis & particle counting was performed using the help of the ImageJ software (available for free on NIH.gov). Griseofulvin for this experiment was purchased from Sigma-Aldrich. All water used was nano pure water obtained from an in-house filtration system.

Calibration of the TL slurry instrument was performed to provide an estimate of the limits of detection, shown in Figure 3. All calibration measurements shown were acquired for crystalline griseofulvin in water, chosen as a simple model for a dissolution test of an amorphous solid dispersion. Based on the observations in the figure, lower limits of detection (LoD) were found to be 140 ppm ($3 \cdot \text{sb} - \text{criterion}$), with the low values arising primarily from the noise of the PMT. The TL measurements have shown that it is sensitive to low levels of crystallinity content. The main challenge with the TL measurement is that there are many impacts where there are no crystals to crush due to the stochastic nature of the measurement. Instead of measuring the integrated signal from TL as has been done in previous TL studies, the authors have opted to measure the number of impacts with crystallinity present, as shown in **Figure 2.17**. As you can see, as the concentration of crystalline griseofulvin increases in the solutions, there is a corresponding number of impacts with a positive signal above a set threshold. Initial data shown here has not yet been optimized, such that better limits of detection can be expected.

The results from TL detection are compared with measurements probed by SHG microscopy in **Figure 2.16**. The SHG results were produced by image analysis of multiple SHG images each with a field of view of $\sim 2\text{mm}^2$. Examples of SHG images in the flow cell can be found in Figure 3. The initial SHG measurements were taken in transmission imaging mode on the Formulatrix SONICC instrument. Future measurements will be performed in a line scanning configuration. The advantage of going to a line scan configuration is that it is more conducive to particle counting experiments. We have previously shown that particle counting experiments are better for low levels of crystallinity and give better signal to noise than signal integration. Information that will be extracted from the line scan experiment will be the particle width (most APIs crystallize in spherical shapes and even those that don't still look spherical when first

crystallizing) and the number of particles per unit of time. The flow rate will be known a priori such that the percent crystallinity per unit of volume can be determined. This will allow for quantification of the rates of crystallization and more precise knowledge of the kinetics of crystallization of ASDs within biomimetic media. An additional benefit of doing a line scan over a full image of the field of view, is that since the particle field is static, particles could be missed because of the movement of the slow axis. With the line scanning, the laser beam will be repeatedly moved across the line at 8 kHz, which is sufficiently fast to feel confident that no particles are slipping past without being detected.

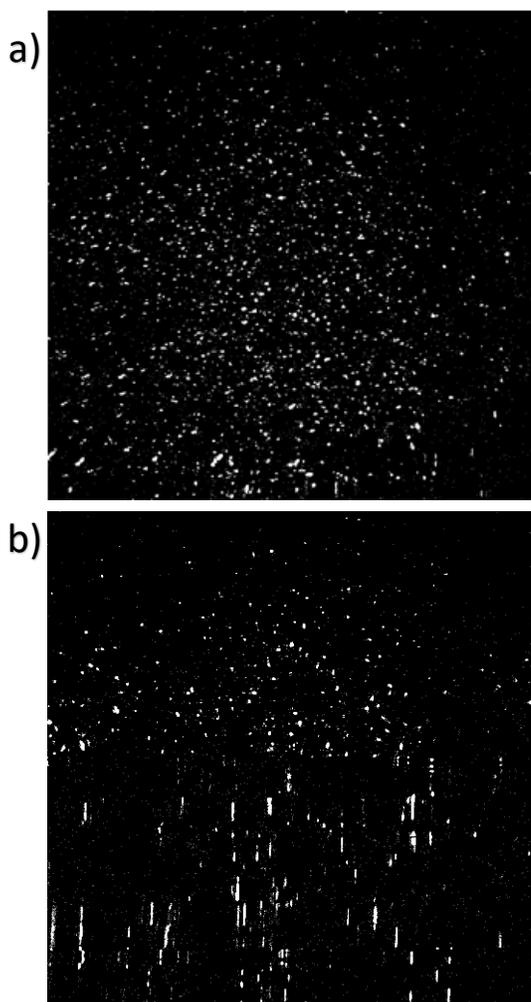


Figure 2.16: A) SHG image of the flow cell, image taken while pump is off. B) SHG image showing blurring particles due to flow of liquid during imaging.

From these combined results, TL may be well-suited for rapid screening during dissolution testing to determine whether an ASD will fail or pass. While the large measurement variance from sample stochastics complicates quantitative analysis by TL as conceived herein, the high selectivity still suggests a role for outlier detection. Whereas the SHG measurements could be used for in depth analysis of the crystallization kinetics to further understanding of supersaturation within the GI tract and how trace crystallinity affects the supersaturation.

The similarities in symmetry properties between piezoelectricity and SHG suggest compatibility of TL for detection of trace residual crystallinity in homochiral crystals. Since the molecular building blocks of homochiral crystals are inherently noncentrosymmetric, the space groups in which homochiral molecules crystallize must necessarily also be limited to the noncentrosymmetric subset. Most of these crystal forms are symmetry-allowed for both SHG and piezoelectricity. It is estimated that 80% (Babu & Nangia, 2011) of new small molecule active pharmaceutical ingredients are homochiral, and therefore compatible with both TL and SHG analyses.

In addition to the work done on TL & SHG of slurries, follow up work on the original TL project was performed. That is summed up in the following two paragraphs. Triboluminescence was found to work well on model amorphous solid dispersions; being able to accurately recover the percent crystallinity and provide a limit of detection 2 orders of magnitude lower than other instruments commonly used to assess crystallinity content in powders. An ASD of griseofulvin in HPMCAS (hydroxypropylmethyl cellulose accinate succinate) was prepared by our collaborators at Merck & Co. The sample was evaluated with both TL, SHG, and PXRD which confirmed that there was no crystallinity in the ASD. The sample was then stressed under elevated relative humidity (RH) and temperature. The RH was increased using a saturated solution of NaCl. The temperature was elevated and held at 65°C for 48hr at which point SHG and TL measurements were performed on the material. SHG image analysis allowed for quantitation of the crystalline content, which was determined to be 0.11%. The trace crystallinity was proven to be detectable by TL analysis although with relatively large shot-to-shot variability.

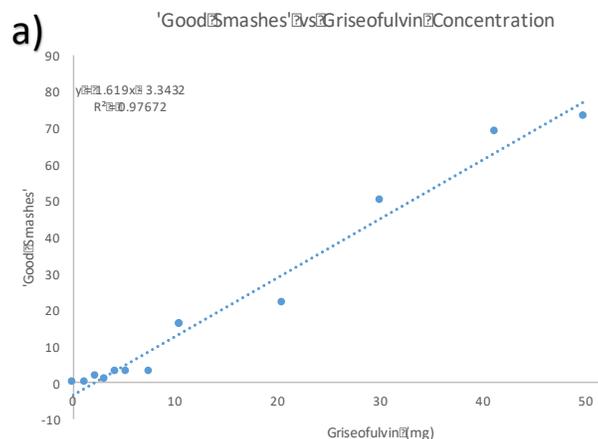


Figure 2.17: Number of impact events with signal above the noise threshold with respect to concentration of the solution. Solution contains griseofulvin in water.

Analysis of crystallization kinetics with both Triboluminescence and SHG have been shown to be great improvements over current methods of analysis for dissolution testing. Both TL and SHG have lower limits of detection than commonly used methods, and they offer increased speed of the measurements. The increased speed offers more information on the underlying kinetics of supersaturation and crystallization within biomimetic media (which is relatively hard to measure using traditional methods). The speed of the measurements and simplicity of the instrument make it applicable to process analytical technologies. The crystal size and polymorphic form dependence of TL parallel SHG measurements, both consistent with a mechanism dominated by the piezoelectric properties of the sample. The symmetry demands of piezoelectricity also underpin the origin of the selectivity of TL to the crystalline API fraction. The capabilities of TL suggest applications for outlier detection during process monitoring of amorphous solid dispersions.

2.9 Other Applications of Triboluminescence

In collaboration with the Cooks lab, TL was briefly explored as a novel ionization method. The rationale being that if TL really was going through the piezoelectric mechanism, the energy required to produce dielectric breakdown is also enough to generate ions. The TL instrument was brought to a mass spec with a “sniffer” attached. The heated inlet was brought near the TL smashing area. MS peaks were recorded following the solenoid impact. However, similar MS

peaks were observed when powdered sample was manually placed into the MS inlet. It was concluded that the hammer was likely hitting hard enough to force the sample into the air and into the inlet. However, further studies into this area may prove promising.

2.10 Acknowledgements

I would like to acknowledge the help of the group members. In particular, Scott Griffin, Julia White, Noah Hudacek, Sidnie Bienz, Joshua Wagner, and our collaborators at Merck. I would also like to thank the Jonathan Amy Facility for Chemical Instrumentation for their extensive help on design and fabrication of the instrument.

CHAPTER 3. ITERATIVE NON-NEGATIVE MATRIX FACTORIZATION

3.1 Introduction

Deconvolution of complex waveforms is a common problem in measurement science of digital signals, with applications spanning super-resolution microscopy, fluorescence lifetime analysis, spectroscopy, and image generation from single-photon/particle counting measurements. (Smith et al.) In these applications, deconvolution substantially improves resolution in space, time, and/or frequency. As such, a diverse portfolio of algorithms exists to enable routine deconvolution of digital signals. One common strategy is to perform a digital convolution with a modified form of an inverse function to undo the impact of the original instrumental convolution. The Wiener filter is a common example, in which the inverse filter is weighted by the frequency-dependent signal to noise in the Fourier domain. (Taylor et al.) Richardson-Lucy deconvolution is another common iterative method for deblurring an image from a known point spread function. (Richardson) Other techniques used for deconvolution include system identification (Vajda et al., 1988), constrained optimization (Madden et al., 1996), the use of cubic spline basis functions (Verotta, 1993), maximum entropy (Skilling & Bryan, 1984), and genetic algorithms (Coello). More recently, digital filters derived from linear discriminant analysis (LDA) have been shown to enable digital deconvolution without explicit need for independently assessing the noise characteristics of the instrument. (Sullivan et al.) While generally successful, this previous suite of techniques is collectively only applicable for systems in which the impulse response function (IRF) of the instrument is known *a priori*. Unfortunately, this foreknowledge is not always trivial to obtain. Independent measurements from model systems of single photons/particles/wavelengths may not be available or easily accessible. Furthermore, drift in instrumental parameters can result in time-dependent changes in the IRF between the calibration and measurement steps, potentially invalidating deconvolution approaches requiring foreknowledge of the IRF.

To address this limitation, there is a clear need in developing blind deconvolution strategies, in which the measurements themselves provide sufficient information to support recovery of both the IRF and the deconvolved output in a single integrated analysis. Blind Richardson-Lucy (BRL) deconvolution, commonly used in image analysis, assumes a general functional form of the convolution kernel (i.e., point spread function) and relies on an iterative maximum-likelihood

approach for deconvolution.(Fish et al.) Alternatively, a method for using least-squares deconvolution with Laguerre expansions can be applied to deconvolution of exponentially decaying functions by interpreting the time-dependent evolution as a linear combination of Laguerre polynomial functions.(Liu et al.) While Laguerre expansion method is well suited to fluorescence lifetime recovery by nature of the exponential components within the Laguerre polynomials, both Laguerre and Richardson-Lucy approaches are still relatively sensitive to the particular selection of inputs used to form the basis set for the deconvolution. As such, significant assumptions regarding the general functional form for the IRF are typically required prior to implementation, limiting extension to broader classes of unsupervised deconvolution challenges in which substantial a priori information of the IRF is not available. Furthermore, the degree to which the results accurately converge to recover the underlying IRF and deconvolved data hinge on the similarity between the IRF and the generating basis set functions.

The challenges and potential benefits associated with blind deconvolution are particularly prominent in single particle measurements with complex IRFs, such as often arise in ion and photon counting. The interplay between optics and electronics can result in nontrivial IRFs that are difficult to predict a priori, and correspondingly difficult to universally describe by the approaches in the preceding paragraph. However, the rewards for success in deconvolution are numerous. Photon counting (PC) is well established as a successful strategy for sensitive detection of low intensity light and is widely used in applications ranging from time resolved fluorescence measurements to quantum computing at wavelengths spanning the infrared through the X-ray (Natarajan et al., 2012). PC typically operates in Geiger mode, in which a photon-induced voltage transient exceeds some critical threshold and registers as a digital count. The principle advantage of photon counting is the elimination of electronic white noise and noise contributions from the shot to shot variance in the detector gain. However, PC generally exhibits a limited dynamic range stemming from pulse-pile-up, in which the timeframe for adjacent photon events arises fast relative to the duration of the single-photon IRF. Conceptually, this dynamic range limitation arises from the convolution of the instantaneous photon arrival event with the IRF. (Muir et al.) As a result, the nature of the IRF and the ability to successfully perform deconvolution directly impact the temporal resolution, sensitivity, and dynamic range of photon counting.

Complications from a complex IRF in photon counting are clearly illustrated in a recent work based on triboluminescence (TL) detection of crystalline drugs within nominally amorphous

pharmaceutical materials, such as amorphous solid dispersions (ASDs)(Casey J. Smith et al.). TL describes the emission of light during mechanical action on the crystalline fraction of an ASD. ASDs are samples of interest, as they are most commonly used as pharmaceutical formulations where the active pharmaceutical ingredient (API) is kinetically trapped in an amorphous state by a polymer matrix, most commonly via hot melt extrusion or spray drying. Residual crystallinity has the potential to decrease the stability of ASDs. Therefore, a sensitive, crystal selective method, such as TL, is desirable for early detection of drug crystallization to identify quality concerns during the early stage of formulation exploration. While previous studies have demonstrated the utility of TL in this application-space, fundamental understanding of the origin of the TL signal production will provide a foundation for assessing the potential broader scope of use.(Casey J. Smith et al., 2018 Chang2016) Deconvolution of the TL signal will allow for a greater fundamental understanding of the mechanism underpinning TL measurements.

The temporal duration of these bursts complicates studies to elucidate the fundamental dominant mechanisms driving signal production; the detection electronics must be slow enough to capture the entire burst envelope, but fast enough to enable meaningful recovery of the characteristic timescales of the individual single crystal TL bursts. Two likely mechanisms were proposed for TL production from organic crystals; i) TL production from piezoelectric activity within the crystal, resulting in dielectric breakdown and plasma luminescence, and ii) solid-state electron/hole recombination resulting in light emission. In principle, the characteristic time frame for light emission should be quite different for these two mechanisms, providing a means for discriminating between them. In practice, this analysis is complicated by practical considerations associated with experimental constraints. A single TL burst envelope often contains emission of hundreds of photons generated by many crystals over a few microseconds to milliseconds. Detection of single-photon events arising within a burst envelope was promoted through intentional impedance mismatch, in which the terminal impedance at the oscilloscope card was set to a value higher than the impedance of the transmission line. While this approach greatly improves the signal to noise for quantification courtesy of Ohm's Law and extends the temporal envelope of single photon events for ease in digital detection, it also resulted in a complex IRF containing substantial ringing from electronic reflections arising from the mismatched impedance.

In the work described herein, a digital filtering approach integrating non-negative matrix factorization (NMF) for blind deconvolution is developed theoretically and subsequently applied

to deconvolve the time-trace of TL events. The use of NMF exploits the fundamental physics associated with the signal production. The number of photons arriving at a given position in time must strictly be a positive integer. Similarly, photon absorption at the detector results in a cascade of electrons producing dominant voltage transients identical in sign each time. This strategy builds on an alternating least-squares formulation of NMF, in which conventionally a data matrix V is expressed as the product of two matrices WH containing only non-negative entries. Using an initial guess for one matrix (W) allows identification of the entries of the second (H), which minimize the sum of squared deviations between the data and the matrix product. Using only the non-negative entries for H , the process is repeated to solve for the entries of W). In this manner, H and W are solved iteratively until non-negativity in each is achieved. In order to perform deconvolution of single-photon events from a time-trace, this general strategy of NMF was modified in two notable ways: i) the NMF mathematical framework was adapted to explicitly incorporate convolution and deconvolution operations, and ii) NMF analysis was performed over only small sections of data as a rolling filter to greatly reduce computation cost in the analysis of large data sets. This non-negative matrix factorization filter (NMF-F) approach was evaluated on test data with known ground-truth results, then implemented to provide insights into the fundamental origins of the light-emitting processes underpinning TL measurements of the model pharmaceutical materials doxycycline hyclate, an antibiotic of the tetracycline class, and griseofulvin, an anti-fungal agent.

3.2 Methodology

The NMF algorithm and related multivariate curve resolution methods (de Juan et al.) are founded on both the alternating least squares approach originally proposed by Paatero and Tapper (1994) (Paatero & Tapper) and an adaptation of a convolution model originally proposed by Chambers *et al.* The objective function in this case corresponds to a maximum likelihood estimate (MLE) based on minimizing the sum of squared deviations between the measured time-trace and the recovered fit. Implicit in MLE is the assumption of normally distributed, uncorrelated instrument noise, consistent with the experimental expectations of our measurements. In the present case, the intensity of the detected light must be a positive number, serving to constrain a fit to recover the intensity.

In convolutional NMF applied to a time-trace, the detected voltage is represented by the vector $\bar{\mathbf{d}}_{meas} (m \times 1)$, which is the convolution of the ground truth impulse train, $\bar{\mathbf{d}}_{true} (m \times 1)$, with the $\mathbf{irf} (n \times 1)$ and addition of noise. The convolution operation was recast as a product with an $m \times m$ matrix \mathbf{P} by convolving the IRF with an identity matrix. Additional noise is indicated by the vector $\bar{\mathbf{e}}$, m is the number of digitization events, and n is the size of the \mathbf{irf} . It may appear counter-intuitive to expand the initial $(n \times 1)\mathbf{irf}$ vector to a $(m \times m)$ matrix $\mathbf{P} = \mathbf{irf} \otimes \mathbf{I}_m$ through a convolution with an $(m \times m)$ identity matrix, considering NMF typically is used to reduce the inherent dimensionality of the data set. However, this expansion was performed for mathematical convenience in casting the convolution operation into a form directly compatible with matrix evaluation and inversion.

$$\bar{\mathbf{d}}_{meas} = \mathbf{irf} \otimes \bar{\mathbf{d}}_{true} + \bar{\mathbf{e}} = \mathbf{irf} \otimes (\mathbf{I}_m \bar{\mathbf{d}}_{true}) + \bar{\mathbf{e}} = \mathbf{P} \bar{\mathbf{d}}_{true} + \bar{\mathbf{e}} \quad (\text{Eq 3.1})$$

If the IRF were known *a priori*, the data trace $\bar{\mathbf{d}}_{meas}$ could be fit assuming (incorrectly) a photon arrival at each position in time through Equation 2. Such a fit recovers the amplitude of the impulse at each position in $\bar{\mathbf{d}}_{fit}$, which will initially exactly recover the time-trace $\bar{\mathbf{d}}_{meas}$ including errors from noise as the number of equations and parameters are exactly matched. Assuming the rate of photon arrivals is significantly less than the digitization rate, roughly half of the amplitudes of the initial fit will be positive and half negative. All amplitudes less than zero (or more generally, less than a prespecified threshold) can be considered nonphysical and removed (by removing a column from the fitting parameters in \mathbf{P} in the evaluation by the nonnegativity criterion, as illustrated in **Figure 3.1**. This process can be repeated until only positive (above threshold) amplitudes are recovered.

$$\bar{\mathbf{d}}_{fit} = (\mathbf{P}^T \mathbf{P})^{-1} \mathbf{P}^T \bar{\mathbf{d}}_{meas} \quad (\text{Eq 3.2})$$

In practice, this approach requires foreknowledge of the IRF. If, instead, we could somehow know *a priori* the positions of all the impulses corresponding to photon arrivals given by $\bar{\mathbf{d}}_{fit}$ but not the IRF, the problem could be inverted to solve instead for the IRF. This process is illustrated in Equation 3. In the absence of measurement noise, the measured data trace can be recovered by the convolution (illustrated by the tensor product \otimes) of the impulse train with the IRF. Introducing an identity matrix (\mathbf{I}) in the convolution allows for matrix inversion and recovery of the best-fit of the IRF based on analysis of the entire time-trace, as shown in Equation 3.

$$\mathbf{irf} = [(\bar{\mathbf{d}}_{fit} \otimes \mathbf{I}_n)^T (\bar{\mathbf{d}}_{fit} \otimes \mathbf{I}_n)]^{-1} (\bar{\mathbf{d}}_{fit} \otimes \mathbf{I}_n)^T \bar{\mathbf{d}}_{meas} \quad (\text{Eq 3.3})$$

In this work, the algorithm is made to iterate through both the recovery of the photon events in $\bar{\mathbf{d}}_{fit}$ and the IRF. First, the IRF is set to a guess value that we assume is true for the initial recovery of the photon events. Then, the parameters recovered from the first half of the iteration are used in a second least squares minimization for recovery of the IRF, which is then reused iteratively in the recovery of $\bar{\mathbf{d}}_{fit}$ until convergence is achieved. After convergence with non-negativity imposed on the IRF, another round of iterative optimization is performed to allow the recovery of weak but nonzero ringing features by relaxing the non-negativity constraint on the final IRF. The overall workflow of the NMF algorithm is illustrated in **Figure 3.1**.

In principle, the expressions in Equations 2 and 3 can be evaluated on the complete data set as indicated in the equations. In practice, the dimensionality of long data-traces complicates direct inversion. Inversion of large data sets belongs to the class of NP-Hard problems in most cases. (Vavasis) To help alleviate this problem, NMF was implemented as a rolling filter (NMF-F) on a small subset of the data, with the analysis performed over a sliding window across the data vector. The $\bar{\mathbf{d}}_{meas}$ vector mentioned in Equation 1 was defined to be a 100-point window of the data, which was on the order of the length of the IRF. NMF analysis was performed using Equation 2 to identify the locations within the window corresponding to initiation of photon events (assuming a known IRF). This approach is qualitatively similar to a linear fitting digital filter (e.g., Savitsky-Golay), but with the NMF constraints imposed. (Savitzky & Golay) Following this rolling analysis, the entire data set was used in evaluating the optimal IRF using Equation 3, followed by iteration between these two steps until convergence was achieved.

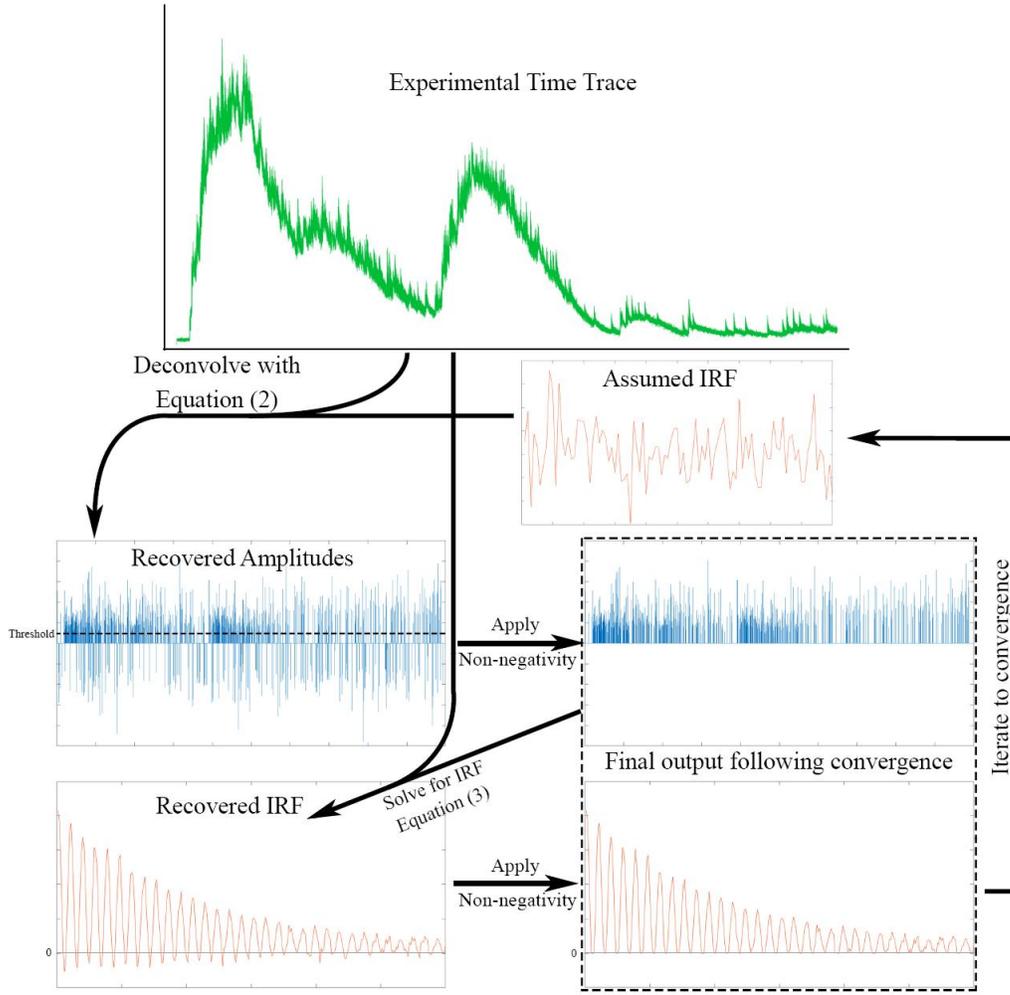


Figure 3.1: Illustration of the NMF-F algorithm workflow.

3.3 Algorithm

The NMF-F algorithm used in this publication integrates capabilities from both MLE NMF and convolutional NMF to perform a rolling NMF on a convolution operation rather than multiplication. To integrate these disparate approaches, we use the identity operation in matrix multiplication, combined with the associative property of convolution and multiplication to generate the final expression shown in equation 4.

$$\bar{\mathbf{d}}_{meas} = \bar{\mathbf{d}}_{true} \otimes \mathbf{irf} = (\bar{\mathbf{d}}_{true} \otimes \mathbf{I})\mathbf{irf} \quad (\text{Eq. 3.4})$$

Introduction of the IRF to the convolution in equation 4 allows for the recovery of the best fit of the IRF through Equation 3. Convolution of $\bar{\mathbf{d}}_{fit}$ with the identity matrix in this manner enables isolation of the IRF for mathematical inversion. Once the IRF is determined, local implementation on short sections of data as a digital filter greatly reduces the computational time. Where a smaller section of the data, $[\mathbf{d}_i : \mathbf{d}_{i+n}]$, which is now $(1 \times n)$ where n is now the desired filter size. This new $[\mathbf{d}_i : \mathbf{d}_{i+n}]$ vector is then evaluated iteratively through the entire data set. Performing the analysis as a rolling filter in this way substantially reduces the time and computational resources required, enabling implementation on simple benchtop computers.

3.3.1 Baseline Correction

-Convolve data with a high pass filter

3.3.2 Matrix Factorization

The general strategy for implementing NMF-F is illustrated in **Figure 3.1**. The NMF-F is evaluated for the i -th element over a range of data equal to the length of the IRF, n . A photon event is allowed to arise in every position within that range when evaluating the amplitude of the i -th element. This process is repeated for every i -th element of the measured data to generate the deconvolved output and recover the IRF.

3.3.3 Initialization

- Set threshold for negativity constraint (accounting for noise).
- Generate initial IRF guess (e.g., uniformly distributed random numbers; $n=100$)

3.3.4 Solving for the $\mathbf{d}_{fit,i}$ element of the \mathbf{d}_{fit} Vector

- Generate the matrix $P(n \times n)$ by $P = I_n \otimes irf$, where P is initially the convolution of an identity matrix of dimension n with the IRF, truncated to produce an $n \times n$ matrix.
- For the i -th element of $d_{measured,i}$ select a section of data of length n and invert the equation $[\mathbf{d}_i : \mathbf{d}_{i+n}]_{measured} = [\mathbf{d}_i : \mathbf{d}_{i+n}]_{fit} P_{(n \times n)}$ to determine the best fit values of $[\mathbf{d}_i : \mathbf{d}_{i+n}]_{fit}$

- Identify the elements of $[d_i : d_{i+n}]_{fit}$ below threshold and remove the corresponding columns of P to generate $P'_{(n \times m)}$ where $m < n$.
- Express $[d_i : d_{i+n}]_{measured} = [d_i : d_{i+n}]_{fit} P'_{(n \times m)}$ and invert to recover an updated $[d_i : d_{i+n}]_{fit}$
- Iterate until no below threshold events are recovered for the amplitudes in $[d_i : d_{i+n}]_{fit}$
- Retain the value of $d_{fit,i}$. If $d_{fit,i}$ is nonzero, subtract $d_{fit,i} \otimes irf$ from $[d_i : d_{i+n}]_{measured}$
- Repeat for all i positions of the measured data.

3.3.5 Solving for the IRF

- Express $d_{measured}$ in the following form: $d_{measured} = d_{fit} \otimes irf = (d_{fit} \otimes I_{n \times n}) irf$
- Multiply by the transpose and invert $(d_{fit} \otimes I)$ to recover the least-squares best fit values for irf
- If imposing non-negativity to the IRF set below threshold entries to zero.

3.3.6 Iterate to Convergence

- Using the updated IRF, solve for d_{fit} and update.
- Using the updated d_{fit} , solve for the IRF and update.

3.3.7 Autocorrelation

- Take the fast Fourier transform (FFT) of d_{fit} , such that $\phi = FFT(d_{fit})$.
- Multiply $\phi \cdot \phi^*$ element-wise.
- Take the inverse FFT of $\phi \cdot \phi^*$.

3.4 Experimental

Triboluminescence measurements of the were obtained using an instrument described in detail in a preceding publication (Casey J. Smith et al.). In brief, kinetic energy impulses were delivered to the sample via a solenoid driven impactor, which mechanically compressed the sample to induce triboluminescence. A single impact delivered ~ 0.5 J over a 7 mm^2 area. The light emitted by the sample was collected and focused onto a photomultiplier tube (PMT, H12310-40,

Hamamatsu, Japan) which was connected to a PCI-express digital oscilloscope card (ATS-9350, AlazarTech, Pointe-Claire, QC, Canada) to record the temporal response of the detector. An Arduino micro-controller was used to trigger data collection and control movement of the sample stage. An impedance mismatch was intentionally introduced between the transmission line (50Ω) and the digitizing electronics ($1M\Omega$). This mismatch provided a signal to noise advantages through two independent mechanisms: i) increasing the integrated voltage through the increased resistance of the digitizer according to Ohm's Law, $V = IR$, and ii) simplifying digitization of the transients by increasing the temporal duration of each voltage transient according to the RC time constant associated with discharge of the transmission line capacitance through the resistor. The increase in voltage manifested as a series of reflections separated by the round-trip transit time through the transmission line of ~ 100 ns, followed by eventual decoherence to a constant baseline offset. The resulting complex waveform produced by an impulsive photon absorption event still retains essentially the same the high frequency information content of the original current transient from the detector, but with different frequencies shifted in phase to produce the observed impulse response function.

Simulations were conducted to quantitatively evaluate the performance of the NMF-F for systems with known ground truth results. A Poisson probability density function (pdf) was used for describing the probability of detecting a photon during an individual impact event. Because the duration of the burst was much longer than the digitization rate, the maximum mean of the Poisson distribution was still less than one at all time-points within the burst profile. For each detected photon, the amplitude was determined by a lognormal pdf describing the variance in gain within the dynode string of the PMT used (Kissick et al.). Once the ground truth was established for the locations and amplitudes of the voltage transients, each voltage transient was then convolved with an impulse response function measured for a single photon event. Finally, normally distributed white noise with a mean of 0 and a standard deviation an order of magnitude lower than the maximum photon amplitude was added to make the simulated time trace. Normally distributed white noise was added after convolution. These data were used to evaluate the algorithm's ability to accurately recover photon events with increasing probabilities of pulse pile-up, in which a new photon event is initiated within the temporal envelope of a preceding IRF.

Comparisons were performed between the NMF-F approach and conventional NMF of the entire data set *en masse*. In conventional NMF, computational time generally increases non-

polynomially with data size (is an NP-hard problem in most cases), which results in large computational time savings from any reduction in the set size. As a consequence of this scaling, attempts to perform conventional NMF on the nearly 2.5 million data points from a single impact event resulted in impractically long computational times, despite performing calculations with sparsity constraints on the Purdue Radon Supercomputing Cluster (45 nodes, One Hyper-Threaded Quad-Core Xeon E3-1284L/8 cores/32 GB RAM per node). By comparison, use of the sliding filter design in NMF-F enabled analysis in ~18 min for pooled analysis of 35 impact events (87.5 million data points) on a Dell XP15 laptop with a 2.6 GHz processor and 16 GB of RAM.

Doxycycline Hyclate was used as received from Sigma-Aldrich (St. Louis, MO).

Griseofulvin was obtained from Sigma-Aldrich (St. Louis, MO), and hydroxypropylmethylcellulose acetate succinate (HPMCAS) was purchased as AQOAT from Shin Etsu (Tokyo, Japan). The spray dry formulation comprised of griseofulvin (20% w/w) and HPMCAS-MF (80% w/w). The solid components were dissolved in Acetone (4% w/v). The solution was spray dried on a 4M8-TriX spray dryer (ProCepT NV, Zelzate, Belgium) with a two-fluid nozzle diameter of 0.6 mm. The dried particles were carried by air into a cyclone and then into a collection vessel. The processing air flow rate was set to be 0.40 m³/min. The atomization air and the solution feed flow rates were taken to be 3.0 L/min and 6 mL/min, respectively. The spray drying conditions were set to reach inlet temperature of 91°C and outlet temperature of 50°C. The cyclone differential pressure was set at 30 mbar and tangential air flow rate was chosen to be 0.1 m³/min while the cooling air was shut off. The spray dried products were further dried under vacuum at ambient temperature overnight. The griseofulvin ASD was confirmed to be amorphous using powder x-ray diffraction (PXRD), second harmonic generation (SHG) microscopy, and TL following preparation. The ASD was then stressed for 48 hours at 65°C and at elevated relative humidity using a saturated NaCl solution to induce crystallization.

3.5 Application to Triboluminescence

A representative time-trace for a TL measurement is shown in **Figure 3.2**. Several key aspects of the measurement are noteworthy. First, at least three separate burst-events are detectable from the time-trace, the first of which at ~ 0.1 ms being the largest. Second, the intentional impedance mismatch introduces a relatively long exponential decay on the detected signals arising from burst events. Thirdly and less obviously at this scale, the time-trace includes high-frequency

spikes arising from the absorption of individual photons. While the impedance mismatch was advantageous for detection of the individual transients, the resulting temporal convolution substantially complicated efforts to characterize the intrinsic burst profile of individual TL events and inform models for describing the mechanisms of TL production.

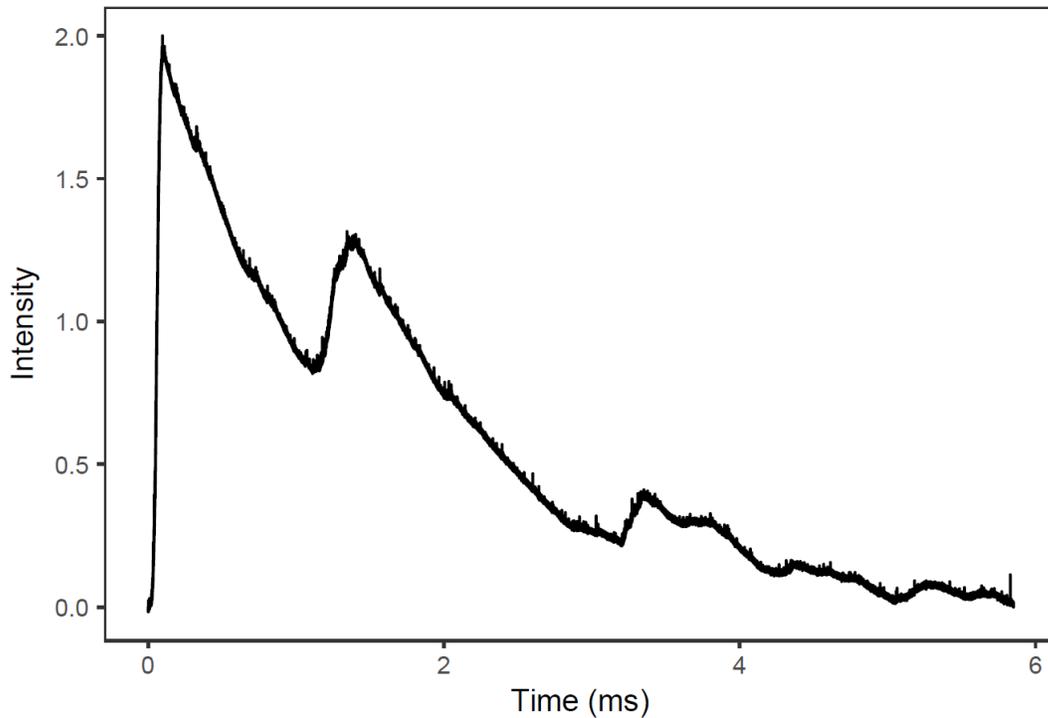


Figure 3.2: Example signal transient pulled from a larger set of TL measurements which illustrate the complexity and difficulty in distinguishing individual photon events. Transients on a microsecond time scale appear as impulses over the millisecond scale of the emission event.

The merits and limitations of the NMF-F were initially evaluated using the simulated data as shown in **Figure 3.3**, for which ground-truth results were known *a priori* for algorithm evaluation. Simulated time-traces were generated by convolving a noise-free IRF extending over 100 pixels with a pulse-train described by a lognormal distribution of amplitudes, followed by addition of Gaussian noise to the post-convolution time-traces. NMF-F was compared to standard blind Richardson-Lucy (BRL) deconvolution and to deconvolution using a Wiener filter, which requires advanced foreknowledge of the IRF. The Wiener deconvolution performed with a noise-free exact IRF approaches the theoretical limiting behavior expected by any blind deconvolution

algorithm, including NMF-F. The signal transients recovered by NMF-F analysis are shown in **Figure 3.3a** (blue trace-top) along with a comparison with the simulations (black trace-bottom). The IRF recovered using the NMF-F algorithm is shown in **Figure 3.3b** (blue dashed-middle), together with the initial guess (red-bottom) used for the IRF to initiate the NMF-F. While the NMF-F algorithm converges well for random initial guess inputs for the IRF (*vide infra*), the BRL algorithm used as a comparator required initial guesses for the IRF with significant similarities to the ground truth result for computational convergence. The recovered photon events (blue-top) by NMF-F were compared with the ground truth positions (black-bottom) in **Figure 3.3c**, in which the ground truth results were plotted as negative transients to help visualize peak positions. Receiver operating characteristic (ROC) plots were calculated in **Figure 3.3d** to assess the ability of the algorithm for reliable recovery of single-photon events relative to blind deconvolution by BRL and by the Weiner filter using the exact IRF. This process was performed for a mean number of photons per digitization timepoint, $\lambda = 0.1$, which is the Poisson distribution mean value. For example, $\lambda = 0.1$ corresponds to a photon originating every 10th digitization event on average. As shown in the ROC plot in **Figure 3.3d**, the area under the curve is close to unity for NMF-F, suggesting the potential for low false positive and false negative rates following NMF-F deconvolution. Most significantly, blind NMF-F substantially outperforms the comparator BRL algorithm, approaching behaviors anticipated by algorithms with exact foreknowledge of the IRF.

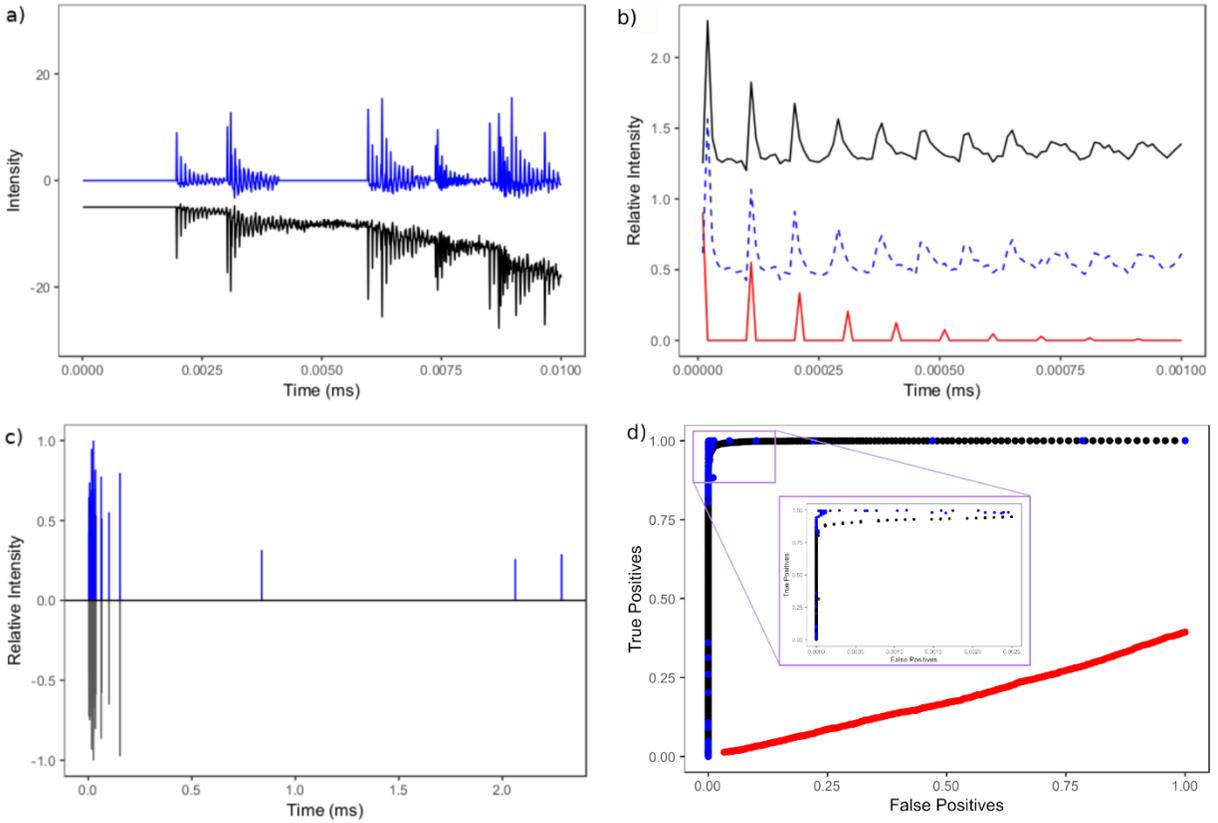


Figure 3.3: Using simulated data zoomed in to microsecond time scales and normalized results for ease of comparison with the ground truth. A) shows the recovered transient (blue-top) compared to the original simulated transient (black-bottom). Where the black trace is displayed with negative amplitude for ease of comparison. In B), the ground truth (black-top, offset by 1 Norm.U.), recovered (blue dashed, offset by 0.5 Norm.U.), and initial guess (red-bottom) IRFs are compared and show good agreement. In C), the recovered photon events (blue-top) are compared with the ground truth events (black-bottom) and show nearly complete recovery with minimal error. The digital filter's accuracy and ability to recover photons in high density data is highlighted in D), by the ROC plots made with the same data and analyzed using NMF-F (blue), Wiener deconvolution (black), and BRL (red). Both the Wiener deconvolution and BRL were performed using the true IRF, while the NMF-F analysis was performed with a good guess from previous experiments. The AUC for NMF-F is 0.987 and 0.998 for the Wiener deconvolution. While the NMF-F performed better at lower false positive rates, it performs worse as the false positive rate increases.

The insensitivity of the NMF-F approach to the initial guess for the impulse response function was assessed by repeating the analysis for a host of different initial IRF guesses. Initial waveforms for the assumed IRF included single impulses, pulse trains with incorrect periods, and random normally distributed waveforms. Each individual implementation of NMF-F blind deconvolution with the most challenged case of fully randomized input, results in convergence recovering both the IRF and the deconvolved time-trace with a success rate of 82.5%. The remainder of which converges to IRFs containing many impulses that more resemble the Gaussian noise. As such, implementation with just 5 initial random guesses results in a 99.92% chance for successful convergence with no *a priori* knowledge. Of course, this success rate of individual implementations can approach unity with initial guesses based on prior knowledge. Fortunately, identification of the outliers was straightforward, as they greatly overestimated/underestimated the total number of detected photons relative to the set closely matching the ground truth IRF or produces an IRF that isn't realistic. An example with a random guess for the IRF is shown in **Figure 3.4**, indicating both the initial IRF, the final convergent solution for the IRF generated by the NMF-F and BRL, the ground truth IRF, and the fraction of single photon events properly identified by the NMF-F (~97%). Given the complexity and high frequency content of the ground truth IRF, the recovery of reasonable results for random initial guess values is quite remarkable and attests to the overall stability of the NMF-F computational method. No comparable convergence could be achieved using the BRL algorithm. It is worth emphasizing that analysis with a random guess requires no significant prior knowledge or assumptions of either the photon positions or the instrumental IRF.

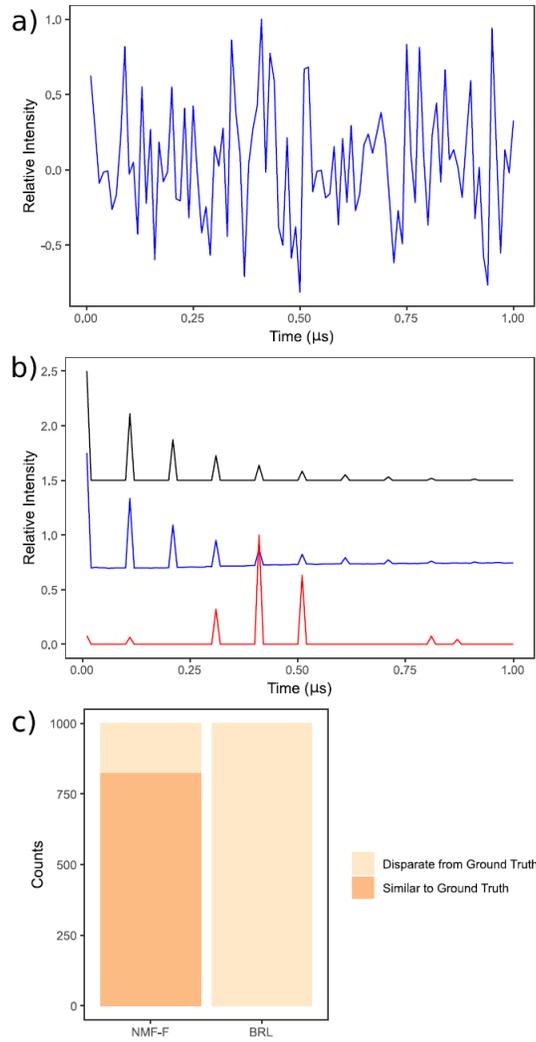


Figure 3.4: An example of the algorithm’s ability to recover an IRF similar to the ground truth.

(a) Representative initial random guess used in the NMF-F and the BRL algorithm. (b)

Recovered IRFs for the NMF-F (blue-middle) and the BRL (red-bottom) algorithms are

compared with the ground truth (black-top), which are offset for ease of comparison. (c) Bar chart highlighting the difference in ability to recover an IRF similar to the ground truth of the NMF-F and BRL algorithms. NMF-F converged to a similar IRF to the ground truth 82.5% of 1000 attempts, while BRL failed in all attempts.

Next, the NMF-F was applied to the measured triboluminescence signal shown in **Figure 3.2**. The results of the deconvolution are shown in **Figure 3.5**, in which the photon arrival times are compared with the original data. The expanded portion of **Figure 3.5** allows identification of individual photon-events at the beginning of a burst generated from a triboluminescence event. From inspection of the results, the deconvolution algorithm notably increases the localization of photon events and suppresses noise from convolution with the IRF, from the baseline offset, and from the white electronic noise. In **Figure 3.6**, the NMF-F algorithm was applied to the ASD sample, where crystals were grown in elevated humidity and temperature.

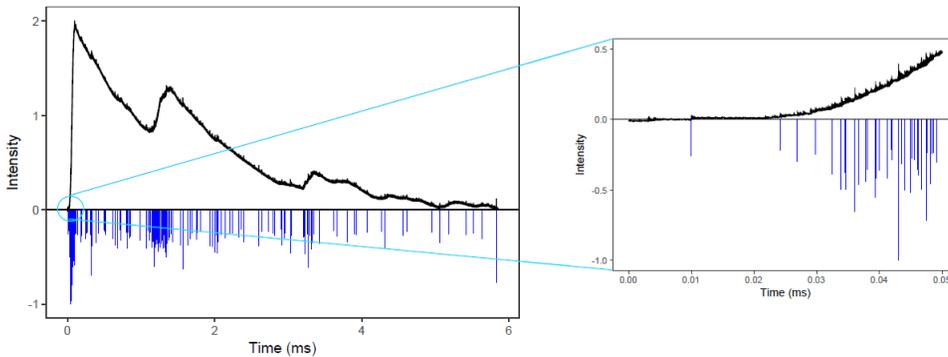


Figure 3.5: Same TL signal transient from Figure 3.2 with the recovered photon events (blue – lower). The inset illustrates how many photons are within the initial rise caused by a TL burst event.

Following deconvolution, autocorrelation analysis of the measured TL time-traces was performed to yield information on the characteristic time-scale for the individual burst events, and in turn insights into the mechanisms of TL in pharmaceutical powders.(X. Y. Dow et al.) The results of the autocorrelation analysis are shown in **Figure 3.7**, revealing a characteristic timescale for the burst events on the order of $\sim 50 \mu\text{s}$. The importance of deconvolution is also illustrated in **Figure 3.7**, in which the autocorrelation analysis of the raw time-trace as received produced bias by a factor of ~ 10 in the evaluation of the characteristic timescale. The bias seen in the autocorrelation of the raw data $\bar{c}_{meas}(\tau)$ can be explained by Equation 5. The autocorrelation is affected by both the inherent characteristics of the TL process $\bar{c}_s(\tau)$ and the IRF $\bar{c}_{IRF}(\tau)$, where in this case the IRF influence dominates.

$$\bar{c}_{measured} = \bar{c}_s(\tau) \otimes \bar{c}_{IRF}(\tau) \quad (\text{Eq. 3.5})$$

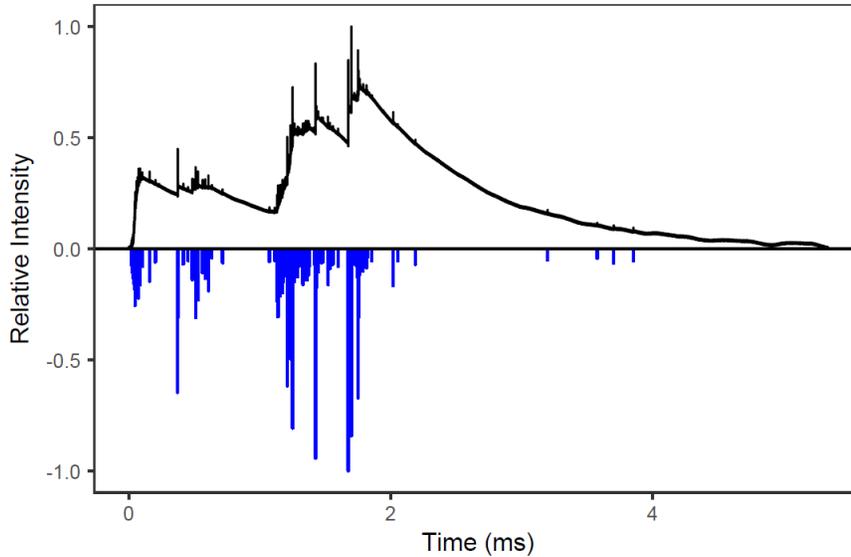


Figure 3.6: Amorphous solid dispersion of griseofulvin in HPMCAS-MF, stressed to induce crystallinity (black – upper) and the recovered photon events (blue - lower, inverted for clarity).

Inspection of the deconvolved autocorrelogram for TL bursts provides several insights into the possible mechanisms driving the observed signal production. Notably, longer timescale recurrences in the autocorrelogram that might arise from ringing events were not observed, consistent with TL arising stochastically in $\sim 50 \mu\text{s}$ bursts. High-speed imaging was performed to assess the piston speed at the point of impact, and from that the timescale for traversal through the sample. The maximum piston speed was $\sim 2 \text{ m/s}$ corresponding to the piston taking $\sim 0.5 \text{ ms}$ to travel through the 1 mm thick sample. This timescale is significantly longer than the individual burst events from autocorrelation analysis, but comparable to the entire envelope of bursts produced by a single impact event. The $50 \mu\text{s}$ time scale is in qualitative agreement with the $\sim 20 \mu\text{s}$ temporal envelopes reported in laser induced breakdown spectroscopy experiments, where emission of excited gas-phase AIO caused by a spark created by a laser pulse was measured over time in different bath gases. (Piehler et al.)

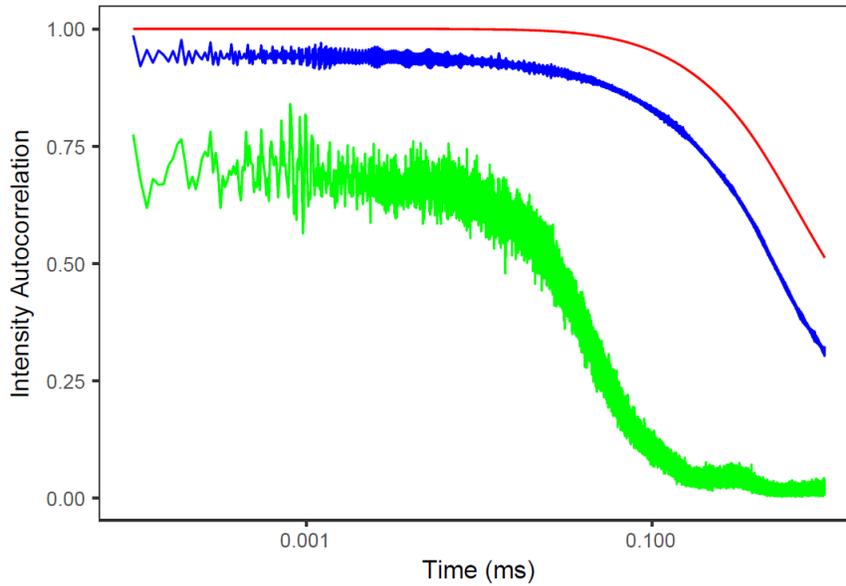


Figure 3.7: Results of autocorrelation on the raw TL data (red-top), the measured IRF (blue-middle), and the deconvolved photon events (green-bottom). The x-axis is on a log scale.

Consistent with this collective set of results, we can propose the following mechanism for the observed TL. During the $\sim 500 \mu\text{s}$ traversal of the piston through the powdered sample, individual crystalline particles independently produced TL. The temporal profiles of these individual TL bursts were consistent with localized plasma production from dielectric breakdown, presumably from rapid changes in capacitance upon lattice collapse. In some TL events, multiple bursts were observed, consistent with the collapse of several different independent particles during the $\sim 500 \mu\text{s}$ impact event. The routine observation of the most pronounced TL burst amplitudes relatively early in the impact window is consistent with the larger particles collapsing at earlier times, subsequently followed by bursts associated with smaller crystals as shown by the secondary peaks in **Figure 3.2** (Casey J. Smith et al.).

If individual TL bursts arise from the stochastic collapse of individual particles as suggested by these measurements, a TL mechanism based on capacitive discharge should enable quantitative connections to be drawn between the crystal size and the amplitude of the individual TL burst events. Such connections would raise the intriguing possibility of probing the crystal size distribution from the histogram of intensities for the individual TL bursts. This histogram would be dependent on the size distribution of individual crystals, which is distinct and complementary

to the conglomerate particle size distribution most easily accessible from alternative analysis approaches.

3.6 Conclusions

An NMF-F digital deconvolution algorithm was developed to recover the individual photon arrival events and applied for the analysis of TL observations. In addition to reliably recovering arrival times in simulations with known ground-truth results, the NMF-F algorithm simultaneously recovered reliable estimates for the ground-truth impulse response function. Adaptation of NMF as a digital filter enabled analysis of high-density data traces that were prohibitively time-consuming for conventional NMF analysis. Random initial guesses for the IRF converged to results similar to the ground truth 82.5% of the time, recovering 97% of the impulse events with a false positive rate of 0.06% with no required prior knowledge of the IRF.

Autocorrelation analysis was performed on the deconvolved impulse train to evaluate several hypotheses regarding the fundamental mechanism of triboluminescence. (Becker et al.) The result provided the characteristic timescales of triboluminescence emission events, which was obscured in the original data transients from convolution by the impulse response function. The recovered timescale for TL bursts was consistent with a mechanism for TL arising from single crystals fracturing and producing plasma discharge as they were impacted by the piston. This proposed mechanism leads to the possibility of probing crystal size distribution through TL measurements.

3.7 Acknowledgements

The author would also like to acknowledge support from the Jonathan Amy Facility for Chemical Instrumentation for their assistance in the design and fabrication of the prototype instrument used in these studies as well as the Merck & Co., Inc., Kenilworth, NJ, USA, Pharmaceutical Sciences Innovation Leadership Team, Timothy Rhodes, Roy Helmy, Yong Liu, Zack Zhiqiang Guo, and Maya Lipert for assistance in gaining funding and / or for valuable technical advice.

CHAPTER 4. CALIBRATION-FREE SHG IMAGE ANALYSIS

4.1 Introduction

Modern drug discovery efforts fueled by the growth of high throughput screening and combinatorial chemistry as tools in the pharmaceutical industry have led to more drug candidates with greater chemical complexity and correspondingly lower aqueous solubilities.(Babu & Nangia, 2011; Thayer, 2010) Poorly water-soluble active pharmaceutical ingredients (APIs) lead to reduced bioavailability, which in turn results in an inadequately effective drug formulation.(Babu & Nangia, 2011) Consequently, the preparation of amorphous forms of an API has become increasingly favorable due to their higher apparent solubilities and faster dissolution rates.(Taylor & Zhang, 2016; Yu, 2001) However, amorphous APIs are typically metastable forms with a thermodynamic driving force to crystallize over a range of timescales dependent on the composition of the formulation and the external environment. A common strategy used to prevent crystallinity is the addition of crystallization-inhibiting polymers that contributes to the stabilization of amorphous formulations, resulting in an amorphous solid dispersion (ASD).(Baghel et al., 2016; Williams et al., 2013; Yu, 2001) However, even trace residual crystallinity can provide nuclei for crystal formation during storage and subsequently following introduction to the body, with corresponding reduction in bioavailability.(Wanapun et al., 2010; Wanapun et al., 2011) Therefore, the ability to detect and characterize crystallinity within drug products is critical for the assessment of the safety, stability, and efficacy of ASD drug products. Several methods are currently available for analysis of residual crystallinity within nominally amorphous formulations, including infrared and Raman spectroscopy,(Berry et al., 2008; Bugay, 2001; Ng et al., 2013; Taylor & Zografi, 1998; Yang et al., 2014) powder X-ray diffraction (PXRD),(Berendt & Munson, 2012; Shah et al., 2006) solid-state nuclear magnetic resonance (ssNMR),(Berendt et al., 2006; Johari et al., 1990; Offerdahl et al., 2005; Sheth et al., 2005) and differential scanning calorimetry (DSC).(Schmitt et al., 2015; Zimper et al., 2010) All of the methods mentioned are often limited to detecting total crystallinity no lower than about 1% (w/w%) under ideal conditions (i.e., high drug load and limited excipient interference) utilizing benchtop instrumentation, which may not be sufficient as even trace amounts of crystals can impact efficacy.

Thus, it is important that the analytical tool used for the characterization of crystallinity in ASDs is selective and as sensitive as possible.

Nonlinear optical (NLO) imaging has proven to be a successful method in the pharmaceutical field for the sensitive detection of trace crystallization in amorphous systems. (Chowdhury et al., 2016; Kestur et al., 2012) Second harmonic generation (SHG) is a coherent process, in which the incoming light is converted to light of twice the frequency. The symmetry properties of SHG only permit this second-order process in ordered noncentrosymmetric systems; amorphous or disordered systems do not generate coherent SHG. Crystals of homochiral molecules generally fall into space groups that are symmetry-allowed for SHG, providing highly selective and sensitive detection of homochiral crystalline APIs. (Hsu et al., 2013; Kissick et al., 2011; Newman et al., 2015; Schmitt et al., 2016; Toth et al., 2015) Limits of detection for chiral crystallinity by SHG have been reported in the part-per-trillion regime in melt-quenched films and in the ppm regime for powders. (Correa-Soto et al., 2017; Kestur et al., 2012) More recently, SHG has been used to monitor changes in crystallinity in amorphous solid dispersions in spray-dried dispersions. (Novakovic et al., 2017) In that work, the integrated SHG intensity was found to scale with the fraction of crystallinity arising during storage at high temperature and humidity, with limits of detection for crystallinity down to ~0.01%. This capability greatly reduced the timescale required for early identification of conditions leading to crystal formation. Complementary nonlinear optical methods including coherent anti-Stokes Raman (CARS) imaging involving interactions from multiple beams have also been demonstrated for compositional analysis of pharmaceutical materials. (Hartshorn et al., 2013; Pack et al., 2017) In general, CARS lacks the symmetry-dictated selectivity of even-ordered nonlinear optical processes, although multimodal imaging integrating both SHG and CARS has the potential to recover such information. (Pack et al., 2017)

Despite successes in detection, precise quantification of residual crystallinity by SHG, continuously spanning a wide dynamic range of crystalline content remains challenging. To date, quantification of crystallinity has relied almost exclusively on the integrated SHG intensity produced within a field of view. This approach has significant limitations for developing predictive models for crystallization kinetics, namely, at low levels of crystallinity, the noise levels are nearly on par with the level of signal. Quantification of SHG by the integrated SHG intensity also typically requires the need for calibration standards, (Kestur et al., 2012) in which measurements

of the pure substance provide a reference point for connecting the measured SHG intensity back to a percent crystallinity. However, access to reliable standards is not always trivial. For many APIs, several different crystal forms are potentially kinetically accessible during crystal formation. Given the high sensitivity of SHG to crystal form,(Toth et al., 2015) standards prepared based on the assumption of one polymorphic form would be inapplicable for measurements obtained with a different crystal form. In addition, the integrated SHG can be dependent on the crystal size distribution, such that sets of measurements acquired with identical % crystallinities can produce quite different integrated SHG activities. Algorithms have been developed to address bias from the particle size distribution, but represent yet another potential complication for routine quantification of crystallinity.

In one recent departure from the signal integration approach for quantification, Schmitt et al. directly recovered the particle size distribution via image analysis for poorly soluble particulates present in commercially available formulations of Abraxane™.(Chowdhury et al., 2016) SHG-active particles as large as 120 μm in diameter were observed within the commercial lyophilized powders. In comparison to signal integration, this particle counting approach has the potential to further lower the limits of detection by SHG (by analogy with photon counting). Differences in particle size and orientation greatly influence the SHG intensity, but not the integrated area in the image. Despite the numerous advantages of particle counting in the limit of low crystallinity (<1%), the dynamic range is quite limited; as the overall crystallinity increases, so too does the probability of observing spatially overlapping crystals. This effect is analogous to paralysis in photon counting, in which overlap in the temporal transients ultimately sets the upper limit on the accessible dynamic range. Unlike photon counting, in which transients of closely separated photons in time are counted as a single event, particle counting can be biased instead from overlap in space between adjacent particles. Unfortunately, the bias introduces a measurement gap, in which neither technique (particle counting or signal integration) yields an optimal, unbiased signal to noise ratio.

In this work, the two methods (signal integration and particle counting) are bridged through the derivation of a statistical model based on analytical expressions for the optimization of signal to noise. Specifically, a particle-counting algorithm capable of unbiased extension of the counting regime to overlap with the regime optimized for signal integration was developed. Based on the relative variance of the recovered volume, an analytical expression was derived for the percent crystallinity for which both bias-corrected particle-counting and signal integration produce an

identical average signal to noise ratio. Crystalline fractions higher than this point are optimally determined by signal integration and fractions below by bias-corrected particle-counting. The average SHG-activity per unit volume measured in the particle-counting regime served as an internal calibration standard for quantifying crystallinity in the signal integration regime, removing the need for external standards. Monte Carlo simulations were performed to assess the merits of the approach. Measurements were performed on evacetrapiB, a cholesteryl ester transfer protein inhibitor, formulated tablets (Sullivan et al., 2014) to experimentally demonstrate the bridging of the two regimes and autocalibration.

4.2 Mathematical Theory

4.2.1 Particle Counting Regime

In the particle-counting regime, in which particles are well-separated and non-overlapping, volume determination from image analysis provides a signal to noise advantage relative to integration of the SHG intensity. This advantage in particle-counting comes from the suppression of dark counts and the neglect of variance in the SHG intensity due to particle orientation and size. However, at higher % crystallinity, particle-counting approaches begin to introduce bias from “pulse-pile-up”, in which overlap between adjacent particles within the depth of field results in counting as a single particle rather than as a pair.(Molnar et al., 2016; Muir et al., 2012) Analogous to single-channel measurements with pulse-pileup, methods to correct for this bias can extend the dynamic range of particle-counting, but at the expense of signal to noise.(Molnar et al., 2016) Segmentation algorithms can separate overlapping particles that are relatively large (i.e., span many pixels), but are not reliable for particles only a few pixels in dimension as studied herein.(Closser et al., 2013) The primary objective of the theoretical framework outlined below is to extend the dynamic range of particle-counting by correcting for bias from particle-particle overlap, and identifying the point at which integration of the SHG intensity offers a signal to noise advantage over bias-corrected particle-counting. Furthermore, the SHG intensities measured in the particle-counting regime can directly serve to calibrate the intensity-based volume assessment without the need for an external reference sample.

In the absence of overlap between spherical particles and in the limit of a large depth of field, the volume of an individual particle can be determined from the cross-sectional area, A , by

$V_i = \frac{4}{3\sqrt{\pi}} A_i^{3/2}$. The total volume is given by the summation of the volumes of all n particles present in a given field of view. The mean of the total volume, $\mu_{V_{Total}}$, can be estimated by the mean particle area, μ_A , and the mean number of particles, μ_n .

$$\mu_{V_{Total}} = \mu_n \mu_V \cong \mu_n \frac{4}{3\sqrt{\pi}} \mu_A^{3/2} \quad (\text{Eq 4.1})$$

The approximation in Equation (4.1) neglects computational bias arising from the nonlinear relationship between the variances in particle size and the recovered volume but should be reasonably reliable for narrow distributions in particle sizes. The relative variance in the total volume in the limit of low volumes can be related back to the Poisson distributed number of particles and the distribution in particle sizes. Summing the contributions from both the variance in the particle radius and the Poisson number of particles yields the following expression for the relative variance, $\sigma_{V_{measured}}^2$, in the measured volume in the limit of no overlapping particles, normalized by the mean measured volume, $\mu_{V_{measured}}$.

$$\left(\frac{\sigma_{V_{measured}}}{\mu_{V_{measured}}} \right)^2 = \frac{1}{\mu_n} \left[1 + 9 \left(\frac{\sigma_r}{\mu_r} \right)^2 \right] \quad (\text{Eq 4.2})$$

As the density of particles increases, the probability of particle-particle overlap becomes non-negligible and the mean ground truth area μ_A will no longer equal the mean measured area by particle-counting, $\mu_{A_{measured}}$. Considering only the probability of two-particle overlap, the measured volume $V_{measured} = \sum_i \frac{4}{3\sqrt{\pi}} A_i^{3/2}$, over-estimates the ground truth total volume V_{Total} . This bias in the volume can be removed by incorporating a correction term C , together with knowledge of the difference between the observed number of particles, N , in the presence of some overlap from naïve particle-counting and the underlying ground truth number of particles, n .

$$\mu_{V_{re-cov}} = \frac{4}{3\sqrt{\pi}} \mu_{A_{measured}}^{3/2} (2\mu_N - \mu_n)(1 - 2C) \quad (\text{Eq 4.3})$$

The analytical expression for the correction term C was derived by determining the mean calculated volume and the true underlying volume.

$$C = \frac{\frac{4\pi}{3}(2\mu_r^3)}{\frac{4}{3\sqrt{\pi}} \left[2\mu_r^2(\pi - \mu_\theta + \frac{1}{2}\sin(2\mu_\theta)) \right]^{\frac{3}{2}}} \cong 0.842 \quad (\text{Eq 4.4})$$

Recovery of a reasonable estimate for the mean number of underlying particles μ_n is more challenging. An approximate value for μ_n can be calculated from foreknowledge of the average particle area μ_A (e.g., measured at low concentration) and the field of view (*FOV*) through the dimensionless parameter $\bar{A}_0 = (\mu_A / \text{FOV})$.

$$\mu_n \cong \frac{1 - \sqrt{1 - 8\bar{A}_0\mu_N}}{4\bar{A}_0} \quad (\text{Eq 4.5})$$

From Equation (4.5), it should be clear that the estimation of μ_n is only applicable in the regime in which $\bar{A}_0\mu_N < \frac{1}{8}$. Allowing for a reasonable buffer, the particle-counting analysis algorithm presented herein should be applicable provided that $\bar{A}_0\mu_N < 0.1$. Experimentally, $\bar{A}_0\mu_N \cong 0.1$ corresponds to approximately 10% of the pixels within the field of view occupied by particles. In this regime, the probability of three or more particle overlap events becomes significant and are not accounted for in the two-particle corrections considered herein. Computationally, the factor of $\sqrt{1 - 8\bar{A}_0\mu_N}$ cannot be evaluated for $\bar{A}_0\mu_N > \frac{1}{8}$, for which the square root becomes negative. Fortunately, integration of the SHG intensity provides an alternative approach for quantification of crystallinity in the regime $\bar{A}_0\mu_N > 0.1$, enabling further extension of the dynamic range by bridging the counting and averaging regimes.

The relative variance in the recovered volume when incorporating the two-particle correction is given by the following expression.

$$\left(\frac{\sigma_{V_{\text{reco}}}}{\mu_{V_{\text{reco}}}} \right)^2 \cong \left[\frac{\mu_n - 2\mu_p + 4C^2\mu_p}{(\mu_n - 2\mu_p + 2C\mu_p)^2} \right] \left[1 + 9 \left(\frac{\sigma_r}{\mu_r} \right)^2 \right] \quad (\text{Eq 4.6})$$

In Equation (4.6), μ_p is the mean number of “dimer” particles, in which two actual particles are treated as a single particle by naïve shape-independent particle-counting algorithms, and is given by $\mu_p = \mu_n - \mu_N$. As expected, the relative variance from Equation (4.6) simplifies to

that in Equation (4.2) in the limit of μ_p approaching zero, corresponding to negligible particle-particle overlap.

For practical purposes of quantification, the recovered volume is converted into the volume fraction (i.e., the crystalline fraction relative to the total interrogated volume). Conversion from the volume fraction to the more commonly used mass fraction requires knowledge of the packing density of the sample, which can be independently determined.

$$\text{Volume Fraction} = V_{\text{recov}} \sqrt{\frac{\pi}{A_0 \times FOV^3}} \quad (\text{Eq 4.7})$$

4.2.2 Integrated Intensity Regime

The volumes recovered from particle-counting can be integrated with the corresponding intensity information for autocalibration of intensity-based volume assessments. In brief, the average SHG intensity per unit volume can be estimated by integrating the SHG intensity of all the counted particles and dividing by the recovered volume of crystalline material. Once this calibration has been performed using images that are amenable to particle-counting analysis, the integrated intensity can be converted to a crystalline volume in regions where particle-counting is inapplicable. Equation (4.7) rescales the recovered crystalline volume from the SHG intensity measurements to be in a volume fraction of the total field of view. In this manner, the total crystallinity can be determined without the need for preparation of calibration standards of known crystallinity.

From previous theoretical and experimental treatments, SHG microscopy measurements can be considered in three separate regimes depending on the mean particle size, μ_r , relative to the beam waist, W_0 , and the effective interaction length / depth of field, l_{eff}^2 .

Regime I. $\mu_r < W_0$. The mean particle size is less than the width of the beam.

If the mean particle size is less than the width of the beam, then the entire crystal falls within one focal volume. The intensity of that particle scales quadratically with the volume in this limit, and therefore with the sixth power of the particle radius.

$$I_i = K \left[\frac{4}{3} \pi r_i^3 \right]^2 \Omega_i \quad (\text{Eq 4.8})$$

The total mean intensity, $\mu_{I_{total}}$, is given by summation over all the individual average crystallite contributions.

$$\mu_{I_{total}} \cong K \left(\frac{16\pi^2}{9} \right) \mu_n \mu_r^6 \mu_\Omega \quad (\text{Eq 4.9})$$

Given the linear scaling, the relative variance in the recovered volume is identical to the relative variance in the total intensity in Equation 4.9. For an assembly of many randomly oriented particles the total intensity is directly proportional to the total volume. This assumption holds in the limit of large n , such that the average intensity is given by the product of the total number of particles and the average intensity per particle. In this limit, the total SHG intensity scales linearly with the total number of particles, and therefore linearly with the total volume.

$$I_{total} \cong K'' n \mu_V \quad (\text{Eq 4.10})$$

Measurements acquired in this limit correspond to the observation of puncta typically only one or two pixels in dimension.

Regime II. $W_0 < \mu_r < l_{eff}$

In this limit, the mean particle size is larger than the beam waist, but smaller than the effective interaction length (either dictated by the depth of field or the scattering length, l_s). This limit is quite common for pharmaceutical materials acquired with low numerical aperture objectives and long depths of field, such as in the present study. If it is assumed for simplicity that the single pixel area is approximately equal to the cross-sectional area of the beam, the volume of material contributing to the SHG in a given pixel is given by the pixel area and the particle diameter.

$$I_{pixel} = K \left[A_{pixel} 2r_i \right]^2 \Omega_i \quad (\text{Eq 4.11})$$

The total intensity of the i th particle is then given by summing the intensities of every pixel in the cross-section of the particle. Defining the area in terms of pixels allows the total intensity of the i th particle to be estimated by the following expression, in which K' incorporates the squared area of a pixel in its definition.

$$I_i = 4K' (\pi r_i^2) r_i^2 \Omega_i \quad (\text{Eq 4.12})$$

$$I_{total} = \sum_{i=1}^n 4K' (\pi r_i^2) r_i^2 \Omega_i \quad (\text{Eq 4.13})$$

The recovered total volume of the ensemble of particles is given by $V_i = \frac{4}{3} \pi r_i^3$, which combined with Equation 4.12 yields the following expression.

$$V_i = \frac{4}{3} \pi \left(\frac{I_i}{4\pi K' \Omega_i} \right)^{3/4} \quad (\text{Eq 4.14})$$

Although the intensity of a single particle scales nonlinearly with the volume according to Equation 4.14, it also scales linearly with the total number of particles. In the limit of a large number of particles, the intensity is then approximately linearly proportional to the total volume.

$$\mu_{V_{\text{re cov}}} \cong k \mu_{I_{\text{total}}} \quad (\text{Eq 4.15})$$

Regime III. $l_{\text{eff}} < \mu_r$

In the limit of large single-crystalline domains compared to the depth of field and/or scattering length, the intensity measured by a single pixel in a beam-scattering system will be a constant. As such, the total volume will be directly proportional to the number of pixels that intersect the crystal. The total intensity of a single particle within a single z-slice is given by the cross-sectional area multiplied by the effective interaction length l_{eff} .

$$I_{i,z} = K \left[l_{\text{eff}} A_{\text{pix}} \right]^2 \Omega_{i,z} \quad (\text{Eq 4.16})$$

In Equation 4.16, A_{pix} is the cross-sectional area of the beam, which is assumed to be on the order of a single pixel as in Regime 2 (an error in this approximation will primarily result in a simple rescaling, which is implicitly handled by the proportionality constant K). In this regime, the particle is much larger than the beam dimensions, such that the particle contributes to the volume determination over many z-slices. Assuming negligible losses due to scattering and neglecting perturbations from phase-matching, the total integrated intensity will be equal to the intensity per focal volume multiplied by the number of focal volumes required to map the entire crystal. In this regime, the total intensity of a given particle will be directly proportional to the total volume of the particle.

In both intensity-based expressions, the orientation term represents a new source of measurement variance that is absent in the determinations based on image analysis. Unfortunately, the precise analytical form for the probability density function describing the orientation-dependence of the SHG intensity is nontrivially related to the inherent tensor properties of the crystals. For simplicity, we will assume the intensity-dependence can be approximated by an exponential pdf, the relative variance is straightforward to evaluate.

$$\mu_{\Omega} = \frac{1}{\Omega_0} \int_0^{\infty} \Omega e^{-\frac{\Omega}{\Omega_0}} d\Omega = \Omega_0 \quad (\text{Eq 4.17})$$

$$\left(\frac{\sigma_{\Omega}}{\mu_{\Omega}} \right)^2 \cong \left(\frac{\frac{1}{\Omega_0} \int_0^{\infty} \Omega e^{-\frac{\Omega}{\Omega_0}} d\Omega - \mu_{\Omega}^2}{\mu_{\Omega}^2} \right) = \left(\frac{2\mu_{\Omega}^2 - \mu_{\Omega}^2}{\mu_{\Omega}^2} \right) = 1 \quad (\text{Eq 4.18})$$

We can now compare this naïve result with the relative variances for different limiting assumed crystal tensors. Simulations were performed to numerically evaluate the pdf and relative variance due to crystal orientation under several different limiting conditions. In brief, the distributions in total SHG intensity with linearly polarized incident light were generated under several different limiting conditions and compared with the proposed exponential pdf. Overall reasonably good agreement was observed between the simulated histograms and the assumed exponential pdf, with relative variances within spitting distance of unity.

Random numbers for the three Euler twist and azimuthal rotation angles ψ and ϕ were drawn from a uniform distribution. For the polar tilt angle θ , random numbers were produced according to a sinusoidal probability density function using the inverse cumulative density function approach, in which uniformly distributed numbers for the CDF were inverted to produce randomly distributed numbers in θ . For each set of $>10^4$ random Euler angles for X-polarized incident light, the intensity of the corresponding coparallel and cross-polarized SHG was calculated based on the squared magnitude of the laboratory-frame Jones tensor elements for that particular crystal orientation ($|\chi_{xxx}|^2$ and $|\chi_{yxx}|^2$, respectively). The total intensity was generated by the summation of the co-parallel and cross-polarized contributions. As such, the simulations yield signals proportional to the anticipated hyper-Rayleigh scattered intensity collected in the forward direction.

4.3 Experimental Details

Monte Carlo simulations were performed to assess the accuracy of the equations outlined in the theory section. A script written in MATLAB, version R2014a, modeled the amorphous solid dispersion by generating a set number of particles with random positions in a 512 x 512 pixel image. Particles were simulated as spheres to the best approximation with pixels. Particle radii

were randomly selected from a lognormal distribution, using $\mu = 0.693$ and $\sigma = 0.336$ in units of pixels. These numbers were selected to produce images similar to the particle size distribution produced in the experimental data. The intensity of each particle was dictated by two factors: i) an orientational term accounting for the difference in SHG intensity for different crystal orientations, and ii) a term related to crystal size. For the orientational term, the intensity was rescaled by a random number from an exponential distribution, consistent with previous observations of the modeled and measured distribution in intensities for BaTiO₃ nanoparticles (Kestur et al., 2012). Building on previous SHG microscopy analyses (Kestur et al., 2012), the size dependence of the intensity scales differently depending on the size of the particle relative to the depth of field and beam waist. In Regime I, corresponding to particles smaller than the beam waist, the SHG intensity scales with the squared volume of the particle. In Regime II, in which particles are larger than the beam waist but smaller than the depth of field, the individual pixels were scaled quadratically with particle radius. For particles larger than the depth of field, consistent with Regime III, the SHG intensity is independent of particle size and was not rescaled.

The average single particle area relative to the field of view, \bar{A}_0 , is a key factor in the analysis. For the simulations, the value of \bar{A}_0 was determined to be 5.9×10^{-5} (pix²/pix²). To determine \bar{A}_0 for the experimental data, particle areas were averaged from images that showed very little particle overlap namely, thirteen fields of view for the blank sample, thirteen fields of view for the 0.05% crystallinity by mass sample, and 3 fields of view for the 0.10% crystallinity by mass sample. See Materials section for details on the preparation of these samples. \bar{A}_0 for the experimental data was determined to be 4.2×10^{-5} .

Evacetrapib was obtained from Eli Lilly and Company (Indianapolis, IN). Evacetrapib tablets (unit formula shown in Table 1) prepared identically to the final dosage form. The actual tablets which contain 50% drug in the sodium salt form were used to make the standards. No crystal forms are known for the sodium salt. The free-form (protonated) of the crystalline drug is then spiked into the ground tablet powder by light mixing. Tablets were confirmed to contain amorphous sodium salt after grinding and prior to mixing by PXRD and FTIR analysis. Spiked standards of crystalline drug substance in the drug product were prepared by adding known amounts of crystalline drug substance to the ground evacetrapib tablets powder followed by additional mixing in the mortar and pestle. The crystalline evacetrapib used for spiking was

roughly 98-100% crystalline. Apart from the API, none of the intragranular or extragranular excipients produces substantial SHG. Although microcrystalline cellulose produces some weak background SHG from locally ordered domains (Schneider et al., 2012), all the bulk crystalline materials adopt centrosymmetric and SHG-inactive space groups.

Table 1. Unit Formula for Evacetrapib Tablets

	Amount (mg/tablet)	Function
Crystalline Evacetrapib	0	API
Intragranular Ingredients		
Spray Dried Dispersion	260.0	Active dispersion
Evacetrapib Free Acid	--	Crystalline active
Hypromellose	--	Dispersion polymer
Croscarmellose Sodium	13.00	Disintegrant
Microcrystalline Cellulose	97.50	Filler
Sodium Chloride	84.50	Disintegrant
Sodium Bicarbonate	84.50	Disintegrant
Colloidal Silicon Dioxide	2.600	Glidant
Magnesium Stearate	0.8125	Lubricant
Extragranular Ingredients		
Croscarmellose Sodium	13.00	Disintegrant
Microcrystalline Cellulose	93.275	Filler
Magnesium Stearate	0.8125	Lubricant
Tablet Core Weight	650.0	
Film Coating Ingredients		
Color Mixture	26.00	Film Coat
Coated Tablet Weight	676.0	

Infrared spectra of the powders (~20 mg) were obtained in absorbance mode using a Nicolet 6700 FT-IR (ThermoFisher Scientific, Waltham, Massachusetts, USA) equipped with a globar infrared beam source, KBr beam splitter, DTGS detector, and a Smart OMNI sampler attenuated total reflectance (ATR) module with a germanium crystal and low-pressure sample tower. The scan range was set from 500 to 4,000 cm^{-1} with 2 cm^{-1} resolution, and 64 scans were co-added. All spectra were acquired and altered using OMNIC 8.3 software package. An automatic baseline correction and ATR correction were applied to the spectra when plotted and all spectra were normalized to the peak at approximately 1560-1570 wavenumbers cm^{-1} .

The powder samples were packed into a 25 mm diameter cavity steel containment cell ensuring a flat level powder surface was obtained. The powder X-ray diffraction (PXRD) data

were obtained using CuK α radiation with a Bruker D4 Endeavor powder diffractometer (Bruker AXS, Madison, Wisconsin, USA) operating at 40 kV and 40 mA. The X-ray measurements were conducted in Bragg-Brentano geometry using a scan range of 4–40° 2 θ and 6–14° 2 θ at scan time per step (step size 0.028°) of 0.5 and 10 seconds, respectively. The diffraction data was baseline corrected utilizing the Bruker software.

SHG images were acquired using a modified SONICC (second order nonlinear imaging of chiral crystals) instrument from Formulatrix (Formulatrix, Bedford, Massachusetts, USA), following custom modifications performed in-house to enable epi-detection of powders. In brief, a polarizing dichroic mirror was added to the beam path of the Formulatrix system in the backward-collected direction, which was focused onto the entrance of a multimodal liquid light guide (Newport, Irvine, California, USA) terminating in an additional narrow bandpass 530 nm dichroic mirror (Semrock, Rochester, New York, USA) designed to separate the SHG from two-photon excited ultraviolet fluorescence (TPE-UVF). Use of a polarizing dichroic mirror retained direct compatibility with two photon excited ultraviolet fluorescence (TPE-UVF), which requires 532 nm incident light to pass through the same dichroic mirror used to reflect the epi-detected SHG of the same wavelength. The 1064 nm fundamental beam was generated by a Fianium (Fianium, Southampton, UK) fiber laser with a pulse repetition rate of 50 MHz, and pulse width of ~150 fs. An image was created by scanning the incident beam across the sample to generate a 512 \times 512 pixel image. The laser was focused onto the sample using a 10X objective with a depth of field of 200 μ m. For each sample investigated, approximately 10 mg of powder was placed in a DSC bottom pan set into wells of a 96 well plate.

The particle-counting analysis was performed using the “Analyze Particles” built-in algorithm in ImageJ (Brown et al., 2003) to generate lists of particle numbers, areas, and average intensities. Reduction of background was accomplished by applying size and intensity thresholds for counting. The size threshold used excluded anything less than two pixels and the intensity threshold used excluded any pixel with intensity less than three counts.

4.4 Simulations

Analysis of the simulation data shown in **Figure 4.1** demonstrated quantitative agreement between the ground truth number of particles (black line) present in an image and the number recovered by the statistical model (blue diamonds). This high degree of correlation was maintained

while spanning the regimes of both low (<1%) and high crystallinity (>1%), for which representative images are provided within the figure. The expression in Equation (4) significantly extends the range over which the true underlying mean number of particles can be recovered prior to introducing bias. For reference, particle-counting performed without bias-correction indicated clear deviation between the naively measured versus bias-corrected crystallinity, introducing significant systematic errors for ~1000 particles (corresponding to ~6% of the area occupied by particles with $\bar{A}_0 = 5.9 \times 10^{-5}$).

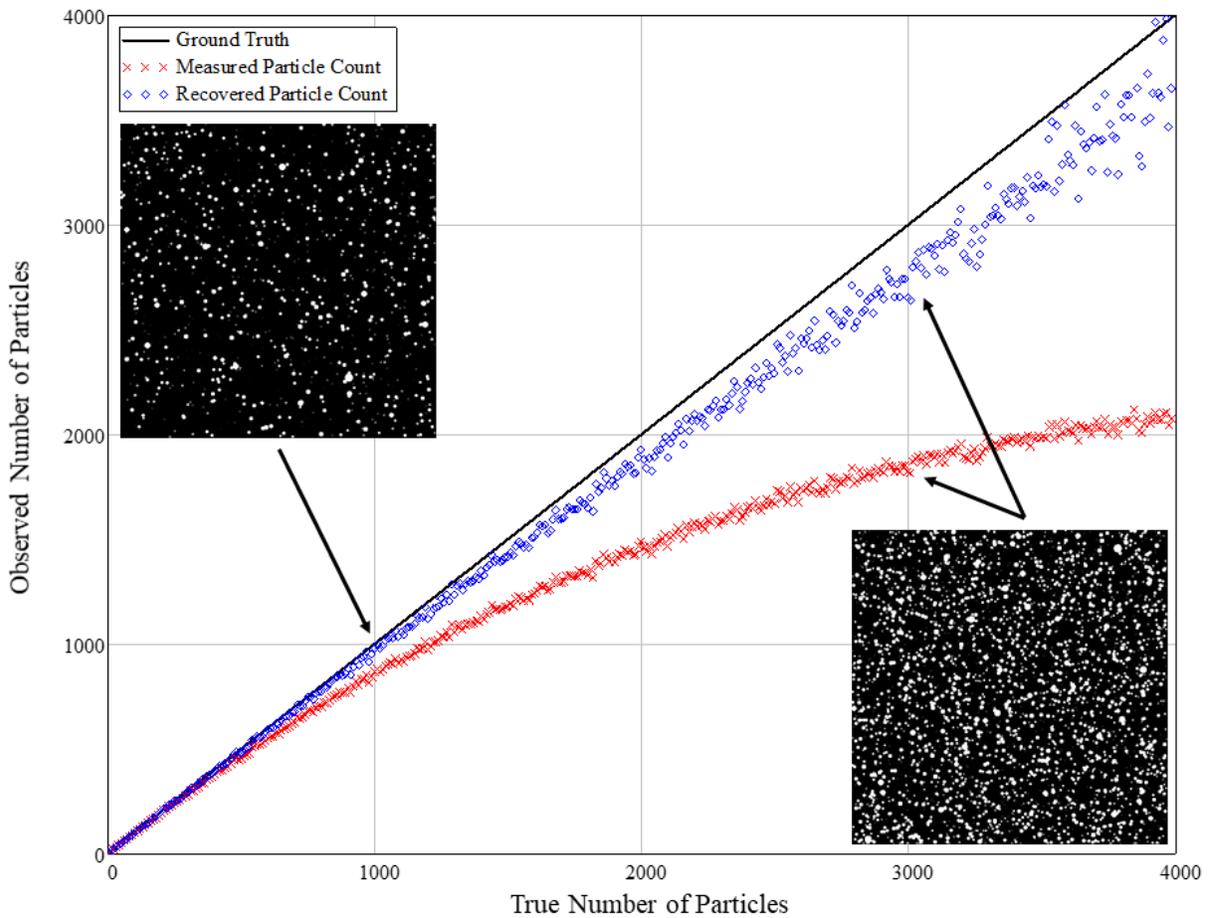


Figure 4.1: Red data points are the raw data measured from particle-counting simulations. Each point represents a simulated image. True particle counts range from 1 particle to 4000 particles per image. Blue data points are the estimated true particle count μ_n as estimated by Equation 4.4. Inset images are simulated images of 1000 and 3000 particles per image.

The most compelling outcome is the reliability in the volume recovery. Both the volumes measured prior to bias-correction (red x's) and the volumes recovered (blue diamonds) are plotted alongside the ground truth particle volumes in **Figure 4.2**. Notably, the volumes calculated using naïve analysis techniques resulted in significant bias in the crystallinity that was largely removed by the correction described in the Theory section. For particle densities higher than ~ 2100 particles per image, the probability of $\overline{A_0\mu_N} > 0.1$ becomes significant, such that Equation (4.4) is no longer applicable. In this case, the probability of three or more particles overlapping is too great for the corrected counting algorithm to reliably operate. The bias introduced by the bias-corrected counting algorithm is small throughout this entire operational regime. However, the increase in relative variance at high crystalline fractions is indicative of the anticipated reduction in the SNR associated with larger correction factors.

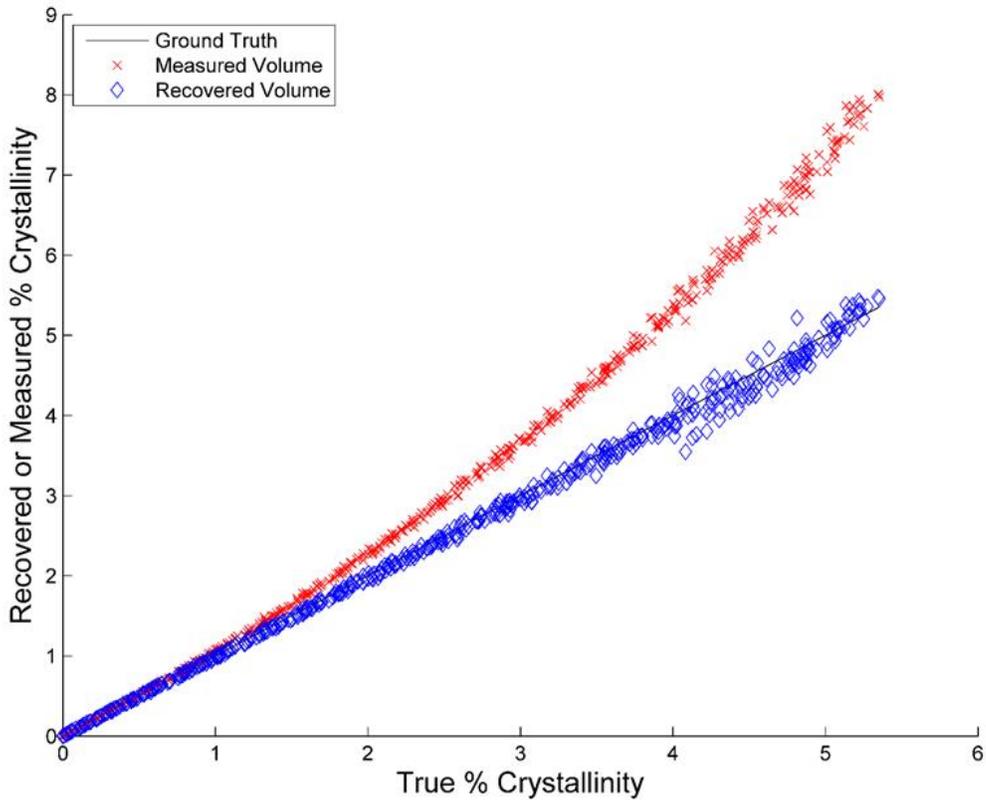


Figure 4.2: Each point represents the volume determination of a simulated image. Red x's are the raw measured volume. Blue diamonds are the recovered volumes using Equation 4.3. The solid black line is the ground truth value.

A theoretical comparison of the predicted SNR from the analytical modeling for the counting and integrating approaches is provided in **Figure 4.3**. In the hypothetical scenario of particles of identical size and brightness, the signal to noise ratio is dictated solely by the Poisson statistics describing the number of particles in each field of view, in which the variance is equal to the mean. All results in **Figure 4.3** are rescaled by this theoretical limiting result in order to better highlight the differences. The upper limit (dash-dot) corresponds to the SNR produced considering only the Poisson distributed number of particles and the inherent variance in their cross-sectional areas. As the number of particles increases, rescaling by the bias-correction extends the dynamic range, but reduces the SNR achievable by particle-counting. By comparison, the relative SNR from integration in both Regimes II and III is independent of the mean number of particles. In Regime III, the reduction in SNR arises from the additional variability in intensity due to crystal orientation that does not directly impact particle-counting approaches. In Regime II, additional variance in intensity from differences in particle thickness further reduces the SNR. In the present case, the microscope has an effective field of view of 1.4 mm x 1.4 mm and the particles probed had an average diameter of 10.4 μm . In the commonly encountered Regime II, the SNR of particle-counting is increased by >3.5-fold compared to intensity integration in the low crystallinity limit, corresponding to a >10-fold reduction in measurement time for a comparable SNR. Over the entire accessible range of volumes considered in the simulations in all particle size regimes, particle-counting resulted in higher SNR for the recovered volume relative to intensity integration, although the advantage is reduced as the particle density increases.

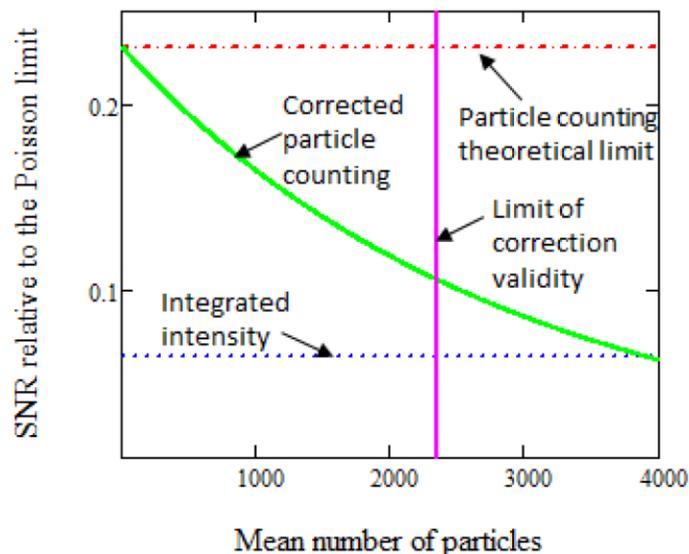


Figure 4.3: Comparison of the predicted signal to noise ratios for volume determination by particle-counting versus integrating the SHG intensity, normalized by the inherent variance in the number of particles from Poisson statistics (i.e., variance equal to the mean number of particles). All calculations were performed for a mean particle area relative to the field of view $\bar{A}_0 = 5.9 \times 10^{-5}$. The vertical bar indicates the upper limit for the particle-counting correction.

In addition to the improvement in the signal to noise in the low crystallinity regime, the particle counting approach described herein has the distinct advantage of minimizing potential complications associated with bias from crystal polymorphism. All crystals of homochiral molecules by necessity must adopt noncentrosymmetric space groups. With the rare exception octahedral symmetry, all other space groups of homochiral molecules are symmetry-allowed for SHG. However, the relative brightness of the SHG-activity can vary significantly depending on the crystal form. If crystals initially adopt one crystal form, then convert to a different form (e.g., as per the Ostwald rule of stages), measurements based on the SHG-activity would result in a bias from the change in crystal form. In contrast, the particle counting approach described herein is based exclusively on the cross-sectional area irrespective of the integrated SHG intensity, removing bias from variance in crystal form.

4.5 Application to Crystalline Volume Determination within a formulated ASD

Physical mixtures of crystalline evacetrapib drug substance in an amorphous solid dispersion (ASD) formulated tablet were made to assess the merits of the proposed auto-calibration algorithms, the SHG images of which are shown in **Figure 4.4**. Thirteen fields of view were collected for the lowest two concentrations of the physical mixture (the 0% and 0.05%) and three fields of view were collected for the rest of the concentrations (0.1%, 0.25%, 0.5%, 1%, 2.5%, 5%, and 10% crystallinity by weight). The higher number of measurements performed at lower concentrations was chosen to reduce the Poisson uncertainties from the finite number of particles present at those concentrations. The corrections for the particle count, the total volume, and the volume fraction were performed using the particle-counting correction algorithm. The limits of detection were calculated for both the integrated intensity and the bias-corrected particle-counting method, they were found to be 0.4% and 0.17% (1700 ppm) respectively.

In the present study, the size distribution of crystals was consistent with Regime II based on measurements of both the beam waist and depth of field. In brief, the beam waist was smaller than a single pixel, such that all particles counted spanning ≥ 3 pixels were large relative to the beam waist. The depth of field was measured to be 200 μm , which was significantly greater than the upper limit of the size distribution observed by particle counting. This regime is the most likely to be encountered in analyses of pharmaceutical materials; particles much smaller than the focal volume produce relatively weak SHG and appear as puncta in images, while particles larger than the depth of field are of a size compatible with numerous alternative conventional analyses. In Regime II, the effective interaction length is dictated by the crystal size, such that changes in the crystal size distribution can potentially influence the autocalibration. Fortunately, knowledge of the crystal size distribution from microscopy measurements provides a route to analytical model and correct for such effects. (Chowdhury et al., 2016)

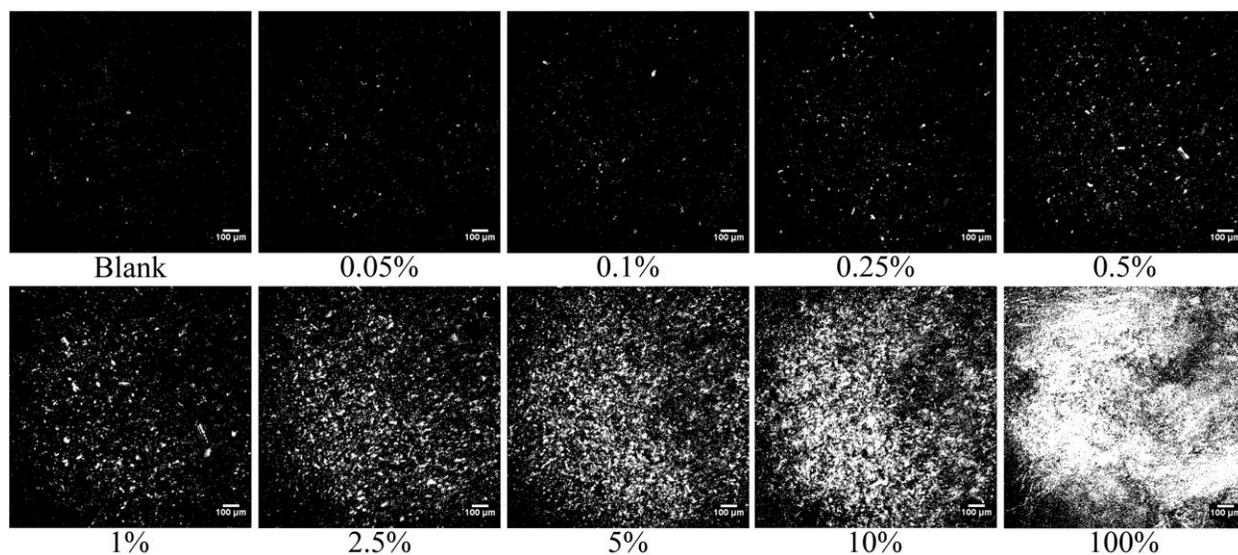


Figure 4.4: Shows representative images of the different physical mixtures of spiked crystalline evacetrapib in the ground ASD tablets. Some of the excipients are SHG-active as seen in the blank. All concentrations except the 100% were used for the analysis. The 100% image is included here for reference only.

The limits of detection by SHG were significantly higher than those observed in previous studies of model systems, dictated largely by weak but nonzero background contributions from the excipients. The most likely source of the background SHG is the microcrystalline cellulose, from which weak SHG-activity has been reported previously. (Toth et al., 2015) Native cellulose as generated biologically has been shown to be SHG-active. (Brown et al., 2003) While the cellulosic materials commonly used as excipients in pharmaceutical formulations generally undergo substantial treatments prior to use, residual weak SHG from a sparse population of locally ordered domains may potentially contribute to the observed background. Analysis of SHG active areas by complementary methods such as Raman spectroscopy or x-ray diffraction, previously successfully demonstrated by Schmitt et al. and Newman et al. respectively, may provide routes for more definitively determining the chemical composition of the SHG-active domains arising within the excipients.

Despite the advantages of both noise reduction and the absence of independent calibration standards, several application spaces are not yet well suited for the particle-counting autocalibration approach described herein because of the assumptions made in the model. One assumption of the model is that the particle size distribution does not change substantially over time, with increase in crystallinity arising primarily from interconversion on a per-particle basis.

In cases in which the sizes of individual crystals evolve during the analysis, (e.g., through molecular diffusion, crystal growth, and/or Ostwald ripening), additional parameters related to the anticipated evolution of the size distribution could be integrated into the approach, but are not described herein. Furthermore, the model is currently limited exclusively to spheroidal particles. In practice, many APIs crystallize in needle-like high aspect ratio crystal, for which the relationship between cross-sectional area and volume is not as concise as described herein. Extension of the model to crystals with a greater diversity of morphologies is still in progress.

It should be noted that the method described recovers the volume fraction by image analysis, but not the mass fraction. Independent knowledge of the void volume in the probed powder would enable conversion of the volume fraction to the mass fraction. In the present case, the correction for void volume was performed using the known mass fraction as a calibrant. In practice, gravimetric or volumetric methods could be used to determine the void volume and recover the mass fraction.

The observed ~2.5-fold improvement in the limit of detection is in excellent agreement with the 3.5-fold enhancement predicted by the statistical modeling. The particle sizes encountered in this study were consistent with Regime II, in which the particles were large relative to the beam diameter, but small relative to the long depth of field of the 10x objective. The agreement between theory & experiment supports the reliability of both the statistical modeling and the simulations.

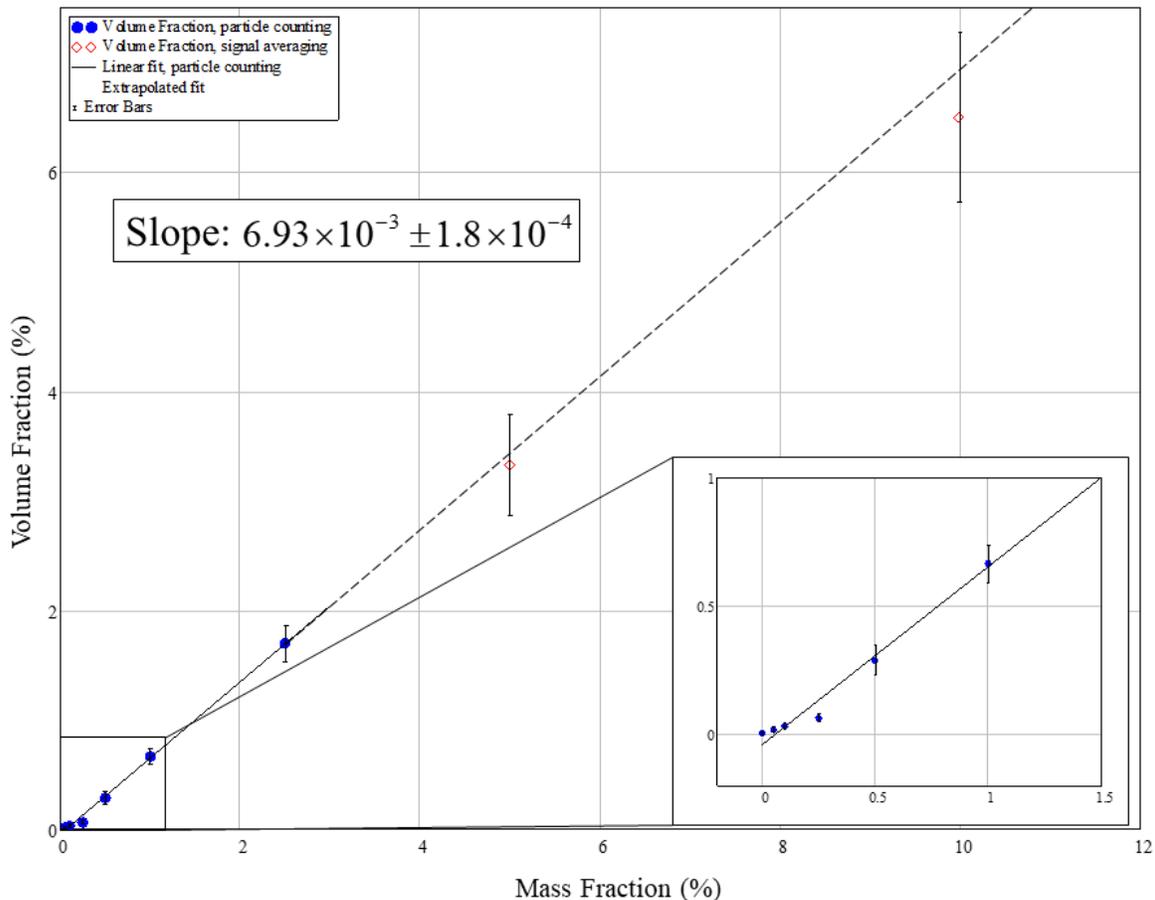


Figure 4.5: Image Analysis of SHG micrographs allowed determination of crystalline fraction of evacetrapib in formulated tablets. The recovered volume obtained by particle counting is plotted in blue (0.05% mass fraction to 2.5%). The linear least squares fit is calculated using only the solid blue data set. The red diamond data set shows the recovered volume fraction from the intensity measurements obtained by autocalibration performed from the measurements indicated in the solid blue points, Equation 4.6.

In addition to removing bias and increasing SNR in the low crystallinity limit, the particle-counting algorithm has the distinct advantage of enabling autocalibration. In the low crystallinity regime, the volume fraction is determined independently of the SHG activity. As such, the SHG activity per unit volume determined in the low crystallinity limit can serve to calibrate extension to higher crystallinity beyond the point at which particle-counting remains reliable. Results from the counting algorithm allowed for the determination of both $\overline{A_0}$ and the SHG per unit volume, which were in turn used for autocalibration to extend the dynamic range beyond the particle-counting regime. For the experimental data of this study $\overline{A_0}$ was determined to equal 4.2×10^{-5}

and the SHG intensity per unit volume was found to equal 3 counts per cubic pixel. In the high crystallinity limit, the total SHG intensity is divided by the determined SHG per unit volume to give the equivalent volume, and that volume is rescaled to the volume fraction using Equation (4.6). In **Figure 4.5** the high crystallinity points in red (5% and 10%) agree quite well with the extrapolated calibration curve determined by the low crystallinity blue points. Two important caveats are noteworthy: i) uncertainties in the slope will generally compound upon extrapolation, and ii) the average SHG/V calculated in this manner is ensemble-averaged over the entire field of view and may not be representative of all positions within the image for non-uniform illumination (e.g., in the presence of vignetting).

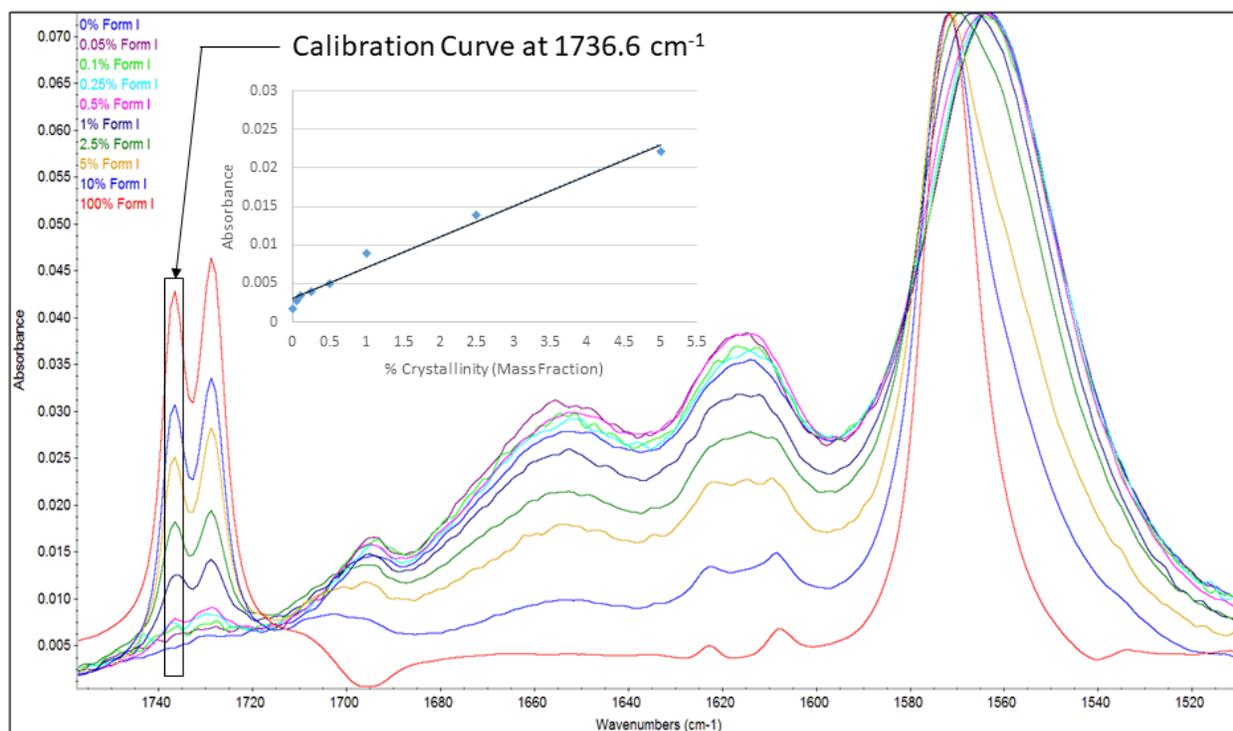


Figure 4.6: The FT-IR absorbance of different concentrations of evacetrapib. The 1736.6 cm^{-1} peak was used for quantification (shown in the black box). Limit of detection using this technique was determined to be 0.8%.

The same samples were analyzed by a suite of complementary methods for comparison with the SHG microscopy approach. The data from FT-IR absorption spectroscopy are shown in **Figure 4.6**. The peak at 1736.6 cm^{-1} (highlighted with a black box in **Figure 4.6**) was chosen for quantification of the crystalline fraction. A calibration curve gives the limit of detection at 0.8%

using the peak height. The same samples were also analyzed using PXRD, the baseline corrected data is shown in **Figure 4.7**. The strongest peak at 12.05° was used to make a calibration curve based on peak height. The inset in **Figure 4.7** shows the calibration curve that was used to determine the limit of detection. A comparable limit of detection of 0.9% was obtained using the peak height for quantification of crystallinity. In all cases considered, the proposed particle-counting approach provides a clear advantage for lowering the limits of detection.

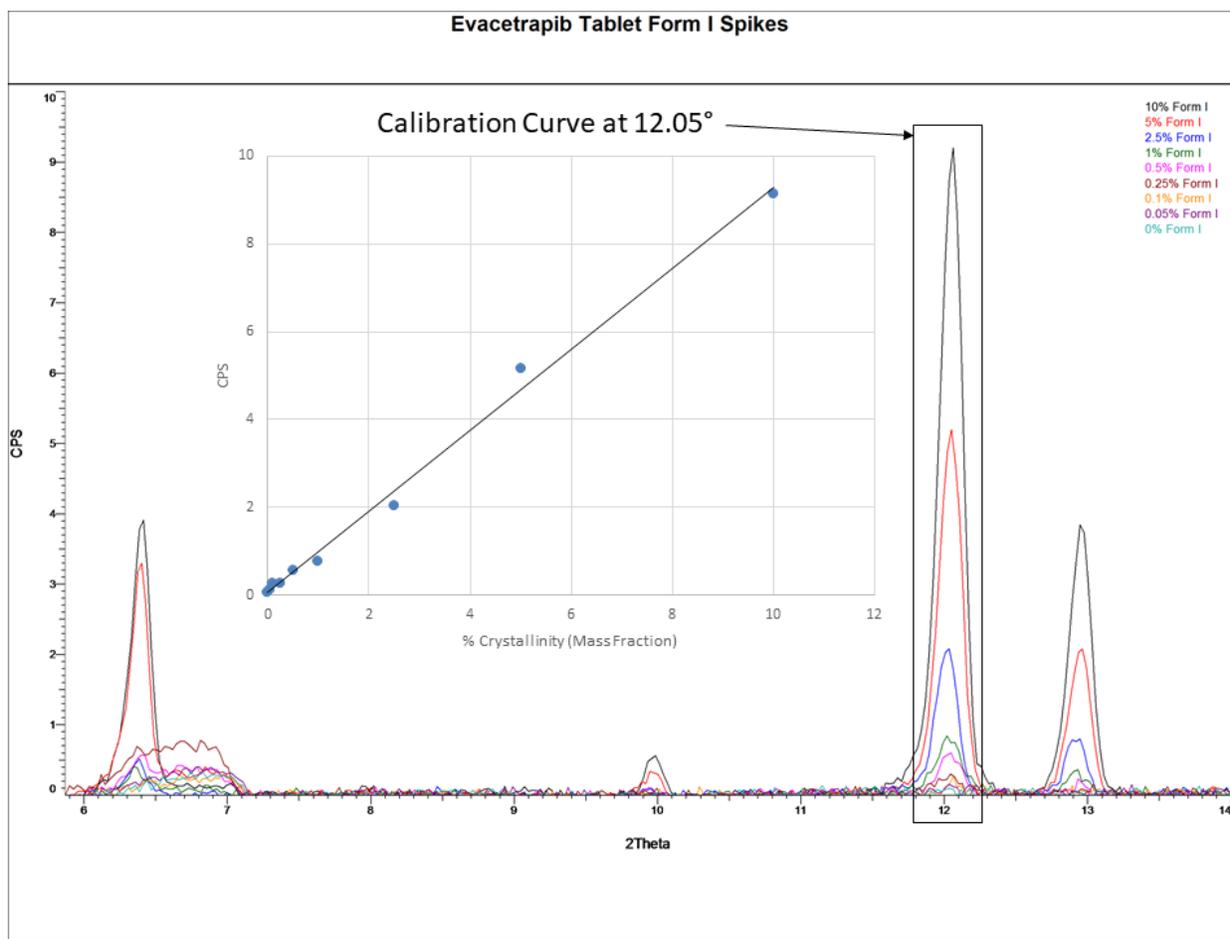


Figure 4.7: The baseline corrected spectra from PXRD for various concentrations of crystalline evacetrapib in drug formulation. The area of the most intense peak, located at 12.05° 2θ , was used for evaluation of the data. Limit of detection was determined to be 0.9%.

4.6 Conclusions

A statistical model for removing bias due to particle overlap and for improving the SNR of particle-counting in the regime of low crystallinity was presented and validated both theoretically

and experimentally. The analytical expressions derived herein provide a correction term for volume determinations from particle-counting, in addition to allowing for the calculation of total crystallinity based on SHG intensities from the particle-counting regime without needing a known calibration sample. It is further possible now to define a transition point based on the relative variance in volume such that the crystalline volume is optimally evaluated by particle-counting below this point and by the integrated SHG intensity above it. Data calculated by applying this model to simulations of crystalline samples were shown to reliably recover the true particle count and crystallinity volume, notably extending the dynamic range of particle-counting methods by approximately an order of magnitude. The model was further validated using experimental data and was shown to successfully recover the particle count, total volume, and the volume fraction from SHG images of an amorphous solid dispersion of evacetrapib. The approach described here further offers the advantage of autocalibration by enabling the use of the SHG intensity per unit volume measured by particle-counting in the low crystallinity regime for the evaluation of the volume in the high crystallinity regime. This statistical model significantly improves the accessibility of quantitative crystallinity information from nonlinear optical imaging, laying the foundation for broader adoption in pharmaceutical analyses.

4.7 Acknowledgements

The authors gratefully acknowledge funding through the Lilly Research Awards Program, grant #107309, and from the NSF Chemical Measurement and Imaging (CMI) Grant Opportunity for Academic Liaison with Industry (GOALI) award, grant #209682. We also collectively acknowledge Formulatrix for extended loan of a SONICC instrument and support of its modification for powders analysis.

CHAPTER 5. STOCHASTIC DIFFERENTIAL SCANNING CALORIMETRY

5.1 Introduction

In pharmaceutical drug development, drug substances and formulations with long-term physical and chemical stability ensure full optimized of a drug when administered to the patient. During the time consuming process of formulations development (estimated >3 months and \$1 million) it is thus critical to employ methods that inform on possible avenues of reduced bioavailability.(Strovel et al., 2016) Identifying possible failure points early can reduce the changes of late-stage failure, which can be costly and time consuming to remediate. For active pharmaceutical ingredients administered as crystalline formulations, phase transitions to alternate crystal forms can have deleterious effects on aqueous dissolution rates, affecting both oral and parenteral bioavailability.(Cheney et al., 2010) Administering these drugs as amorphous solid dispersions or as liquid formulations does not altogether avoid the problem, as spontaneous crystallization can occur during storage or even *in vivo* under certain conditions.(Andronis et al., 1997; Baghel et al., 2016; Jackson et al., 2015; Konno & Taylor, 2008; Newman et al., 2012; Rumondor et al., 2009; Shamblin & Zografis, 1999) Phase transitions can dictate both direct physical depletion of bioavailable drugs through transitions to non-soluble crystal forms, as well as susceptibility to chemical depletion through transitions to more reactive crystalline, liquid or glassy states. There is therefore a need in the pharmaceutical industry to thoroughly map the phase space of potential drug candidates, excipients, and mixtures.

Several thermal methods exist for characterizing phase transitions within pharmaceutical powders. Differential thermal analysis applies equal heat to reference and sample pans and measures the temperature differential between the two pans to detect thermal events associated with phase transitions. Thermal gravimetric analysis (TGA) continuously measures the mass of a sample while varying temperature. TGA is especially useful for dehydration and decomposition. Differential scanning calorimetry (DSC) enables the detection of phase transitions as a function of the energy differential between reference and sample pans heated to the same temperature. DSC is a widely used technique for phase transition characterization because of its sensitivity, ease-of-use, and short measurement time.(Clas et al., 1999)

Despite these advantages, DSC probes the energy transfer into the full ensemble of the sample, averaging over all particles in the formulation. Phase transitions for ensembles of crystals are predicted to be stochastic, with nucleation being the rate-limiting step on a per-particle basis.(Khaliullin et al., 2011; Oxtoby, 1998; Sleutel et al., 2014) If phase transformation is stochastic, individual transient events may be obscured or missed in ensemble-averaged analysis. Additionally, single-particle measurements can enable the accurate modelling of kinetics for phase transitions, allowing the distinction between different mechanisms of transition (i.e., concerted or continuous).(Liu et al., 2014)

Optical microscopy measurements of changes in gross morphology coupled with DSC can recover information on single-particle phase transformations. Reading and coworkers have developed an image analysis algorithm dubbed thermal analysis by structural characterization (TASC) to detect phase transitions and the spatial distribution of phases from images acquired using hot-stage microscopy.(Alhijaj et al., 2018; Alhijaj et al., 2015; Alhijaj et al., 2017; Reading, 2017; Reading et al., 2014) TASC measures the subtle difference between an original image and subsequent images to quantify morphological changes within a region-of-interest as temperature is varied. The algorithm also accounts for the possibility of translation of the region-of-interest within the field-of-view (FoV). TASC is a fast and inexpensive method for characterizing the spatial distribution of phase transformations. However, TASC methods rely exclusively on morphological changes in the optical response, which are only intrinsically related to a subset of all possible phase transformations. Solid / solid phase transformations are largely inaccessible by TASC, and glass / solid phase transformations can be challenging to reliably detect without perturbing the sample. Desolvation (including dehydration) or polymorph transitions in the solid state are typically not accompanied by obvious changes in the gross morphology of a particle but can profoundly influence physio-chemical characteristics that are accompanied by thermal transients in DSC measurements. In addition, structural re-arrangements between polymorphs with similar bulk free energies are quite challenging to detect by either conventional DSC or bright-field microscopy. Coupling DSC with concurrent imaging capabilities that are more directly tied to crystal form would help address these ambiguities associated with conventional bright-field microscopy and DSC.

In this work, we demonstrated the integration of second harmonic generation (SHG) microscopy with DSC analysis to monitor phase transitions on a per-particle basis. This

multimodal approach enabled simultaneous single-particle analysis for distinguishing the impact of crystal size and shape on phase transitions in a single experiment, complementing the ensemble-averaged information obtained from DSC. Furthermore, the strict symmetry requirements for SHG allow the mapping of complex desolvation and crystal form transitions with sensitivity in the ppm regime.(Kestur et al., 2012) Previous work has utilized the exquisite sensitivity of SHG microscopy to internal structure of the lattice in non-centrosymmetric crystals as a standalone tool for polymorph discrimination, calibration-free quantification of trace crystallinity, and protein crystal centering at X-ray beamlines.(Chowdhury et al., 2014; Scarborough et al., 2017; Casey Jake Smith et al., 2018) Raman spectroscopy, terahertz spectroscopy, and SHG have all previously been used to characterize crystallinity, polymorphism, and phase transitions.(Francis et al., 2018; Galland et al., 2009; Heinz et al., 2009; Simon et al., 2015; Strachan et al., 2007; Strachan et al., 2005; Zeitler et al., 2007; Zeitler et al., 2006) The merits of integrated SHG and DSC measurements are explored for a model system to then explore the complex phase transitions in trehalose dihydrate.

5.2 Experimental

5.2.1 SHG Microscopy

The experimental apparatus is depicted in **Figure 5.1** and consists of a home-built SHG microscope integrated with an optical DSC stage (Linkam, DSC450). A tunable 80 MHz, Ti:sapphire, femtosecond laser (Spectra-Physics, Mai Tai) was used for the incident light source. The fundamental beam was raster-scanned across the sample using a resonant scanning mirror at 8.8 kHz (EOPC) for the fast-scan axis and a galvanometer mirror (Cambridge-Tech) for the slow-scan axis. A 4x, 0.1 NA objective (Nikon) was used to focus the beam onto the sample, and the SHG signal was collected in the epi direction through the same objective used for delivery of the fundamental beam. The laser was tuned to 800 nm with a power of 80–120 mW at the sample. Two long-pass dichroic mirrors (Chroma, 650DCXR) and a band-pass filter (Chroma HQ400/20M-2P) were used to isolate the 400 nm SHG signal before it was detected by a photomultiplier tube (PMT) (Hamamatsu, H7422P-40 MOD). Responses of the PMT were digitized synchronously with the laser pulses by using a digital oscilloscope card (Alazar Tech, ATS9350) and mapped onto 512×512 images via custom software written in-house

(MATLAB).(Muir et al., 2014) The SHG videos were recorded at 17 frames per second for urea samples and at eight frames per second for trehalose samples.

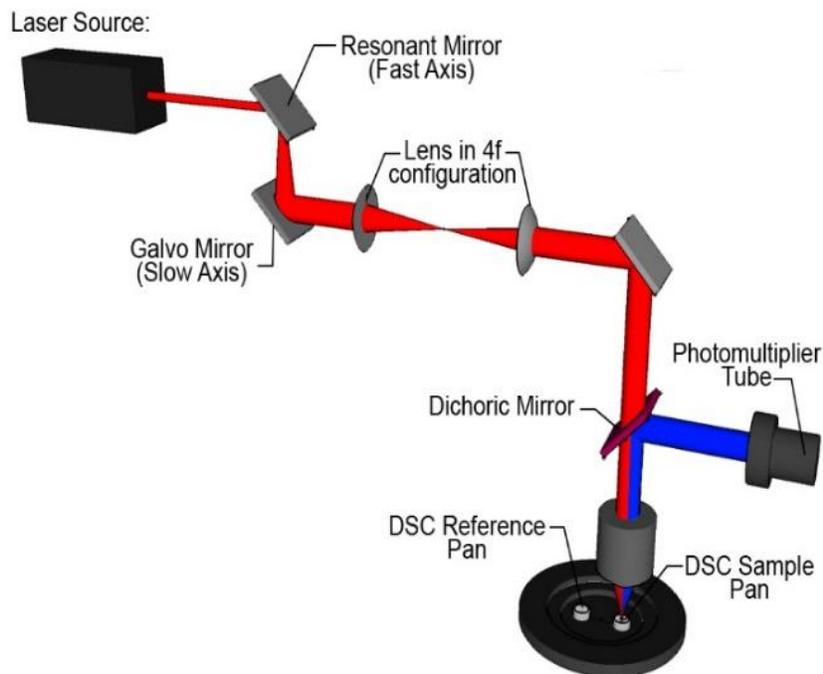


Figure 5.1: Schematic of the integrated SHG-DSC microscope. A pulsed laser beam (800 nm, 80 MHz) was scanned with a galvo-resonant scan pair, which was 4f coupled to the back of an objective and focused onto a DSC sample pan. The SHG signal was collected in the epi direction, isolated with a dichroic mirror and detected with a photomultiplier tube.

5.2.2 DSC measurements

Stochastic differential scanning calorimetry (SDSC) measurements were acquired by integration of the home-built SHG microscope with a Linkam Optical DSC450 stage. The DSC temperature ramp range was 125–145 °C for the urea samples and 80–250 °C for the trehalose samples with ramp rates of 10 °C/min and 20 °C/min, respectively. Before the data collection, an isothermal hold time of two minutes for both urea and trehalose ensured a reliable starting temperature for all experiments to reduce the effects of initial temperature variation. Temperature and input power were measured by the DSC at a rate of five Hz. Standard aluminum sample pans were used in all DSC experiments, and were not sealed to allow for optical access of the sample. The DSC was calibrated using an indium standard in an open pan.

5.2.3 Data analysis

ImageJ (NIH) was used to perform single-particle analysis on the images acquired with the integrated SHG-DSC microscope. The single-particle SHG areas were measured by calculating the fraction of pixels above a threshold within the region of interest.

5.2.4 Sample preparation

Supersaturated solutions of urea (Sigma-Aldrich) were prepared in DI water solution and rapidly recrystallized in an ice bath. The water was decanted, and the crystals were left to dry overnight. Crystals were then passed through a 63 μm mesh into the aluminum DSC sample pan.

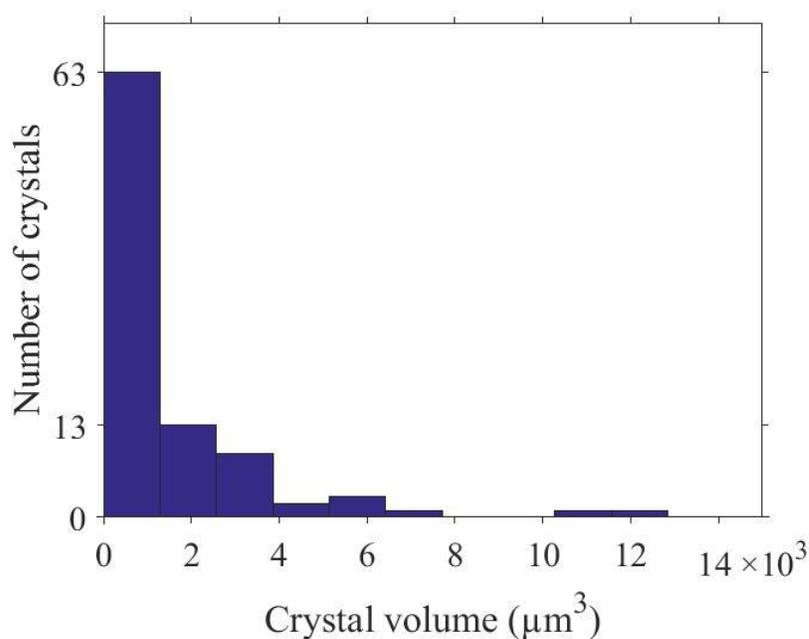


Figure 5.2: Distribution of crystalline urea volumes measured via image analysis on the SHG microscope.

D-(+)-trehalose dihydrate purchased from Sigma Life Sciences was used as received.

5.3 Application to Urea Crystals

Initial proof of concept studies for SDSC were performed using urea, which undergoes a simple single-stage solid/liquid phase transformation. **Figure 5.2** shows the DSC trace for the

melting transition of a urea sample (black line) overlaid with single-particle SHG area (colored lines). The SHG areas of each particle were normalized to the initial SHG area. The DSC trace indicates that the melting transition occurred at ~ 135 °C. This temperature is slightly higher than the literature value of 133 °C. This difference is attributed to a non-negligible heat transfer time to the sample. The need for optical access required the use of an unencapsulated sample pan, resulting in slower heat transfer. (Gaisford et al., 2016) Despite some baseline drift in the per-particle SHG activity from particle motion, the single-particle SHG area shows much sharper transitions from the SHG-active crystalline form to the SHG-inactive molten liquid of urea indicated by the DSC trace, consistent with independent, stochastic melting events. Single-particle analysis yields a mean phase transition duration of 1.0 ± 0.6 s/particle while the DSC melting peak has a full-width-half-maximum of 3.3 s. This difference indicates that the rate of phase transformation was dictated largely by the rate for seed formation of the liquid-state within individual particles, followed by rapid phase transformation within a given particle. **Figure 5.3** shows representative frames from a video acquired with the SHG microscope during the melting of urea. From the SHG results, the nucleation rate at the phase transformation temperature can be estimated. Approximations for the per-particle mass using the density of urea and volume estimated from particle cross-sectional area yields a peak nucleation rate of $1.6 \pm 0.2 \times 10^4$ nuclei per second per milligram (based on the observed rate of 6.13 nuclei/s in a FoV with an estimated 0.39 μg of particulate mass; standard deviation was determined based on Poisson statistics for the number of crystals within the FoV).

The nucleation rate is calculated based on the number of phase transformations observed in the FoV per unit time and the estimated mass of the crystals in the FoV. The estimated mass of the crystals comes from the known density of urea and the estimated volume of the crystals, which is calculated from the crystal cross-section area (assuming rod-like shape). The distribution of crystal volumes in the FoV are shown in **Figure 5.2**. A total mass of 0.39 μg was calculated from the sum of these volumes multiplied by the density of urea (1.32 g/cm^3). The mass combined with the observed nucleation rate of 6.13 nuclei/s yields the nucleation rate of $1.6 \pm 0.2 \times 10^4$ nuclei per second per milligram. The uncertainty in the nucleation rate arises from the number of crystals counted in the FoV, which is dictated by Poisson statistics. The standard deviation, σ of a Poisson distribution with a mean, μ is equal to $\sqrt{\mu}$. We assume the mean of the distribution is approximately equal to the number of crystals counted (58), yielding a standard deviation of ~ 7.6 .

When propagated into nucleation rate, it results in an uncertainty of $\pm 0.2 \times 10^4$ nuclei per second per milligram.

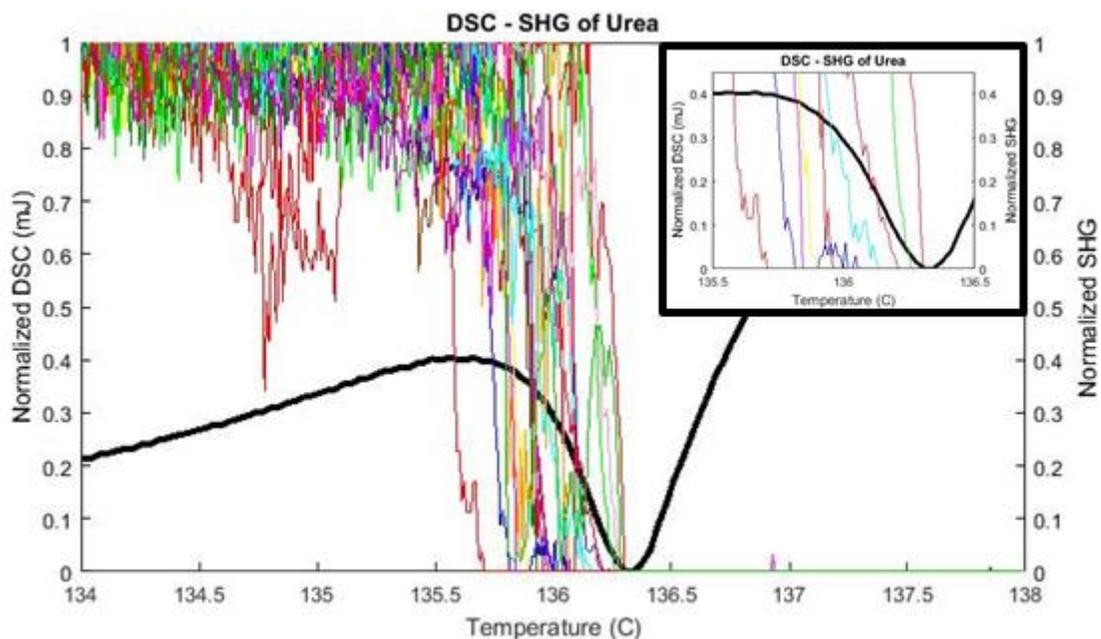


Figure 5.3: SDSC analysis of urea by combined SHG microscopy and DSC. The solid black line is the DSC trace of the melting of urea with endothermic direction being down (left-axis). The thin colored lines show the normalized SHG intensity of individual urea particles (right-axis), acquired concurrently with DSC measurements. DSC and SHG measurements both indicate a structural transformation in the same temperature range. The inset shows a zoom-in of the temperature range 135.5-136.5°C.

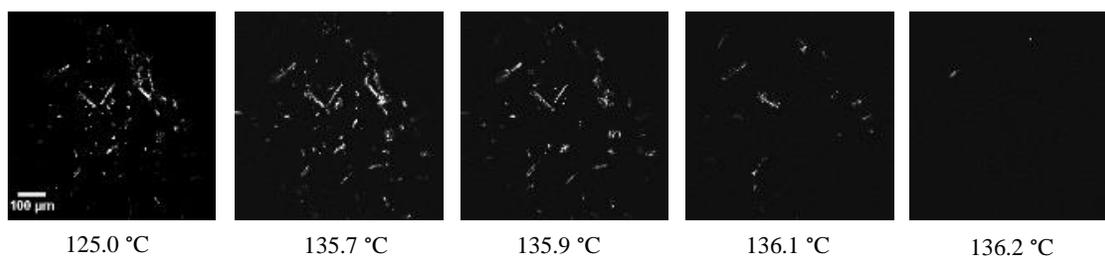


Figure 5.4: SHG microscope images of urea crystals during SDSC analysis. SHG microscopy reveals the stochastic nature of the melting transition and enable analysis on a per-particle basis.

The simultaneous SHG measurements were also used to determine the impulse response function (IRF) of the DSC instrument. As described in a preceding paragraph, the thermal events

induced in the sample produced an instrument response with a temporal delay associated with heat transfer times. The macroscopic DSC endotherm arose from the net collective contributions from many such stochastic events. As such, the DSC observables were given by the convolution of the impulsive phase transformations with the IRF of the DSC system. Assuming a double-exponential IRF consistent with heat flow through a thermal resistor, the measured set of impulsive phase transformations by SHG, \mathbf{x} , can be combined with the recorded DSC trace, \mathbf{y} , to recover the maximum likelihood estimate (MLE) for the impulse response function $\hat{f}_{IRF} = f_{IRF}(\hat{\mathbf{a}})$ described by the set of parameters \mathbf{a} through $\hat{\mathbf{a}} = \arg \min_a \|\mathbf{y} - \mathbf{x} \otimes f_{IRF}(\mathbf{a})\|^2$. The results of an MLE fit of the measured DSC transient using the single-crystal phase transformation data from SHG is shown in **Figure 5.4**, recovering a phase-lag of 3.088 ± 0.013 s and rising and falling exponential time constants of 1.161 ± 0.019 s and 0.435 ± 0.013 s, respectively, for the double-exponential IRF. The phase-lag between the DSC event and the mean of the SHG events is attributed to delays from diffusive heat transfer from the individual particles to the pan and sample mount of the DSC instrument. The two time constants are tentatively attributed to the thermal diffusion times from the sample to the sample stage, and from the stage to the heat sink. (Höhne et al., 1996) Even in a relatively simple melting transition, SHG provides the capability of mapping phase transformations on individual particles with temporal resolution much faster than achievable from conventional ensemble-averaged DSC measurements. The pairing of both methods provides information inaccessible by DSC alone; namely, that phase transformation proceeds through slow nucleation followed by rapid growth on a per particle basis with an intrinsic rate of 1.0 ± 0.6 s/particle (variance is dominated by particle-to-particle stochasticity).

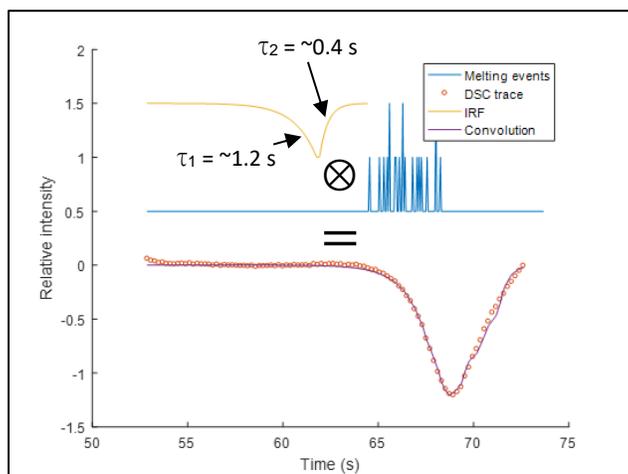


Figure 5.5: Comparison of the experimental DSC trace and the DSC trace generated from melting events. The best-fit double exponential IRF (yellow) is convolved with impulsive melting events (blue) of single-particles observed in SHG images to generate the purple trace, which is in good agreement with the experimental DSC data (red). The time constant and offset of double-exponential IRF were optimized to minimize difference between experimental and generated DSC traces. Best-fit values of 1.161 ± 0.019 s, 0.435 ± 0.013 s, and 3.088 ± 0.013 s were computed for the double-exponential time constants and the offset respectively.

5.4 Application to Trehalose

The performance of SDSC using the integrated SHG-DSC microscope was further tested with trehalose dihydrate, a molecule with complex phase behavior, including the existence of multiple transient crystal forms accessible under varying conditions.(Singh, 2018; Taylor et al., 1998; Lynne S Taylor & Peter York, 1998; Lynne S Taylor & P York, 1998) Since significant differences in crystal structure can exist between polymorphs, hydrates, and solvates, SHG has the potential to aid in distinguishing transitions due to its sensitivity to noncentrosymmetric crystal forms.(Chowdhury et al., 2014) **Figure 5.5** shows the SHG images of the trehalose dihydrate and **Figure 5.6** summarizes the SDSC analysis of trehalose dihydrate. Each colored line represents the percentage of pixels in an individual trehalose crystalline particle that exceeds a threshold.

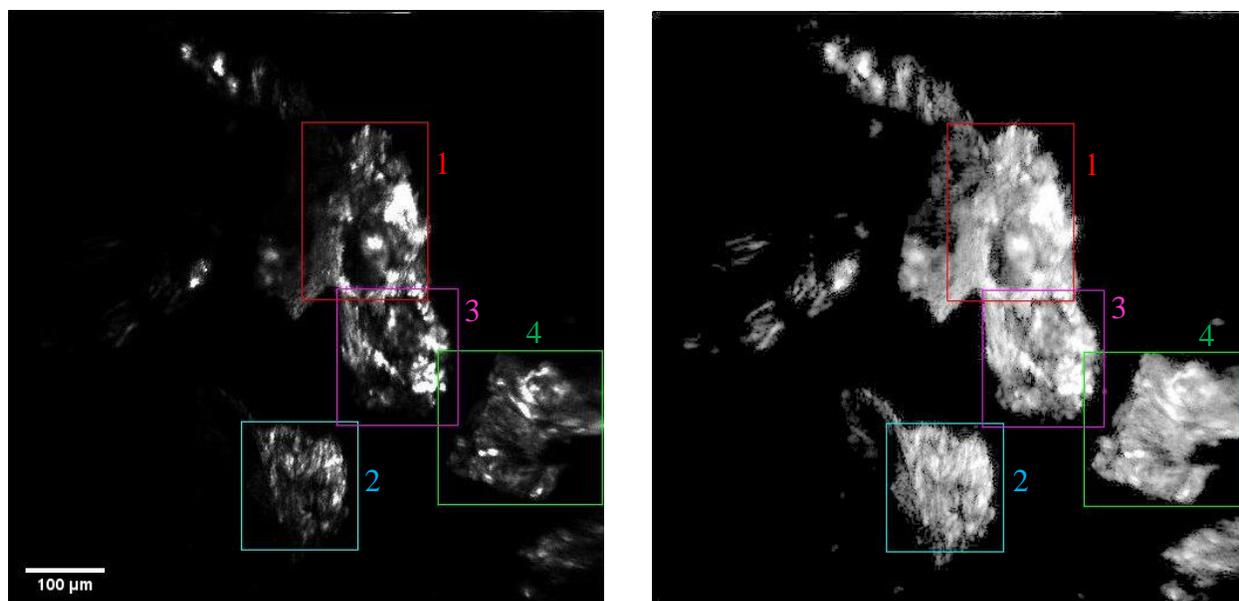


Figure 5.6: The left image is the summed frames over the temperature range of 140–220°C. Each particle included in the analysis is enclosed in a box. The image on the right is the same image but on a log scale to better differentiate background and signal.

Four individual particles were monitored in this study. The black line is the DSC curve, in which dips correspond to endothermic events. For the temperature range of 140–220 °C, the collected SDSC images were summed and the logarithm of this summed image was cropped to the measured region of each particle. Each of the cropped images points to their representative SHG trace. DSC and SHG data were interpreted using previously reported literature values to map the polymorphic transitions of the sample through the temperature ramp. At the initial 80 °C, the weak SHG signal arises from the trehalose dihydrate (T_h) form that was initially added to the DSC pan before the temperature ramp. (Sussich et al., 1998) The composition of this starting material was confirmed by powder X-ray diffraction (pXRD) to be primarily composed T_h , although the unstable anhydrous form (T_a) may also be present at trace levels; previous reports indicate observation of the T_a form following room temperature storage at low relative humidity. The first signs of phase transformation activity occurred in the temperature range of 80–130 °C. The DSC curve exhibits a major endothermic peak in this range. This peak was previously reported to be the result of the dehydration of T_h . (Raimi-Abraham et al., 2014) Additionally, two distinct trends appear in the individual particle SHG data. Three of the trehalose particles increased in SHG activity while the other particle decreased over this same lower temperature range. The disparity

in SHG activity between individual particles can be explained by the two major dehydration pathways available for T_h in this temperature range; T_h can either partially dehydrate with rearrangement to the crystalline T_γ form or completely dehydrate to an amorphous form (T_{am}). (Raimi-Abraham et al., 2014; Sussich et al., 1998) Similarly, the T_α form can also transition to the T_{am} form upon heating. (Raimi-Abraham et al., 2014; Rani et al., 2006; Sussich et al., 1999; Sussich et al., 1998) The T_γ polymorph is theorized to be a mixture of T_h encapsulated in a shell of the crystalline, anhydrous form (T_β), which is SHG-active. (Raimi-Abraham et al., 2014; Sussich et al., 1999) Therefore, we conclude that the three particles that increased in SHG area in this temperature range likely transitioned to the SHG-active T_γ form, while the particle that reduced in SHG activity likely transitioned to the SHG-inactive T_{am} form at ~ 110 °C from either the T_h or T_{am} form.

The SHG activity generally increased for individual particles over the subsequent temperature range of 130–160 °C. The DSC curve exhibited small but reproducible features in this temperature range, the magnitude of which varies between experiments, and has been reported to be due to the dehydration of T_h . (Raimi-Abraham et al., 2014) Measurements of the fraction of SHG-active pixels in individual particles show that the three particles that previously exhibited SHG signal continue to convert over this temperature range. This increase in highly SHG-active fraction can simply be attributed to crystal growth over this temperature range. The fourth particle that previously decreased in SHG area shows a sharp increase in SHG area starting at ~ 160 °C. SHG-inactive T_{am} has been previously reported to transition to SHG-active T_β in the range of 150–200 °C. (Raimi-Abraham et al., 2014; Sussich et al., 1998) Likewise, the T_α form has been reported to be able to transition directly to the T_β when held at a temperature of ~ 150 °C. (Raimi-Abraham et al., 2014; Rani et al., 2006; Sussich et al., 1999; Sussich et al., 1998) However, this T_α transition has not been observed with the faster heating rates (20 °C/min.) used in the present study. Therefore, we conclude that this polymorphic transition from T_{am} to T_β is likely the cause of the sharp peak in SHG from the particle represented by the magenta trace.

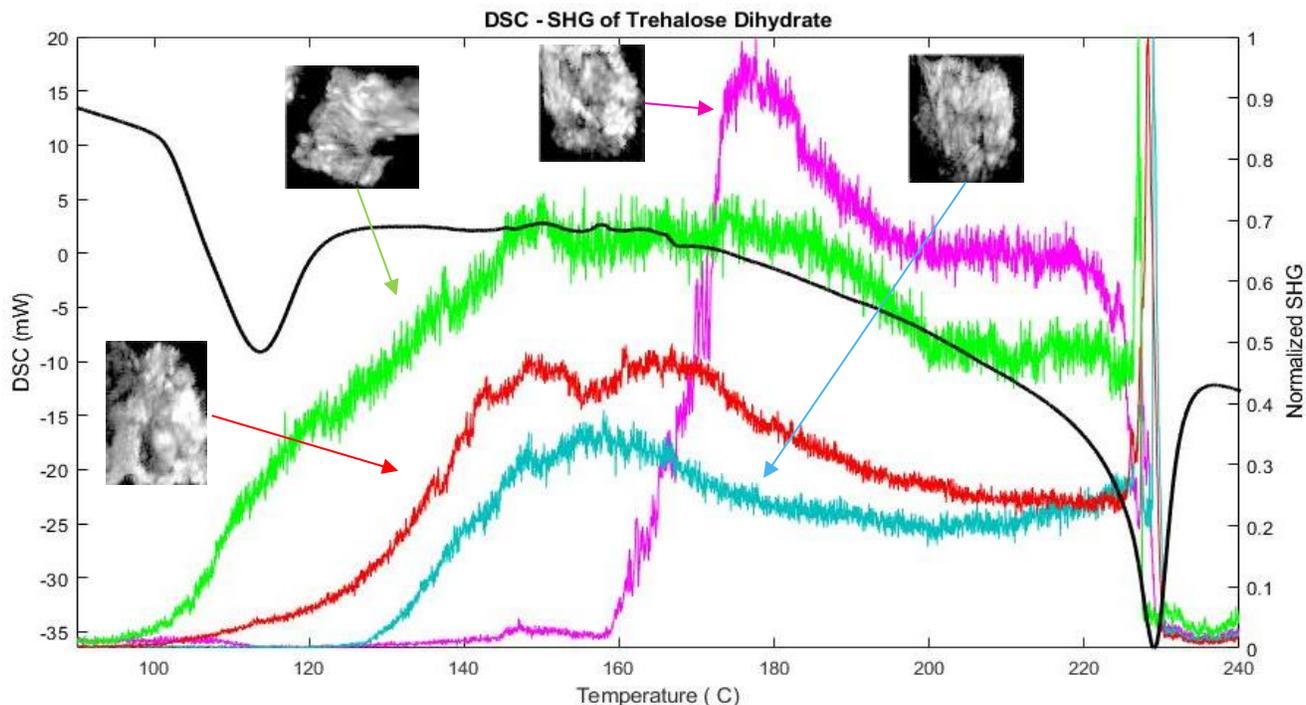


Figure 5.7: SDSC analysis of trehalose by SHG microscopy. The colored lines show the percentage of pixels in a trehalose particle above an SHG intensity threshold overlaid with the black line of the DSC trace with the endothermic direction being down. SDSC images were summed for the temperature range of 140–220°C and the logarithm of the summed image was cropped to the measured region of each particle (each cropped image points to the representative SHG trace).

Following the monotonic increases in transitioning to the T_{β} form from 130–160 °C, decreases in SHG were observed for all four particles in the range 160–200 °C. No peaks are observed in the DSC curve in this range, but the trehalose particles exhibit slowly varying and inhomogeneous loss in SHG activity. Several possible mechanisms for this loss in SHG activity were considered, two of which are described in detail. One possible explanation for loss in SHG activity with large crystal-to-crystal variability may be from interference effects as crystal sizes change during phase transformation. (Ximeng Y Dow et al., 2016) Interference fringes from single crystals were reported previously in SHG microscopy measurements. However, interference effects were deemed unlikely for three reasons: i) interference effects would be anticipated to produce increases or decreases in SHG activity with equal probability, while only decays were observed experimentally, ii) the backwards coherence length is short (~100 nm) relative to the sizes of the crystals, such that average activities would be expected, and iii) the interference is

expected to vary with thickness, such that crystals with variance in dimensions would likely produce fringes over which the intensities are integrated. Alternatively, the decay in SHG signal can be explained by condensation of water vapor released by T_h onto the optical window of the DSC. Condensation was commonly seen in our experiments when working with hydrated samples because open crucible pans were used in the experiment to allow for imaging during DSC scans. Additionally, no purge gas was used to remove water vapor from the system, as addition of purge gas introduced measurement variance in the DSC analyses from variable heat transfer to the bath gas. Previously reported results by TGA indicate gradual water loss up to 200 °C, (Raimi-Abraham et al., 2014) consistent with the observed reduction in SHG signal that is proposed to arise from scattering losses from condensation. Furthermore, condensation is likely to impact some locations within the FoV to a greater extent than others, potentially providing an explanation for the large crystal-to-crystal variability in the loss. Thus, condensation is predicted to be the main cause of this loss in SHG signal.

Following the slow SHG signal loss from condensation, dramatic changes in the SHG-activity arose from 200–240 °C. In this range the DSC curve exhibited an endothermic peak at ~230 °C, and the SHG data displayed a sharp increase and subsequent decrease during the DSC peak. It is clear from the literature that the endothermic DSC peak arises from a melting transition. (Raimi-Abraham et al., 2014; Sussich et al., 1999) The final fall in SHG area is attributed to this melting event. However, the source of the sharp increase in SHG area prior to the melt has no precedent in reported DSC measurements. Two origins for the transient SHG area increase were considered. First, significant mobility in the sample arose during the melting transition, which could alter the positions of the sample crystals relative to the focal plane. Crystals moving into the focal plane could result in an increase in SHG area, as SHG intensity scales with the squared power of incident light. Second, the increase could be explained by a change in the crystal form, degree of crystallinity, or crystal size of trehalose. To evaluate the first mechanism, a custom optic, designed in-house, was added to the beam path to extend the depth-of-field from ~20 μm to ~100 μm (detailed in a manuscript in preparation). The sharp peak in SHG prior to the melt was still observed while imaging with an extended depth-of-field, suggesting that crystal movement into the focal volume is not likely to be the major cause of the increase in SHG area. Alternatively, the increase in SHG area could arise from a change in the crystal form, degree of crystallinity, or crystal size of trehalose. While it is possible that there exists a previously unknown polymorphic

transition immediately prior to the melt of trehalose, this possibility was rejected on the principal of Occam's razor, considering the extensive body of prior work done to characterize the polymorphism of trehalose. A more plausible origin for the spike in SHG area is from the final conversion of the remaining T_h core of T_γ polymorph to T_β , promoted by an increase in the energy in molecular diffusion immediately prior to the melting transition. Additionally, another closely related contribution to the increase in SHG area could be the rapid growth of T_β crystalline domains of trehalose from residual T_{am} due to increase mobility immediately prior to the melt. In summary, the particles observed in this study are theorized to undergo one of two general phase transition pathways: $T_h \rightarrow T_{am} \rightarrow T_\beta \rightarrow \text{melt}$ or $T_h \rightarrow T_\gamma \rightarrow T_\beta \rightarrow \text{melt}$.

Notably, this "cold transformation" phenomenon, in which phase transformation to the SHG-active crystalline form occurs immediately prior to the transition to the liquid form, was clearly observable by SHG but undetectable in the DSC measurements. This latter proposed mechanism has precedent in observations of "cold crystallization" in DSC, in which transient crystallization arises upon transformation from a glassy state to a liquid. (Wunderlich, 1958) In DSC, cold crystallization appears as an exothermic peak immediately preceding an endothermic peak. Cold crystallization has been observed in studies of T_h for the transition from a glassy material to the T_β form. (Sussich et al., 1999) However, to our knowledge, this study represents the first observation of cold transformation between two different solid-state crystalline forms prior to the melt. Importantly, evidence supporting cold transformation was undetectable by DSC alone, presumably obfuscated by the comparatively longer response time and inherent ensemble averaging associated with the DSC measurements.

The final DSC peak in the trehalose dihydrate experiment was analyzed using the IRF recovered from the urea experiment (see **Figure 5.5**). The "cold phase transformation" events and the melting events of the four particles were convolved with the best-fit, double-exponential IRF determined from analysis of the SDSC data for urea. The results of the convolution are shown in **Figure 5.8** as dashed lines. Neither the cold phase transformation events nor the melting events produced curves that matched well with the DSC peak. This perhaps indicates that the four particles observed using SHG microscopy are not representative of the full population of particles in the sample. A more complete set of phase transformation events would possibly produce a curve that would better match the DSC peak.

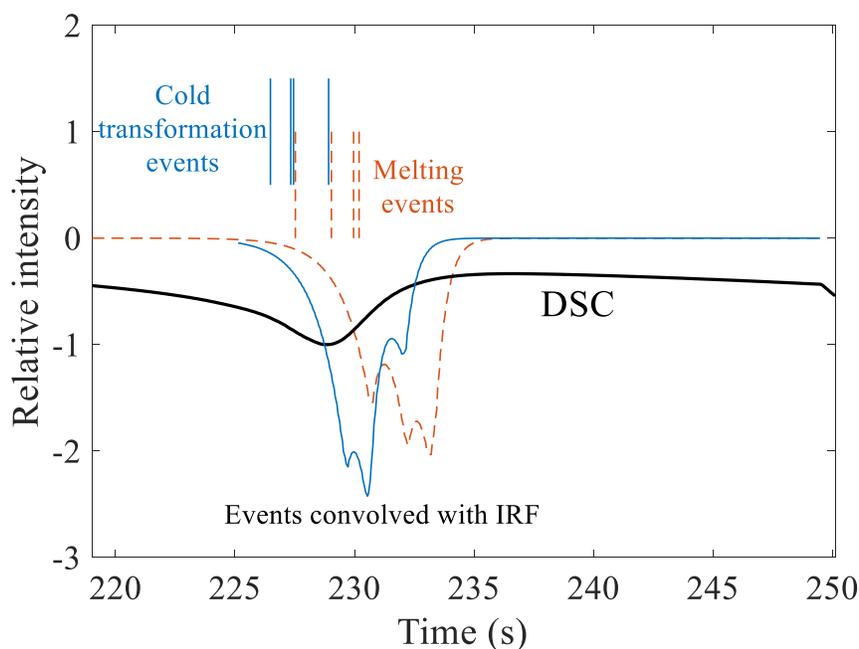


Figure 5.8: The final DSC peak in the trehalose dihydrate experiment was analyzed using the IRF recovered from the urea experiment. Results of the convolution are shown as the solid blue and dashed orange lines. Neither separately nor together do they match the DSC peak.

5.5 Conclusions

SDSC was shown to enable connection of internal structural rearrangements within crystalline materials measured by SHG on single particles to the heat flow recorded by DSC. SHG measurements were sensitive to variations in the molecular packing arrangements during the phase transitions of individual particles that are generally inaccessible in conventional microscopy. Coupling the structural data acquired from SHG imaging with DSC enabled single-particle measurements disentangled from the ensemble-averaged thermal transients recorded with DSC. A proof-of-concept SDSC experiment was performed on urea, which has a single solid/liquid phase transformation. From these measurements, the IRF of the DSC system was determined, showing that the melt proceeds through a slow nucleation step followed by rapid growth on a per particle basis. Next, SDSC was used to characterize the phase behavior of a more complex system, trehalose dihydrate. The dehydration events of trehalose dihydrate were readily observed by SDSC and the higher temporal resolution of SHG measurements, compared to that of the temporal resolution of typical DSC measurements, allowed for detection of a rapid phase transformation not

observable by DSC. Thus, SDSC is proposed as a novel technique for the pharmaceutical drug development pipeline for the characterization of single-particle phase transformations.

CHAPTER 6. GENERATIVE ADVERSARIAL LINEAR DISCRIMINANT ANALYSIS

6.1 Introduction

Measurement science has seen a vast improvement with the advent of modern instrumentation with ever-increasing sampling capacity, giving rise to a large volume, complex raw measured data. It has been emerging as a practical concern of how to wisely make decisions in the era of Big Data in which the speed of data growth has already exceeded Moore's law[2]. To make use of the raw data acquired, a general first step is to conduct a feature extraction to reduce the dimensionality of the raw data. There exist abundant models to perform this task, especially with the advance in artificial neural networks that solves problem magically without solid mathematical foundation. Among the models that are routinely adopted for feature extraction, Linear Discriminant Analysis (LDA) is a common choice, especially among chemists, given its mathematical simplicity and it being a supervised model.

Despite LDA being known for its relative simplicity mathematically, the mathematical underpinnings of the method require a matrix inversion step when evaluating the eigenvalues & eigenvectors. For a matrix to be invertable, that matrix must be square. However, non-square matrices can still have a left or right inverse under certain conditions. Getting a square matrix for use in LDA can be accomplished in two ways, i) perform a preprocessing step to reduce the dimensionality to the number of replicates of data, or ii) increase the number of measurements to match the dimensionality. Neither method is easy, and both have disadvantages. If a preprocessing is performed to reduce the dimensionality of the data set, some information is inevitably lost. And acquiring more measurements to match the dimensionality is often imposes a significant time burden. Another variation on acquiring more measurements, is the generation of additional measurements based on knowledge of the mean & standard deviation within the data set. The contrarian would point out that the generation of additional data only serves to introduce noise. To that point, the contrarian would be correct. Noise is introduced, but LDA can be easily performed without preprocessing or performing additional measurements.

Herein, generative adversarial linear discriminant analysis (GALA) is introduced as a method to perform LDA without preprocessing or collecting a large data set through the generation of data and then performing an iterative generative adversarial attack to reduce the noise associated

with the generation of data. Briefly, generated noisy decoy data is perturbed to classify as one of the LDA classes. LDA is performed again but now with an additional class for the decoy data. The process of generating decoy data, perturbing decoy data, and performing LDA is iterated until convergence. Through this generative adversarial method, LDA is performed without the need of preprocessing or the acquisition of large data sets.

6.2 Mathematical Framework

6.2.1 Overview

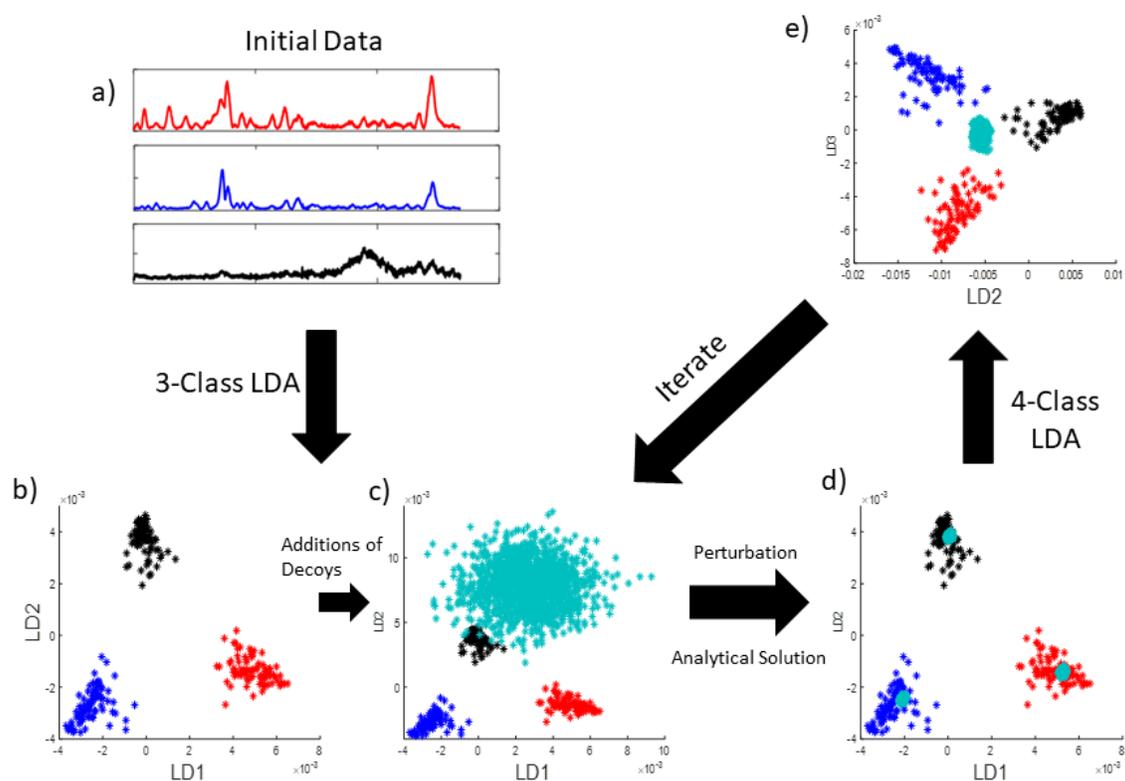


Figure 6.1: Workflow of generative adversarial linear discriminant analysis algorithm. Initial data that falls into three classes is projected onto a reduced dimensional space. Random decoy data are then generated and projected onto this 2D space. The decoy data are then attacked to make them overlap with the initial data on the reduced dimensional space while minimizing the magnitude of the perturbation. LDA is performed again, but now with an additional class to separate out the decoy data. The process is iterated.

Figure 1 shows the overall workflow for Generative Adversarial LDA (GA-LDA). In b) LDA is performed using additional digitally generated spectra to make the matrix invertible. This step adds noise, but also removes the need to either collect more measurements or perform a preprocessing step. Next, a generative adversarial iteration (c through e) is performed to remove noise and improve the loading plots for LDA. C) shows the projection of decoy data following generation from a uniform random distribution. D) The decoy data is then perturbed to classify as one of the original classes while retaining the original spectral features of the decoy data. e) LDA is performed again, but now with an additional class for the decoy data. The generation of decoy data, perturbation, and LDA is performed until convergence (c through e).

6.2.2 Performing LDA

LDA was performed by solving the equation $J\mathbf{W} = \mathbf{S}_w^{-1}\mathbf{S}_b\mathbf{W}$. Where J is the resolution, \mathbf{W} is a matrix of the optimal projection vectors to maximize resolution, \mathbf{S}_w is the within class variance matrix, and \mathbf{S}_b is the between class variance matrix. LDA eigenvectors for the data set are shown in the **Figure 6.5**. The dimensionality of the data set was reduced from 1340 to 2. To make \mathbf{S}_w a square matrix (and therefore, invertible), additional spectra were generated via knowledge of the mean and standard deviation of each wavelength for each class. Enough spectra were generated to make each class have an equivalent number of spectra to spectral channels.

Typically, the within-class variance is made invertible by performing some preprocessing to the data set. Examples of preprocessing include thresholding and principle component analysis.

6.2.3 Generating Decoy Data

Decoy data were generated to visually represent the initial loading plots. This was performed by generating spectra from a uniform distribution with an upper bound of 6×10^{-3} and a lower bound of 0. These bounds were chosen because they reflect the upper bound and lower bound of the measured Raman spectra. Decoy data were then projected onto the reduced dimensional space using the eigenvectors from LDA. This did not place the decoy data in the center of the graph, but closest to the class that represented the background signal. The placement of the decoy data shows that the initial loading plots do indeed contain some spectral features that allow for discrimination

between class 1 and class 2. Even though upon inspection of the eigenvectors it is difficult to visually detect these spectral features.

6.2.4 Launching an Attack

The central objective of the adversarial attack is to identify the perturbation, δ , that optimally alters the classification of an initial spectrum to a target group, subject to constraints imposed by a cost function. For an initial spectrum, \mathbf{x}_s , the perturbed spectrum, \mathbf{x}' , is given by $\mathbf{x}' = \mathbf{x}_s + \delta$. The general strategy in optimization of δ is illustrated in **Figure 6.2**, which is intended to serve as a graphical depiction of an attack implemented in a reduced-dimensional space. For example, these reduced dimensions may represent the principal axes produced by linear discriminant or principal component analyses. Each wavelength channel in the original spectral space results in a “nudge” to collectively contribute to the position of the spectrum in the reduced dimensional space. While the primary spectral features (indicated by the thin black arrows) combined to dictate the general position within the reduced dimensional space, randomness within the noise (indicated by the short red arrows in **Figure 6.2A** and **6.2B**) produces a spread about that mean position.

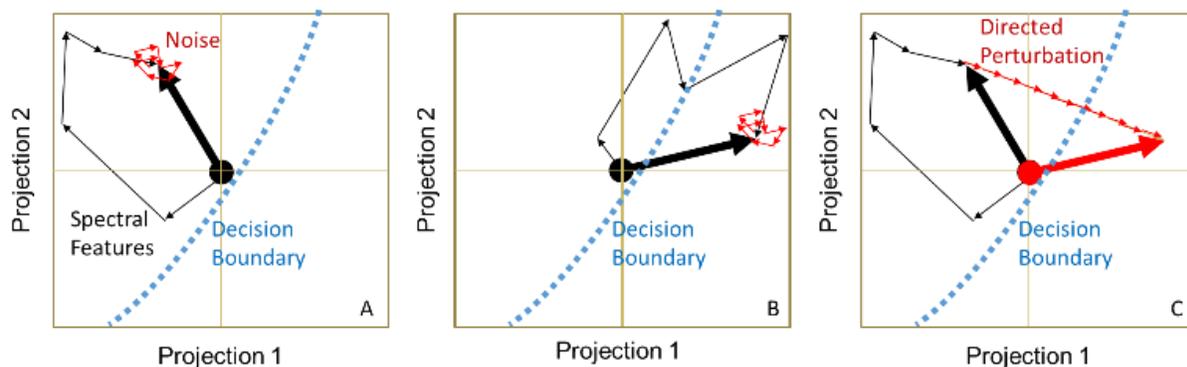


Figure 6.2: Conceptual illustration of adversarial spectroscopy. Major spectral peaks drive the position of spectra in lower-dimensional projections, demonstrated by the set of thin, long black arrows for Classes A and B, respectively. In panel C the addition of patterned perturbations in the vector δ optimally relocates the position from the initial Class A to Class B in this reduced dimensional space.

Additional “nudges” by addition of a perturbation to the original spectrum can relocate the position of the initial spectrum in the reduced dimensional space to one significantly closer to the

target, as illustrated in **Figure 6.2C**. The vector of deviations \mathbf{d} from the initial sample spectrum, \mathbf{x}_s , to the “target”, \mathbf{x}_t , in the reduced dimensional space is given by the following expression:

$$\mathbf{d} = \mathbf{D} \cdot [(\mathbf{x}_s + \boldsymbol{\delta}) - \mathbf{x}_t] \quad (\text{Eq 6.1})$$

Where \mathbf{D} is the matrix of eigenvectors that project the high-dimensional data to a lower-dimensional space (such as PCA or LDA). In the absence of other considerations, the optimal perturbation, $\hat{\boldsymbol{\delta}}$, will be one that maximizes the probability that the perturbed spectrum will be classified as the target. The reduction in dimension associated with \mathbf{D} results in an underdetermined problem, with, in general, an infinite number of selections for $\boldsymbol{\delta}$ producing comparable values for \mathbf{d} . In the present study, three dimensions in LDA-space were considered in the analysis of spectra with 1340 elements, such that \mathbf{D} is a $n \times 1340$ matrix, where n is the number of dimensions the data is reduced to & $n = 2$ initially and $n = 3$ on later iterations. These eigenvectors are visualized in **Figure 6.5**.

Selection of one among the infinite sea of possible perturbations was performed by also considering an additional term in the cost function in Eq. (6.2) to minimize the squared magnitude of the perturbation. The total cost function for the reduced dimensional analysis was given by the sum of the two terms, which collectively minimized the sum of squared deviations to the target in the reduced dimensional space while minimizing the overall squared magnitude of the perturbation, $\boldsymbol{\delta}$, in spectra space, both evaluated as the squared L2 norms.

$$\hat{\boldsymbol{\delta}} = \underset{\boldsymbol{\delta}}{\text{argmin}} \left[\|\mathbf{D} \cdot (\mathbf{x}_s + \boldsymbol{\delta} - \mathbf{x}_t)\|_2^2 + \beta \|\boldsymbol{\delta}\|_2^2 \right] \quad (\text{Eq 6.2})$$

The scaling parameter β in Eq. (6.2) allows for empirical adjustment of the importance given to proximity to the target relative to minimizing perturbation of the major spectral features. The first term in the cost-function is designed to “fool the classifier” by minimizing the distance to the target in the reduced dimensional space, while the second term is targeted to “fool the human” by minimizing the perturbation in spectral space. In the present study, a value of $\beta = 1$ was used for simplicity.

The optimal perturbation was determined analytically by rewriting the cost function as a function of the perturbation, see Equation 6.3.

$$F(\bar{\boldsymbol{\delta}}) = \|\mathbf{D} \cdot (\bar{\mathbf{x}}_s + \bar{\boldsymbol{\delta}} - \bar{\mathbf{x}}_t)\|_2^2 + \beta \|\bar{\boldsymbol{\delta}}\|_2^2 \quad (\text{Eq 6.3})$$

Then solve for the optimal perturbation by setting the partial derivative $\frac{F}{\delta}$ to zero. With $\frac{F}{\delta} = 0$, $\bar{\delta}$ can then be solved as:

$$\hat{\delta}_i = \frac{(\mathbf{D}^T \mathbf{D}(\bar{x}_s - \bar{x}_t))_i}{\text{Diag}(\mathbf{D}^T \mathbf{D})_i + \beta} \quad (\text{Eq 6.4})$$

6.2.5 Iterative Algorithm

Algorithm 1: Generative Adversarial Linear Discriminant Analysis (GALDA) algorithm

Input: Excel Spreadsheet
Output: Optimized LDA loading
Data: Spectroscopic data set

```

/* Data preprocessing */
data ← read from excel
training set, validation set ← data
/* GALDA with training set */
spectral mean and standard deviation of each class ← training set
perform 3-class LDA
while exit condition not met do
    generate Decoys
    perturb Decoys such that they classify to desired classes
    append perturbed Decoys to training data as a 4th class
    perform 4-class LDA
    if LDA loading smoothed to desired extent then
        | break while loop
    else
        | keep iterating
    end
end
/* validate GALDA Optimized LDA loadings with validation data */

```

Figure 6.3: Pseudo code to describe the algorithm.

6.3 Experimental – Raman Data Set

Pure clopidogrel bisulfate Form I and Form II were produced in-house at Dr. Reddy’s Laboratories and were used as received. Both the Form I and Form II particles were spherical with similar particle size distributions (diameter: ~25 μm). 252 Raman spectra were collected from the clopidogrel samples and separated into three classes (84 spectra per class) – Form I, Form II, and background (glass slide). The ground truth identity of these samples was known *a priori*.

Raman spectra were acquired using a custom Raman microscope, built in-house and described in detail previously (Zhang et al., 2018). In brief, a continuous wave diode laser (Toptica, 785nm wavelength) coupled into a Raman probe (InPhotonics, RPS785/24) was collimated by a

fused silica lens and directed through an X-Y scan head composed of two galvanometer scanning mirrors. Two additional fused silica lenses formed a 4f configuration to deliver a collimated beam on the back of a 10x objective (Nikon). The Raman signal from the sample was collected through the same objective and descanned back through the same beam path into the Raman probe. A notch filter was built in the Raman probe to reject the laser signal. Raman spectra were acquired using an Acton SP-300i spectrometer with a 100x1340 CCD array, and controlled by a computer running WinSpec32 software. The laser power measured at the sample was ~30 mW. The exposure time was 0.5 s per spectral frame. To achieve higher signal to noise ratio for high quality training data for classification, 30 consecutive frames were averaged for each spectrum acquired over a spot size of ~2-3 μm diameter within the field of view. A Savitzky-Golay filter was applied to smooth the spectra (Ehrentreich & Summchen, 2001), and a rolling ball filter was used to remove the fluorescence background (Liland et al., 2010). Finally, the spectra were normalized to their integrated intensities, i.e., the area under the curves.

The optimal $\hat{\delta}$ was recovered by solving the cost function either by via partial derivatives or through use of a genetic algorithm, with comparable results obtained by each. The script to perform partial derivatives calculations and the genetic algorithm were written in-house in MatLab2018b. Code for the partial derivatives calculation can be found in Appendix A. Results from the genetic algorithm were not included because they gave similar results to the partial derivatives method and required significantly more time and computational power to complete.

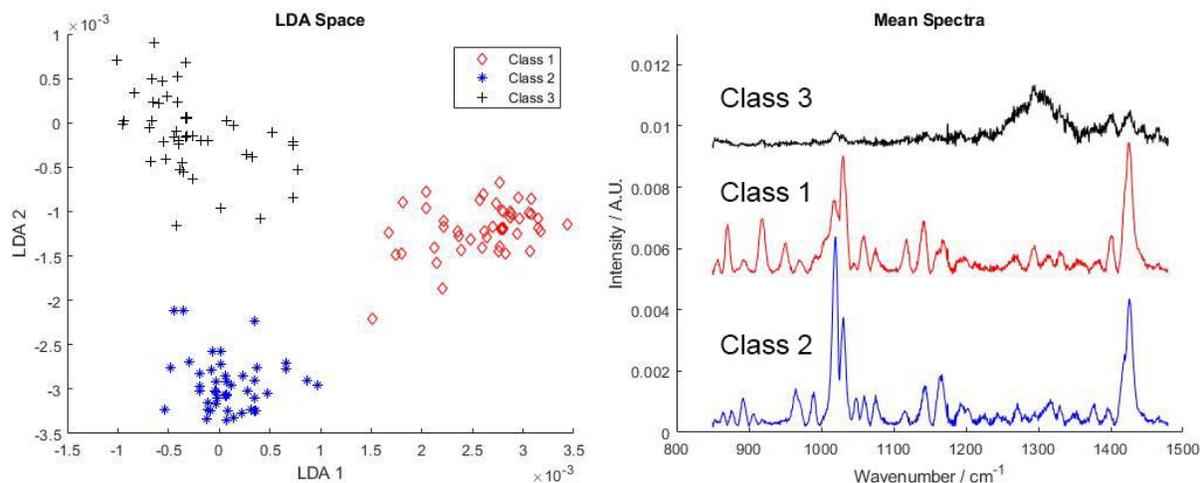


Figure 6.4: Projection of experimental Raman spectra in LDA-space (on the left), with the corresponding mean spectrum for each class (on the right), offset for clarity. The mean spectrum for each class is an average of 84 measurements.

6.4 Performing an attack on LDA

The mean spectra, average of 84 measurements, for three classes are shown in **Figure 6.4**. The spectra corresponding to the background were classified as Class 3 (top spectra, black). The spectra belonging to the two polymorphs of clopidogrel bisulfate were classified as Classes 1 and 2 (red, middle and blue, bottom trace respectively). The recorded spectra for Class 3, identified as background, shared one major feature of note, that being a large rolling peak around 1280 cm^{-1} . Spectra collected for Classes 1 and 2 showed clearly notable differences distinguishable by the relative peak intensities of the major features at ~ 1019 and 1030 cm^{-1} , along with numerous minor peaks present in one or the other of the two sets of spectra. Linear discriminant analysis provided clear separation between the different spectral classes upon dimension reduction.

LDA was performed on the data set of Raman spectra. In order to solve for the optimal projections on LDA space, the inverse of the within-class variance matrix was matrix-multiplied by the between-class variance matrix. Additional spectra were generated to enable direct matrix inversion via knowledge of the mean and standard deviation of each wavelength for each class. These additional spectra allowed for the eigenvectors to be calculated to project the Raman spectra onto the lower dimensional LDA space. The stability of the LDA projection was tested via a cross-validation method. 20% of the spectra were withheld from calculation of the LDA eigenvectors, and then projected onto the reduced dimensional space. 100% of all withheld cross-validation

spectra were classified correctly. The eigenvectors from LDA to project Raman spectra onto the lower-dimensional space are seen in **Figure 6.5**. Inspection of **Figure 6.5** shows that some of the spectral features of the Class 1 and 2 are present in the eigenvectors. Of note, however, is the high frequency content present in both eigenvectors. It will be shown later in the manuscript, that the high frequency content is the main handle that the adversarial attack uses to induce misclassifications and can be improved by the iterative generative method shown herein.

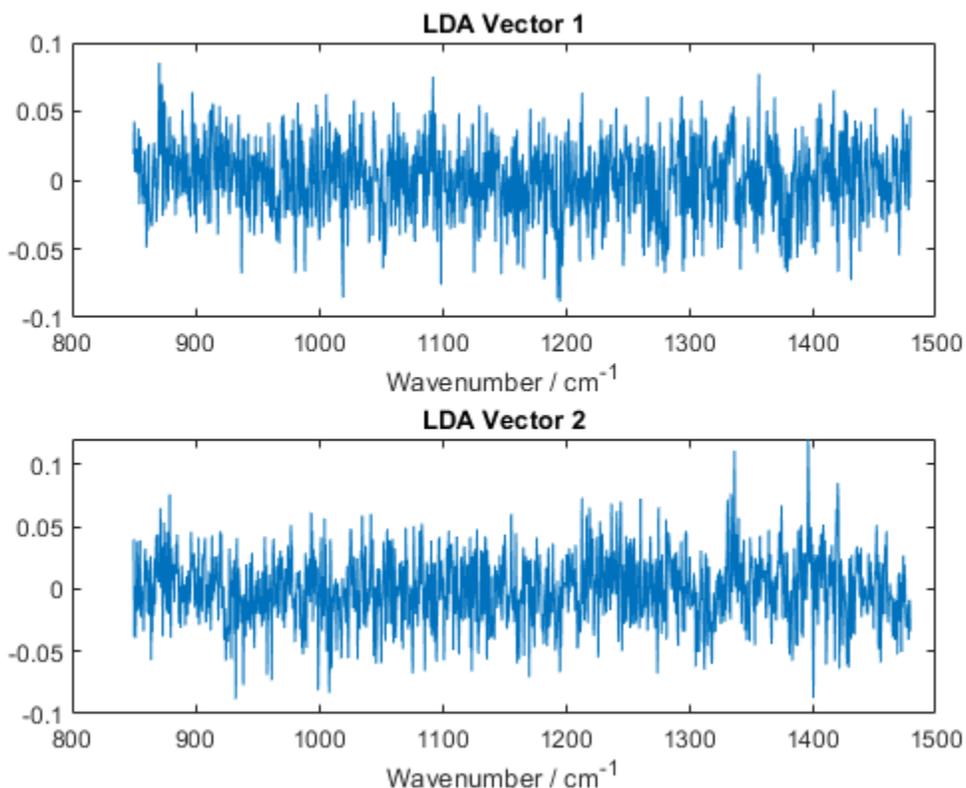


Figure 6.5: The eigenvectors recovered from LDA to project Raman spectra onto the lower-dimensional LDA space. The underlying spectral features are present and support reliable classification in validation data sets but are largely obscured by noise from the use of finite training data.

Using this data set of 252 spectra, the results of “attacks” in the reduced-dimensional LDA space are shown in **Figure 6.6-6.10**. Each perturbation was designed to displace the initial spectrum toward the target class within the reduced dimensional space. Because of the reduction in dimension, the direction of perturbation δ is underdetermined, such that the optimal perturbation was calculated based on Equation 6.4. In brief, the Euclidean distance to the target in the LDA-

space was minimized while also minimizing the overall magnitude of the perturbation in spectral-space. This approach resulted in patterns of perturbation shown in **Figures 6.6-6.10** which are representative of all attacks performed. The attack in **Figure 6.6** was designed to move a spectrum from Class 1 (middle right, red) to Class 2 (bottom left, blue). The green x's represent the optimized perturbation, along the path from initial to target classification represented by $|\Delta|$, where $\Delta=0$ corresponds to the initial, unperturbed spectrum and $\Delta=1$ corresponds to the perturbed spectrum positioned at the mean of the target class in LDA-space. The purple x's represent the region of greatest uncertainty in classification, which is defined as the classification transitions from 95% confidence as the initial class to 95% confidence as the target class. This region of interest is highlighted in **Figure 6.7**. In brief, ratios of probabilities based on maximum-likelihood estimation were evaluated in the reduced dimensional space using on the mean and standard deviation along each of the linear discriminant coordinates for the initial and target classification.

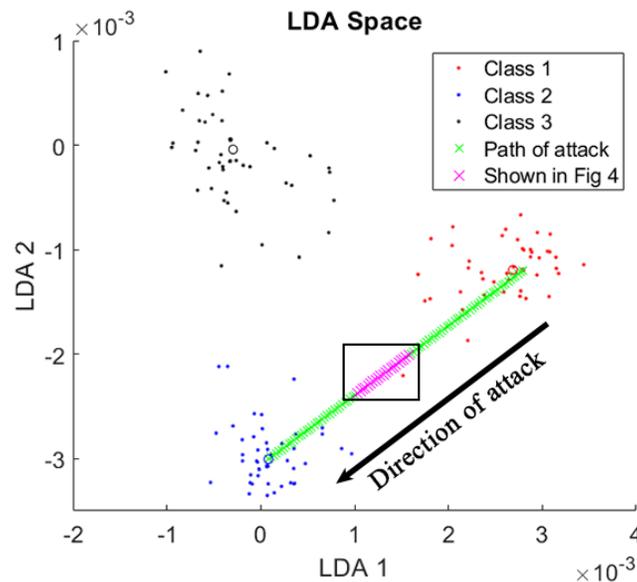


Figure 6.6: Demonstration of an incremental attack from Class 1 (red) to Class 2 (blue). Shown by the x's progressively moving towards the mean of Class 2, indicated by hollow blue circle.

The probability of the attack belonging to each class was calculated at each point, x. The probability of the purple x's in the box belonging to each class is shown in Figure 6.7. The purple x's denote the region of greatest uncertainty.

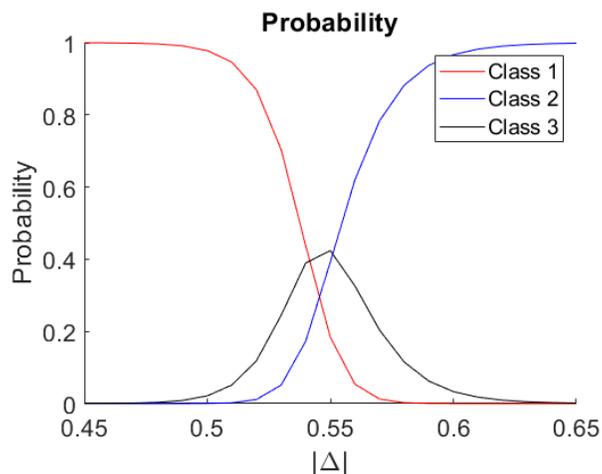


Figure 6.7: Zoom in of the probability of the perturbation belonging to each class in the region of greatest uncertainty, as shown by the boxed, purple x's in figure 6.6. $|\Delta|$ represents the degree of displacement toward the mean of the target class.

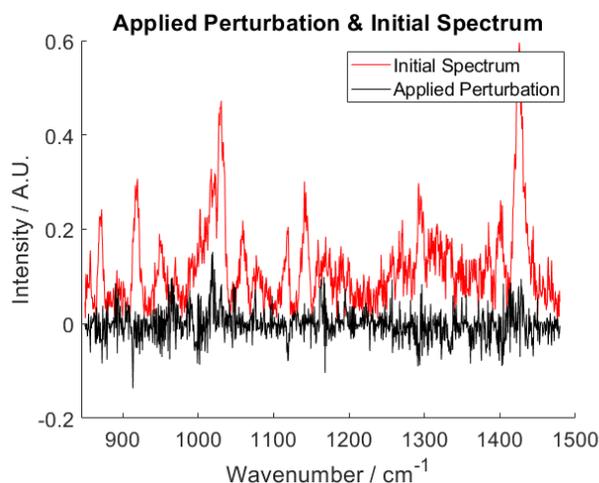


Figure 6.8: Comparison of the initial spectrum (red, top) and the applied perturbation (black, bottom). No offset applied.

Interestingly, the perturbations that optimally produce changes in classification are not due to changes in the major spectral features, as one might initially anticipate. Indeed, it is visually difficult to detect the perturbations in spectral space that squarely relocate the spectra to the target classes in the reduced dimensional space ($\geq 95\%$ confidence). **Figure 6.8** shows the perturbation applied (black data trace) in order to achieve a misclassification. In **Figure 6.9**, inspection of the perturbed spectra for reclassification from Class 1 to Class 2 suggests strong visual similarities to the initial classification, despite high statistical confidence in classification as the target following

perturbation. These trends are repeated in considerations of perturbations to induce misclassification from Classes 3 to 2 shown in **Figure 6.10** as well as from Classes 1 to 3 as shown in the **Figure 6.11**, where additional examples of spectroscopic attacks are detailed. All examples of perturbed spectra were misclassified with a confidence interval of greater than 95% as the target class. The dominance of high frequency content in the perturbation as seen in **Figure 6.8** highlights the significance of the variance of the signal in both the initial and target spectra, which appears to drive much of the spectral power in inducing misclassification. This outcome is somewhat surprising, as no constraints other than proximity to target and minimization of magnitude were imposed on the cost function given in Equation (6.2).

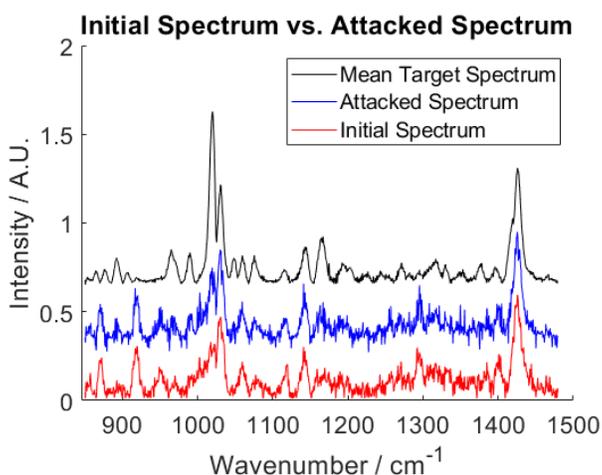


Figure 6.9: Comparison of the initial spectrum (red, bottom) from Class 1 with the attacked spectrum (blue, middle) classified as Class 2, and the mean target spectrum (black, top) of Class 2.

The relative magnitudes of the perturbations required to confidently induce misclassification are surprisingly small. Analysis of the median magnitudes of the initial spectrum and the applied perturbation corresponding to >95% confidence in misclassification show that the applied perturbation was ~12% the magnitude of the initial spectrum, and it is enough to unequivocally alter the spectral classification. The graphical depiction in **Figure 6.2** may provide some insights regarding the absence of similarity between the perturbation with either the initial or target spectra. From inspection of the figure, the most direct path from the initial to the target in the reduced dimensional space will generally not pass through the origin. As such, reductions

of the major peaks in the initial spectrum and growth of the major peaks for the target spectrum will not generally correspond to the optimal perturbation.

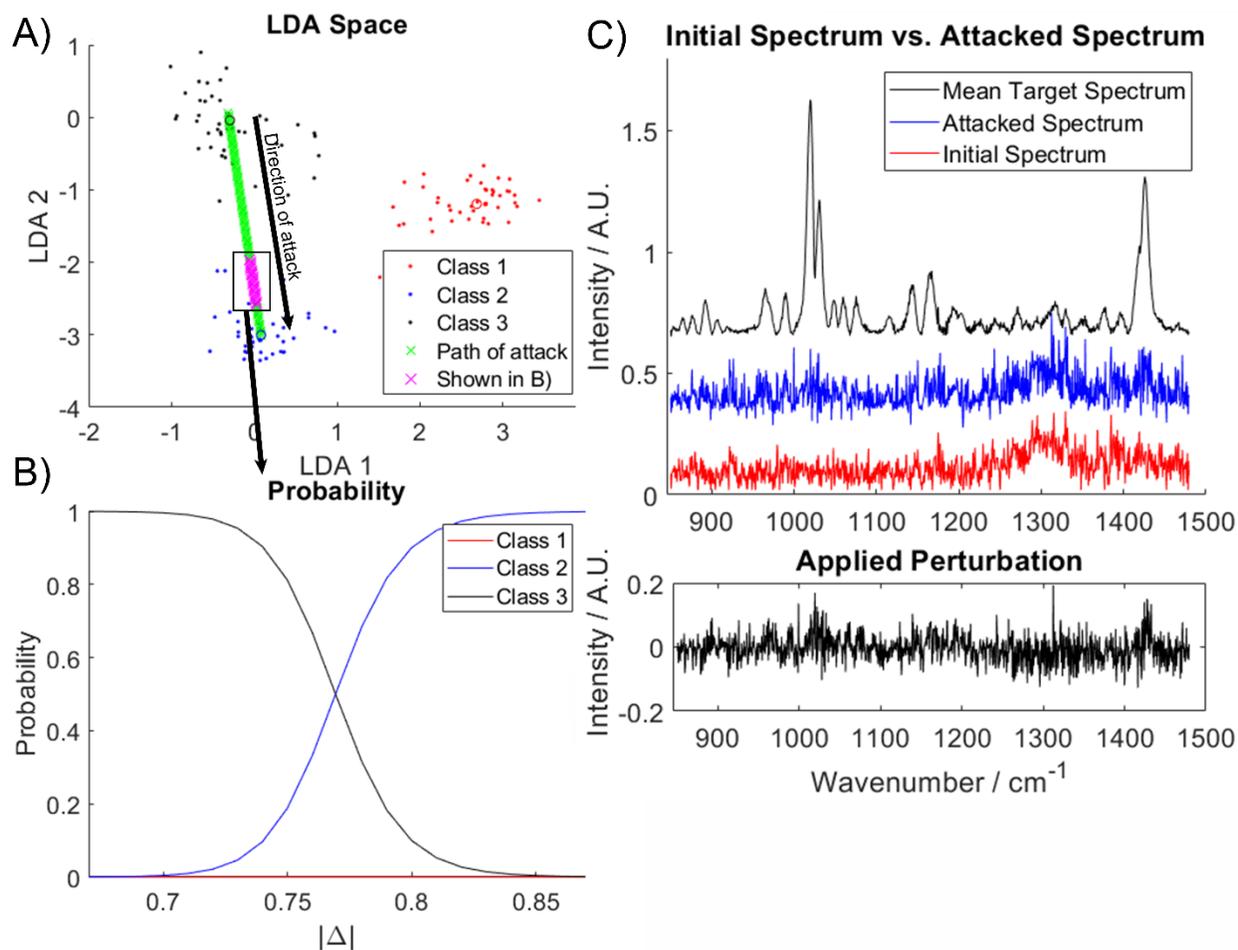


Figure 6.10: Additional example of an adversarial attack, in which a Class 3 spectrum is perturbed such that it is misclassified as Class 2. A) The direction of the perturbation and the final position in LDA space are shown in the reduced dimensional space. B) The region of greatest uncertainty in classification is highlighted by the boxed, purple x's in panel A. $|\Delta| = 0$ is the initial unperturbed spectrum, and $|\Delta| = 1$ is the mean position of class 2 in LDA-space. C) The attacked spectrum (middle, blue) is classified as Class 2 with the initial spectrum (bottom, red) from Class 1 and the target mean spectrum from Class 2 (top, black). D) Optimized perturbation used to induce misclassification.

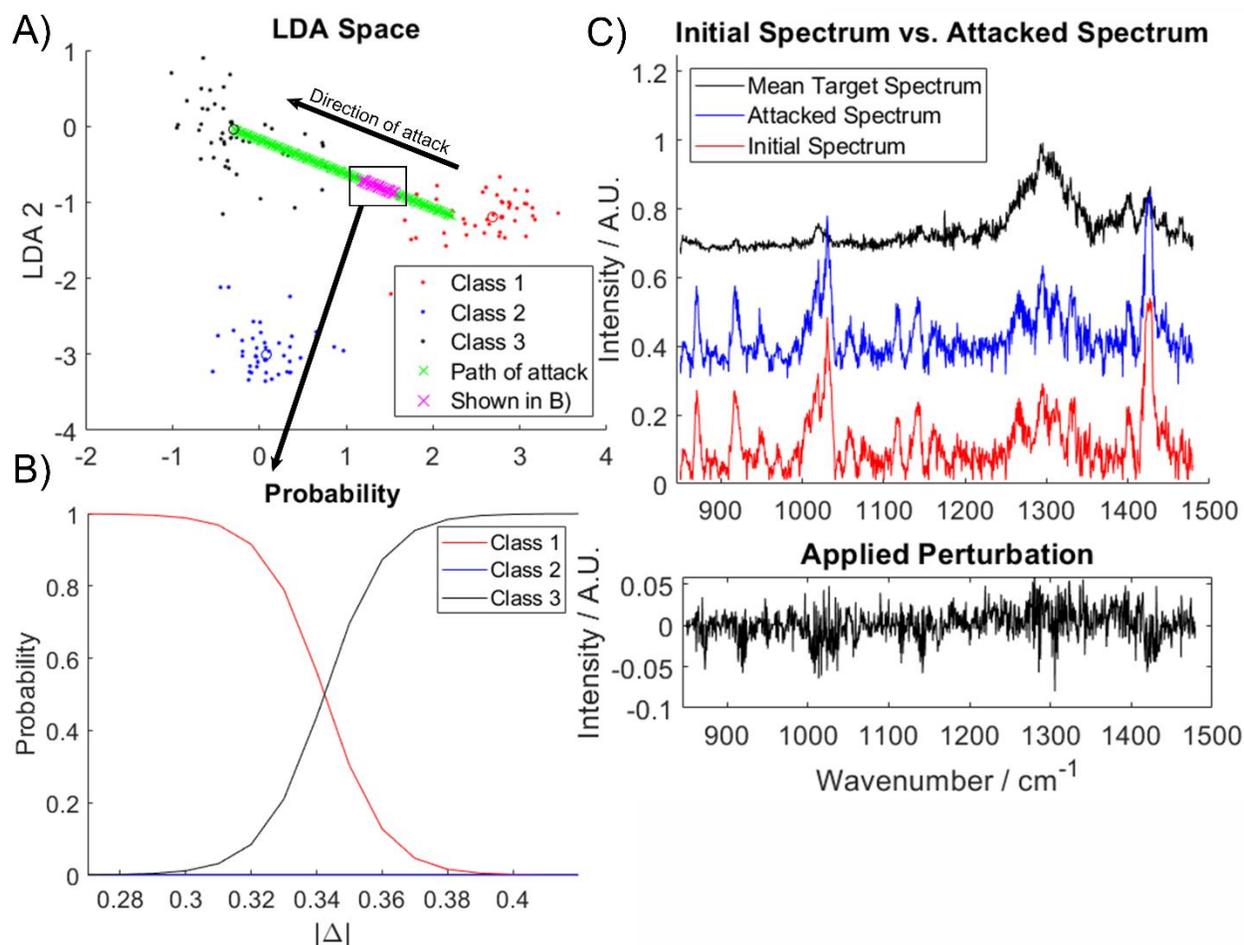


Figure 6.11: Demonstration of an attack on an initial spectrum from class 1 to induce a misclassification as Class 3. A) The direction of the perturbation and the final position in LDA space are shown in the reduced dimensional space. B) The region of greatest uncertainty in classification is highlighted by the boxed, purple x's in panel A. $|\Delta| = 0$ is the initial unperturbed spectrum, and $|\Delta| = 1$ is the mean position of class 3 in LDA space. At $|\Delta| = 0.38$ there is a >95% probability that the attacked spectrum belongs to Class 3 based on the position in LDA-space. C) Comparison of the initial spectrum (bottom, red) with the attacked spectrum at $|\Delta| = 0.38$, which bears clear visual similarity to the initial unperturbed spectrum. D) Optimized perturbation used to induce misclassification.

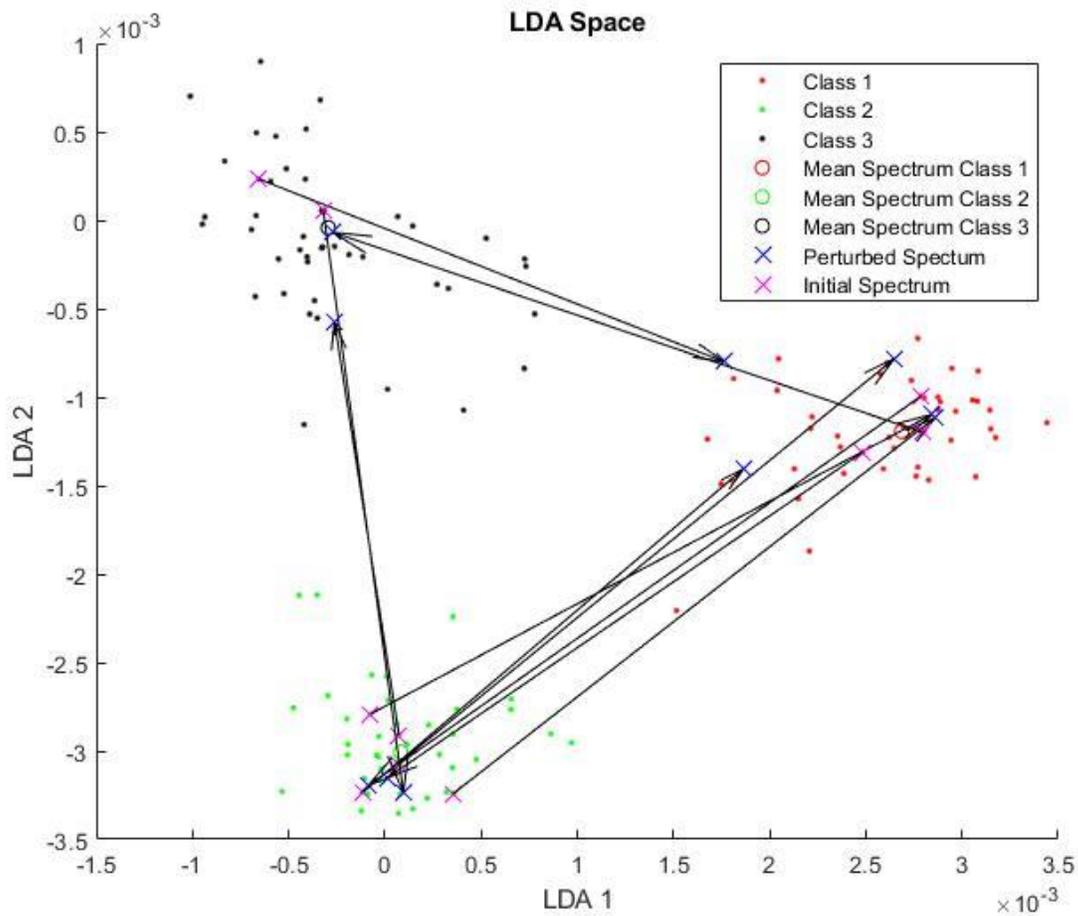


Figure 6.12: Attack strategies similar to those detailed previously were performed for an additional 10 initial spectra, selected at random, for randomly selected target classification and spectra. All results are quantitatively similar to those detailed previously.

These results highlight the growing challenges in ensuring statistical validity in regulatory, business, and legal decisions derived from data-intensive measurements that may be subject to incidental or intentional perturbation. As demonstrated herein, adversarial attacks on spectral information can profoundly influence decision outcomes, even when the perturbations are challenging to visually identify upon manual inspection. These results demonstrate that subtle digital alteration of files used in the otherwise benign operation of background subtraction can result in dramatically different outcomes in decision-making based on spectral analysis. Such relatively subtle changes spread over the entire spectrum would generally be challenging to discriminate from random noise. These results also demonstrate the vulnerability of dimension reduction techniques to instrument drift and cosmic rays among other common detector noise

sources. As the volume of data integrated for decision-making increases along with the corresponding degree of pre-processing required for data mining, chemical decision-making is only poised to be increasingly susceptible to manipulation through adversarial perturbations.

Attacks such as those demonstrated herein may also lead to future studies for improving chemical classifiers through generative adversarial strategies analogous. Just as the use of GANs has improved the reliability and stability of ANNs, so too may adversarial attacks in LDA and PCA provide similar future improvements in generative architectures. By identification of vulnerabilities, generative adversarial strategies may enable improvements in reliability and stability in conventional linear spectral analysis pipelines.

6.5 Iterative Generative Approach to Improving Loading Plots

It has been shown that the adversarial portion of the GALDA approach works well. Now the adversarial portion will be applied to the algorithm. The LDA was initialized as described in the experimental section. The initial loading plots for this LDA are seen in **Figure 6.13**.

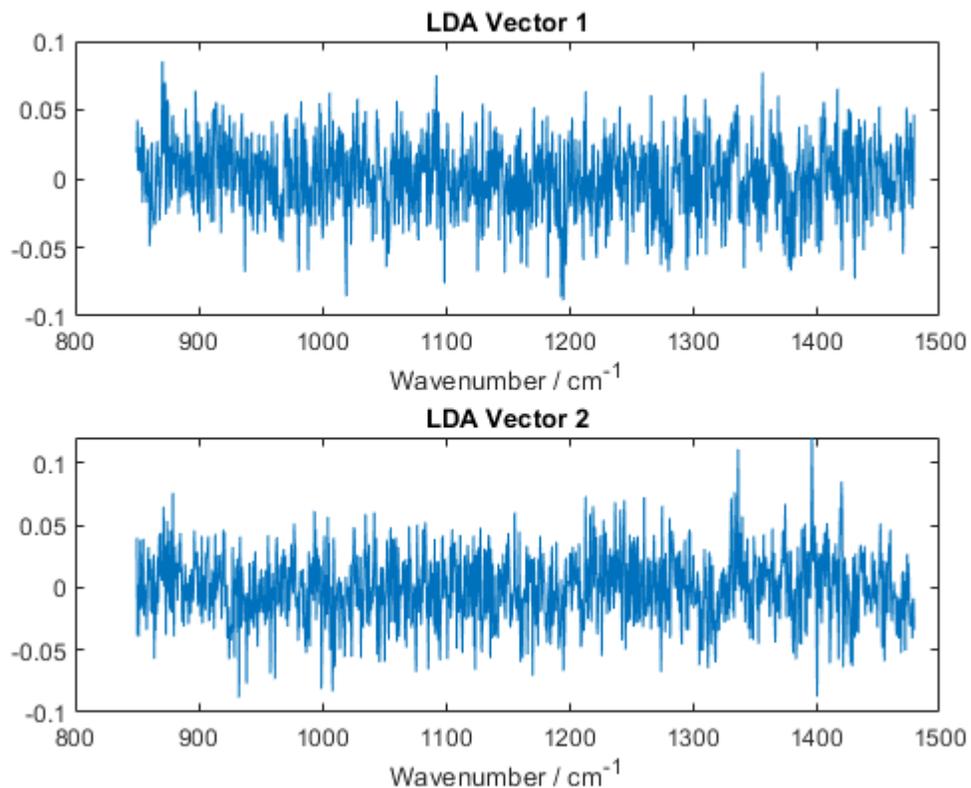


Figure 6.13: Initial loading plots after performing LDA on the raw data set with additional spectra generated from knowledge of the mean and standard deviation.

Inspection of the two loading plots in **Figure 6.13** shows that there are some spectral features in the loading plots. However, the amplitude of these features is small and are nearly obscured by the high levels of noise. Even with the high levels of noise in the loading plots, the data are still well separated as seen in **Figure 6.14**.

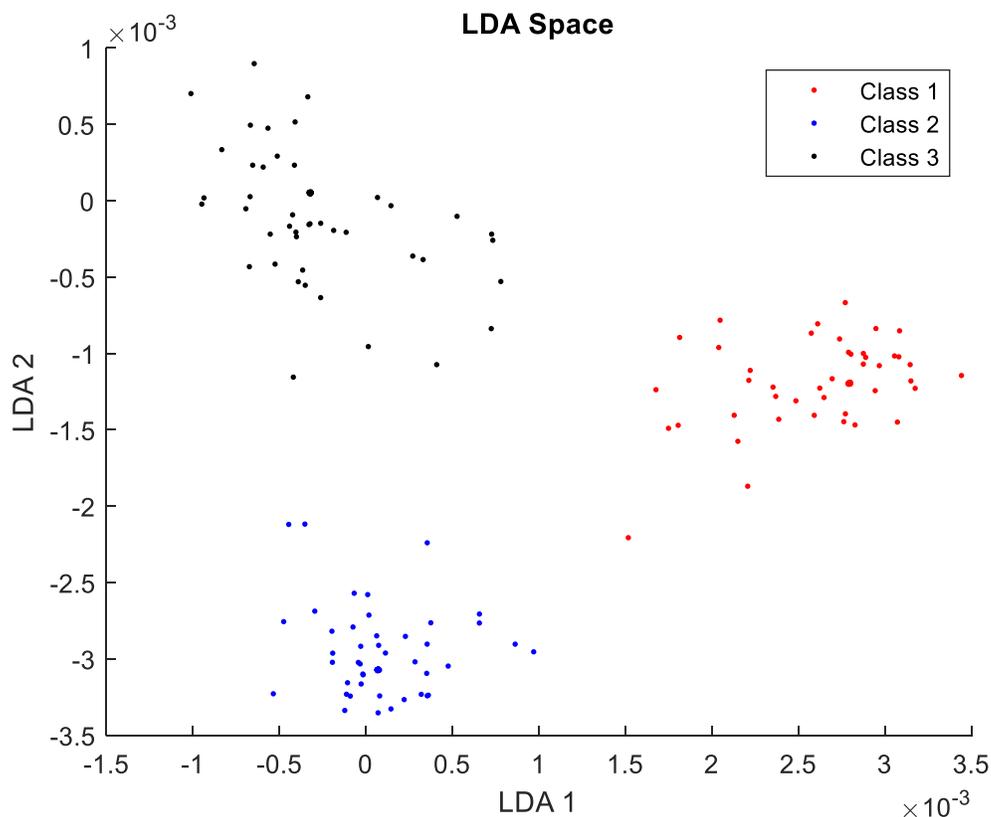


Figure 6.14: Initial placement in LDA space using the two eigenvectors shown in **Figure 6.13**.

The next step to improve the loading plots, is to generate decoy data that is noisy. By including decoy data that is noisy and setting it as an additional class within the LDA space, the LDA will have to focus on the spectral features to retain the same level of resolution. Decoy data was generated using a uniform distribution, the upper bound being 6×10^{-3} and the lower bound being 0. These bounds mimicked the max and minimum values of the measured Raman spectra. The length of decoy data was the same as the measured Raman spectra, e.g. 1340.

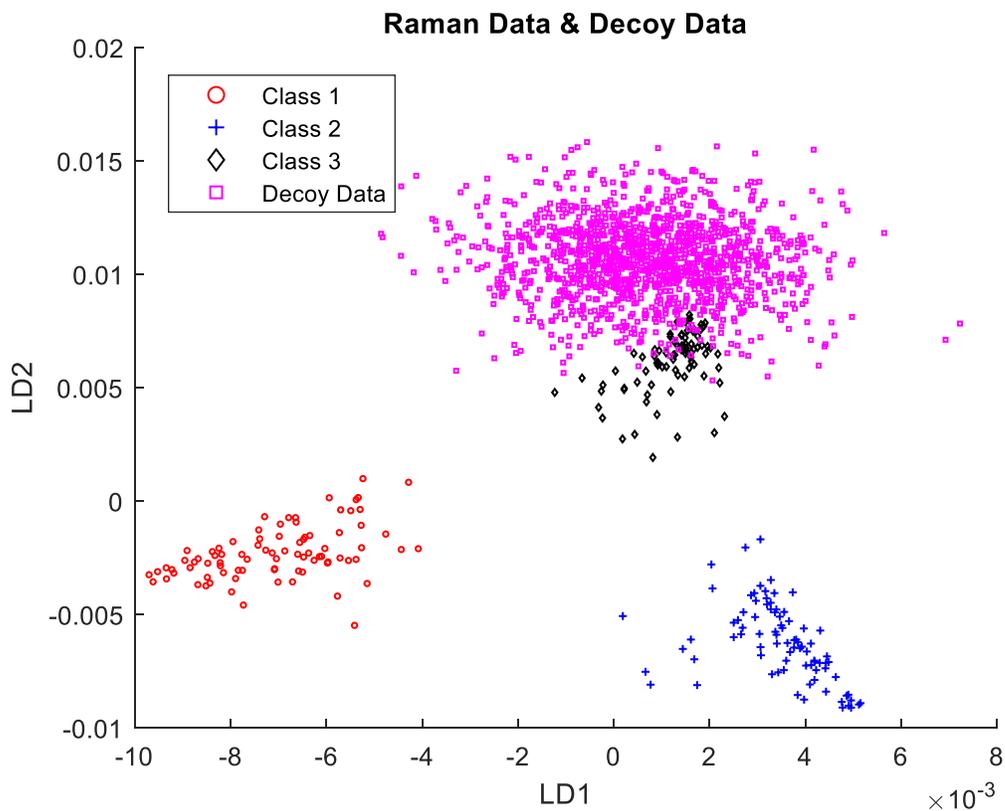


Figure 6.15: Decoy data projected onto the reduced LDA-space as purple squares. Note that the decoy data cluster around the background data which primarily consists of noise.

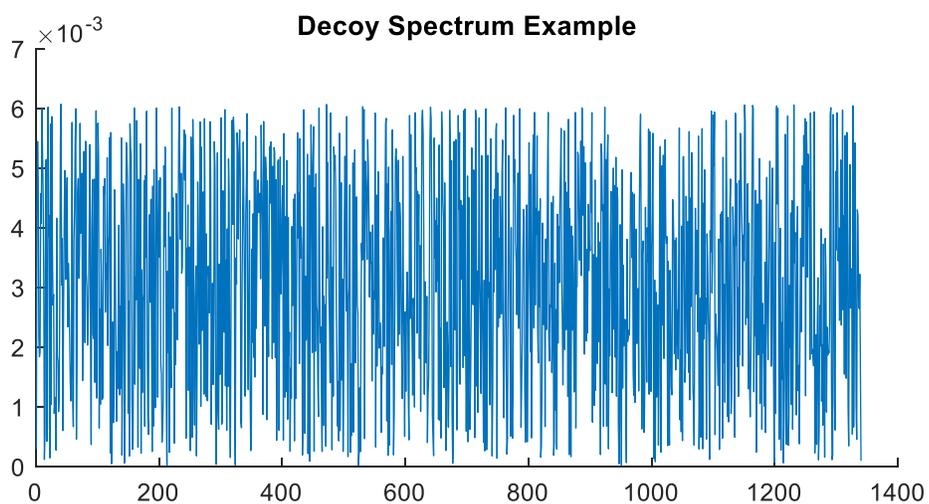


Figure 6.16: Example of a generated Decoy Spectrum. Purposefully noisy and drawn from a random uniform distribution.

It is important to note that the decoy data cluster around the class 3 or the background data in LDA-Space. This is likely because the background data have the highest levels of noise in the measurements as would be suspected. Each iteration 1340 decoy spectra are generated and added to previous iteration's decoy data. This way, improvements are cumulative and build upon previous iterations.

Following generation of decoy data, each individual decoy spectrum is perturbed such that its position in LDA-space is at the mean of a chosen class. The decoy spectra are divided up and an equal number is perturbed to each of the three classes.

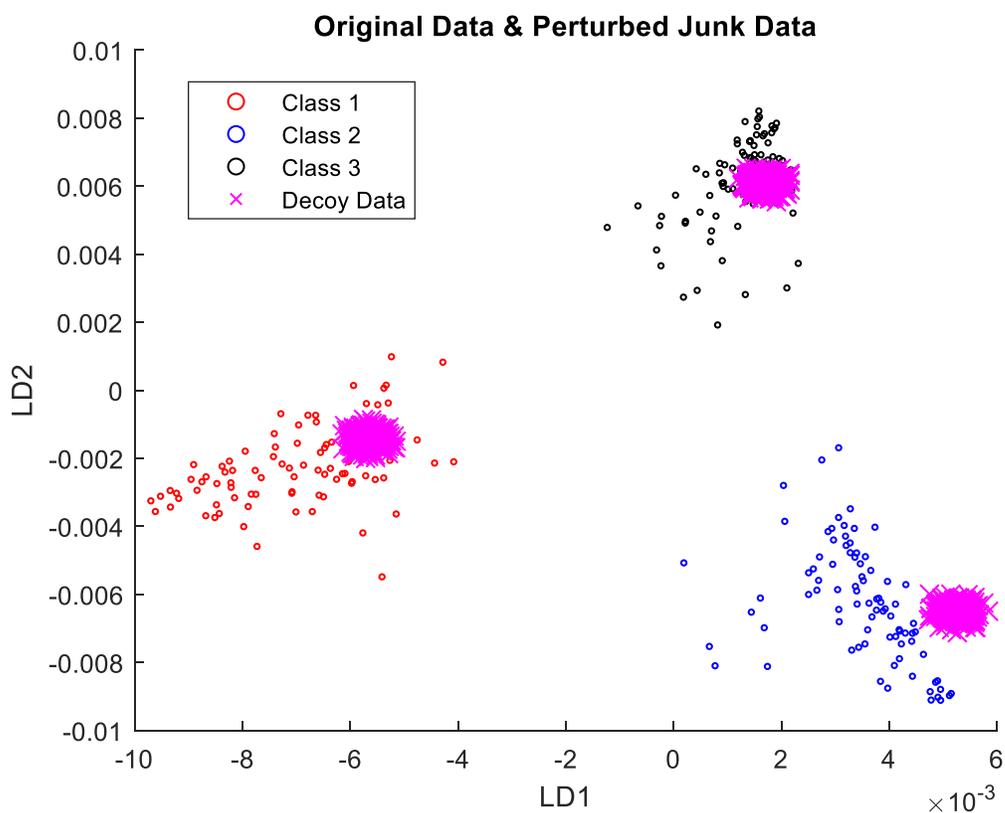


Figure 6.17: The location of the decoy data following perturbation to the mean of each class in LDA-space. The optimal perturbation was calculated using Equation 7.4.

Following the perturbation of the decoy data, LDA was performed again on the whole data set (Raman spectra & decoy spectra). This time designating all decoy data as a 4th class. LDA needs to find the optimal eigenvectors that give the best resolution for all classes. By designating the perturbed decoy data which now overlaps each of the 3 classes as a 4th class, LDA is forced to

rely on the underlying spectral features and not just the noise as it has previously done in the first iteration.

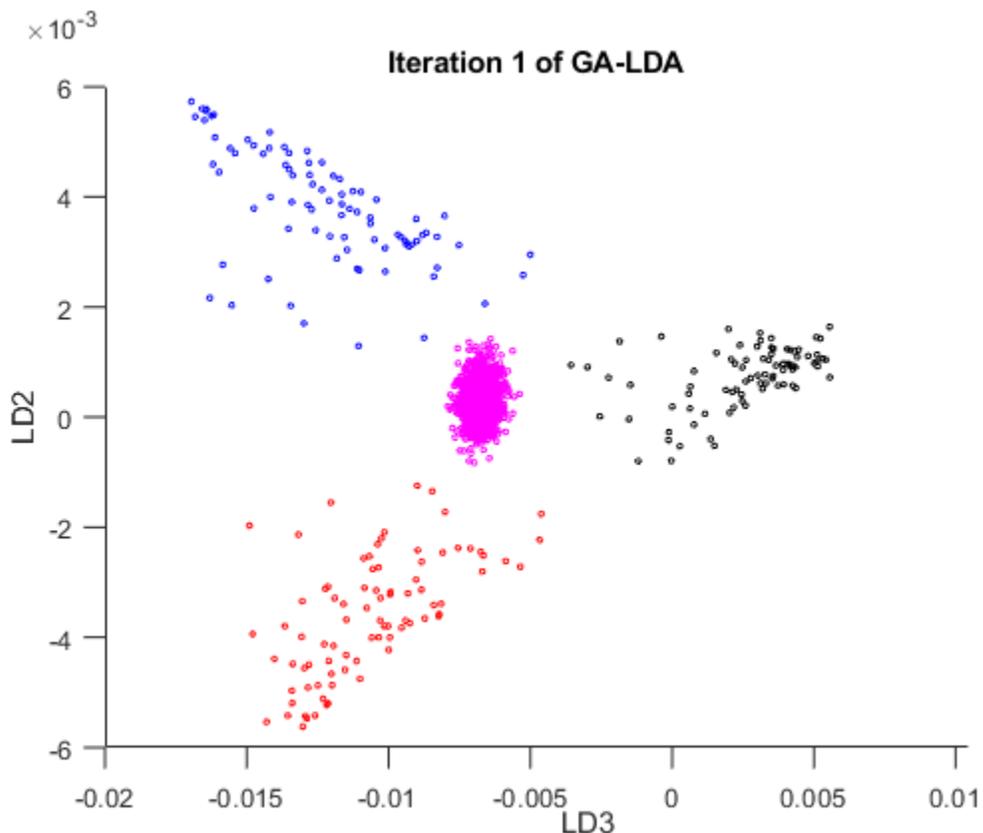


Figure 6.18: LDA-space after the first iteration of GA-LDA. The 2nd and 3rd LDA axis are shown in this plot. Decoy data are shown in purple, with the original Raman spectra shown in blue, black, and red.

Figure 6.18 shows that the data are still separated according to each class. However, the resolution on a whole has decreased. Each of the three original classes are significantly closer to one other than on the original LDA projection. This decrease in resolution is an expected result of improvement of the loading plots. Originally, LDA was able to artificially create a higher resolution through the noise levels found in the original loading plots, see **Figure 6.14**. Each of the three classes share many features in spectral space, which is why the lower resolution is probably closer to the real resolving power of LDA upon this data set. Upon closer inspection, the decoy data are very well separated from the rest of the data set via the 1st LDA axis. **Figure 6.19**

shows the 3D scatter plot, clearly demonstrating the high level of separation of the decoy data from the Raman spectra.

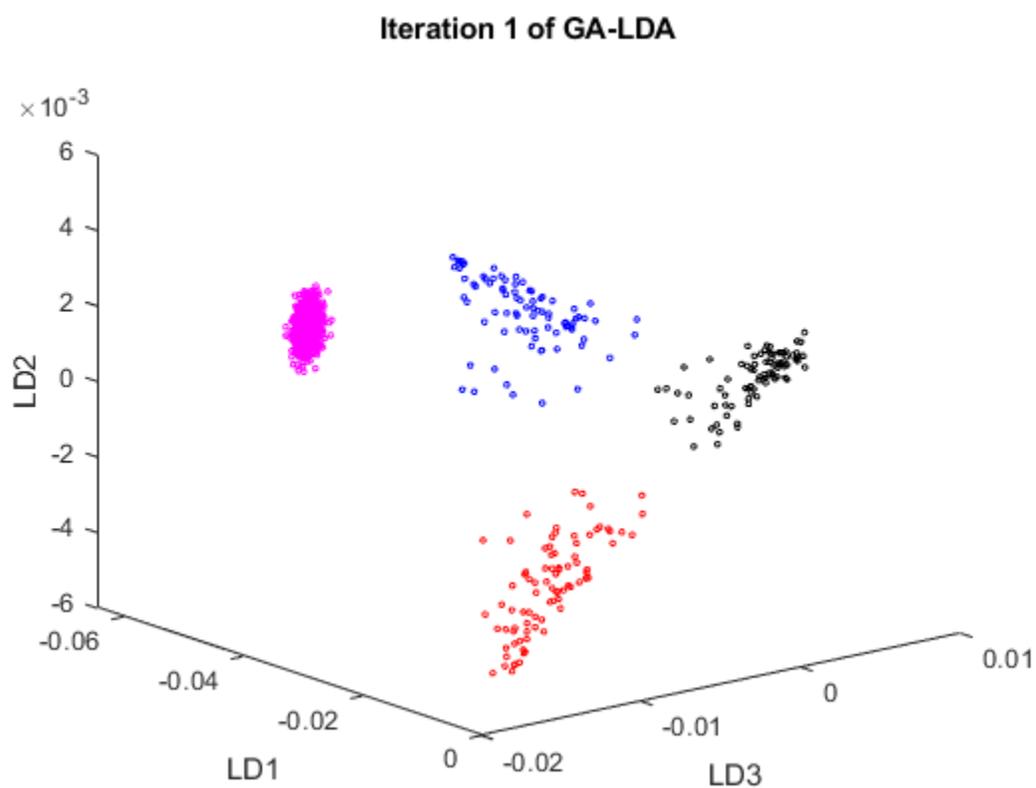


Figure 6.19: Another view of the first iteration of GA-LDA. This time showing all three LDA axis to demonstrate the high degree of separation of decoy data from measured Raman spectra.

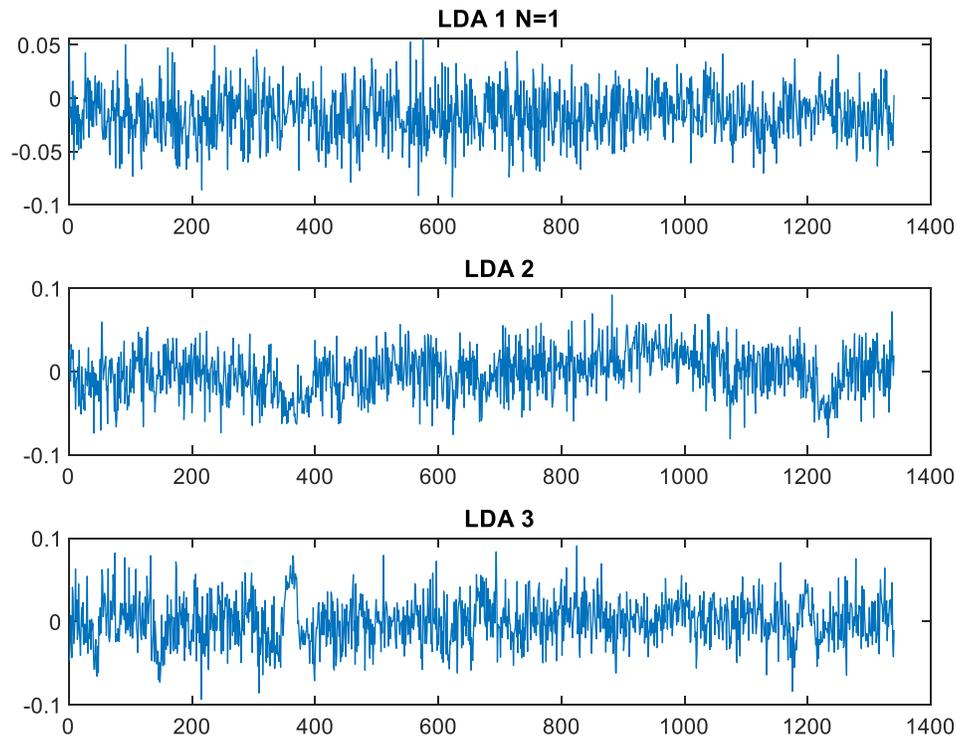


Figure 6.20: The loading plots following the first iteration of the GA-LDA. Note the improvement over the original loading plots when LDA was first performed.

Obvious improvements to the loading plots can be seen after just the first iteration. Spectral features can start to be seen in each of the 3 eigenvectors. Further improvements can be expected with more iterations of the algorithm. 20 iterations of the algorithm were performed in approximately one minute using a standard desktop computer. Results of the 20 iterations can be seen in the following figures.

Iteration 20 of GA-LDA

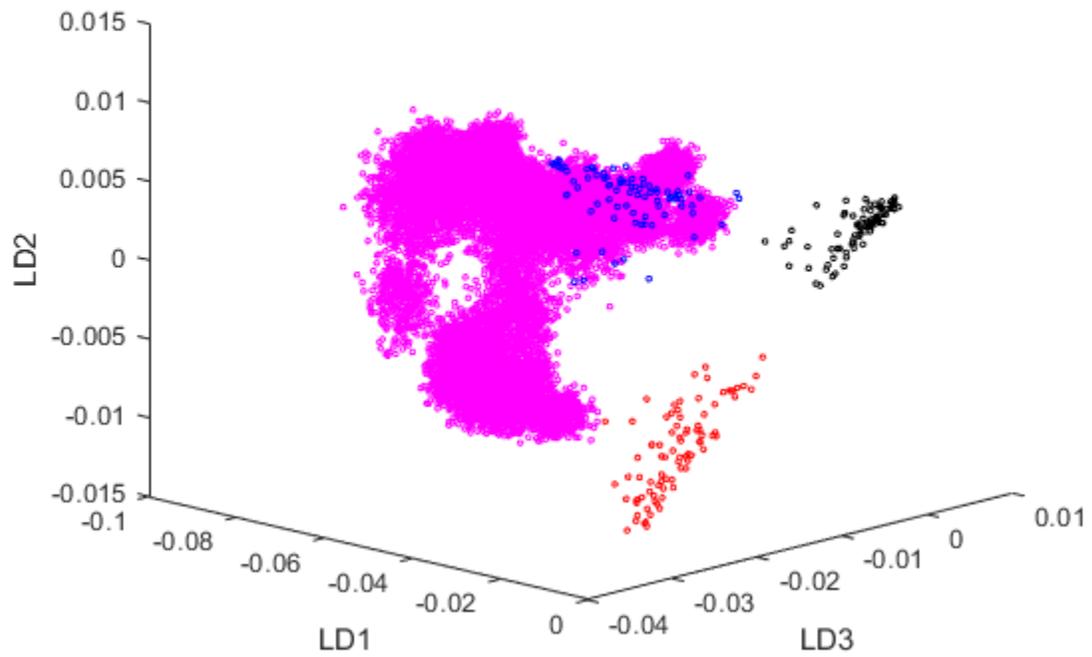


Figure 6.21: 3D LDA-space after the 20th iteration of GA-LDA. Decoy data are shown in purple, with the original Raman spectra shown in blue, black, and red. The first eigenvector provides a large discrimination power to separate the decoy data from the Raman spectra.

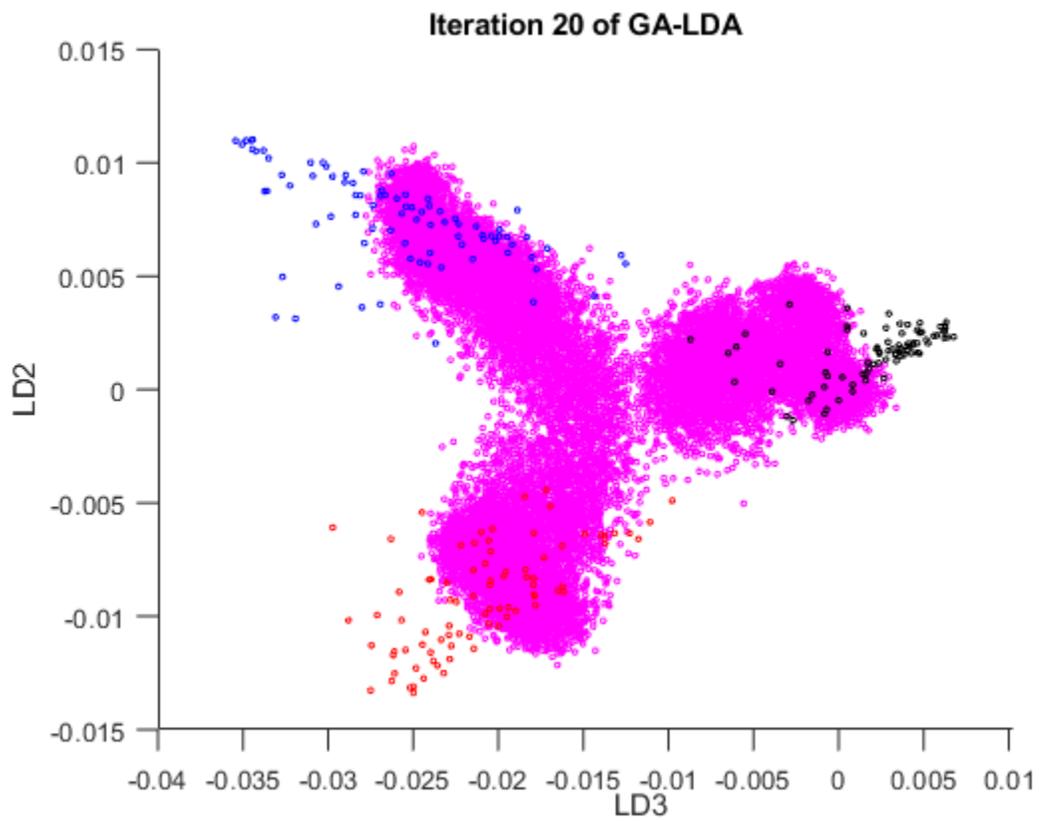


Figure 6.22: Another view of the decoy data in purple and the original Raman spectra in blue, black, and red. The 2nd and 3rd eigenvectors best show the resolution in the Raman spectra.

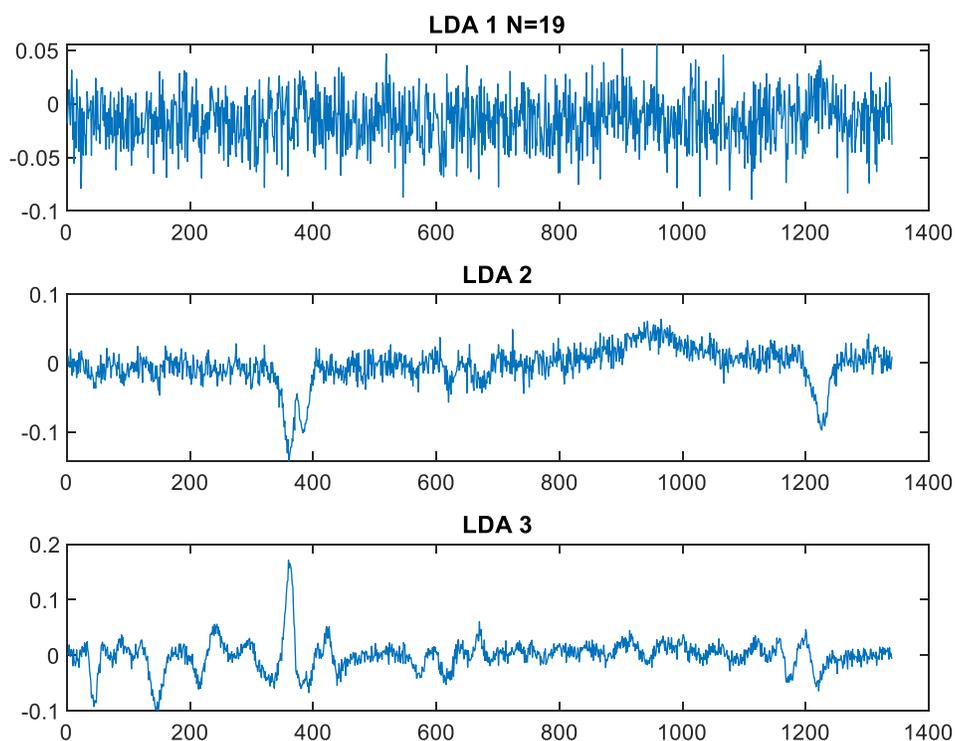


Figure 6.23: The loading plots (eigenvectors) after 20 iterations of the GA-LDA algorithm. Significant improvements have been made over the original LDA loading plots. The 2nd and 3rd eigenvectors now show well defined spectral features that correspond to the spectral features found in each of the 3 original classes of Raman spectra.

Inspection of the loading plots (eigenvectors) following 20 iterations of the GA-LDA algorithm show significant improvements over the original LDA loading plots. The 2nd and 3rd eigenvectors now contain well defined spectral features that correspond to features found in the measured Raman spectra. The 2nd and 3rd eigenvectors show evidence of smoothing, e.g. reduction of noise. Further smoothing can be obtained by performing more iterations; however, spectral features appear to not improve significantly after the 20th iteration of the algorithm. The 1st eigenvector remains noisy. The eigenvectors are ranked in order of which ones provide the largest degree of separation. One might ask why the noisy eigenvector provides the largest resolution. When the decoy data (which is primarily made up of noise) is considered, the noisy eigenvector makes sense. This 1st eigenvector provides the large separation between the decoy data and the measured Raman spectra.

6.6 Conclusions

The GA-LDA algorithm shows a definitive improvement to the loading plots, both reducing the noise and increasing the spectral features found within the eigenvectors. The time and resources required are relatively low such that this technique could be easily adopted to most applications where LDA analysis is used for decision making.

Additionally, GA-LDA provides insight into how machine learning techniques might be applied to chemical problems. GA-LDA is a linear approach that is relatively straightforward to understand and could provide a gateway for researchers to move into non-linear techniques (such as neural networks and generative adversarial networks) for more complicated problems.

APPENDIX A. CODE

CHAPTER 3 CODE: NON-NEGATIVE MATRIX FACTORIZATION FILTER

```
clear all
% close all

PerformAutocorrelation = 0;
Reducing = 0;
datatype = 1; %0 for .bin and 1 for xls
plot_irf = 0;
plot_events = 0;
%% Load data
% loads in data from the alazar cards, rescales to voltage, and splits the
% data into individual hits from tribo data collection
if datatype == 0
    Filename = '\\10.164.16.153\Data\Scott
Griffin\deconvolution\DoxycyclineHyclate\08-Aug-
2017_16_12_28Channel_B_Raw_Data.bin'; % '\\10.164.16.153\Data\Scott
Griffin\deconvolution\DoxycyclineHyclate\08-Aug-
2017_16_12_28Channel_B_Raw_Data.bin';
    fid = fopen(Filename);
    filter = 100; %size of the filter you want to use
    Range = 200; %input range for alazar card
    Buffers = 35; %amount of concatenated hits
    corrected_split_data = zeros(2450000,Buffers); %preallocation for speed
    reduction = 1; %for if you want to reduce the data for speed, not really
    necessary anymore
    Data = fread(fid,'uint16');
    %% Convert into voltage and make positive
    Rescaled_Data = ((Data / 2^15) * -Range) + Range; %converts to voltage, flips
to positive for ease of visualization, and shifts baseline towards 0
    split_data = reshape(Rescaled_Data,[],Buffers); %splits up the data into
the individual impact events
    for i = 1:Buffers;
        baseline = mean(split_data(1:40000,i)); %finds the mean of the first
millisecond of data in each trace
        corrected_split_data(:,i) = (split_data(:,i)-baseline); %subtracts the
baseline mean from the corresponding trace and saves them
    end
end

% loads in data from csv files
if datatype == 1
    Filename = '\\10.164.16.153\Data\Scott Griffin\deconvolution\ROC
plot\TimeTraces&GroundTruth\TimeTrace9.csv';
    Rescaled_Data = csvread(Filename,0,0);
    Filename2 = '\\10.164.16.153\Data\Scott Griffin\deconvolution\ROC
plot\TimeTraces&GroundTruth\GroundTruthTimeTrace9.csv';
    GroundTruthData = csvread(Filename2,0,0);
    filter = 100;
    reduction = 1;
    Buffers = 1;
    corrected_split_data = Rescaled_Data;
```

```

end
IRFs = xlsread('\\10.164.16.153\Data\Scott
Griffin\deconvolution\Gaus_IRF.xlsx');
%% This section cuts out the important part of the data for the smashing event
for i = 1:Buffers
    maximums(:,i) = max(corrected_split_data(:,i)); % finds max values for each
hit
end

for i = 1:Buffers % This loop finds the beginning and end of each event
    temp3 = find(corrected_split_data(:,i) == maximums(i)); %pulls out the
correct max value for the impact event
    if size(temp3) > 1
        value1(i) = 0;
        value2(i) = 0;
    elseif size(temp3) == 1
        value1(i) = find(corrected_split_data(:,i) == maximums(i));
        value2(i) = find(corrected_split_data(:,i) == maximums(i));
        while corrected_split_data(value1(i),i) > 0
            if corrected_split_data(value1(i),i) > 0
                value1(i) = value1(i) - 1; % stored values for the beginning of
events
            end
        end
        while corrected_split_data(value2(i),i) > 0
            if corrected_split_data(value2(i),i) > 0
                value2(i) = value2(i) + 1; % stored values for the end of events
            end
        end
    end
end
recovered_events = cell(1,Buffers);
recovered_irf = cell(1,Buffers);
for n = 1:Buffers
    if value1(n) == 0
        n = n + 1;
    end
    Reduced_Data = corrected_split_data;
%% Create baseline correcting high-pass filter and apply it
    x_hp = linspace(-3,3);
    pdf = normpdf(x_hp,0,1);
    pdf_normalized = -pdf/sum(pdf);
    impulse = zeros(1,100);
    for i = 1:length(x_hp)
        if i == 50
            impulse(1,i) = abs(sum(pdf_normalized));
        else
            impulse(1,i) = 0;
        end
    end
    final_filter = pdf_normalized + impulse;
    Reduced_Data = conv(final_filter,Reduced_Data); %applies the filter
    Reduced_Data = Reduced_Data(50:length(Reduced_Data)-50); %gets rid of the
extra points added by the convolution
    %% Fit
    filtersize = filter/reduction; %size of data transient
    irf_results = zeros(filtersize,1); %preallocate memory for speed

```



```

threshold      NegIndex = find(C_prime<Threshold); %discard below the

                if isempty(NegIndex)
                    More = 0;
                else
                    P_prime = P_prime(:,Index);
                end
            end
            conc_prime = P_prime\data_prime;
        %% Puts the photon amplitudes together with the arrival times
        Ampl_prime = zeros(filtersize,1);
        for j = 1:size(P_prime,2)
            index = find(P_prime(:,j)==decay_guess(1));
            Ampl_prime(index)= conc_prime(j);
        end
        if isempty(P_prime)
            else
                data(startpoint:endpoint) = data(startpoint:endpoint) -
(P_prime(:,1)*Ampl_prime(1));
            end
            data_fit(startpoint,1) = Ampl_prime(1); %creates the final
array of recovered photon amplitudes and arrival times
            new_start = find(Ampl_prime ~= 0);
            new_start = new_start(new_start(:) >= 2);
            if isempty(new_start) == 1 || new_start(end) <= 1
                startpoint = startpoint + filtersize;
                endpoint = endpoint + filtersize;
            else
                startpoint = new_start(1) + startpoint - 1 ;
                endpoint = endpoint + new_start(1) - 1;
            end
        end
    end

    I = eye(filtersize,filtersize); %identity matrix for calculating
the weighted convolution matrix
    M = conv2(data_fit,I); % new 'P' matrix for recovering the IRF
    M = M(1:length(M)-(length(I)-1),:); %needs to be cut down to be the
proper length
    %% calculates the irf and decides if the fit is good enough
    irf = M\Reduced_Data;
    irf = irf/sum(irf);
    if counter >= 1
        Past_Diff = Diff;
        Past_nonzero = nonzero;
    end
    Diff = norm(irf - decay_guess); %Euclidian distance between the
recovered IRF and the guess used for the convergence condition
    nonzero = length(find(data_fit>Threshold)); %count how many photon
events were recovered
    if counter >= 1
        comparison_Diff = Diff - Past_Diff;
        if comparison_Diff < 0
            condition1 = 1;
        else
            condition1 = 0;
        end
    end
end
end

```

```

        if condition1 == 1
            decay_guess = irf;
        else
            keepgoing = 0;
            final_Diff = Past_Diff;
            final_nonzero = Past_nonzero;
            [irf2,final_nonzero2,data_fit_final] =
NNMF_IRF_testing_part2(irf,Reduced_Data,Threshold); %Sends the results to a
final iteration of the algorithm
            %% 'annealing' step where a non-converging result has noise
added to help it converge by moving it away from the local minima
            if final_nonzero2 < 1500 || sum(isnan(irf2)) > 0
                if sum(isnan(irf2)) > 0
                    decay_guess = IRFs(k,:)'/5 + normrnd(0,1,100,1);
                    keepgoing = 1;
                else
                    keepgoing = 1;
                    decay_guess = irf2/5 + normrnd(0,1,100,1);
                end
            else
                keepgoing = 0;
            end
        end
        counter = counter + 1
    end
    results(:,k) = vertcat(MuGuess,final_Diff,final_nonzero2);
    data_results(:,k) = data_fit_final;
    irf_results(:,k) = irf2;
end
minimums = min(results,[],2);
location = find(results(2,:) == minimums(2));
recovered_events{: ,n} = data_results(:,location);
recovered_irf{: ,n} = irf_results(:,location);
end

%% autocorrelation or recovered data
if PerformAutocorrelation == 1
    FT1=fft(data_results(:,location)); %FT along vib mirror time axis
    FTFT1=gmultiply(FT1,conj(FT1)); %FTFT* = square mag FT
    clear FT1;
    AC1=real(ifft(FTFT1)); %iFT(FTFT*)=AC
    clear FTFT1;
    AC_deconvolved = AC1(1:floor(length(data)/2));

    figure3 = figure;
    plot(AC_deconvolved,'LineWidth',2);
    xlabel({'Data'},'FontSize',14);
    ylabel({'Autocorrelogram'},'FontSize',14);
    title({'Autocorrelogram of deconvolved data'},'FontSize',14);
end

```

CHAPTER 4 CODE: CALIBRATION FREE SHG IMAGE ANALYSIS

```

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Instructions %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%       This program is designed to correct bias in particle counting
%       experiments. ImageJ should be used before this program. This program

```

```

% requires the summary and results from ImageJ to work.
%
% ImageJ instructions:
% 1- Make a stack using the brightest frame from the z-scan measurement
% 2- Make a duplicate of your stack (Ctrl-Shift-D)
% 3- Perform an intensity threshold on the duplicate stack (Ctrl-Shift-T)
% 4- Analyze -> Set Measurements: have the following selected
%     a) Area
%     b) Mean gray value
%     c) Stack position
%     d) Redirect to: Stack (or whatever the name of the original is)
% 5- Analyze -> Analyze Particles
%     a) Size Threshold (3 or less pixels) or (0.0003 in^2) *sometimes pixel
%         units don't work on ImageJ
%     b) Display results
%     c) Summarize
%     d) Include holes
% 6- Save Summary in an excel file (Name it "Summary.xlsx")
% 7- Save Results in a text file (excel file is also acceptable) (Name it
%     "Results.xlsx" or "Results.txt")
% 8- Make sure these two files are in the same folder
% 9- Determine the average particle size using ImageJ (Place that data in
%     this script at the appropriate spot)
% 10-
%
% Known bugs:
% 1- When the slice names of the stack of images are numbers and not
%     text. The program will not run properly. To remedy this, go to
%     line 135 and uncomment the line below it and comment out the
%     original line
% 2- The results file must only have 4 columns (with column 4 being
%     the slice information (which image it comes from) for that
%     particle. Check line 103 (xlsx) or 91 (txt).

filepath = '\\10.164.16.153\Data\Scott
Griffin\AbbVie\6_29_2018\plateID_31\'; %don't forget last '\\'

filename_results = '1786_Results.xlsx'; %text file or xlsx?
Results_is_text_file = 0; %0 = no, 1 = yes

filename_summary = '1786_Summary.xlsx';

filename_output = 'Matlab Output.xlsx';

Average_Particle_Size = 14.55; %Average Area
FOV = 512^2; % Field of View: this needs to be in the same units as the average
area

%this functionality currently doesn't work
experimentaldata = 0; % make this 0 if not experimental data or if you don't
want to use the blank correction

%% Open Files
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% End of User Input Section %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

```

```

%%%%%%%%%% Do not change any of the following inputs %%%%%%%%%%%
%

file_results = strcat(filepath, filename_results);
file_summary = strcat(filepath, filename_summary);
file_output = strcat(filepath, filename_output);

A0 = Average_Particle_Size / FOV; %Units must match!
C = 0.842;

%experimental data
blankavecount = 1;
blankavevolume = 2.104 * 10^-5;

if experimentaldata == 0;
    blankavecount = 0;
    blankavevolume = 0;
end

%Open up the text file (with the individual particle data)

if Results_is_text_file == 1;
    fid = fopen(file_results);
    text_file_data = textscan(fid, '%f %f %f %f');
    fclose('all');
    particle_data_temp = text_file_data{2};
    intensity_data = text_file_data{3};
    particle_position = text_file_data{4};

    particle_data = zeros(length(particle_data_temp),1);
    for i = 1:length(particle_data_temp);
        particle_data(i,1) = particle_data_temp(i);
    end
end

if Results_is_text_file == 0;
    Results = xlsread(file_results);
    particle_data = Results(:,2);
    intensity_data = Results(:,3);
    particle_position = Results(:,4);
end

[summary, summarytxt, summaryraw] = xlsread(file_summary);

%column 1 measured count
%column 2 measured total area
%column 3 measured average area
%column 4 measured %Area
%column 5 measured mean intensity

datasize = size(summary);
outputsize = datasize + 1;
outputsize = num2str(outputsize(1,1));
datasize2 = length(particle_data);
datasize = num2str(datasize(1,1));

```

```

range = strcat('B1:', 'B', datasize);

counts_per_image = xlsread(file_summary, range);
%% Particle Data manipulations
% Calculate Volumes for each image
end_point = str2double(datasize); %check this value of plus one not sure why it
is here

radius = (particle_data./pi).^(1/2);
volumes = (4/3) .* pi .* radius.^3;
areas = pi .* radius.^2;
volume = zeros(end_point,1);
area = zeros(end_point,1);

%% Corrections to recover n

N = zeros(end_point,1);

for j = 1:end_point;
    N(j) = summary(j,1);
    %N(j) = summary(j,2); % for when the first column in summaryraw
    %is numbers and not text
end

n_correction = zeros(end_point,1);
for k = 1:end_point;
    n_correction(k) = ((1 - (1 - 8*A0*N(k)).^5)/(4*A0)) - blankavecount;
end

%Add in an if functionality that prevents the previous equation from
%working if A0*N = .1 or smaller. This will prevent the equation from
%breaking down prematurely.

%% Sum up the intensities and volumes of the particles
intensity = zeros(end_point,1);

for state = 1:end_point;
    for m = 1:datasize2;
        if particle_position(m) == state;
            volume(state) = volume(state) + volumes(m);
            intensity(state) = intensity(state) +
(intensity_data(m)*particle_data(m));
        end
    end
end

%% v correction

v_correction = zeros(end_point,1);

for p = 1:end_point;
    v_correction(p) = volume(p) * (1 - 8*A0*(n_correction(p) - 1) * (1 -
C)); %the 4 was changed to an 8 and now it works

```

```

    v_correction(p) = v_correction(p) - blankavevolume;
end

%% Volume Fraction

volume_fraction = zeros(end_point,1);

for p = 1:end_point;
    volume_fraction(p) = v_correction(p) * sqrt(pi/(A0 * FOV^3));
end

%% Output Files

output = strcat('H2:H',outputsize);
output2 = strcat('I2:I',outputsize);
output3 = strcat('J2:J',outputsize);
output4 = strcat('L2:L',outputsize);
output5 = strcat('N2:N',outputsize);
output7 = strcat('A2:F',outputsize);

xlswrite(file_output, summaryraw, output7);
xlswrite(file_output, n_correction, output);
xlswrite(file_output, volume, output2);
xlswrite(file_output, v_correction, output3);
xlswrite(file_output, volume_fraction, output4);
xlswrite(file_output, intensity, output5);

Titles = [{'Particle Count'}, {'Total Area'}, {'Mean Area'}, {'% Area'}, {'Mean
Int'}, {''}, {'Recov #'}, {'Meas Vol'}, {'Recov Vol'}, {''}, {'Vol Fraction'},
{''}, (Molnar et al.)];
xlswrite(file_output, Titles, 'Sheet1', 'B1:N1');

```

CHAPTER 6 CODE: ADVERSARIAL SPECTROSCOPY

```

%% Generative Adversarial Spectroscopy

%Written by Casey Smith
    %1st edition - 5/14/2019
    %Updated - 11/15/2019: Target changed from mean of target class, to
    %individual spectrum of target class. See line 129.
    %Updated - 11/19/2019: Updated LDAfunc to LDAfunc2 which now outputs
    %eigenvalues too

%Overview:
    %Perform LDA on a data set
    %Generate junk data
    %Cause junk data to misclassify to various classes
    %Perform LDA again, but now with the junk data as an additional class
    %Iterate (until impatience?)

clearvars;

```

```

plot_figures = 0;

%% Import data set
%This script is designed for spectra to be saved in an excel file, with
%each classes saved on different sheets.

%Retrieve info on sheets (names & number)
InputFilename = 'ProcesseSpec_All.xlsx';
[~, InputFileSheets] = xlsfinfo(InputFilename); % Get info about inp ut spread
sheet for automatic class counting.
SheetNum = numel(InputFileSheets); % Automatically determine the number of
worksheets in input file.

for i = 1:SheetNum
    data{i} = xlsread(InputFilename, InputFileSheets{i}); %Save each sheet in
a different cell of the variable 'data'
    datasize(i,:) = size(data{i});
end

%% LDA Prep
%We need to generate enough additional spectra to make the system
%overdetermined. E.g. 84 spectra with 1340 channels would require an
%additional 1256 spectra to be overdetermined.

for i=1:SheetNum
    meanSpec{i} = mean(data{i},2); %determine the mean spectra of each class
    stdSpec{i} = std(data{i},0,2); %determine the standard deviation at each
wavelength for each class
    sizeInfo{i} = size(data{i}); %save the matrix dimensions of each class to
determine how many spectra need to be generated.
    SpecLength = max(sizeInfo{i});
end

%generate new spectra based off of statistics
for i=1:SheetNum
    SpecOut = zeros(max(sizeInfo{i}));
    %for j=1:max(sizeInfo{i})
    for j=1:3000
        New_Spec = normrnd(meanSpec{i}, stdSpec{i});
        SpecOut(:,j) = New_Spec;
    end
    GeneratedData{i} = horzcat(data{i},SpecOut);
    clear SpecOut New_Spec j k
end

%% Perform LDA

[LDA_Projection, Vecs] = LDAfunc(GeneratedData); %LDA_Projection is all of the
data (including generated data)

%perform projection of just original data
if plot_figures == 1
for i=1:SheetNum
    LDA_Projection2{i} = Vecs*data{i};
    LDA_Projection2{i} = LDA_Projection2{i}';
end

```

```

end
figure;
hold on
clr{1}='ro';
clr{2}='b+';
clr{3}='kd';
clr{4}='ms';
for i=1:SheetNum
    X= LDA_Projection2{i}(:,1);
    Y= LDA_Projection2{i}(:,2);
    scatter(X,Y,10,clr{i});
end
legend('Class 1', 'Class 2', 'Class 3')
xlabel('LD1');
ylabel('LD2');
%title('Original Data Only')
hold off
end
%% Generate Junk Data
loop = 1;
while loop < 21

SavedVecs{loop} = Vecs;

for i = 1:SheetNum
    JunkLowerBound(i) = min(min(data{i})); %determine the lowest value for the
uniform distribution
    JunkUpperBound(i) = max(max(data{i})); %determine the highest value for the
uniform distribution
end

JunkData = random('uniform', min(JunkLowerBound)*.7, max(JunkUpperBound)*.7,
[SpecLength]);
JunkLDA = transpose(Vecs*JunkData);

if plot_figures == 1
%perform projection of original data & junk data
for i=1:SheetNum
    LDA_Projection2{i} = Vecs*data{i};
    LDA_Projection2{i} = LDA_Projection2{i}';
end
figure;
hold on

for i=1:SheetNum
    X= LDA_Projection2{i}(:,1);
    Y= LDA_Projection2{i}(:,2);
    scatter(X,Y,4,clr{i});
end
xlabel('LD1');
ylabel('LD2');
%title('Original Data & Junk Data')
scatter(JunkLDA(:,1),JunkLDA(:,2),4,clr{4});
hold off
end

```

```

%% Classify Junk data to various classes

targetClass = randi(3,[length(JunkData), 1]);
targetSpec = randi(min(min(datasize)), [length(JunkData), 1]);

GenStatus = 1;
parfor i = 1:length(JunkData)
    %JunkData_Perturbed(:,i) = GA_Func(JunkData(:,i),targetClass(i),Vecs,
meanSpec);
    JunkData_Perturbed(:,i) = JunkData(:,i) - ASolution(Vecs', JunkData(:,i),
meanSpec{targetClass(i)});
    %JunkData_Perturbed(:,i) = JunkData(:,i) - ASolution(Vecs', JunkData(:,i),
data{targetClass(i)}(:,targetSpec(i)));
    GenStatus = GenStatus + 1
end

if loop == 1
    GeneratedData{4} = JunkData_Perturbed;
else
    GeneratedData{4} = horzcat(GeneratedData{4},JunkData_Perturbed);
end

if plot_figures == 1
%perform projection of perturbed data & original data
PerturbedLDA = transpose(Vecs*JunkData_Perturbed);
PerturbedLDA = PerturbedLDA';

figure;
hold on
clr{1}='r';
clr{2}='b';
clr{3}='k';
%clr{4}='v';
for i=1:SheetNum
    X= LDA_Projection2{i}(:,1);
    Y= LDA_Projection2{i}(:,2);
    scatter(X,Y,4,clr{i});
end
xlabel('LD1');
ylabel('LD2');
title('Original Data & Perturbed Junk Data')
scatter(PerturbedLDA(1,:),PerturbedLDA(2,:),100,'mx');
hold off
end

%perform LDA again
[LDA_Projection, Vecs] = LDAfunc(GeneratedData);

loop = loop+1

if plot_figures == 2 %For plotting in 3D for a higher level LDA
%Plot the 4 classes - original data (1-3), and the newly reassigned
%JunkData Class

LDA_Projection2{1} = (Vecs*data{1})';

```

```

LDA_Projection2{2} = (Vecs*data{2})';
LDA_Projection2{3} = (Vecs*data{3})';
%LDA_Projection2{4} = (Vecs*JunkData_Perturbed)';
LDA_Projection2{4} = (Vecs*GeneratedData{4})';

figure;
hold on
clr{1}='r';
clr{2}='b';
clr{3}='k';
clr{4}='m';
for i=1:4
    X= LDA_Projection2{i}(:,1);
    Y= LDA_Projection2{i}(:,2);
    Z= LDA_Projection2{i}(:,3);
    scatter3(X,Y,Z,4,clr{i});
    %scatter(Y,Z,4,clr{i});
end
xlabel('LD2');
ylabel('LD3');
zlabel('LD1');
title('Iteration 20 of GA-LDA')
hold off
end

end

%% Plot SavedVecs over time
if plot_figures == 2
for i = 2:length(SavedVecs)
    figure1 = figure;
    subplot(3,1,1)
    plot(SavedVecs{1,i}(1,:))
    title(strcat('LDA 1 N=', num2str(i-1)))
    subplot(3,1,2)
    plot(SavedVecs{1,i}(2,:))
    title('LDA 2')
    subplot(3,1,3)
    plot(SavedVecs{1,i}(3,:))
    title('LDA 3')
    %print(figure1, strcat('C:\Users\CaseyJSmith\Dropbox\Research\Adversarial
Attacks\GAS\Gif2\',num2str(i)), '-dpng', '-r300');
    %close all
end
end
%% Plot some of the junkdata vs junkdata perturbed

% chosen_spectra = 6; %1 to 1340
% figure
% subplot(3,1,1)
% plot(JunkData(:,chosen_spectra))
% title('JunkData')
% subplot(3,1,2)
% plot(JunkData_Perturbed(:,chosen_spectra),clr{targetClass(chosen_spectra)})
% text = strcat('Perturbed JunkData ', num2str(chosen_spectra), ' to Class ',
num2str(targetClass(chosen_spectra)));

```

```

% title(text)
% subplot(3,1,3)
% plot(JunkData(:,chosen_spectra) - JunkData_Perturbed(:,chosen_spectra));
% title('Difference of original & perturbed')

%% Plot the outcome - used early on for proof of concept

% targetclass = 1;
% chosen_spectra = data{3}(:,1);
% sigma = ASolution(Vecs', chosen_spectra, meanSpec{targetclass});
%
% figure
% hold on
% plot(sigma)
% title('Best Mask')
% hold off
%
% figure
% hold on
% final_spec = chosen_spectra - sigma; %the final attacked spectra
% title('Attacked Spectra(r) vs. Original Spectra(k)')
% offset = 0.003; %for clarity of the graph
% plot((final_spec+offset),'r')
% plot(chosen_spectra,'k')
%
% %plot LDA Space
% figure
% hold on
% LDA_attacked = transpose(Vecs*final_spec);
% LDA_initial = transpose(Vecs*chosen_spectra);
% for i = 1:3
%     LDA{i} = transpose(Vecs*data{i});
%     LDA_meanSpec{i} = transpose(Vecs*meanSpec{i});
% end
% title('LDA Space')
% xlabel('LDA 1')
% ylabel('LDA 2')
% plot(LDA{1}(:,1),LDA{1}(:,2),'r.') % class 1
% plot(LDA{2}(:,1),LDA{2}(:,2),'g.') % class 2
% plot(LDA{3}(:,1),LDA{3}(:,2),'k.') % class 3
% plot(LDA_meanSpec{1}(:,1),LDA_meanSpec{1}(:,2),'or') % class 1 mean
% plot(LDA_meanSpec{2}(:,1),LDA_meanSpec{2}(:,2),'og') % class 2 mean
% plot(LDA_meanSpec{3}(:,1),LDA_meanSpec{3}(:,2),'ok') % class 3 mean
% plot(LDA_attacked(:,1),LDA_attacked(:,2),'xb','MarkerSize',10) % attack
% plot(LDA_initial(:,1),LDA_initial(:,2),'xm','MarkerSize',10) % initial
%
% legend('Class 1', 'Class 2', 'Class 3', 'Mean Spectra Class 1', 'Mean Spectra
Class 2', 'Mean Spectra Class 3', 'Attacked Spectra', 'Initial Spectra')

```

ASolution – Analytical Solution of optimal perturbation

```

function [sigma] = ASolution(Vecs, victimSpectra, meanSpec_targetClass)
%Analytical Solution: Returns the analytical perturbation for an adversarial
%attack
% Ziyi Cao wrote the code and confirmed the math

```

```

% Garth did the math behind this

% updated by Youlin on 2019-12-19: to allow for multiple number of classes by
matrix operation
sigma = Vecs*Vecs'*(victimSpectra-meanSpec_targetClass) ./ (sum(Vecs.^2,2) +
1);

end

```

LDAfunc – A function that performs LDA

```

function [LDA_Projection,EigenVectors] = LDAfunc(Data)
%LDA: Input the data (each class in it's own cell) and function will output
%the eigenvectors.
% Detailed explanation goes here

SheetNum = numel(Data);
N = zeros(SheetNum, 1); % Number of spectra in each class.
for i=1:SheetNum
    [SpecLength, N(i)] = size(Data{i}); % SpecLength is the number of wavelengths
recorded in each spectrum.
end

MeanSpec = zeros(SpecLength, SheetNum); % First index - Wavelengths; Second
index - SheetNum. It is a bit counter-intuitive, but it makes sense if you write
the entire matrix down.
VarSpec = zeros(SheetNum, SpecLength, SpecLength);
for i=1:SheetNum
    MeanSpec(:, i) = mean(Data{i}, 2);
    VarSpec(i, :, :) = cov(Data{i}'); %generates the within-class variance
matrix for each i class, using Matlab built-in covariance function.
end

%Determine the total within-class variance matrix
S_w = squeeze(sum(VarSpec, 1)); % squeeze() function is used here to get rid of
the first (SheetNum) dimension in VarSpec.

%Determine the total weighted mean spectrum

Sum_spec_tot=sum(MeanSpec.*repmat(N', SpecLength, 1), 2); % Using A.*repmat(B)
can achieve multiplication by row or by column.
Num_spec_tot=sum(N);
Mu_tot = Sum_spec_tot / Num_spec_tot; %this algorithm agrees with MathCad

S_B = zeros(SpecLength);
for i = 1:SheetNum
    S_B = S_B + (MeanSpec(:, i) - Mu_tot)*(MeanSpec(:, i) - Mu_tot)'*(N(i));
end
S_B = S_B/Num_spec_tot;

%DETERMINE THE PROJECTIONS THAT MAXIMIZE THE FISHER LINEAR DISCRIMINANT
%[Vecs,Vals] = eigs(S_B/S_w, SheetNum-1); %note - eig() returns eigenvals as
the diagonal

```

```

%[Vecs,Vals] = eigs(S_w\S_B, SheetNum-1); %note - eig() returns eigenvals as
the diagonal
[Vecs,Vals] = eigs(inv(S_w)*S_B,SheetNum-1);
% NOTE: Newer Matlab version encourages usage of "Matrix Division" instead of
multiplication with inv(). This is because inv() function in Matlab is slow,
buggy, and unreliable due to historical reasons. In most cases inv() should
work just fine. If this can cause confusion, please revert S_B/S_w back to
inv(S_w)*S_B.

%elements of a matrix in ascending order, with the most positively valued
eigenvalues given last. Note - inv(X) returns the inverse matrix of X.

%Next transform the original data by projection onto the linear
discriminant coordinates. We'll constrain the plotting to the first two
(assuming at least three classes of data).

%Project the raw data onto the selected principal components within
'transform' and write to an output file

for i=1:SheetNum
    LDA_subset{i}= Vecs'*Data{i};
    LDA_subset_trans{i}=LDA_subset{i}'; %Transpose operation arranges the
    %data as XY pairs for 2D plotting and XYZ triplets for 3D plotting.
end

LDA_Projection = LDA_subset_trans;
Eigenvectors = Vecs';
end

```

APPENDIX B. BUILD LIST FOR TL INSTRUMENT

#	Part	Vendor	ID
1	PMT	Hamamatsu	H10721-210
Many	Microscope Slides (cheap)	Fisher Scientific	S67112A
1	Oscilloscope	Picotech	PP958 - 300 Series Model 3203D
1	Flange Nut (Thompson)	Motion Industries (Lafayette based)	MTS10X2M
1	Lead Screw (Thompson)	Motion Industries (Lafayette based)	SPT10X2M
1	Coupling	Misumu	CPCX16-5-5
1	Stepper Motor	Adafruit	324
1	Stepper Shield	Adafruit	1438
2	Deep Groove Ball Bearing	VBX Ball Bearing	627-2Z
1	Arduino Board	Adafruit	
Several	Base microscope slides	Ted Pella	260226
Various	Rods		
2	Lenses	Thorlabs	Appropriate Optics. . .

PUBLICATIONS

1. Smith, C. J.; Griffin, S. R.; Eakins, G.; Deng, F.; White, J. K.; Thirunahari, S.; Ramakrishnan, S.; Sangupta, A.; Zhang, S.; Novak, J.; Liu, Z.; Rhodes, T.; Simpson, G. “Triboluminescence from Pharmaceutical Formulations”; *Anal. Chem.*, **2018**, 90 (11), pp 6893-6898
2. Smith, C. J.; Dinh, J.; Schmitt, P.; Stroud, P. A.; Hinds, J.; Johnson, M.; Simpson, G. J.; “Calibration-Free SHG Image Analysis for Quantification of Trace Crystallinity within Final Dosage Forms of Amorphous Solid Dispersions”; *Applied Spectroscopy* **2018**. <https://doi.org/10.1177/0003702818786506>.
3. Griffin, S. R.; Biechele-Speziale, J. A.; Smith, C. J.; You-Dow, X.; White, J. K.; Zhang, S.; Novak, J.; Liu, Z.; Simpson, G. J.; “Iterative Non-Negative Matrix Factorization Filter for Blind Deconvolution in Photon/Ion Counting”; *Anal. Chem.*, **2019**, 91 (8), pp 5286-5294. DOI: 10.1021/acs.analchem.9b00120.
4. Song, Z.; Sarkar, S.; Vogt, A.; Danzer, G.; Smith, C. J.; Gualtieri, E.; Simpson, G. J.; “Kinetic Modeling of Accelerated Stability Testing Enabled by Second Harmonic Generation Microscopy”; *Anal. Chem.*, **2018**, 90 (7), pp 4406-4413
5. Sherman, A. M.; Geiger, A. C.; Smith, C. J.; Taylor, L.; Hinds, J.; Stroud, P. A.; Simpson, G. J.; “Stochastic Differential Scanning Calorimetry by Nonlinear Optical Microscopy”; *Anal. Chem.*, **2020**, 92, 1, 1171-1178

REFERENCES

- Aitipamula, S., Banerjee, R., Bansal, A. K., Biradha, K., Cheney, M. L., Choudhury, A. R., Desiraju, G. R., Dikundwar, A. G., Dubey, R., Duggirala, N., Ghogale, P. P., Ghosh, S., Goswami, P. K., Goud, N. R., Jetti, R., Karpinski, P., Kaushik, P., Kumar, D., Kumar, V., Moulton, B., Mukherjee, A., Mukherjee, G., Myerson, A. S., Puri, V., Ramanan, A., Rajamannar, T., Reddy, C. M., Rodriguez-Hornedo, N., Rogers, R. D., Row, T. N. G., Sanphui, P., Shan, N., Shete, G., Singh, A., Sun, C. Q. C., Swift, J. A., Thaimattam, R., Thakur, T. S., Thaper, R. K., Thomas, S. P., Tothadi, S., Vangala, V. R., Variankaval, N., Vishweshwar, P., Weyna, D. R., & Zaworotko, M. J. (2012). Polymorphs, Salts, and Cocrystals: What's in a Name? *Crystal Growth & Design*, *12*(5), 2147-2152. <https://doi.org/10.1021/cg3002948>
- Alhijjaj, M., Belton, P., Fabian, L., Wellner, N., Reading, M., & Qi, S. (2018). Novel Thermal Imaging Method for Rapid Screening of Drug–Polymer Miscibility for Solid Dispersion Based Formulation Development. *Molecular pharmaceuticals*, *15*(12), 5625-5636.
- Alhijjaj, M., Reading, M., Belton, P., & Qi, S. (2015). Thermal analysis by structural characterization as a method for assessing heterogeneity in complex solid pharmaceutical dosage forms. *Analytical chemistry*, *87*(21), 10848-10855.
- Alhijjaj, M., Yassin, S., Reading, M., Zeitler, J. A., Belton, P., & Qi, S. (2017). Characterization of heterogeneity and spatial distribution of phases in complex solid dispersions by thermal analysis by structural characterization and X-ray micro computed tomography. *Pharmaceutical research*, *34*(5), 971-989.
- Andronis, V., Yoshioka, M., & Zograf, G. (1997). Effects of sorbed water on the crystallization of indomethacin from the amorphous state. *Journal of Pharmaceutical Sciences*, *86*(3), 346-351. <https://doi.org/10.1021/js9602711>
- Babu, N. J., & Nangia, A. (2011). Solubility Advantage of Amorphous Drugs and Pharmaceutical Cocrystals. *Crystal Growth & Design*, *11*(7), 2662-2679. <https://doi.org/10.1021/cg200492w>
- Baghel, S., Cathcart, H., & O'Reilly, N. J. (2016). Polymeric Amorphous Solid Dispersions: A Review of Amorphization, Crystallization, Stabilization, Solid-State Characterization, and Aqueous Solubilization of Biopharmaceutical Classification System Class II Drugs. *Journal of Pharmaceutical Sciences*, *105*(9), 2527-2544. <https://doi.org/10.1016/j.xphs.2015.10.008>
- Bauer, J., Spanton, S., Henry, R., Quick, J., Dziki, W., Porter, W., & Morris, J. (2001). Ritonavir: An extraordinary example of conformational polymorphism. *Pharmaceutical Research*, *18*(6), 859-866. <https://doi.org/10.1023/a:1011052932607>
- Becker, W., Bergmann, A., Hink, M. A., König, K., Benndorf, K., & Biskup, C. (2004). Fluorescence lifetime imaging by time-correlated single-photon counting. *Microscopy Research and Technique*, *63*(1), 58-66. <https://doi.org/10.1002/jemt.10421>
- Berendt, R. T., & Munson, E. J. (2012). Effect of enantiomeric ratio and preparation method on proline crystal form. *Crystengcomm*, *14*(7), 2479-2488. <https://doi.org/10.1039/c2ce06445f>
- Berendt, R. T., Sperger, D. M., Isbester, P. K., & Munson, E. J. (2006). Solid-state NMR spectroscopy in pharmaceutical research and analysis. *Trac-Trends in Analytical Chemistry*, *25*(10), 977-984. <https://doi.org/10.1016/j.trac.2006.07.006>

- Berry, D. J., Seaton, C. C., Clegg, W., Harrington, R. W., Coles, S. J., Horton, P. N., Hursthouse, M. B., Storey, R., Jones, W., Friscic, T., & Blagden, N. (2008). Applying hot-stage microscopy to co-crystal screening: A study of nicotinamide with seven active pharmaceutical ingredients. *Crystal Growth & Design*, 8(5), 1697-1712. <https://doi.org/10.1021/cg800035w>
- Brown, R. M., Millard, A. C., & Campagnola, P. J. (2003). Macromolecular structure of cellulose studied by second-harmonic generation imaging microscopy. *Optics Letters*, 28(22), 2207-2209.
- Bugay, D. E. (2001). Characterization of the solid-state: spectroscopic techniques. *Advanced Drug Delivery Reviews*, 48(1), 43-65. [https://doi.org/10.1016/s0169-409x\(01\)00101-6](https://doi.org/10.1016/s0169-409x(01)00101-6)
- Byrn, S. R., Pfeiffer, R. R., Stephenson, G., Grant, D. J. W., & Gleason, W. B. (1994). SOLID-STATE PHARMACEUTICAL CHEMISTRY. *Chemistry of Materials*, 6(8), 1148-1158. <https://doi.org/10.1021/cm00044a013>
- Cheney, M. L., Shan, N., Healey, E. R., Hanna, M., Wojtas, L., Zaworotko, M. J., Sava, V., Song, S., & Sanchez-Ramos, J. R. (2010). Effects of Crystal Form on Solubility and Pharmacokinetics: A Crystal Engineering Case Study of Lamotrigine. *Crystal Growth & Design*, 10(1), 394-405. <https://doi.org/10.1021/cg901010v>
- Chowdhury, A. U., Dettmar, C. M., Sullivan, S. Z., Zhang, S., Jacobs, K. T., Kissick, D. J., Maltais, T., Hedderich, H. G., Bishop, P. A., & Simpson, G. J. (2014). Kinetic trapping of metastable amino acid polymorphs. *Journal of the American Chemical Society*, 136(6), 2404-2412.
- Chowdhury, A. U., Zhang, S. J., & Simpson, G. J. (2016). Powders Analysis by Second Harmonic Generation Microscopy. *Analytical Chemistry*, 88(7), 3853-3863. <https://doi.org/10.1021/acs.analchem.5b04942>
- Clas, S.-D., Dalton, C. R., & Hancock, B. C. (1999). Differential scanning calorimetry: applications in drug development. *Pharmaceutical science & technology today*, 2(8), 311-320.
- Closser, R. G., Gualtieri, E. J., Newman, J. A., & Simpson, G. J. (2013). Characterization of salt interferences in second-harmonic generation detection of protein crystals. *Journal of Applied Crystallography*, 46, 1903-1906. <https://doi.org/10.1107/s0021889813027581>
- Coello, C. A. C. (2002). Theoretical and numerical constraint-handling techniques used with evolutionary algorithms: a survey of the state of the art. *Computer Methods in Applied Mechanics and Engineering*, 191(11-12), 1245-1287. [https://doi.org/10.1016/s0045-7825\(01\)00323-1](https://doi.org/10.1016/s0045-7825(01)00323-1)
- Correa-Soto, C., Trasi, N. S., Schmitt, P. D., Su, Y. C., Liu, Z., Miller, E., Variankaval, N., Marsac, P. J., Simpson, G. J., & Taylor, L. S. (2017). Second harmonic generation microscopy as a tool for the early detection of crystallization in spray dried dispersions. *Journal of Pharmaceutical and Biomedical Analysis*, 146, 86-95. <https://doi.org/10.1016/j.jpba.2017.07.066>
- de Juan, A., Jaumot, J., & Tauler, R. A. (2014). Multivariate Curve Resolution (MCR). Solving the mixture analysis problem. *Analytical Methods*, 6(14), 4964-4976. <https://doi.org/10.1039/c4ay00571f>

- Dow, X. Y., Dettmar, C. M., DeWalt, E. L., Newman, J. A., Dow, A. R., Roy-Chowdhury, S., Coe, J. D., Kupitz, C., Fromme, P., & Simpson, G. J. (2016). Second harmonic generation correlation spectroscopy for characterizing translationally diffusing protein nanocrystals. *Acta Crystallographica Section D-Structural Biology*, 72, 849-859. <https://doi.org/10.1107/s205979831600841x>
- Dow, X. Y., DeWalt, E. L., Sullivan, S. Z., Schmitt, P. D., Ulcickas, J. R., & Simpson, G. J. (2016). Imaging the nonlinear susceptibility tensor of collagen by nonlinear optical stokes ellipsometry. *Biophysical journal*, 111(7), 1361-1374.
- Ehrentreich, F., & Summchen, L. (2001). Spike removal and denoising of Raman spectra by wavelet transform methods. *Analytical Chemistry*, 73(17), 4364-4373. <https://doi.org/10.1021/ac0013756>
- Elkhabaz, A., Sarkar, S., Dinh, J. K., Simpson, G. J., & Taylor, L. S. (2018). Variation in Supersaturation and Phase Behavior of Ezetimibe Amorphous Solid Dispersions upon Dissolution in Different Biorelevant Media. *Molecular Pharmaceutics*, 15(1), 193-206. <https://doi.org/10.1021/acs.molpharmaceut.7b00814>
- Fei, R. X., Li, W. B., Li, J., & Yang, L. (2015). Giant piezoelectricity of monolayer group IV monochalcogenides: SnSe, SnS, GeSe, and GeS. *Applied Physics Letters*, 107(17), Article 173104. <https://doi.org/10.1063/1.4934750>
- Fish, D. A., Brinicombe, A. M., Pike, E. R., & Walker, J. G. (1995). BLIND DECONVOLUTION BY MEANS OF THE RICHARDSON-LUCY ALGORITHM. *Journal of the Optical Society of America a-Optics Image Science and Vision*, 12(1), 58-65. <https://doi.org/10.1364/josaa.12.000058>
- Francis, A. T., Nguyen, T. T., Lamm, M. S., Teller, R., Forster, S. P., Xu, W., Rhodes, T., Smith, R. L., Kuiper, J., & Su, Y. (2018). In Situ Stimulated Raman Scattering (SRS) Microscopy Study of the Dissolution of Sustained-Release Implant Formulation. *Molecular pharmaceutics*, 15(12), 5793-5801.
- Gaisford, S., Kett, V., & Haines, P. (2016). *Principles of thermal analysis and calorimetry*. Royal society of chemistry.
- Galland, A., Dupray, V., Berton, B., Morin-Grognet, S., Sanselme, M., Atmani, H., & Coquerel, G. (2009). Spotting conglomerates by second harmonic generation. *Crystal Growth and Design*, 9(6), 2713-2718.
- Guedes, F. L., de Oliveira, B. G., Hernandez, M. Z., De Simone, C. A., Veiga, F. J., Maria do Carmo, A., Pitta, I. R., Galdino, S. L., & Neto, P. J. R. (2011). Solid dispersions of imidazolidinedione by PEG and PVP polymers with potential antischistosomal activities. *Aaps Pharmscitech*, 12(1), 401-410.
- Hardy, G. E., Kaska, W. C., Chandra, B., & Zink, J. I. (1981). Triboluminescence-structure relationships in polymorphs of hexaphenylcarbodiphosphorane and anthranilic acid, molecular crystals, and salts. *Journal of the American Chemical Society*, 103(5), 1074-1079.
- Hartshorn, C. M., Lee, Y. J., Camp, C. H., Jr., Liu, Z., Heddleston, J., Canfield, N., Rhodes, T. A., Walker, A. R. H., Marsac, P. J., & Cicerone, M. T. (2013). Multicomponent Chemical Imaging of Pharmaceutical Solid Dosage Forms with Broadband CARS Microscopy. *Analytical Chemistry*, 85(17), 8102-8111. <https://doi.org/10.1021/ac400671p>
- He, Y., & Ho, C. (2015). Amorphous Solid Dispersions: Utilization and Challenges in Drug Discovery and Development. *Journal of Pharmaceutical Sciences*, 104(10), 3237-3258. <https://doi.org/10.1002/jps.24541>

- Heinz, A., Strachan, C. J., Gordon, K. C., & Rades, T. (2009). Analysis of solid-state transformations of pharmaceutical compounds using vibrational spectroscopy. *Journal of Pharmacy and Pharmacology*, 61(8), 971-988.
- Hsu, H. Y., Toth, S., Simpson, G. J., Taylor, L. S., & Harris, M. T. (2013). Effect of substrates on naproxen-polyvinylpyrrolidone solid dispersions formed via the drop printing technique. *Journal of Pharmaceutical Sciences*, 102(2), 638-648. <https://doi.org/10.1002/jps.23397>
- Höhne, G. W. H., Hemminger, W., & Flammersheim, H.-J. (1996). Theoretical fundamentals of differential scanning calorimeters. In *Differential Scanning Calorimetry* (pp. 21-40). Springer.
- Jackson, M. J., Kestur, U. S., Hussain, M. A., & Taylor, L. S. (2015). Dissolution of danazol amorphous solid dispersions: supersaturation and phase behavior as a function of drug loading and polymer type. *Molecular pharmaceuticals*, 13(1), 223-231.
- Johari, G. P., Ram, S., Astl, G., & Mayer, E. (1990). CHARACTERIZING AMORPHOUS AND MICROCRYSTALLINE SOLIDS BY CALORIMETRY. *Journal of Non-Crystalline Solids*, 116(2-3), 282-285. [https://doi.org/10.1016/0022-3093\(90\)90703-o](https://doi.org/10.1016/0022-3093(90)90703-o)
- Kestur, U. S., Wanapun, D., Toth, S. J., Wegiel, L. A., Simpson, G. J., & Taylor, L. S. (2012). Nonlinear optical imaging for sensitive detection of crystals in bulk amorphous powders. *Journal of pharmaceutical sciences*, 101(11), 4201-4213.
- Khaliullin, R. Z., Eshet, H., Kühne, T. D., Behler, J., & Parrinello, M. (2011). Nucleation mechanism for the direct graphite-to-diamond phase transition. *Nature materials*, 10(9), 693.
- Kissick, D. J., Muir, R. D., & Simpson, G. J. (2010). Statistical Treatment of Photon/Electron Counting: Extending the Linear Dynamic Range from the Dark Count Rate to Saturation. *Analytical Chemistry*, 82(24), 10129-10134. <https://doi.org/10.1021/ac102219c>
- Kissick, D. J., Wanapun, D., & Simpson, G. J. (2011). Second-Order Nonlinear Optical Imaging of Chiral Crystals. In R. G. Cooks & E. S. Yeung (Eds.), *Annual Review of Analytical Chemistry*, Vol 4 (Vol. 4, pp. 419-437). <https://doi.org/10.1146/annurev.anchem.111808.073722>
- Konno, H., & Taylor, L. S. (2008). Ability of different polymers to inhibit the crystallization of amorphous felodipine in the presence of moisture. *Pharmaceutical research*, 25(4), 969-978.
- Law, D., Schmitt, E. A., Marsh, K. C., Everitt, E. A., Wang, W., Fort, J. J., Krill, S. L., & Qiu, Y. (2004). Ritonavir-PEG 8000 amorphous solid dispersions: in vitro and in vivo evaluations. *Journal of pharmaceutical sciences*, 93(3), 563-570.
- Liland, K. H., Almoy, T., & Mevik, B. H. (2010). Optimal Choice of Baseline Correction for Multivariate Calibration of Spectra. *Applied Spectroscopy*, 64(9), 1007-1016. <https://doi.org/10.1366/000370210792434350>
- Liu, H., Kwon, O.-H., Tang, J., & Zewail, A. H. (2014). 4D Imaging and Diffraction Dynamics of Single-Particle Phase Transition in Heterogeneous Ensembles. *Nano Letters*, 14(2), 946-954. <https://doi.org/10.1021/nl404354g>
- Madden, F. N., Godfrey, K. R., Chappell, M. J., Hovorka, R., & Bates, R. A. (1996). A comparison of six deconvolution techniques. *Journal of Pharmacokinetics and Biopharmaceutics*, 24(3), 283-299. <https://doi.org/10.1007/bf02353672>

- Molnar, C., Jermyn, I. H., Kato, Z., Rahkama, V., Ostling, P., Mikkonen, P., Pietiainen, V., & Horvath, P. (2016). Accurate Morphology Preserving Segmentation of Overlapping Cells based on Active Contours. *Scientific Reports*, 6, Article 32412. <https://doi.org/10.1038/srep32412>
- Muir, R. D., Kissick, D. J., & Simpson, G. J. (2012). Statistical connection of binomial photon counting and photon averaging in high dynamic range beam-scanning microscopy. *Optics Express*, 20(9). <https://doi.org/10.1364/oe.20.010406>
- Muir, R. D., Sullivan, S. Z., Oglesbee, R. A., & Simpson, G. J. (2014). Synchronous digitization for high dynamic range lock-in amplification in beam-scanning microscopy. *Review of Scientific Instruments*, 85(3), 033703.
- Natarajan, C. M., Tanner, M. G., & Hadfield, R. H. (2012). Superconducting nanowire single-photon detectors: physics and applications. *Superconductor Science & Technology*, 25(6), Article 063001. <https://doi.org/10.1088/0953-2048/25/6/063001>
- Newman, A., Knipp, G., & Zografi, G. (2012). Assessing the performance of amorphous solid dispersions. *Journal of pharmaceutical sciences*, 101(4), 1355-1377.
- Newman, J. A., Schmitt, P. D., Toth, S. J., Deng, F. Y., Zhang, S. J., & Simpson, G. J. (2015). Parts per Million Powder X-ray Diffraction. *Analytical Chemistry*, 87(21), 10950-10955. <https://doi.org/10.1021/acs.analchem.5b02758>
- Ng, Y. C., Yang, Z., McAuley, W. J., & Qi, S. (2013). Stabilisation of amorphous drugs under high humidity using pharmaceutical thin films. *European Journal of Pharmaceutics and Biopharmaceutics*, 84(3), 555-565. <https://doi.org/10.1016/j.ejpb.2013.01.008>
- Novakovic, D., Saarinen, J., Rojalin, T., Antikainen, O., Fraser-Miller, S. J., Laaksonen, T., Peltonen, L., Isomaki, A., & Strachan, C. J. (2017). Multimodal Nonlinear Optical Imaging for Sensitive Detection of Multiple Pharmaceutical Solid-State Forms and Surface Transformations. *Analytical Chemistry*, 89(21), 11460-11467. <https://doi.org/10.1021/acs.analchem.7b02639>
- Offerdahl, T. J., Salsbury, J. S., Dong, Z. D., Grant, D. J. W., Schroeder, S. A., Prakash, I., Gorman, E. M., Barich, D. H., & Munson, E. J. (2005). Quantitation of crystalline and amorphous forms of anhydrous neotame using C-13 CPMAS NMR spectroscopy. *Journal of Pharmaceutical Sciences*, 94(12), 2591-2605. <https://doi.org/10.1002/jps.20469>
- Oxtoby, D. W. (1998). Nucleation of first-order phase transitions. *Accounts of chemical research*, 31(2), 91-97.
- Paatero, P., & Tapper, U. (1994). POSITIVE MATRIX FACTORIZATION - A NONNEGATIVE FACTOR MODEL WITH OPTIMAL UTILIZATION OF ERROR-ESTIMATES OF DATA VALUES. *Environmetrics*, 5(2), 111-126. <https://doi.org/10.1002/env.3170050203>
- Pack, B. W., Babayan, Y., Schrad, M. A., Stroud, P. A., Sperry, D. C., White, K. C., & Aburub, A. (2017). Development of an in vivo-relevant drug product performance method for an amorphous solid dispersion. *Journal of Pharmaceutical and Biomedical Analysis*, 142, 307-314. <https://doi.org/10.1016/j.jpba.2017.05.029>
- Piehler, T. N., DeLucia, F. C., Munson, C. A., Homan, B. E., Miziolek, A. W., & McNesby, K. L. (2005). Temporal evolution of the laser-induced breakdown spectroscopy spectrum of aluminum metal in different bath gases. *Applied Optics*, 44(18), 3654-3660. <https://doi.org/10.1364/ao.44.003654>

- Raimi-Abraham, B. T., Moffat, J. G., Belton, P. S., Barker, S. A., & Craig, D. Q. M. (2014). Generation and Characterization of Standardized Forms of Trehalose Dihydrate and Their Associated Solid-State Behavior [Article]. *Crystal Growth & Design*, 14(10), 4955-4967. <https://doi.org/10.1021/cg5004885>
- Rani, M., Govindarajan, R., Surana, R., & Suryanarayanan, R. (2006). Structure in dehydrated trehalose dihydrate - Evaluation of the concept of partial crystallinity [Article]. *Pharmaceutical research*, 23(10), 2356-2367. <https://doi.org/10.1007/s11095-006-9058-6>
- Reading, M. (2017). Thermal Analysis by Structural Characterization (TASC): Structural and Thermo-Rheological Information from Hot Stage Microscopy. *Microscopy Today*, 25(5), 18-23.
- Reading, M., Morton, M., Antonijevic, M., Grandy, D., Hourston, D., & Lacey, A. (2014). New methods of thermal analysis and chemical mapping on a micro and nano scale by combining microscopy with image analysis. *Microscopy: advances in scientific research and education. Formatex Research Center*, 1083-1089.
- Richardson, W. H. (1972). BAYESIAN-BASED ITERATIVE METHOD OF IMAGE RESTORATION. *Journal of the Optical Society of America*, 62(1), 55-+. <https://doi.org/10.1364/josa.62.000055>
- Rumondor, A. C., Stanford, L. A., & Taylor, L. S. (2009). Effects of polymer type and storage relative humidity on the kinetics of felodipine crystallization from amorphous solid dispersions. *Pharmaceutical research*, 26(12), 2599.
- Savitzky, A., & Golay, M. J. E. (1964). SMOOTHING + DIFFERENTIATION OF DATA BY SIMPLIFIED LEAST SQUARES PROCEDURES. *Analytical Chemistry*, 36(8), 1627-&. <https://doi.org/10.1021/ac60214a047>
- Scarborough, N. M., Godaliyadda, G., Ye, D. H., Kissick, D. J., Zhang, S., Newman, J. A., Sheedlo, M. J., Chowdhury, A., Fischetti, R. F., & Das, C. (2017). Synchrotron X-Ray Diffraction Dynamic Sampling for Protein Crystal Centering. *Electronic Imaging*, 2017(17), 6-9.
- Schmitt, P. D., DeWalt, E. L., Dow, X. Y., & Simpson, G. J. (2016). Rapid Discrimination of Polymorphic Crystal Forms by Nonlinear Optical Stokes Ellipsometric Microscopy. *Analytical Chemistry*, 88(11), 5760-5768. <https://doi.org/10.1021/acs.analchem.6b00057>
- Schmitt, P. D., Trasi, N. S., Taylor, L. S., & Simpson, G. J. (2015). Finding the Needle in the Haystack: Characterization of Trace Crystallinity in a Commercial Formulation of Paclitaxel Protein-Bound Particles by Raman Spectroscopy Enabled by Second Harmonic Generation Microscopy. *Molecular Pharmaceutics*, 12(7), 2378-2383. <https://doi.org/10.1021/acs.molpharmaceut.5b00065>
- Schneider, C. A., Rasband, W. S., & Eliceiri, K. W. (2012). NIH Image to ImageJ: 25 years of image analysis. *Nature Methods*, 9(7), 671-675. <https://doi.org/10.1038/nmeth.2089>
- Selfridge, A. R. (1985). Approximate material properties in isotropic materials. *IEEE transactions on sonics and ultrasonics*, 32(3), 381-394.
- Shah, B., Kakumanu, V. K., & Bansal, A. K. (2006). Analytical techniques for quantification of amorphous/crystalline phases in pharmaceutical solids. *Journal of Pharmaceutical Sciences*, 95(8), 1641-1665. <https://doi.org/10.1002/jps.20644>
- Shamblin, S. L., & Zograf, G. (1999). The effects of absorbed water on the properties of amorphous mixtures containing sucrose. *Pharmaceutical research*, 16(7), 1119-1124.

- Sheth, A. R., Lubach, J. W., Munson, E. J., Muller, F. X., & Grant, D. J. W. (2005). Mechanochromism of piroxicam accompanied by intermolecular proton transfer probed by spectroscopic methods and solid-phase changes. *Journal of the American Chemical Society*, 127(18), 6641-6651. <https://doi.org/10.1021/ja045823t>
- Simon, F., Clevers, S., Dupray, V., & Coquerel, G. (2015). Relevance of the second harmonic generation to characterize crystalline samples. *Chemical Engineering & Technology*, 38(6), 971-983.
- Singh, S. K. (2018). Sucrose and Trehalose in Therapeutic Protein Formulations. In *Challenges in Protein Product Development* (pp. 63-95). Springer.
- Skilling, J., & Bryan, R. K. (1984). MAXIMUM-ENTROPY IMAGE-RECONSTRUCTION - GENERAL ALGORITHM. *Monthly Notices of the Royal Astronomical Society*, 211(1), 111-&. <https://doi.org/10.1093/mnras/211.1.111>
- Sleutel, M., Lutsko, J., Van Driessche, A. E., Durán-Olivencia, M. A., & Maes, D. (2014). Observing classical nucleation theory at work by monitoring phase transitions with molecular precision. *Nature communications*, 5, 5598.
- Smith, C. J., Dinh, J., Schmitt, P. D., Stroud, P. A., Hinds, J., Johnson, M., & Simpson, G. (2018). EXPRESS: Calibration-Free Second Harmonic Generation (SHG) Image Analysis for Quantification of Trace Crystallinity within Final Dosage Forms of Amorphous Solid Dispersions. *Applied spectroscopy*, 0003702818786506.
- Smith, C. J., Dinh, J., Schmitt, P. D., Stroud, P. A., Hinds, J., Johnson, M. J., & Simpson, G. J. (2018). Calibration-Free Second Harmonic Generation (SHG) Image Analysis for Quantification of Trace Crystallinity Within Final Dosage Forms of Amorphous Solid Dispersions. *Applied Spectroscopy*, 72(11), 1594-1605. <https://doi.org/10.1177/0003702818786506>
- Smith, C. J., Griffin, S. R., Eakins, G. S., Deng, F., White, J. K., Thirunahari, S., Ramakrishnan, S., Sangupta, A., Zhang, S., Novak, J., Liu, Z., Rhodes, T., & Simpson, G. J. (2018). Triboluminescence from Pharmaceutical Formulations (vol 90, pg 6893, 2018). *Analytical Chemistry*, 90(21), 13131-13131. <https://doi.org/10.1021/acs.analchem.8b04287>
- Smith, D. A., McKenzie, G., Jones, A. C., & Smith, T. A. (2017). Analysis of time-correlated single photon counting data: a comparative evaluation of deterministic and probabilistic approaches. *Methods and Applications in Fluorescence*, 5(4), Article 042001. <https://doi.org/10.1088/2050-6120/aa8055>
- Steed, K. M., & Steed, J. W. (2015). Packing Problems: High Z' Crystal Structures and Their Relationship to Cocrystals, Inclusion Compounds, and Polymorphism. *Chemical Reviews*, 115(8), 2895-2933. <https://doi.org/10.1021/cr500564z>
- Strachan, C. J., Rades, T., Gordon, K. C., & Rantanen, J. (2007). Raman spectroscopy for quantitative analysis of pharmaceutical solids. *Journal of Pharmacy and Pharmacology*, 59(2), 179-192.
- Strachan, C. J., Taday, P. F., Newnham, D. A., Gordon, K. C., Zeitler, J. A., Pepper, M., & Rades, T. (2005). Using terahertz pulsed spectroscopy to quantify pharmaceutical polymorphism and crystallinity. *Journal of pharmaceutical sciences*, 94(4), 837-846.
- Strovel, J., Sittampalam, S., Coussens, N. P., Hughes, M., Inglese, J., Kurtz, A., Andalibi, A., Patton, L., Austin, C., & Baltezor, M. (2016). Early drug discovery and development guidelines: for academic researchers, collaborators, and start-up companies.

- Sullivan, S. Z., Schmitt, P. D., Muir, R. D., DeWalt, E. L., & Simpson, G. J. (2014). Digital Deconvolution Filter Derived from Linear Discriminant Analysis and Application for Multiphoton Fluorescence Microscopy. *Analytical Chemistry*, 86(7), 3508-3516. <https://doi.org/10.1021/ac404150d>
- Sussich, F., Princivalle, F., & Cesàro, A. (1999). The interplay of the rate of water removal in the dehydration of α , α -trehalose. *Carbohydrate research*, 322(1-2), 113-119.
- Sussich, F., Urbani, R., Princivalle, F., & Cesaro, A. (1998). Polymorphic amorphous and crystalline forms of trehalose [Article]. *Journal of the American Chemical Society*, 120(31), 7893-7899. <https://doi.org/10.1021/ja9800479>
- Sweeting, L. M. (2001). Triboluminescence with and without Air. *Chemistry of Materials*, 13(3), 854-870. <https://doi.org/10.1021/cm0006087>
- Taylor, J. R. B., Chan, J. J. M., & Thomas, G. (2015). Wavelet-based blind deconvolution of near-field ultrasound scans. *Iet Image Processing*, 9(8), 672-679. <https://doi.org/10.1049/iet-ipr.2013.0791>
- Taylor, L. S., Williams, A. C., & York, P. (1998). Particle size dependent molecular rearrangements during the dehydration of trehalose dihydrate-in situ FT-Raman spectroscopy. *Pharmaceutical research*, 15(8), 1207-1214.
- Taylor, L. S., & York, P. (1998). Characterization of the phase transitions of trehalose dihydrate on heating and subsequent dehydration. *Journal of pharmaceutical sciences*, 87(3), 347-355.
- Taylor, L. S., & York, P. (1998). Effect of particle size and temperature on the dehydration kinetics of trehalose dihydrate. *International journal of pharmaceuticals*, 167(1-2), 215-221.
- Taylor, L. S., & Zhang, G. G. Z. (2016). Physical chemistry of supersaturated solutions and implications for oral absorption. *Advanced Drug Delivery Reviews*, 101, 122-142. <https://doi.org/10.1016/j.addr.2016.03.006>
- Taylor, L. S., & Zograf, G. (1998). The quantitative analysis of crystallinity using FT-Raman spectroscopy. *Pharmaceutical Research*, 15(5), 755-761. <https://doi.org/10.1023/a:1011979221685>
- Thayer, A. M. (2010). Finding solutions. *Chemical & Engineering News*, 88(22), 13-18.
- Toth, S. J., Madden, J. T., Taylor, L. S., Marsac, P., & Simpson, G. J. (2012). Selective Imaging of Active Pharmaceutical Ingredients in Powdered Blends with Common Excipients Utilizing Two-Photon Excited Ultraviolet-Fluorescence and Ultraviolet-Second Order Nonlinear Optical Imaging of Chiral Crystals. *Analytical Chemistry*, 84(14), 5869-5875. <https://doi.org/10.1021/ac300917t>
- Toth, S. J., Schmitt, P. D., Snyder, G. R., Trasi, N. S., Sullivan, S. Z., George, I. A., Taylor, L. S., & Simpson, G. J. (2015). Ab Initio Prediction of the Diversity of Second Harmonic Generation from Pharmaceutically Relevant Materials. *Crystal Growth & Design*, 15(2), 581-586. <https://doi.org/10.1021/acs.cgd.5b00020>
- Ulcickas, J. R., & Simpson, G. J. (2020). Polarization dependent second harmonic generation microscopy in turbid media. *Advanced Chemical Microscopy for Life Science and Translational Medicine*,
- Vajda, S., Godfrey, K. R., & Valko, P. (1988). NUMERICAL DECONVOLUTION USING SYSTEM-IDENTIFICATION METHODS. *Journal of Pharmacokinetics and Biopharmaceutics*, 16(1), 85-107. <https://doi.org/10.1007/bf01061863>

- Vavasis, S. A. (2009). ON THE COMPLEXITY OF NONNEGATIVE MATRIX FACTORIZATION. *Siam Journal on Optimization*, 20(3), 1364-1377. <https://doi.org/10.1137/070709967>
- Verotta, D. (1993). 2 CONSTRAINED DECONVOLUTION METHODS USING SPLINE FUNCTIONS. *Journal of Pharmacokinetics and Biopharmaceutics*, 21(5), 609-636. <https://doi.org/10.1007/bf01059117>
- Wanapun, D., Kestur, U. S., Kissick, D. J., Simpson, G. J., & Taylor, L. S. (2010). Selective Detection and Quantitation of Organic Molecule Crystallization by Second Harmonic Generation Microscopy. *Analytical Chemistry*, 82(13), 5425-5432. <https://doi.org/10.1021/ac100564f>
- Wanapun, D., Kestur, U. S., Taylor, L. S., & Simpson, G. J. (2011). Single Particle Nonlinear Optical Imaging of Trace Crystallinity in an Organic Powder. *Analytical Chemistry*, 83(12), 4745-4751. <https://doi.org/10.1021/ac1031397>
- Williams, H. D., Trevaskis, N. L., Charman, S. A., Shanker, R. M., Charman, W. N., Pouton, C. W., & Porter, C. J. H. (2013). Strategies to Address Low Drug Solubility in Discovery and Development. *Pharmacological Reviews*, 65(1), 315-499. <https://doi.org/10.1124/pr.112.005660>
- Wunderlich, B. (1958). Theory of cold crystallization of high polymers. *The Journal of Chemical Physics*, 29(6), 1395-1404.
- Yang, Z., Nollenberger, K., Albers, J., Moffat, J., Craig, D., & Qi, S. (2014). The effect of processing on the surface physical stability of amorphous solid dispersions. *European Journal of Pharmaceutics and Biopharmaceutics*, 88(3), 897-908.
- Yu, L. (2001). Amorphous pharmaceutical solids: preparation, characterization and stabilization. *Advanced drug delivery reviews*, 48(1), 27-42.
- Zeitler, J. A., Kogermann, K., Rantanen, J., Rades, T., Taday, P. F., Pepper, M., Aaltonen, J., & Strachan, C. J. (2007). Drug hydrate systems and dehydration processes studied by terahertz pulsed spectroscopy. *International journal of pharmaceutics*, 334(1-2), 78-84.
- Zeitler, J. A., Newnham, D. A., Taday, P. F., Threlfall, T. L., Lancaster, R. W., Berg, R. W., Strachan, C. J., Pepper, M., Gordon, K. C., & Rades, T. (2006). Characterization of temperature-induced phase transitions in five polymorphic forms of sulfathiazole by terahertz pulsed spectroscopy and differential scanning calorimetry. *Journal of pharmaceutical sciences*, 95(11), 2486-2498.
- Zhang, S. J., Song, Z. T., Godaliyadda, G., Ye, D. H., Chowdhury, A. U., Sengupta, A., Buzzard, G. T., Bouman, C. A., & Simpson, G. J. (2018). Dynamic Sparse Sampling for Confocal Raman Microscopy. *Analytical Chemistry*, 90(7), 4461-4469. <https://doi.org/10.1021/acs.analchem.7b04749>
- Zimper, U., Aaltonen, J., McGoverin, C. M., Gordon, K. C., Krauel-Goellner, K., & Rades, T. (2010). Quantification of process induced disorder in milled samples using different analytical techniques. *Pharmaceutics*, 2(1), 30-49.

VITA

Casey Smith was born in Delta, Utah in 1989. He attended Delta High School and graduated as Valedictorian and Student Body President. Following high school, he served a two-year missionary call for the Church of Jesus Christ of Latter-Day Saints in the Philippines. Casey attended the University of Utah and graduated with a B.S. in Chemistry. In his last year at the University of Utah, he met and married Nalani Artinger. He then accepted an offer to attend graduate school at Purdue University and joined Garth Simpson's research group at Purdue. By the grace of God, Casey and Nalani were blessed with a daughter in 2018. Casey graduated with his Ph.D. in Analytical Chemistry from Purdue University in May 2020.