Inverse Scattering with Chemical Composition Constraints for Spectroscopic Tomography

Luke Pfister[1], Yoram Bresler[1], Rohit Bhargava[1,2] and P. Scott Carney[1]
Department of Electrical and Computer Engineering, University of Illinois Urbana-Champaign, Urbana, IL 61801
Department of Bioengineering, University of Illinois Urbana-Champaign, Urbana, IL 61801
lpfiste2@illinois.edu

Abstract: A dramatic reduction in data required for chemically specific 3-D imaging is achieved through prior constraints on the known constituents of the sample. We solve the inverse scattering problem to determine morphology and composition.

1. Introduction

Vibrational spectroscopy using mid-infrared illumination provides a means to determine chemical species without the need for staining or other sample preparation with exogenous chemicals. However, chemically specific 3-D imaging achieved by 3-D point-scanning mid-infrared spectroscopy is slow and cumbersome. Mid-infrared spectroscopy is made practical by assuming the sample is homogeneous and characterized by a single bulk spectrum (thus reducing the imaging problem to a single dimension), or is a thin section exhibiting only transverse heterogeneities (thus reducing the problem to two spatial and one spectral dimension). Moreover, spatial variations of density and composition within the sample can cause diffraction and scattering that distort the measured spectra. In contrast, imaging methods such as optical coherence tomography (OCT), are based on scattering and encode sample morphology in the spectrum of scattered light to provide non-invasive and non-destructive rapid 3-D imaging. However, OCT yields limited chemical specificity [1]. In this work, we present a new approach that lies in between these two extremes to encode some sample morphology in the spectra while still providing chemically specific images.

In many applications, it is reasonable to assume that the sample consists of a finite number of distinct species of materials. We refer to this as the 'n-species' model. We extend previous work [2] that used a similar model in several ways: we (i) use a non-asymptotic forward model; (ii) incorporate sparsity-driven regularization; (iii) demonstrate material-resolved reconstruction of samples with two spatial dimensions (one transverse and depth, easily extended to three spatial dimensions) from data that is not generated according to the first Born approximation; and (iv) refine the conditions for recovery of a sample consisting of *n*-species from interferometric scattering experiments.

2. Forward Problem

We characterize our object through its linear susceptibility $\eta(\mathbf{r}_{||}, z, k_0)$, where $\mathbf{r}_{||} = (x, y)$ represents the transverse dimensions, z is the axial dimension, and k_0 is wavenumber. Under the n-species approximation we have $\eta(\mathbf{r}_{||}, z, k_0) = \sum_{j=1}^{N_s} p_j(\mathbf{r}_{||}, z) h_j(k_0)$. Here, p_j represents the spatial density and h_j is the spectral response of the j-th chemical species.

The sample is placed in an asymmetric Fourier-transform infrared (FTIR) spectroscopic microscope. Light from a broadband source is focused into a Gaussian beam and focused to a variable distance z_F within the sample. The back-scattered light is detected and the complex field of the scattered field is extracted from the interferometric measurements. The beam is scanned over the transverse positions in the sample.

Under the first Born approximation, the complex scattered field is linear in η, or bilinear in the functions p_j and h_j:

$$S(\mathbf{k}_{||}, k_0; z_F) = \sum_{j=1}^{N_s} \int dz \underbrace{\left(\int d^2 k_{||}' \frac{\exp\left\{ i[k_z(\mathbf{k}_{||} - \mathbf{k}_{||}', k_0) + k_z(\mathbf{k}_{||}', k_0)](z - z_F) \right\}}{k_z(\mathbf{k}_{||}', k_0)} \tilde{g}(\mathbf{k}_{||} - \mathbf{k}_{||}', k_0) \tilde{g}(\mathbf{k}_{||}', k_0) \right)}_{\triangleq A(\mathbf{k}_{||}, k_0, z - z_F)} h_j(k_0) p_j(\mathbf{k}_{||}, z) \quad (1)$$

where $\mathbf{k}_{||}$ are the transverse Fourier coordinates, $k_z(\mathbf{k}_{||}, k_0) = \sqrt{k_0^2 - |\mathbf{k}_{||}|^2}$, and $\tilde{g}(\mathbf{k}_{||}, k_0) = (\sqrt{\pi} \cdot k_0 \cdot \mathrm{NA})^{-1} \exp\left(-\frac{1}{2}(\frac{|\mathbf{k}_{||}|}{k_0 \cdot \mathrm{NA}})^2 \right)$ is the Fourier transform of the Gaussian beam profile in the waist plane. We restrict the integral over $\mathbf{k}_{||}'$ to include only propagating waves, for which $k_z(\mathbf{k}_{||}', k_0)$ is purely real.

We obtain measurements at N_k different values of the wavenumber k_0 at each of N_x transverse positions and N_F

focal planes. The forward model (1) can be discretized and written in block matrix form as

$$
\begin{bmatrix} \mathbf{s}_1 \\ \mathbf{s}_2 \\ \vdots \\ \mathbf{s}_{N_F} \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{D}_1\mathbf{A}_1 & \mathbf{D}_2\mathbf{A}_1 & \dots & \mathbf{D}_{N_s}\mathbf{A}_1 \\ \mathbf{D}_1\mathbf{A}_2 & \mathbf{D}_2\mathbf{A}_2 & \dots & \mathbf{D}_{N_s}\mathbf{A}_2 \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{D}_1\mathbf{A}_{N_F} & \mathbf{D}_2\mathbf{A}_{N_F} & \dots & \mathbf{D}_{N_s}\mathbf{A}_{N_F} \end{bmatrix}}_{\triangleq\ \Phi\ \in\ \mathbb{C}^{N_F N_x N_k \times N_x N_z N_s}} \begin{bmatrix} \mathbf{p}_1 \\ \mathbf{p}_2 \\ \vdots \\ \mathbf{p}_{N_s} \end{bmatrix} \tag{2}
$$

The matrix $\mathbf{A}_i \in \mathbb{C}^{N_x N_k \times N_x N_z}$ is the discretized kernel of the integral (1) when the system is focused to the i-th plane and $\mathbf{s}_i \in \mathbb{C}^{N_x N_k}$ is the corresponding measurement vector. The vectors $\mathbf{p}_j \in \mathbb{C}^{N_x N_z}$, $\mathbf{h}_j \in \mathbb{C}^{N_k}$, are discretized versions of the spatial densities and spectral responses, and the diagonal matrix $\mathbf{D}_j = \mathbf{I}_{N_x} \otimes \mathrm{diag}\left\{\mathbf{h}_j\right\} \in \mathbb{C}^{N_k N_x \times N_k N_x}$, with $\mathbf{I}_N$ denoting the $N \times N$ identity matrix, has repeated copies of $\mathbf{h}_j$ along its diagonal.

3. Inverse Problem

Recovering $\mathbf{h}_j$ and $\mathbf{p}_j$ from measurements given by (2) is a *bilinear* inverse problem. Despite recent progress in understanding bilinear inverse problems, they remain challenging both theoretically and computationally. However, there are scenarios in which it is reasonable to assume either exact knowledge of the chemicals in the sample or that the chemicals are drawn from a "dictionary" $\mathbf{H} \in \mathbb{C}^{N_k \times N_s}$ of possible species that has j-th column $\mathbf{h}_j$. Conditioned on the knowledge of the $\mathbf{h}_j$, recovering the $\mathbf{p}_j$ from data generated by (2) reduces to a linear inverse problem. We restrict our attention to this setting.

Let the matrix $\mathbf{P} \in \mathbb{C}^{N_x N_z \times N_s}$ have j-th column given by $\mathbf{p}_j$. We also define $\bar{\mathbf{p}} = \mathrm{vec}\left(\mathbf{P}\right)$ to be the vector formed by stacking the $\mathbf{p}_j$. We wish to recover $\mathbf{P}$ from measurements $\mathbf{s}$ and knowledge of $\mathbf{H}$. Unfortunately, this problem is ill-posed. The diffraction limit ensures that objects with high transverse spatial frequencies lie in the null space of each $\mathbf{A}_i$, and thus in the null space of Φ. However, with sufficiently many focal planes we can recover any object within the optical passband of the system:

Theorem 1. *Let $\mathfrak{B}$ be the space of signals supported within the optical passband of the system and let $\bar{\mathfrak{B}} = \mathbf{I}_{N_s} \otimes \mathfrak{B}$. Given measurements $\mathbf{s}$, there is a unique vector $\tilde{p} \in \bar{\mathfrak{B}}$ such that $\mathbf{s} = \Phi\tilde{\mathbf{p}}$ whenever the $\mathbf{h}_j$ are linearly independent, $N_F \geq N_s$, and $\left|\mathbf{h}_j\right| \geq 0$ for each j, where the magnitude and inequality are taken elementwise.*

The requirement $N_F \geq N_s$ remains burdensome when the chemicals are drawn from a dictionary with many candidate species. If only a small number of chemicals from the dictionary are present in the sample, recovery is possible from $N_F \leq N_s$ focal planes if a stronger condition on the dictionary is met.

Theorem 2. *Let $\mathbf{H}$ be a dictionary of N_s and let $\mathbf{P}$ have at most $k < N_s$ nonzero columns. We can recover $\mathbf{P}$ from measurements given by (2) if every set of $2k$ columns of $\mathbf{H}$ are linearly independent, and $N_F \geq 2k$, and $\left|\mathbf{h}_j\right| \geq 0$ for each j, where the magnitude and inequality are taken elementwise.*

We recover $\mathbf{P}$ by solving a penalized weighted least squares (PWLS) problem

$$
\min_{\mathbf{P}} \frac{1}{2} \sum_{z_F} \|\mathbf{s}_{z_F} - \sum_{j=1}^{N_s} \mathbf{D}_j \mathbf{A}_{z_F} \mathbf{p}_j\|_2^2 + \lambda\,\psi(\mathbf{P}) \tag{3}
$$

where the regularization functional $\psi : \mathbb{C}^{N_x N_z \times N_s} \to \mathbb{R}$ is necessary to overcome the ill-posedness of reconstruction and encourages solutions to obey a particular signal model. For example, if our object consists of point targets, we can take ψ to be a sparsity-promoting functional such as the ℓ_1 norm, whereas total-variation (TV) regularization is well-suited for the reconstruction of piecewise constant objects. If no additional information regarding the structure of $\mathbf{P}$ is available, we default to Tikhonov regularization by taking ψ to be the squared Frobenius norm.

There are a vast number of ways to solve the PWLS problem (3). For Tikhonov regularization, we use the Conjugate-Gradient (CG) algorithm, whereas if ψ is non-differentiable (as is often the case with sparsity-promoting regularization), we use the Alternating Direction Method of Multipliers (ADMM) [3].

4. Simulations

We simulate reconstruction of samples in two spatial dimensions (one transverse and one depth) with significantly different spatial structure. In the first, we generated scattered measurements from a set of point targets according to the Foldy-Lax model. Data was collected from 3 focal planes distributed evenly in $512 \times 512\mu m$ volume, with a $1\mu m$ pixel size. We obtain 256 wavelength samples between 6 to $11\mu m$ and the simulated system used NA $= 0.5$. Three chemical species are present in the volume: octane, isopropanol, and cyclohexanol. The spectra were found via experimental FTIR measurements of bulk samples. The spectra, sample, and reconstructions for Tikhonov and ℓ_1 regularization are shown in Figure 1. While Tikhonov reconstruction accurately recovered the location of the point targets, many
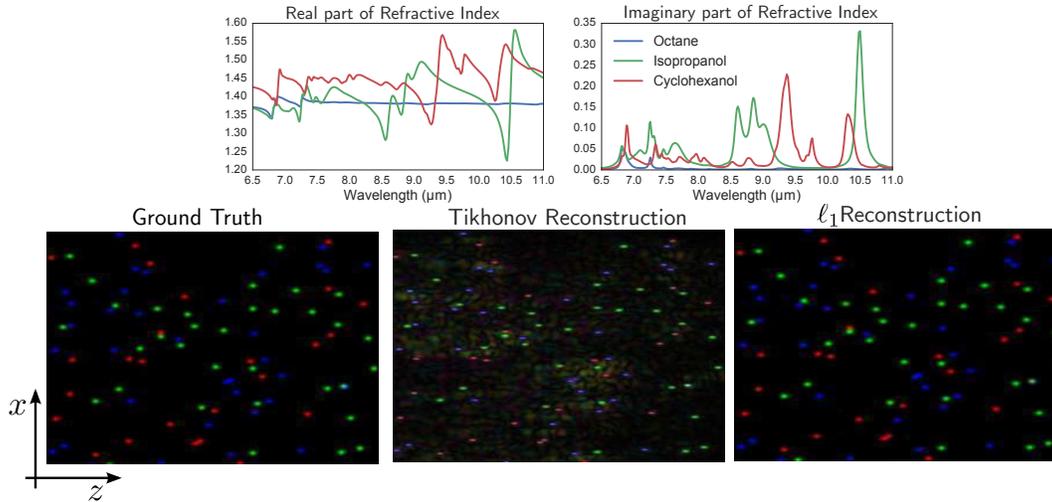
Fig. 1. Reconstruction of point targets. To facilitate display of results, a Gaussian blur is applied to the ground truth and ℓ_1 reconstruction. Blue dots: octane; Green : isopropanol; Red: cyclohexanol
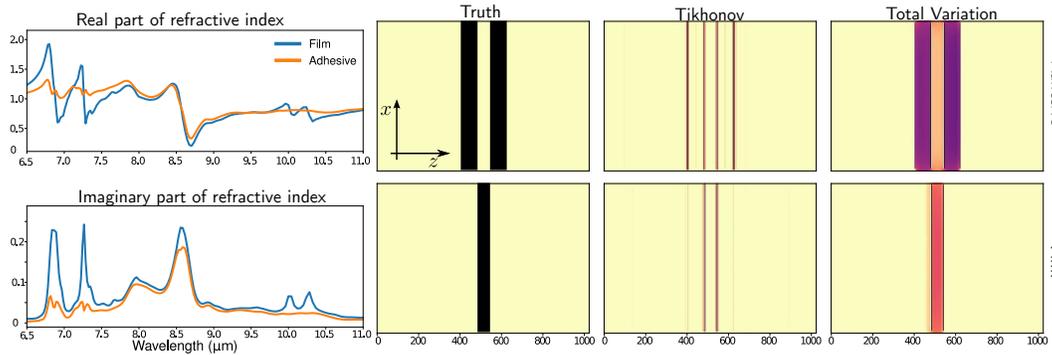

Fig. 2. Reconstruction of tape.

reconstructed locations are the superposition of all three species and artifacts due to multiple scattering remain. By contrast, ℓ_1 regularization eliminates these artifacts and correctly identifies the species at each point.

The second sample models a piece of double sided tape. A 40µm layer of film is located between two 50µm layers of adhesive. Data was generated from the same parameters as in the previous simulation. The spectra for the film and adhesive were found via FTIR. Figure 2 illustrates the spectra, sample, and reconstructions using Tikhonov and total-variation regularization. Tikhonov regularization results in a band-pass filtered version of the original sample, where the maximum and minimum spatial frequencies are dictated by the minimum and maximum illuminating wavelengths, respectively. TV regularization promotes piecewise constant images and is a natural regularizer for this sample. As expected, TV is able to reconstruct the lower spatial frequencies of the sample.

Acknowledgment

References

1. N. Bosschaart, T. G. van Leeuwen, M. C. G. Aalders, and D. J. Faber, "Quantitative comparison of analysis methods for spectroscopic optical coherence tomography," Biomedical Optics Express **4**, 2570 (2013).
2. B. Deutsch, R. Reddy, D. Mayerich, R. Bhargava, and P. S. Carney, "Compositional prior information in computational infrared spectroscopic imaging," J. Opt. Soc. Am. A **32**, 1126–1131 (2015).
3. S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," Foundations and Trends in Machine Learning **3**, 1–122 (2010).