Diffuse Optical Tomography

Ethan Morton and Matt Kubas, *Member, IEEE*

***Abstract*—**

I. Background

Brain imaging is an essential part of healthcare around the world. The currently available options for brain scans are limited and extremely bulky, and some even require cryogenic cooling. Researchers are searching to find new and innovative ways to characterize how the brain works. For this project, we will look into brain activity based on how oxygenated certain regions of the brain are. There already exist systems that can measure this, but the systems are not movable and can take up entire rooms. We will be analyzing a dataset from a brain imager developed by Joseph Culver, Ed Richter, Jason Trobaugh, and Adam Eggebrecht at Washington University in St Louis.

The goal of this equipment is to non-invasively image brain activity and accurately relate specific portions of the brain to specific external stimuli. This system uses Diffuse Optical Tomography (DOT) to image brain activity. This means that the region of interest is imaged by diffusing light through a medium. This is typically done by shining a light through the medium and detecting it on another side in some distinct geometrical array of sources and detectors. The precise structure of the 3D region of interest can be reconstructed through knowledge of the medium and how it reacts to the light chosen. We arrange the sources and detectors into "pairs," where each source will have a distinct and separable effect on the reading of each detector. For the case of our dataset, we have 24 sources and 28 detectors, each supporting two wavelengths, which gives us 1344 pairs. We then numerically break the 3D space of the brain into "voxels" (3D pixels). Each source-detector pair will have a different "sensitivity" to each voxel, which indicates how strongly a reading from that source-detector pair is related to the concentration of oxygenated hemoglobin at that physical location inside the brain. Mathematically, we write this as follows:

$$\mathbf{y} = \mathbf{A} \cdot \mathbf{x} \tag{1}$$

Where $\mathbf{y}$ is a 1344 tall column vector, $\mathbf{x}$ has a height dependent on the number of voxels, and $\mathbf{A}$ is the "sensitivity matrix" which is 1344 tall and has a width equal to the number of voxels. Each physical reading in y represents the superposition of the effects of multiple sources, and each element in $\mathbf{x}$ represents the change in concentration at a given point in the region of interest.

To go "forward" with this problem, e.g., given a simulated $\mathbf{x}$, we need to find equations for the elements of A to find $\mathbf{y}$. This is done by solving the diffusion equation for light diffusing through the brain. Because this project is *not* on the derivation of these relations but rather on the use of these relations, we will paste them in with the reference to the slides Dr. Trobaugh and Dr. Lew provided for this module. The resulting equations are seen below:

$$y = -\frac{v}{D_o}\int \frac{G(\mathbf{r}_s, \mathbf{r'})G(\mathbf{r'}, \mathbf{r}_d)}{G(\mathbf{r}_s, \mathbf{r}_d)}\delta\mu_a(\mathbf{r'})d\mathbf{r'} \tag{2}$$

$$A_{i,j} = -\frac{vh^3}{D_o}\frac{G(\mathbf{r}_s, \mathbf{r'})G(\mathbf{r'}, \mathbf{r}_d)}{G(\mathbf{r}_s, \mathbf{r}_d)} \tag{3}$$

$$x_j = \Delta\mu_a(\mathbf{r}_j) \tag{4}$$

$$G_{inf}(r_s, r) = G(r_s - r) = \frac{1}{4\pi D|r - r_s|}e^{-\mu_{eff}|r-r_s|} \tag{5}$$

We approximate the brain to be a rectangular slab (for simplicity of discovery) which is broken up into voxels (3D pixels), it takes us from an integral, as shown in the equation for y, to a sum, which is why A has discrete elements for each voxel. The summing of these elements represents the integral over "all space" that gives us y. Because of the way the math works out, this is calculated with a set of Green's functions. The equation for Green's function depends exclusively on the position of the light source and the target point in the medium for which the light strength is being calculated.

To detect the brain activity itself, we specifically look for the *change* in light absorption throughout the region, as seen in the equation for x_j. This is because as a region of the brain becomes more active, there will be an increased concentration of oxygenated blood in that region, causing a difference in absorption. Specifically, the absorption curves for oxygenated and non-oxygenated hemoglobin are shown in fig. 1.

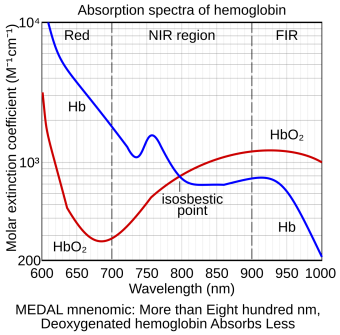

Fig. 1: Absorption spectra of oxygenated and non-oxygenated hemoglobin.

The absorption curves show a few distinct regions of behavior. At wavelengths between 600 and 750, the curves are at very different levels. Between 800 and 900, the levels are relatively similar. These differences are leveraged by including two available wavelengths of light in the sources and detectors: 750nm and 850nm. By using more than one wavelength of

light, we add rows to the A matrix, providing more information and constraints on the matrix.

Methods for simulation and reconstruction are discussed in sections II-A and II-B, but we will discuss the related mathematics here. In general, we must be able to calculate an x from any given y, assuming we are given A. However, A is generally underconstrained and noninvertible, so this becomes an optimization problem rather than a simple inverse problem. To allow for finer control over the effect of inverting A, we add a regularization term that will help prevent solutions with huge spikes. Mathematically, we write this as follows:

$$\text{Inverse Problem} \implies min\{||\mathbf{y} - \mathbf{Ax}||_2^2 + \lambda||\mathbf{x}||_2^2\} \tag{6}$$

$$\mathbf{A}_{reg}^{\#} = (\mathbf{A}^T\mathbf{A} + \lambda\mathbf{I})^{-1}\mathbf{A}^T = \mathbf{A}^T(\mathbf{A}\mathbf{A}^T + \lambda\mathbf{I})^{-1} \tag{7}$$

$$\mathbf{A}_{reg}^{\#} = \mathbf{VFS}^{-1}\mathbf{U}^T \tag{8}$$

$$F_{ii} = \frac{s_{ii}^2}{s_{ii}^2 + \lambda} \tag{9}$$

This gives us a regularized inverse of A, which we can use in place of a true inverse, where λ is a smoothing parameter that is set relative to the most significant singular value, as seen by the final two equations relating the SVD of the regularized inverse to lambda. A higher lambda results in a higher smoothing effect across the volume.

II. Methods

The general mathematics behind the DOT system this report focuses on are described above in section I. For clarity's sake, we will review the system's basic specifications again. The system uses an array of light sources and detectors tuned to two specific wavelengths of light, 750nm and 850nm. The properties of the human brain for these waves of light are listed in fig. 2.

	750 nm		850 nm		Index of Refraction
	μ_a [mm^{-1}]	$\mu_{s'}$ [mm^{-1}]	μ_a [mm^{-1}]	$\mu_{s'}$ [mm^{-1}]	n
Scalp *	0.0170	0.74	0.0190	0.64	1.4
Skull *	0.0116	0.94	0.0139	0.84	1.4
CSF **	0.004	0.3	0.004	0.3	1.4
Gray matter ***	0.0180	0.8359	0.0192	0.6726	1.4
White matter ***	0.0167	1.1908	0.0208	1.0107	1.4

Fig. 2: Table of properties which govern the diffusion of light in brain tissue

These properties μ_a and μ_s' are used to calculate the constants listed throughout section I, using the following formulas:

$$D = \frac{1}{3\mu_a + \mu_s'} \tag{10}$$

$$\mu_{eff} = \sqrt{\frac{\mu_a}{D}} \tag{11}$$

The specific positions of the sources and detectors in the array, which allow us to calculate the Green's function for each pair, are given for a specific arrangement in the dataset. The dataset also includes time-series data of the measurements of each source-detector pair for roughly 450 seconds. This dataset was taken of a patient observing a wedge stimulus moving counter-clockwise in a circle, ticking once sector per second. The circle was divided into 36 sectors. An example experimental setup is shown below in fig. 3.

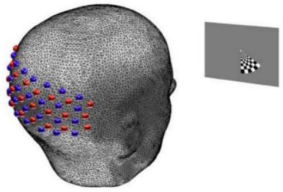

Fig. 3: An example experimental setup of a patient viewing a wedge rotating counter-clockwise in sectors of a circle.

For all simulations, the experimental setup is modeled as a "slab" consisting of voxels. The voxel size is an adjustable parameter but generally kept at 2mm per voxel, with a slab of 60mmx90mmx30mm. The sources and detectors are positioned 1mm above the slab in a grid formation.

A. Visualization

For lower-dimensional data, visualization is generally fairly straightforward. You need only define the axes; plotting software such as Matlab, Gnuplot, etc., will handle everything else. However, for the 4D data in this system, "visualizing" becomes less easily interpretable. For this project, we use both the Matlab VolumeViewer and the Sliceviewer; however, for the animations and figures of this report, we use a simple custom process for looking at a single slice over time. A few examples of the volumetric and sliced techniques are shown below:

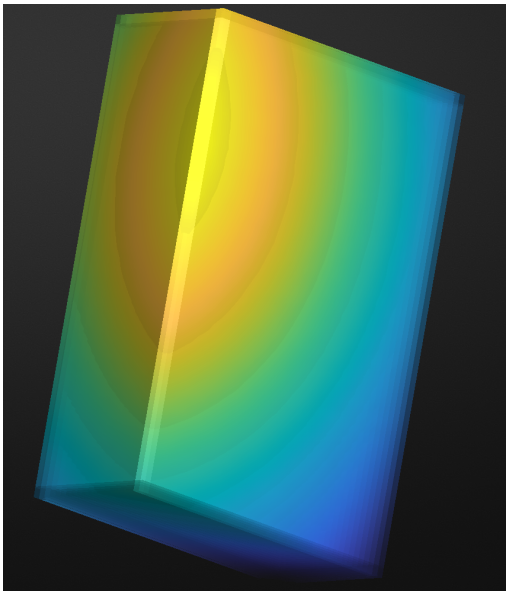

Fig. 4: Volumetric view of a source-detector "banana"

While the volumetric image looks interesting, it doesn't provide as much information about the shape and intensity

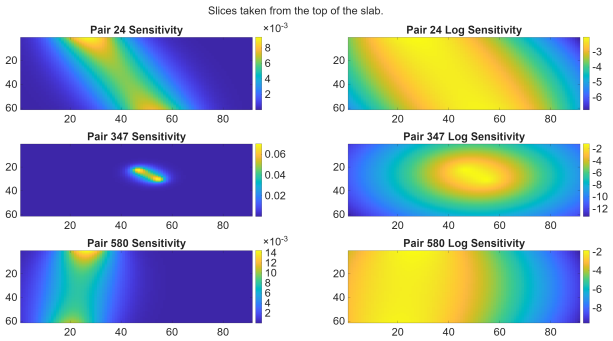

Fig. 5: Slice view of source-detector "bananas"

inside the volume itself. The critical data is often not visually interpretable, such as the variation of a weak signal over time, which would quickly become lost in a sea of pretty colors. For this reason, we use these visualizations mainly as a sanity check for the simulation of our system.

B. Simulation

To test the system's mathematics and verify that the theory works (in theory), we must run some simulations. To do this, we first designate a specific stimulus type and determine its physical response in the brain. That way, we have an accurate representation of the "true" system at work. Thanks to prior knowledge provided by professors Trobaugh and Lew, we know that a good simulation stimulus candidate is a point source moving in a circle inside the volume. We are also given a Hemodynamic Response Function (HRF), which represents the blood oxygenation response to neural activity. This is pictured below in fig. 6.

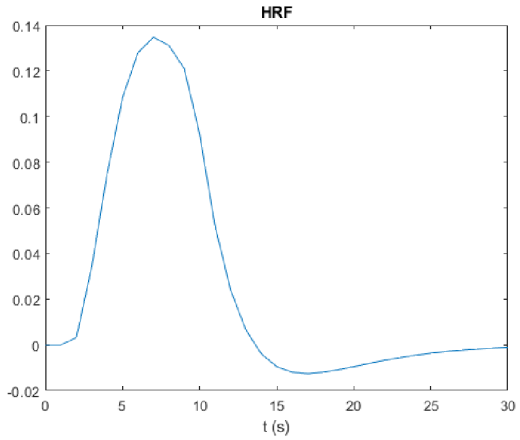

Fig. 6: Graph of the HRF over the course of 30 seconds

We see that it peaks somewhere around 6 seconds and falls off slightly negative, then goes to zero again. To properly simulate the response of the system to the stimulus, we must convolve the HRF with the stimulus to obtain the "true" oxygenation concentration simulation. To set up the stimulus, we define a circle with a radius of half the smallest dimension of the slab, centered in the middle of the slab, and give the stimulus an angular velocity the same as the experimental stimulus in fig. 3. Because the system relies on the *difference* in concentration, as written in eq. (4), we also take the first-order difference of each voxel through time. After convolution and the difference, we have a representative time-series simulation of the change in concentration over time at each voxel. This gives us a matrix with as many rows as voxels and columns as time steps. To obtain the simulated measurements $\mathbf{y}$, we multiply the time-series matrix $\mathbf{X}$ with the matrix A, which will result in another time-series matrix $\mathbf{Y}$, with as many rows as source-detector pairs and as many columns as time steps. To go in the reverse direction and attempt to obtain our original $\mathbf{X}$ matrix, we multiply by the regularized inverse of A, setting λ appropriately in the process. This leaves us with a solution for the original $\mathbf{X}$, which we call $\hat{\mathbf{X}}$.

C. Point Spread

Image resolution can be classified in many ways. In this case, we are measuring the point spread function (PSF) of the reconstructed data using the full-width half maximum (FWHM) classification. For DOT, this describes how differentiable different positions in the brain are and how accurately we can recreate hemoglobin response in the brain from the particular array of sensors. We could achieve higher resolution and vice versa if we used a more densely packed array. In this case, we are not in control of the sensor position, density, or any other physical parameters, we can alter the regularization constant λ. A visualization of how the algorithm for finding the FWHM is shown in fig. 10. The relationship between lambda and the point spread is shown in fig. 7. The graph shows that as the regularization constant gets larger, the FWHM value appears to be an exponential decay in the relationship, but this is just a visual analysis. This means that the more we punish the regularized matrix *A*, the smaller the PSF.

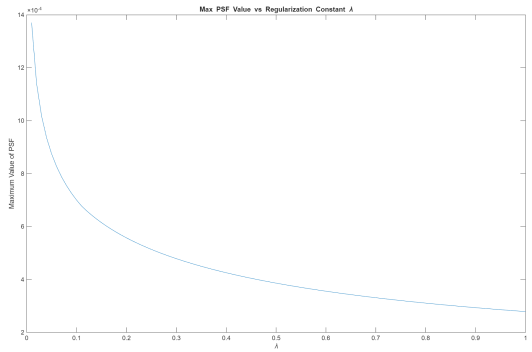

Fig. 7: The relationship between the regularization constant λ and the point spread measured using largest value FWHM (measured in millimeters)

III. Data Processing

The data processing started identically to the visualization method, we calculate the *A* matrix, regularize the *A* matrix to create A_{reg}, and now multiply the regularized *A* matrix by the raw data to get the reconstructed data, this takes the form $\hat{x} = A_{reg}y$ where y is the raw data and $\hat{x}$ is the reconstructed data. This data is inherently noisy, or in other words, a small

signal-to-noise ratio. In an attempt to try something cool, we attempted to de-convolve the HRF from the decoded signal. This was done similarly to the method used in the simulation example by taking the Fourier transform of the decoded signal and the HRF and then dividing the two matrices element-wise. A key thing to note is that both `HRF./x_hat` and `x_hat./HRF` were tried, but only the first one resulted in anything meaningful. We then applied a tight bandpass filter to select the frequency we were explicitly looking for, excluding things like heart rate and high-frequency noise present in the system.

A. Filtering & Noise Reduction

The filtering done in this case was to pinpoint the behavior we expect to see, a counterclockwise rotating behavior mimicking the stimulus the patient is looking at. To remove everything unwanted, we started by trying to de-convolve with the HRF by element dividing or dot dividing in MATLAB (`HRF./x_hat`). This process seemed promising, as the magnitude of the Fourier transform in each case resembled more of what we were looking at. Figure 8 represents the magnitude of the Fourier transform while fig. 9 represents the de-convolved data. Seeing that once we de-convolved with the HRF, that graph looked promising, and thus, we thought it would give great results. Our results showed some behavior, but not exactly what we expected.

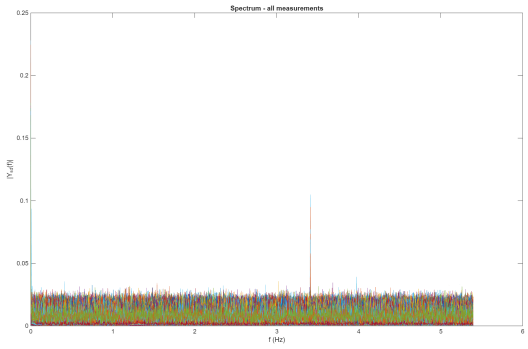

Fig. 8: Magnitude of the discrete Fourier transform of the reconstructed data.

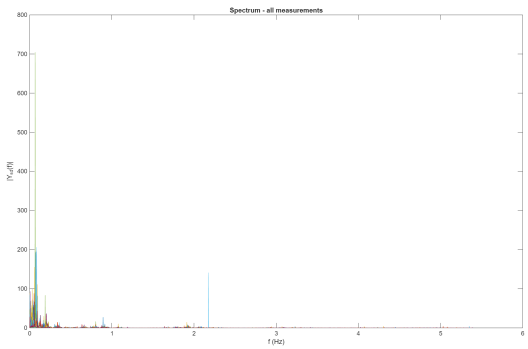

Fig. 9: Magnitude of the discrete Fourier transform of the x_hat data de-convolved with the HRF

We then applied a bandpass filter to the system to create a behavior where the cutoff frequency of both the high-pass and the low-pass were at the same point, $0.02\overline{778}Hz$. This would essentially pick and choose only the behavior we were looking for and eliminate the noise. This was done for both the data that was de-convolved with the HRF and the raw data itself.

Another method we used to attempt to remove noise was to average samples. Since we knew that the data repeated itself roughly 11 times during the trial, we were able to average each repetition together to reduce the amount of noise and get a clearer animation as a result.

IV. Results

A. Simulation

The simulations showed some expected behavior but did not closely resemble the original simulated response. An example of this can be seen in the simulated versus recovered animation for the circular stimulus discussed in section II-B. On the left, we see the original stimulus at the 8th slice in the medium (16mm), and on the right, we see the recovered top slice directly underneath the sources and detectors. For a clearer view of the stimulus, see this animation.

Possible improvements to the visualization process could include setting a threshold to improve the visibility of specific high-interest regions or to more carefully remove the HRF from the resulting $\hat{\mathbf{X}}$, rather than brute-force division.

B. Dataset Analysis

The resulting animations we had achieved were not the exact behavior we were expecting, but there was some resemblance to what we were looking for. The animation for the de-convolved dataset that was averaged can be found at: De-Convolved HRF with BPF at 0.027Hz. The animation has some counterclockwise tendencies but not explicitly. Sometimes, opposite corners would light up, which did not make any sense. This was disappointing, but we did not know what we had done incorrectly.

It became known later that de-convolving with the HRF had not been done and was making our data behave in an uninterpretable manner. Since the Fourier transform of the de-convolved data looked promising, it made it seem like it would work. Regardless, we then removed the de-convolution and were able to obtain the following animation: Reconstructed Data with BPF at 0.027Hz. This animation showed behavior that matched the counterclockwise activation we expected to see. The starting location of the simulation is not exact since we were unsure of when the data sampling began relative to the position of the stimulus.

V. Conclusion

Overall, our results represent a reasonable introduction to DOT and the processing of actual and simulated data. The results show some expected characteristics, such as observed activity shifts in the reconstructed actual data and clear behavior in the recovered simulation data. These behaviors would be good targets for further study of the signal processing and simulation testing involved. Further extensions could include simulation of more complicated stimuli, which necessitate

better signal processing, such as finding ways to account for the unequal sensitivity of each voxel to each pair, or more advanced signal processing of the real data, such as producing cleaner and more concise metrics to analyze the quality of the data and detect signals through the noise.

APPENDIX

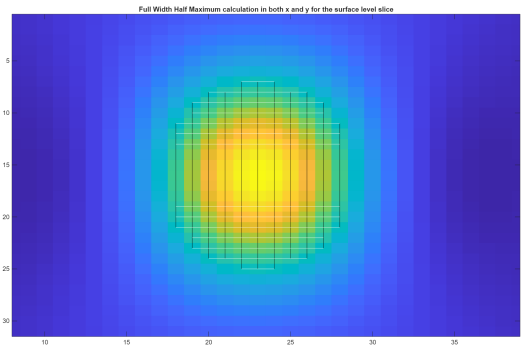

Fig. 10: FWHM calculation visualization for a point source in the center buried halfway in the slab with $\lambda = 0.05$