Optimal Filters for High-Speed Compressive Detection in Spectroscopy

Greg Buzzard

Department of Mathematics
Purdue University

July 31, 2015

Joint work with Brad Lucier, Dor Ben-Amotz (chemistry), David Wilcox (chemistry) and Owen Rehrauer (chemistry) at Purdue.
Partially supported by ONR N00014-13-1-0394.

Motivating question

Question: How can you do spectral analysis using low-power lasers and/or high-throughput sampling?

Applications: Pharmaceutical quality control, chemical threat-detection, forensic analysis, etc.

Raman spectroscopy

- Illuminate a chemical sample with a laser.
- Scattered light produces a frequency spectrum (chemical fingerprint).
- Measure the fingerprint.
- Determine the chemical composition of the sample.

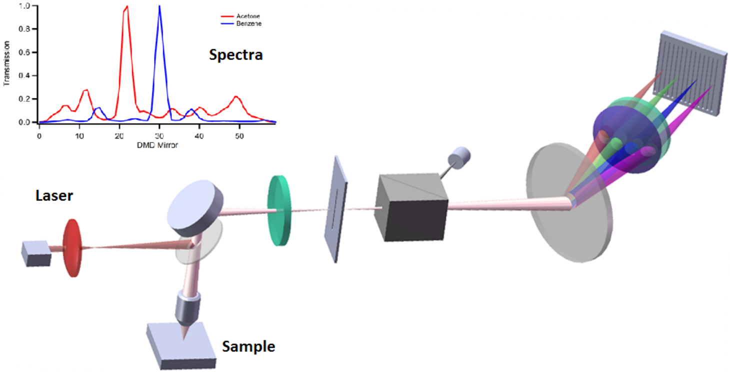

Modeling

Roughly, a photon hits one of the materials in the sample, changes energy in response, and then is diffracted into the appropriate bin and detected. We model this as a discrete Poisson process with some arrival rate p (photons/sec).

Let X be the number of photons counted in a time interval of length τ.
Then

$$X \sim Poisson(p\tau)$$

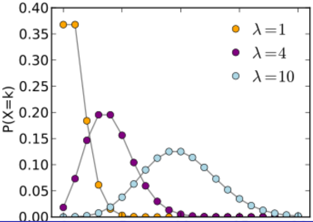

Modeling

For this Poisson X, $\mathrm{E}(X) = \mathrm{Var}(X) = p\tau$. For large τ,

$$X \text{ is roughly } N(p\tau, p\tau), \text{ so}$$
$$X/\tau \text{ is roughly } N(p, p/\tau).$$

I.e., longer exposure yields a more reliable estimate of p.

We get a different count for each energy bin (assumed to be independent). The arrival rate depends on the laser power, the optics and detector, and the chemical composition of the sample.

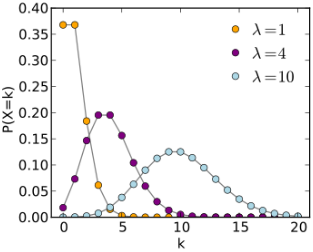

Modeling

Main assumption: Additive spectra

- A sample is composed of some mixture of known chemicals, S_j (with known spectra).
- Each S_j has an associated rate of photon emission, Λ_j (essentially proportional to concentration).
- A photon from S_j lands in energy bin i with probability P_{ij}.
- The photon count in bin i is a Poisson process with rates given by summing over samples.

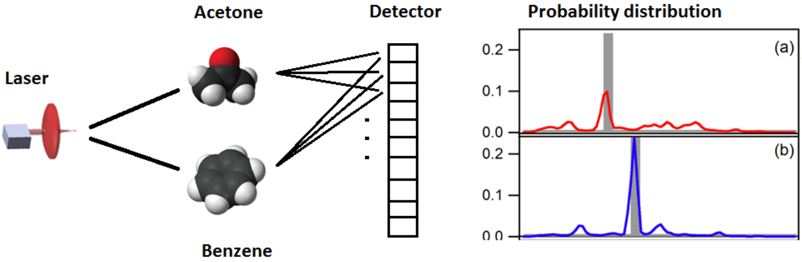

Modeling

Measuring the ith bin for time τ gives a Poisson random variable X_i with mean and variance $\tau(P\Lambda)_i$.

Goal: Take measurements to determine Λ.

This is called spectral unmixing - determine the composition of a sample from spectral measurements.

Noise

Sources of noise:

- The variance associated with X_i is called "shot noise."
- Each detector has noise that scales linearly with time called "dark noise."
- Each detector has noise from each reading called "read noise."

Shot noise is unavoidable - this comes from sampling a Poisson random variable. Dark noise and read noise can be minimized but still must be considered.

Questions: How should we measure the spectrum to determine Λ? What is the resulting uncertainty in our estimate?

Measurements

Standard measurement approach uses charge-coupled devices (CCD, as in cameras): one measurement per frequency bin, all measured simultaneously (assume no dark or read noise).

Get a vector sample X that is Poisson with mean $\tau P \Lambda$, or

$$X = \tau P \Lambda + \epsilon,$$

with $\mathrm{E}(\epsilon) = 0$, $\mathrm{Var}(\epsilon) = \mathrm{diag}(\tau P \Lambda) = \Omega$, and P, τ known.

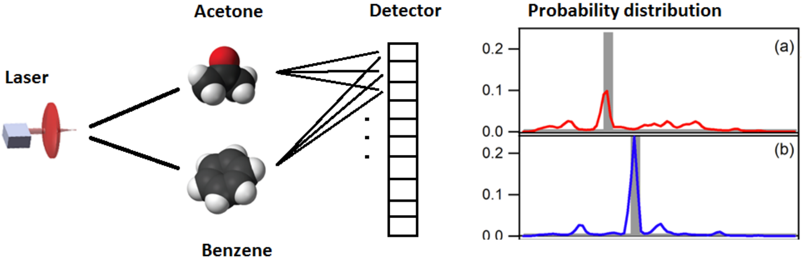

Measurements

Now we have

$$X = \tau P \Lambda + \epsilon,$$

with $\mathrm{E}(\epsilon) = 0$, $\mathrm{Var}(\epsilon) = \mathrm{diag}(\tau P \Lambda) = \Omega$, and P, τ known.

Assuming we know Ω, generalized least squares give the BLUE

$$\hat{\Lambda} = \frac{1}{\tau}(P^T \Omega^{-1} P)^{-1} P^T \Omega^{-1} X.$$

We can use X to estimate Ω or use more sophisticated techniques.

Problem: For small τ, the mean of read noise can be larger than the mean of $X \Rightarrow$ large uncertainty in $\hat{\Lambda}$.

Question: Is there a better way with fewer measurements? I.e., how can we reduce irreducible uncertainty?

Measurements

Now we have

$$X = \tau P \Lambda + \epsilon,$$

with $\mathrm{E}(\epsilon) = 0$, $\mathrm{Var}(\epsilon) = \mathrm{diag}(\tau P \Lambda) = \Omega$, and P, τ known.

Assuming we know Ω, generalized least squares give the BLUE

$$\hat{\Lambda} = \frac{1}{\tau}(P^T \Omega^{-1} P)^{-1} P^T \Omega^{-1} X.$$

We can use X to estimate Ω or use more sophisticated techniques.

Problem: For small τ, the mean of read noise can be larger than the mean of $X \Rightarrow$ large uncertainty in $\hat{\Lambda}$.

Question: Is there a better way with fewer measurements? I.e., how can we reduce irreducible uncertainty?

Filters

Idea: Look for ways to combine photons before counting.

An **optical filter** is a device to collect photons from multiple bins before measuring, possibly also modulating each bin separately. We model this as a column vector with entries between 0 and 1.

One measurement is the dot product of a filter with the vector X of photon counts. We'll need to measure with several filters, but we can't (generally) measure simultaneously, so we can also choose how long to measure with each filter (can assume total time 1).

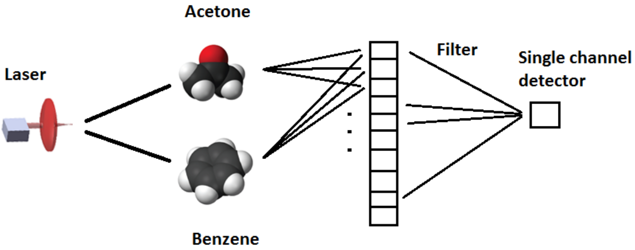

Problem Formulation

Let F be a matrix with each column a filter, T diagonal with the time for filter j in the jjth position, P the matrix with P_{ij} equal to the probability that a photon from sample j lands in bin i, and Λ the vector with Λ_j the rate of photon emission from sample j.

Then X is Poisson with mean $\mu = TF^T P\Lambda$ and covariance matrix $\text{diag}(\mu)$.

Goal: Find F, T and a matrix B to give an "optimal" estimate of Λ using $\hat{\Lambda} = BT^{-1}X$.

Criterion: Minimize $\text{E}(\|\hat{\Lambda} - \Lambda\|^2)$. For $\hat{\Lambda}$ to be unbiased, we need also $BF^T P = I$.

I.e., choose the best filters, times, and recovery method to minimize the 2-norm error in $\hat{\Lambda}$.

Problem Formulation

Let F be a matrix with each column a filter, T diagonal with the time for filter j in the jjth position, P the matrix with P_{ij} equal to the probability that a photon from sample j lands in bin i, and Λ the vector with Λ_j the rate of photon emission from sample j.

Then X is Poisson with mean $\mu = TF^T P\Lambda$ and covariance matrix $\mathrm{diag}(\mu)$.

Goal: Find F, T and a matrix B to give an "optimal" estimate of Λ using $\hat{\Lambda} = BT^{-1}X$.

Criterion: Minimize $\mathrm{E}(\|\hat{\Lambda} - \Lambda\|^2)$. For $\hat{\Lambda}$ to be unbiased, we need also $BF^T P = I$.

I.e., choose the best filters, times, and recovery method to minimize the 2-norm error in $\hat{\Lambda}$.

Problem Formulation

Goal: Find F, T and B to minimize $\mathrm{E}(\|BT^{-1}X - \Lambda\|^2)$ subject to $BF^TP = I$, $0 \leq F_{ij} \leq 1$, T nonnegative diagonal with $\mathrm{tr}(T) = 1$.

Notes:

- Λ affects the solution, but we don't know Λ. Taking Λ all 1's generally works well and gives provable bounds on error.
- If the number of filters equals the number of species, then B is $(F^TP)^{-1}$. If there are more filters than species and T is fixed, then B is determined by GLS.
- This formulation is our reconstruction of the idea of local A-optimality from the field of optimal design of experiments - minimize some measure of size of the ellipsoid of uncertainty associated with $\hat{\Lambda}$.

Problem Formulation

Goal: Find F, T and B to minimize $\mathrm{E}(\|BT^{-1}X - \Lambda\|^2)$ subject to $BF^TP = I$, $0 \leq F_{ij} \leq 1$, T nonnegative diagonal with $\mathrm{tr}(T) = 1$.

Notes:

- Λ affects the solution, but we don't know Λ. Taking Λ all 1's generally works well and gives provable bounds on error.
- If the number of filters equals the number of species, then B is $(F^TP)^{-1}$. If there are more filters than species and T is fixed, then B is determined by GLS.
- This formulation is our reconstruction of the idea of local A-optimality from the field of optimal design of experiments - minimize some measure of size of the ellipsoid of uncertainty associated with $\hat{\Lambda}$.

Problem Formulation

Goal: Find F, T and B to minimize $\mathrm{E}(\|BT^{-1}X - \Lambda\|^2)$ subject to $BF^TP = I$, $0 \leq F_{ij} \leq 1$, T nonnegative diagonal with $\mathrm{tr}(T) = 1$.

Notes:

- Λ affects the solution, but we don't know Λ. Taking Λ all 1's generally works well and gives provable bounds on error.
- If the number of filters equals the number of species, then B is $(F^TP)^{-1}$. If there are more filters than species and T is fixed, then B is determined by GLS.
- This formulation is our reconstruction of the idea of local A-optimality from the field of optimal design of experiments - minimize some measure of size of the ellipsoid of uncertainty associated with $\hat{\Lambda}$.

Problem Formulation

Goal: Find F, T and B to minimize $\mathrm{E}(\|BT^{-1}X - \Lambda\|^2)$ subject to $BF^TP = I$, $0 \le F_{ij} \le 1$, T nonnegative diagonal with $\mathrm{tr}(T) = 1$.

Notes:

- Λ affects the solution, but we don't know Λ. Taking Λ all 1's generally works well and gives provable bounds on error.
- If the number of filters equals the number of species, then B is $(F^TP)^{-1}$. If there are more filters than species and T is fixed, then B is determined by GLS.
- This formulation is our reconstruction of the idea of local A-optimality from the field of optimal design of experiments - minimize some measure of size of the ellipsoid of uncertainty associated with $\hat{\Lambda}$.

Reformulation

Let $A = F^T P$. Then $A\Lambda$ gives the vector of mean counts for measurements of time 1 for each filter.

Using the Poisson assumption (and assuming no noise), then a standard variance calculation gives

$$E(\|BT^{-1}X - \Lambda\|^2) = \sum_{i=1}^{m} \frac{1}{T_{ii}} \|B\mathbf{e}_i\|^2 (A\Lambda)_i.$$

Question: Given A and B, what is the optimal T?

Reformulation

Let $A = F^T P$. Then $A\Lambda$ gives the vector of mean counts for measurements of time 1 for each filter.

Using the Poisson assumption (and assuming no noise), then a standard variance calculation gives

$$E(\|BT^{-1}X - \Lambda\|^2) = \sum_{i=1}^{m} \frac{1}{T_{ii}} \|B\mathbf{e}_i\|^2 (A\Lambda)_i.$$

Question: Given A and B, what is the optimal T?

Optimal T

Problem: Given $a_j \geq 0$, minimize $\displaystyle\sum_j \frac{a_j}{T_j}$

subject to $\sum_j T_j = 1, \; T_j \geq 0$.

Optimal T

Problem: Given $a_j \geq 0$, minimize $\displaystyle\sum_j \frac{a_j}{T_j}$
subject to $\sum_j T_j = 1, \ T_j \geq 0$.

Can use Karush-Kuhn-Tucker (KKT) conditions or Lagrange multipliers, but Cauchy- Schwarz is slicker. Get $T_j = \gamma\sqrt{a_j}$ with γ chosen to get $\sum_j T_j = 1$.

Reformulation

Choosing these T_{jj} for the problem above, we now need to choose F and B to minimize

$$\sum_{j=1}^{m} \|B\mathbf{e}_j\| \sqrt{(A\Lambda)_j},$$

where $A = F^T P$.

- Note that given A and T we can find B, and given A and B we can find T. For fixed A, this gives an iterative method that converges to the optimal B and T for that A.
- All of the uncertainty was integrated away to yield a matrix optimization problem.

Choosing filters

Questions:

- How should we choose filters (or $A = F^T P$) to minimize the sum?
- How many filters should we use?
- Is it better to use arbitrary analog filters (using a spatial light modulator) or can we use binary (0-1) filters (using a digital micromirror device (DMD)?

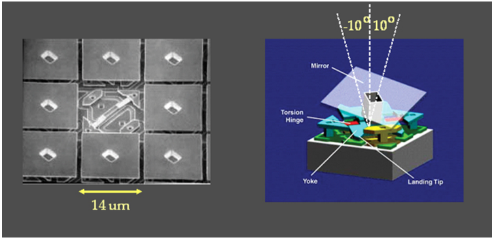

DMD

Choosing filters

Questions:

- How should we choose filters (or $A = F^T P$) to minimize the sum?
- How many filters should we use?
- Is it better to use arbitrary analog filters (using a spatial light modulator) or can we use binary (0-1) filters (using a digital micromirror device (DMD)?

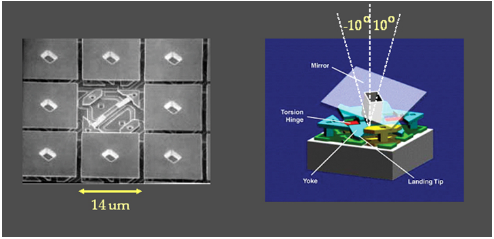

DMD

Choosing filters

Questions:

- How should we choose filters (or $A = F^T P$) to minimize the sum?
- How many filters should we use?
- Is it better to use arbitrary analog filters (using a spatial light modulator) or can we use binary (0-1) filters (using a digital micromirror device (DMD)?

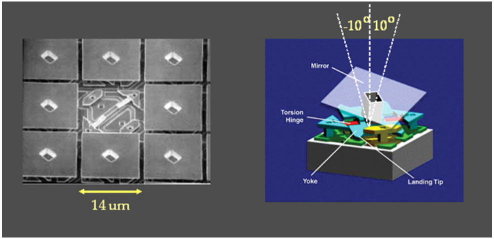

DMD

Nearly binary filters

Choose A and B to minimize

$$\sum_{j=1}^{m} \|B\mathbf{e}_j\| \sqrt{(A\Lambda)_j},$$

subject to $BA = I$, $A = F^T P$ and $0 \leq F_{ij} \leq 1$.

Note that F is (# of bins) $\times$ (# of filters) and that there may be hundreds of bins, so this is a high-dimensional optimization problem.

Observation 1: A is only (# of filters) $\times$ (# of chemicals), which for our application is much lower dimensional than F. So optimize over A that can be written as $F^T P$.

Nearly binary filters

Choose A and B to minimize

$$\sum_{j=1}^{m} \|B\mathbf{e}_j\| \sqrt{(A\Lambda)_j},$$

subject to $BA = I$, $A = F^T P$ and $0 \leq F_{ij} \leq 1$.

Note that F is (# of bins) $\times$ (# of filters) and that there may be hundreds of bins, so this is a high-dimensional optimization problem.

Observation 1: A is only (# of filters) $\times$ (# of chemicals), which for our application is much lower dimensional than F. So optimize over A that can be written as $F^T P$.

Nearly binary filters

Choose A and B to minimize

$$\sum_{j=1}^{m} \|B\mathbf{e}_j\| \sqrt{(A\Lambda)_j},$$

subject to $BA = I$, $A = F^T P$ and $0 \leq F_{ij} \leq 1$.

Observation 2: If we scale row j of A by $r > 0$ and scale column j of B by $1/r$, then still $BA = I$, but

$$\|B\mathbf{e}_j / r\| \sqrt{(rA\Lambda)_j} = \|B\mathbf{e}_j\| \sqrt{(A\Lambda)_j} / \sqrt{r}.$$

Hence for a fixed row of A, there is a maximal scaling that minimizes this term subject to being realizable with a filter.

Nearly binary filters

Choose A and B to minimize

$$\sum_{j=1}^{m} \|B\mathbf{e}_j\| \sqrt{(A\Lambda)_j},$$

subject to $BA = I$, $A = F^T P$ and $0 \leq F_{ij} \leq 1$.

Now instead of A, we can search over rows of A that are unit vectors, $\mathbf{u}$, with all entries nonnegative and find the maximal r so that $r\mathbf{u}^T = F_j^T P$ with $0 \leq F_{ij} \leq 1$. Finding r and F_j given $\mathbf{u}$ is a standard linear programming problem.

This reduces the problem by one degree of freedom for each filter. With two chemicals and 200 bins, this reduces the problem from 400 dimensions to 2.

Nearly binary filters

Choose A and B to minimize

$$\sum_{j=1}^{m} \|B\mathbf{e}_j\| \sqrt{(A\Lambda)_j},$$

subject to $BA = I$, $A = F^T P$ and $0 \leq F_{ij} \leq 1$.

Now instead of A, we can search over rows of A that are unit vectors, $\mathbf{u}$, with all entries nonnegative and find the maximal r so that $r\mathbf{u}^T = F_j^T P$ with $0 \leq F_{ij} \leq 1$. Finding r and F_j given $\mathbf{u}$ is a standard linear programming problem.

This reduces the problem by one degree of freedom for each filter. With two chemicals and 200 bins, this reduces the problem from 400 dimensions to 2.

Nearly binary filters

Using the KKT conditions on this two stage minimization problem, we obtain

Theorem (Nearly Binary Filters)

Let n be the number of chemicals. Then the optimal filters can be chosen to consist of only 0*s and* 1*s, except for possibly $n-1$ entries strictly between* 0 *and* 1 *in each filter.*

In fact, this theorem is obtained by working on only one row of A at a time. In general, the filters that are optimal simultaneously consist of only 0s and 1s (or at least that this can be chosen to be so by adding more filters). In practice, we round to 0 or 1 in order to use DMD.

Nearly binary filters

Using the KKT conditions on this two stage minimization problem, we obtain

Theorem (Nearly Binary Filters)

Let n be the number of chemicals. Then the optimal filters can be chosen to consist of only 0*s and* 1*s, except for possibly $n - 1$ entries strictly between* 0 *and* 1 *in each filter.*

In fact, this theorem is obtained by working on only one row of A at a time. In general, the filters that are optimal simultaneously consist of only 0s and 1s (or at least that this can be chosen to be so by adding more filters). In practice, we round to 0 or 1 in order to use DMD.

Simulation results

Sample spectra and optimal binary (OB) filters obtained by minimizing the objective function.

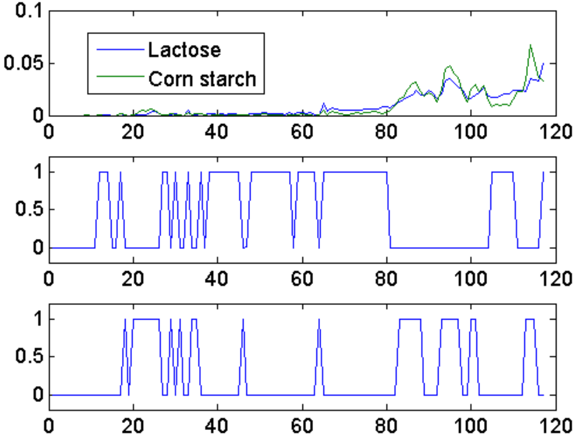

Simulation results

Recovery with OB filters versus CCD, no read noise. Here CCD is uniformly better than OB.

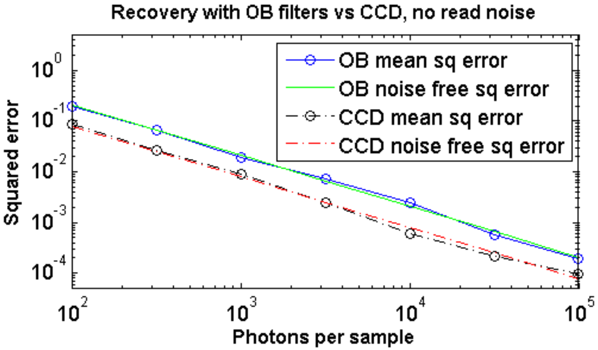

Simulation results

Recovery with OB filters versus CCD, with read noise (modeled as Gaussian with std dev about 7). Now the read noise introduces significant errors for small counts.

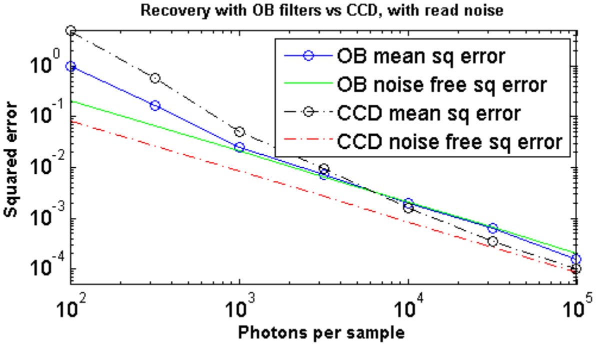

Experimental results

Experimental setup includes a background spectrum (due to optics) in addition to chemical spectra. We model this as another chemical. First we focus on classification: is the sample chemical A or B? With only about 100 photons total, we can reliably distinguish acetone and benzene.

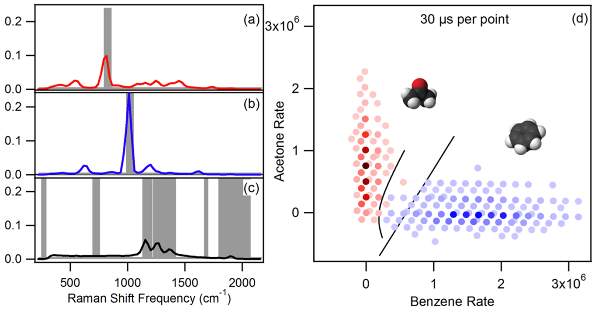

Experimental results

For closely related chemicals, we need longer exposure times.

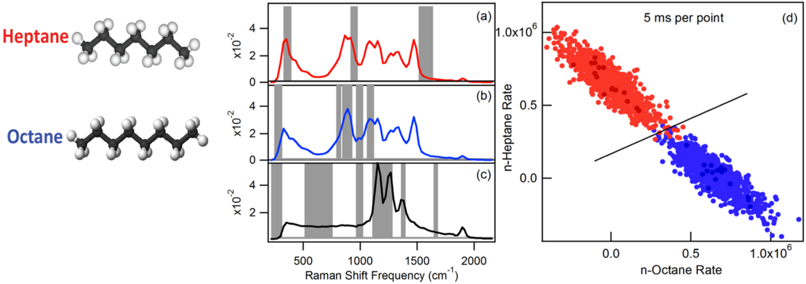

Experiment vs simulation

We can perform simulated experiments by sampling from a Poisson distribution. Comparing these to actual experiments produces very good agreement.

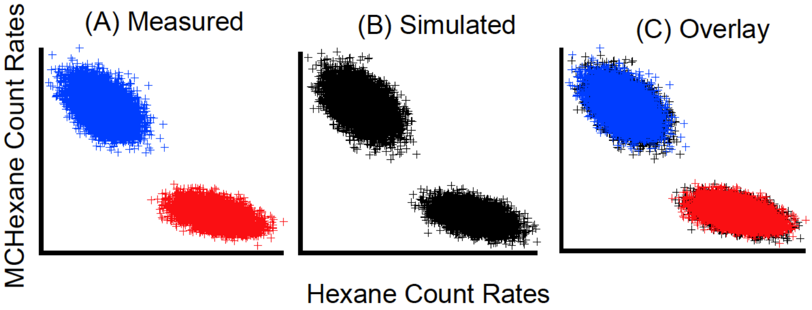

OB vs other filters

Comparing actual experiments using OB filters and two versions of filters commonly used in literature yields favorable results.

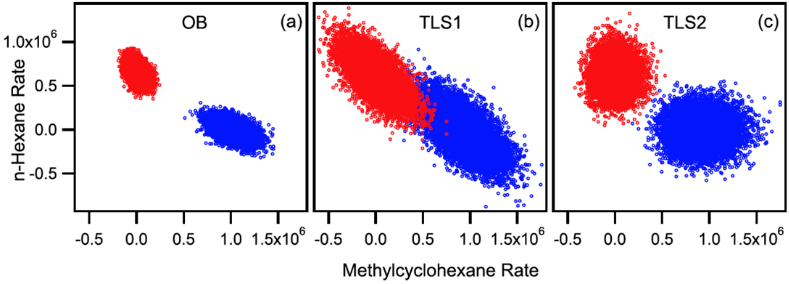

Accuracy as a function of time

Longer exposure time leads to less variance in the estimates.

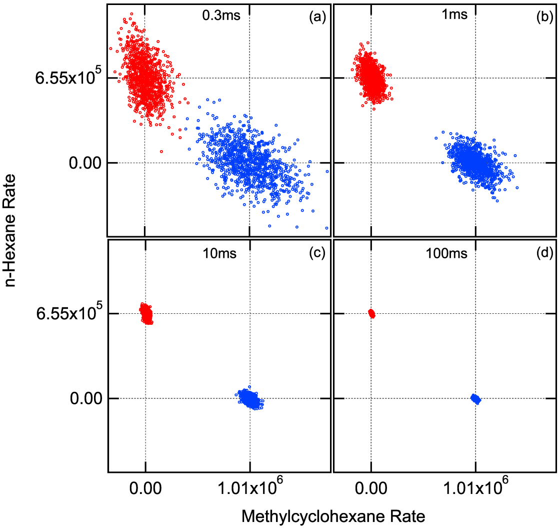

Quantification

The OB filters accurately identify the relative fraction of two chemicals.

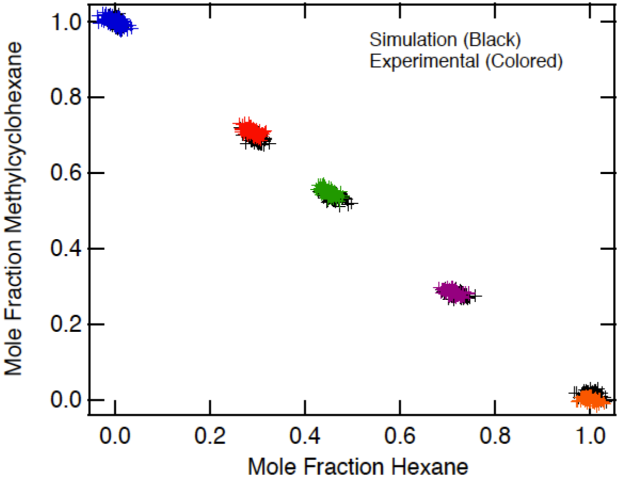

Imaging

Can do raster scanning to create chemical images.

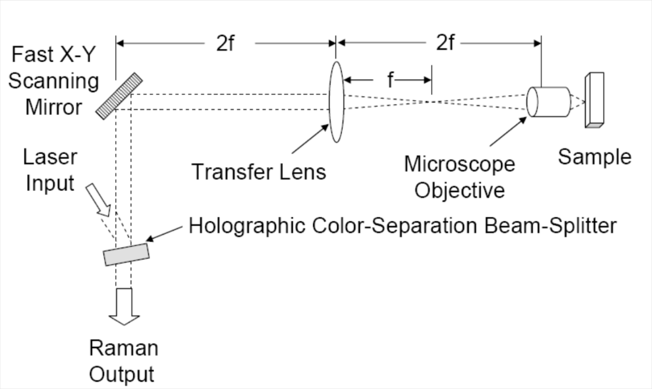

Imaging

Scan of chemical target shows good resolution, even at fast scanning rates.

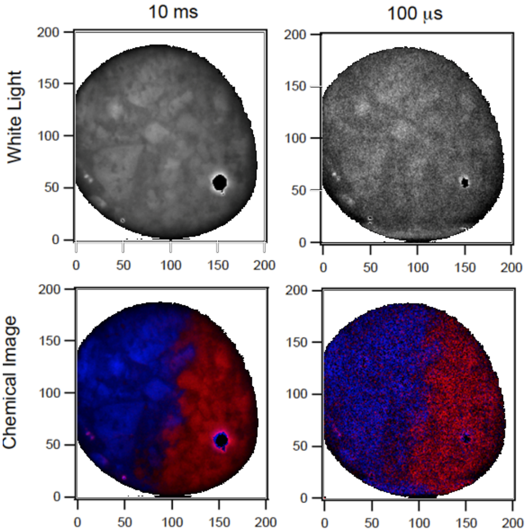

Fluorescence

For some materials, the laser may cause fluorescence - broad spectrum light that can overwhelm the spectrum of interest. How can we detect the signal in this case?

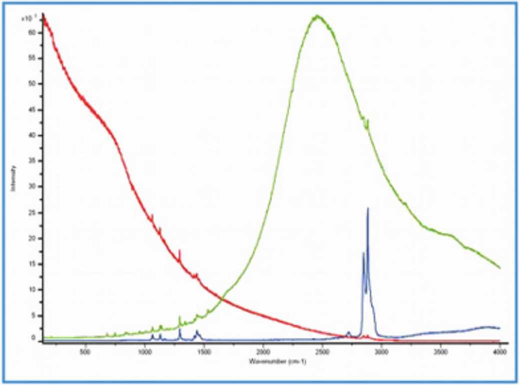

Fluorescence

Idea: use simple functions (polynomials) to model the fluorescence part of the spectrum and include corresponding filters.

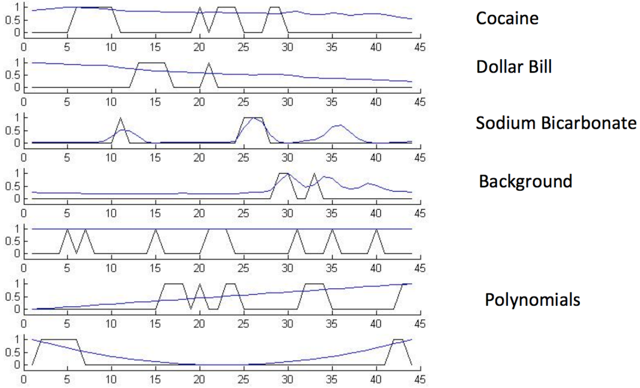

Fluorescence

Applying these filters to a dollar bill laced with baking soda and cocaine shows a clear distinction (I didn't know that chemists could order cocaine, but apparently so, at least in very small quantities!).

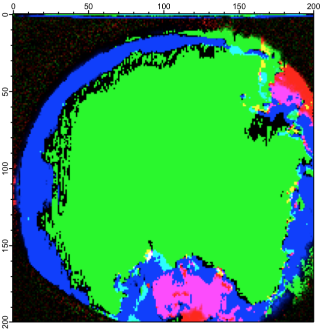

10 ms
integration time

Blue = Sodium Bicarbonate

Red = Cocaine

Green = Background (Dollar Bill)

Future work

- Identify one chemical among many with few filters.
- Adaptive sensing to identify unknown chemicals.
- Apply Bayesian and/or minimax methods to handle initial uncertainty in Λ.

Conclusions

- Physical sciences models are more accurate than biological models!
- OB filters outperform standard methods in low-signal regime.
- Could form the basis for fast/low-power spectroscopy - many applications.
- Use the model to propagate uncertainty from measurements to parameter estimates, and then set up an optimization problem to minimize uncertainty.
- Model is linear in parameters (like many surrogate models): simple in structure, but provides great flexibility and the use of powerful tools for UQ.