CENS
CENTER FOR EMBEDDED NETWORKED SENSING

Images as Sensors:
Using Visible Light Images to Measure Natural Phenomena

Josh Hyman
Advised by: Deborah Estrin and Mark Hansen

Oral Qualification
UCLA Computer Science Department
September 9th, 2008

Problem Statement

- Some natural phenomena can be measured without harming the environment:
  - Temperature
  - Rainfall
  - Humidity
- Others, require destructive instrumentation:
  - CO_2 flux from single plant, meadow, or soil

The site at James Reserve where mosscam is located

Proposal

- Construct a procedure that uses an imager to estimate phenomena that other sensors cannot measure
- Evaluate this procedure in the context of specific ecological applications

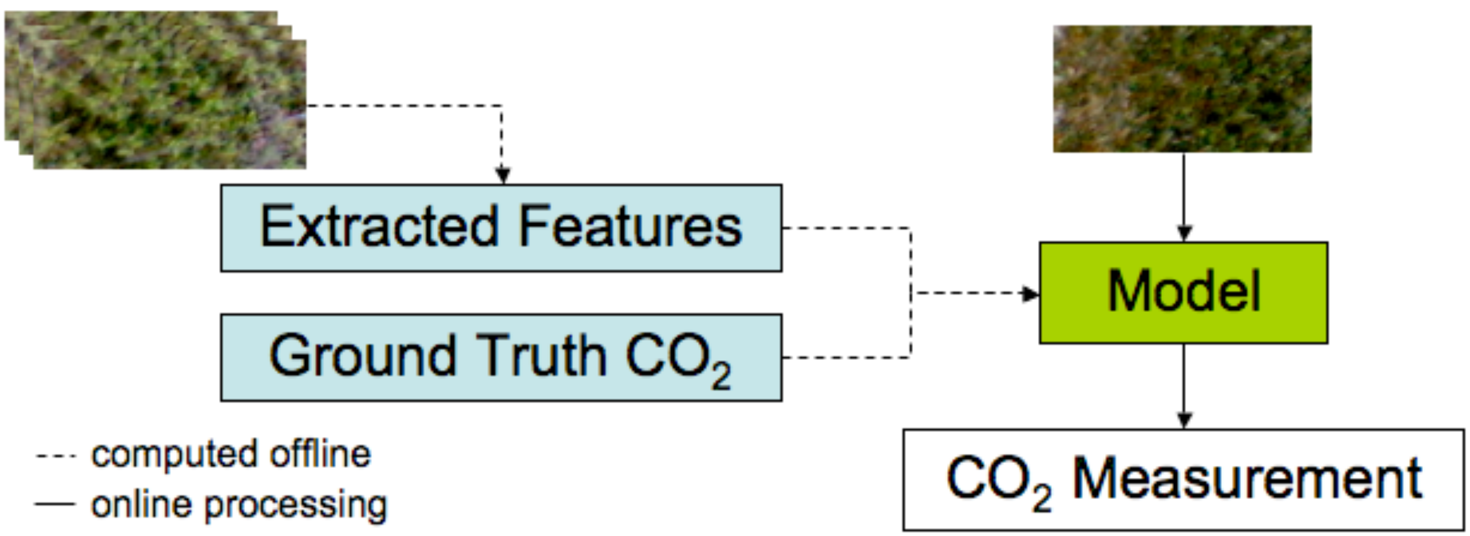

Challenges

- Field lighting conditions are variable
  - Illumination spectra changes hourly during the day
  - Daily spectra change over the course of a year
  - Results in unstable image features

9am

10:30am

12pm

1:30pm

- Approach: Reverse the effect of changing illumination before extracting image features

Challenges

- Not all image features have meaning with respect to the target ecological signal

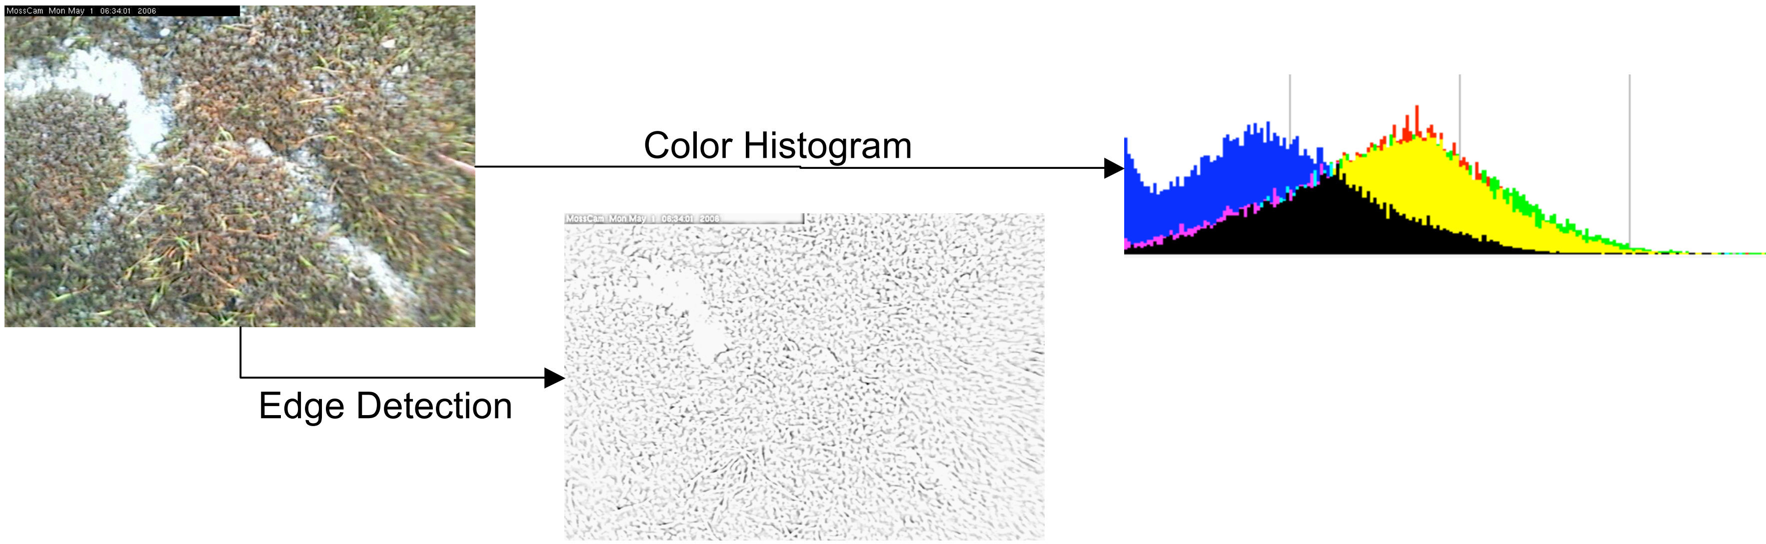

- Approach: Use the spectral reflectance of the subject as the image feature; easily verified in the field in the laboratory

Challenges

- Using distribution (like reflectance spectra) as independent regression inputs results in poor performance

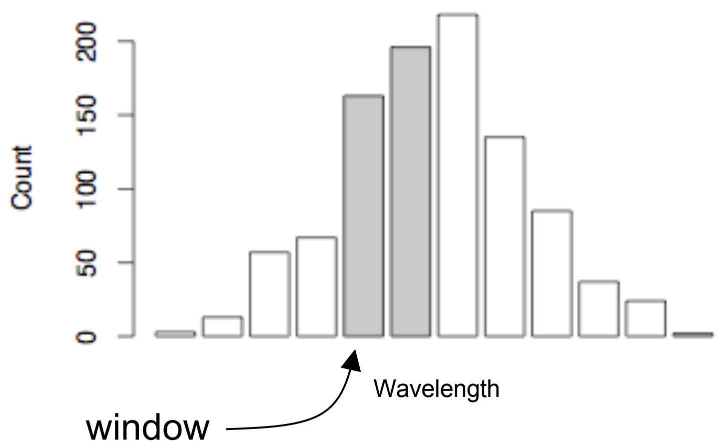

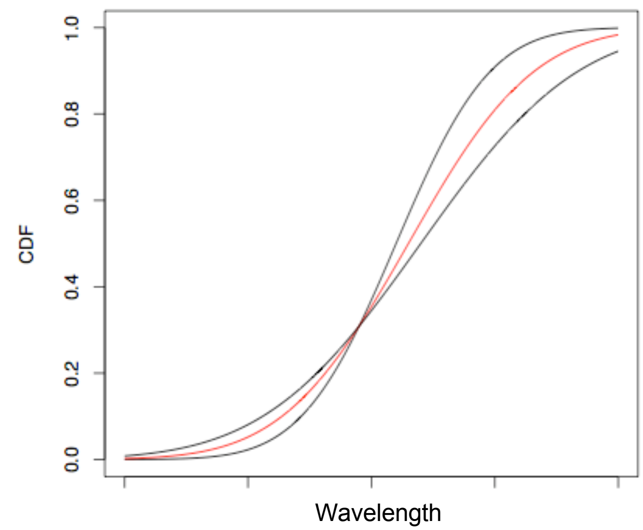

- Approach: Reformulate regression algorithms to use high-dimensional, highly structured data as input

Challenges

- Since the procedure is trained using lab data, and little or no data can be collected in the field, verifying that our estimates are accurate is difficult.

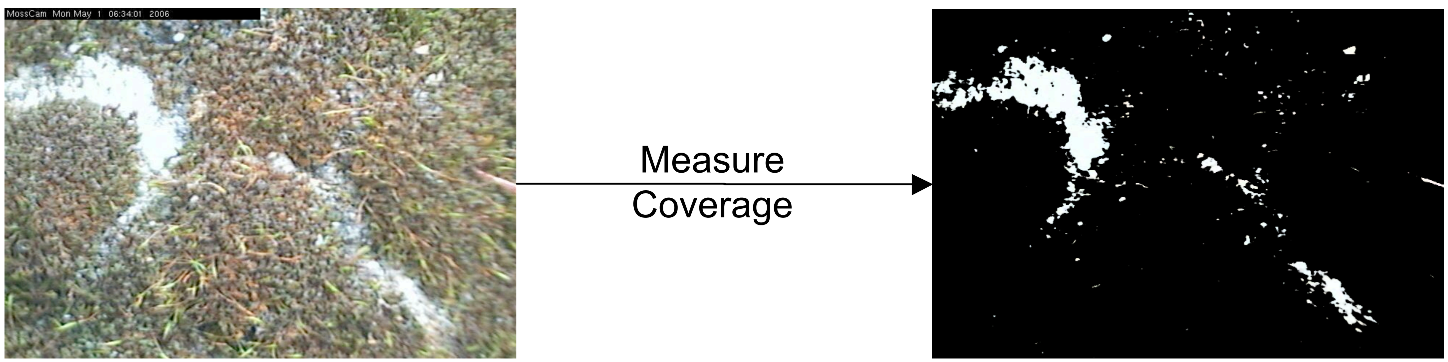

- Approach: use observable system characteristics, not suitable for measurement, to validate accuracy

Planned Contributions

- **Application evaluated image-based sensor toolkit:** A procedure to correlate images to ecological signals of interest using a series best-of-breed computer vision, image processing, and statistical learning algorithms.
- **High-dimensional, highly structured data as regression inputs**
- **Field-robust algorithms and methodology**

Process Overview

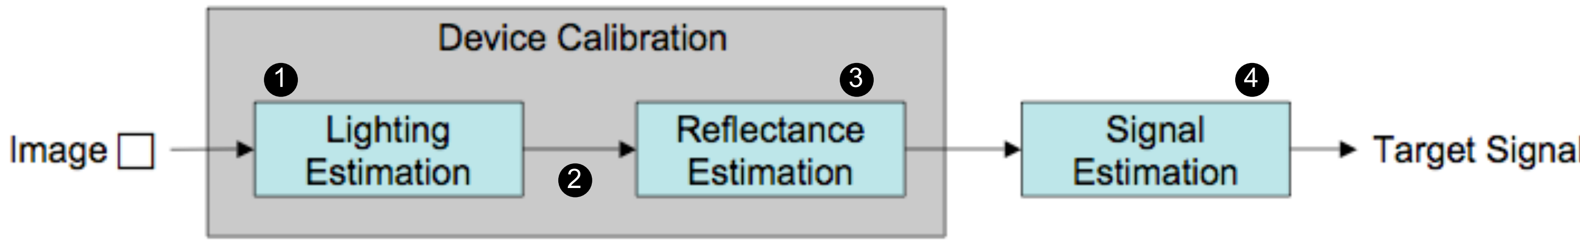

1. Estimate the incident illumination in the scene
2. Transform the image to be under a reference illuminant
3. Estimate the subject's spectral reflectance using color image features
4. Estimate the target signal

State of Current Work

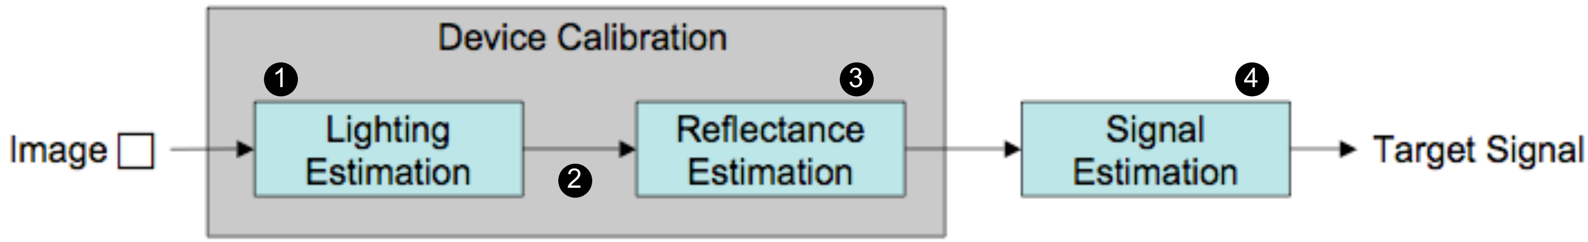

- A system for performing (1), (2), and (3) is operational but requires further tuning and evaluation
- Step (4) is formulated but requires generalization
- End-to-end validation is only formulated

Device Calibration: Physical Model of Image Formation

$$r_k = \int_w E(\lambda)S(\lambda)R_k(\lambda)d\lambda$$

- r_k = response of the kth sensor
- w = bandwidth of device
- $E(\lambda)$ = incident spectral power distribution
- $S(\lambda)$ = subject's relative spectral reflectance
- $R_k(\lambda)$ = the sensitivity of the kth sensor

Device Calibration: Modeled Image Formation

$$r \approx \hat{E}(\lambda)\hat{S}(\lambda)^T R(\lambda)$$

$$r \approx (\mathbf{B_e} w_e)(\mathbf{B_s} w_s)^T \mathbf{R}$$

- Discretize the bandwidth and rewrite in matrix notation
- Model E(λ) and S(λ) using functional PCA
- Results in 6 unknowns (w_e and w_s)

This system is under-constrained.
Estimate the spectra in sequence.

Device Calibration: Estimating Incident Illumination

Color By Correlation (Finlayson et. al.)

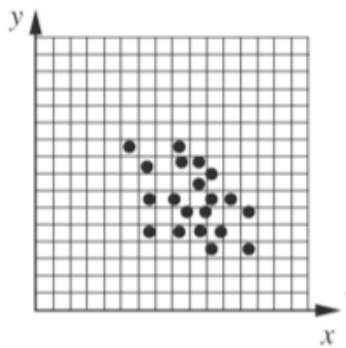

(a)

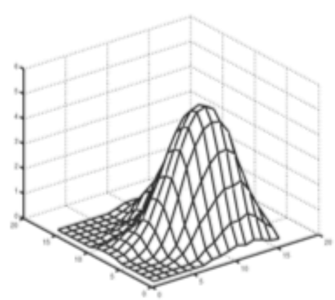

(b)

	ill 1	ill 2	ill 3	ill 4	ill 5	ill 6	ill 7	ill 8
x1, y1	0.1	0.0	0.7	0.1	0.0	0.0	0.0	0.9
x1, y2	0.2	0.1	0.6	0.2	0.0	0.4	0.1	0.9
x1, y3	0.2	0.1	0.5	0.4	0.0	0.4	0.1	0.7
	:	:	:	:	:	:	:	:
	:	:	:	:	:	:	:	:
	:	:	:	:	:	:	:	:
x1, yn	0.6	0.3	0.3	0.9	0.2	0.6	0.1	0.0
x2, y1	0.9	0.3	0.1	0.9	0.4	0.7	0.1	0.0
x2, y2	0.4	0.3	0.1	0.7	0.2	0.5	0.2	0.0
	:	:	:	:	:	:	:	:
	:	:	:	:	:	:	:	:
	:	:	:	:	:	:	:	:
	:	:	:	:	:	:	:	:
	:	:	:	:	:	:	:	:
	:	:	:	:	:	:	:	:
xn, yn	0.0	0.1	0.0	0.1	0.5	0.1	0.6	0.0

(c)

Device Calibration: Lighting Transformation

$$\begin{bmatrix} E_1(\lambda_R)S(\lambda_R) \\ E_1(\lambda_G)S(\lambda_G) \\ E_1(\lambda_B)S(\lambda_B) \end{bmatrix} = T_{light} \begin{bmatrix} E_2(\lambda_R)S(\lambda_R) \\ E_2(\lambda_G)S(\lambda_G) \\ E_2(\lambda_B)S(\lambda_B) \end{bmatrix}$$

$$T_{light} = \begin{bmatrix} E_1(\lambda_R)/E_2(\lambda_R) & 0 & 0 \\ 0 & E_1(\lambda_G)/E_2(\lambda_G) & 0 \\ 0 & 0 & E_1(\lambda_B)/E_2(\lambda_B) \end{bmatrix}$$

- Given some $E(\lambda)$ found previously, we can compute T_{light}
- This assumes that $R(\lambda) = \delta_k \quad k = \{R, G, B\}$
- Though unrealistic, this assumption has been shown to hold for cameras presented with “reasonable” light sources (like daylight)

Device Calibration: Estimating Spectral Reflectance

- Unlike lighting estimation, we have less insight into the relationship between the spectral reflectance model and the image's color coordinates.
- Estimate the 3 model parameters using non-linear regression
- Use chromaticity coordinates as input

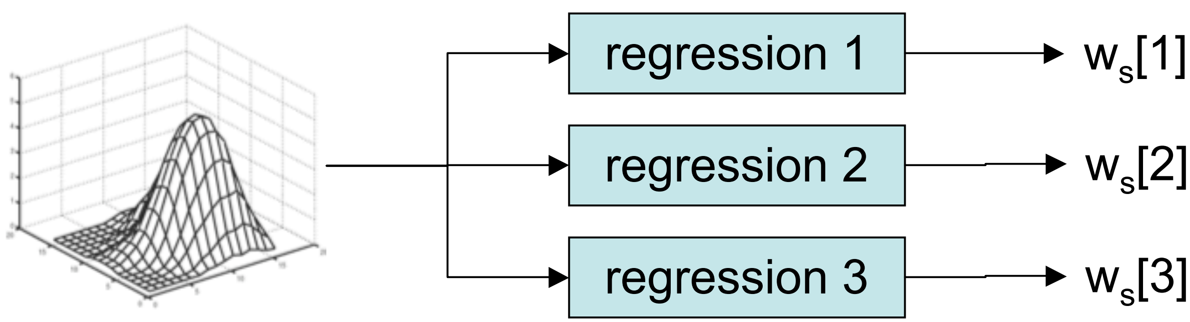

Modeling the Target Signal

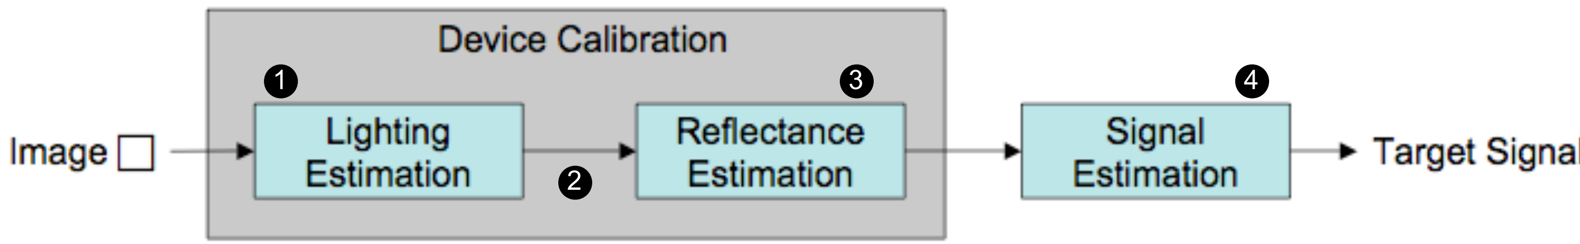

- We've completed device calibration
- Now, we model the target signal (4) using subject's spectral reflectance (the output of device calibration)

Modeling the Target Signal

- We would like to use high-dimensional, highly structured data as regression inputs
- Results in:
  - polynomial explosion in runtime
  - curse of dimensionality

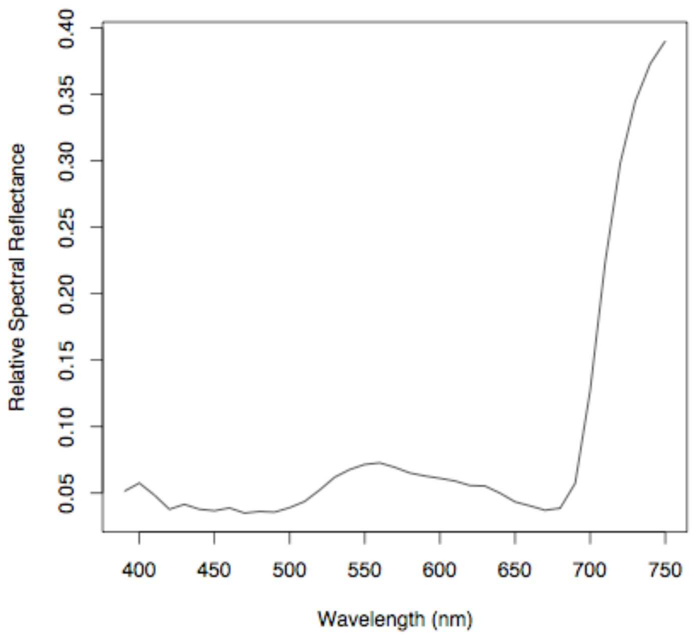

Sample moss relative spectral reflectance to be use in CO_2 modeling

Modeling the Target Signal

- Approximate the data as a mixture of location-scale distributions.
- How does this help?
  - Error function only depends on the location of an example relative to the threshold
  - Requires we keep track of intersections (which is faster than computing all thresholds)

Modeling the Target Signal

- Keeping track of the error is simply a matter of choosing a *threshold CDF* and checking for intersections with other examples.

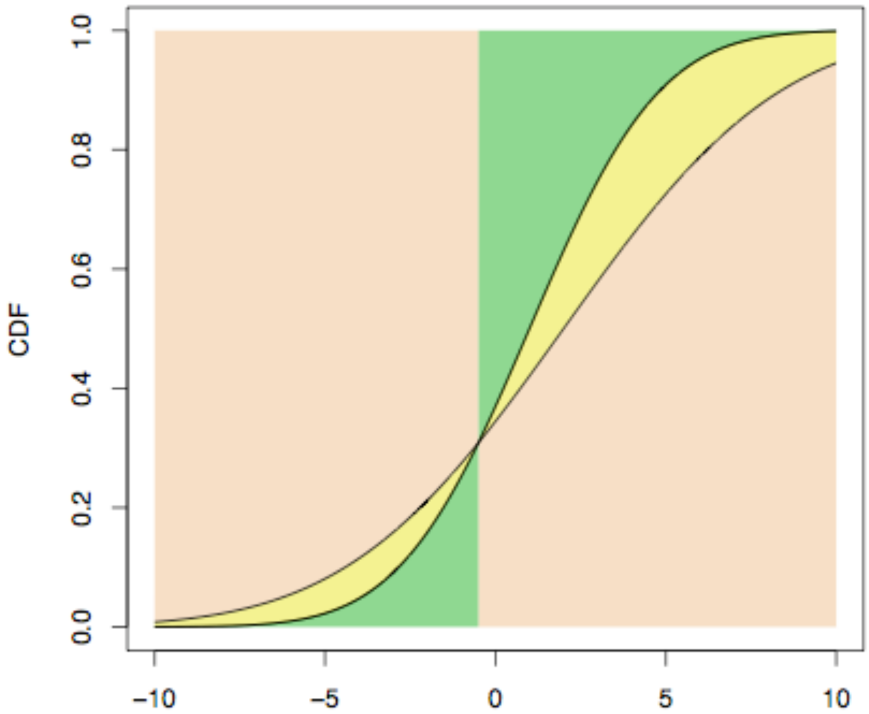

Possible threshold in similar colors

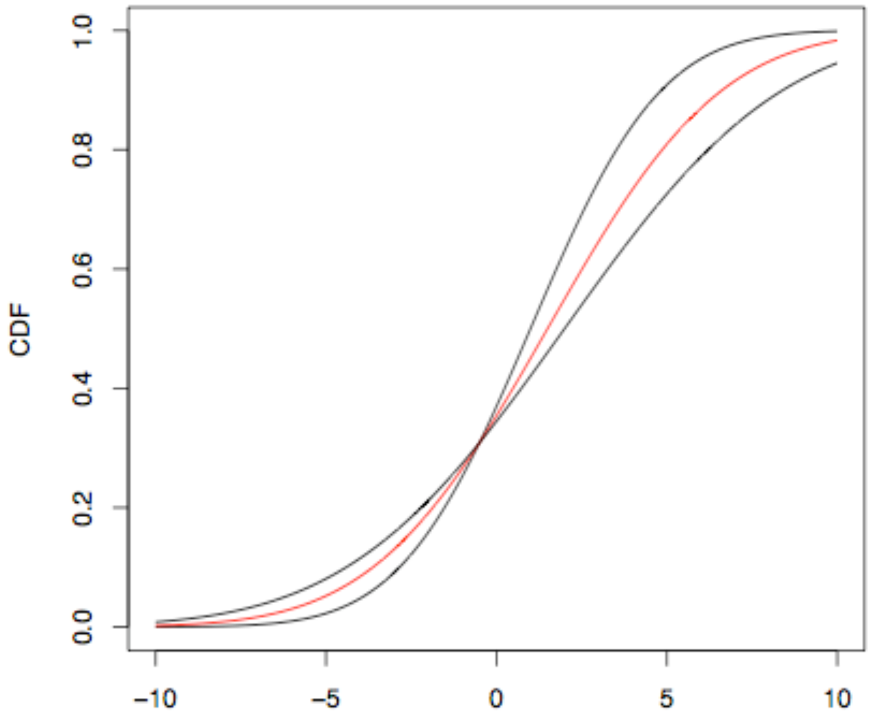

One threshold's (yellow) CDF

Modeling the Target Signal

- Intersections can be found quickly, if the input can be modeled as:
  - A set of distributions with related means and variances
  - A set of distribution mixtures with fixed mean and variance and varying weights
- In these cases we can find the “next” intersection in O(1)

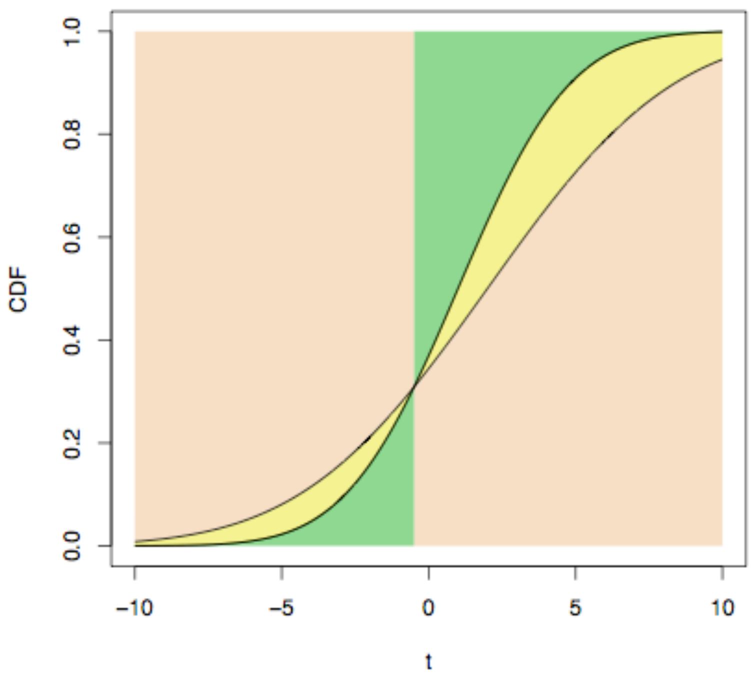

Modeling the Target Signal

- Current regression runtime: $O(nk^2)$
  - n = number of features
  - k = number of examples

- For inputs which meet those criteria, the modified algorithm's runtime: $O(nh + k)$
  - n = number of distribution features
  - h = number of discretized buckets per distribution
  - k = number of examples

Model Validation

- Types of Validation:
  - Magnitude of estimates
  - Temporal correctness
- Common techniques
  - Corroborate estimates with easily observable characteristics
  - Incorporate sensors aside from the imager
  - Compare estimates to results from less controlled experiments

Process Overview (refined)

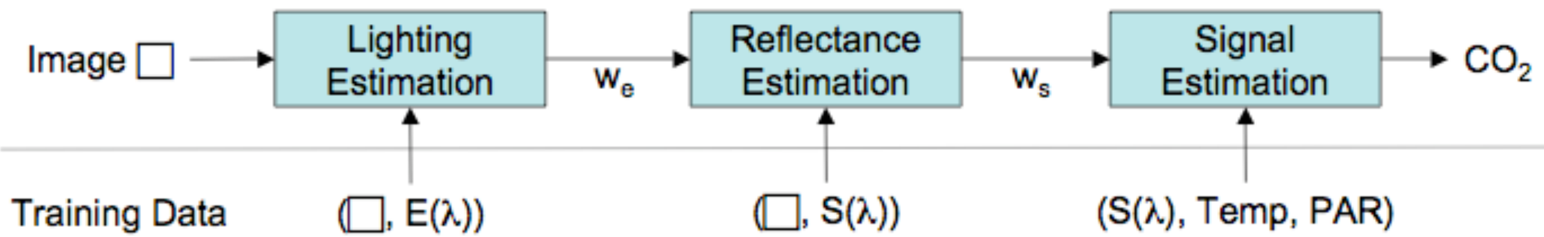

- Training data must be acquired to build this process:
  - The incident lighting in the field, E(λ)
  - The possible reflectance of the subject, S(λ)
  - The "appearance" of the system for measured values of the target signal: S(λ), Temperature, and PAR

Application: Measuring Moss Photosynthesis

Ecologists want to determine the effect of short summer rain events on the moss' ability to survive

- There are no available sensors
- Methods suggested by previously ecological studies have insufficient temporal resolution

Tortula princeps

Photosynthesis begins to occur 5 minutes after being hydrated

Evaluation: Experimental Setup

1. Collect moss from JR
2. Hydrate moss and allow to dry over time
3. Collect samples:
   a. illumination
   b. spectral reflectance
   c. high-quality images
   d. low-quality images

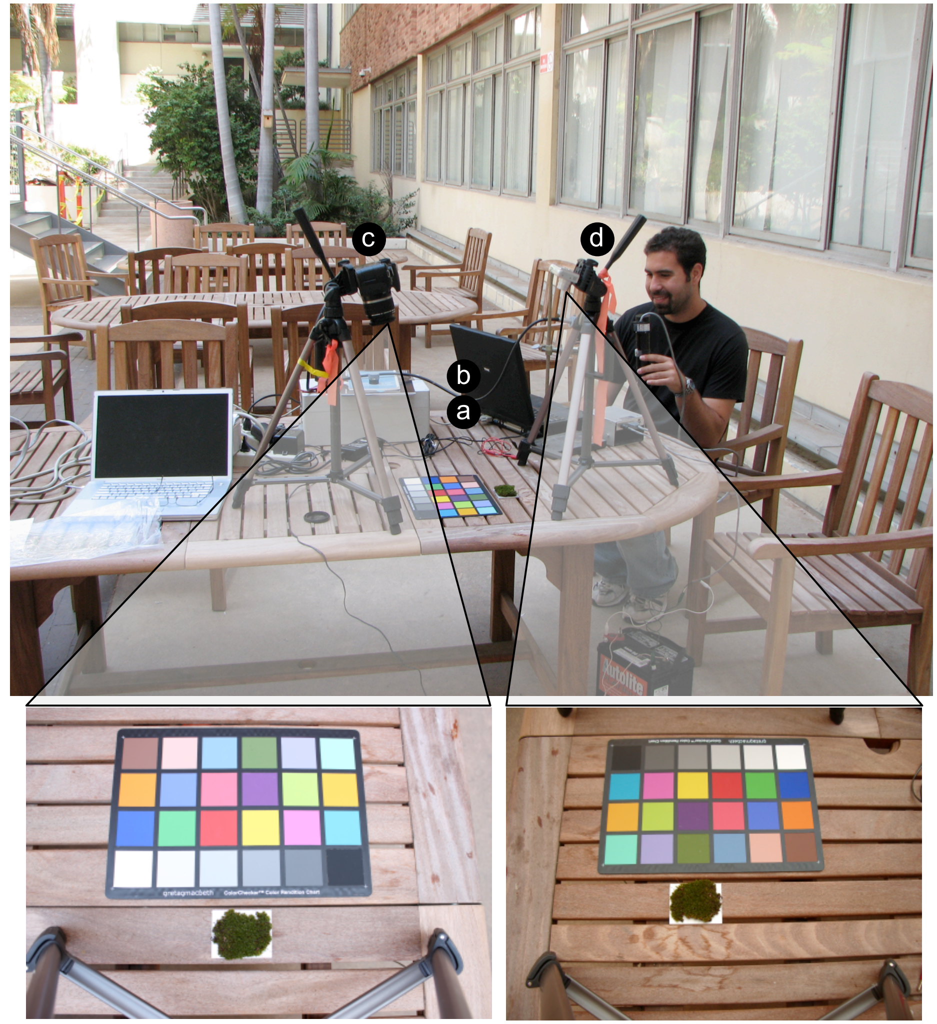

Samples acquired every 15 min for ~6 hrs (23 samples total)

Evaluation: Incident Illumination Modeling

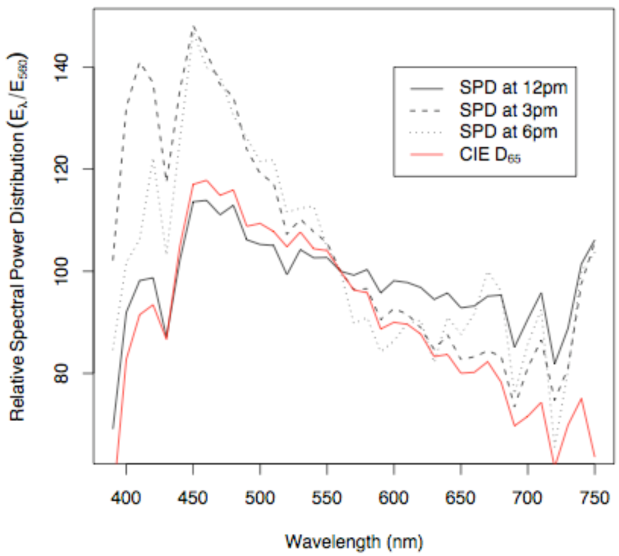

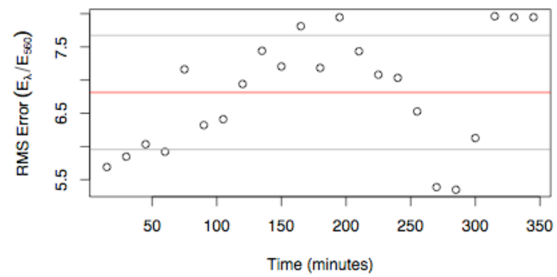

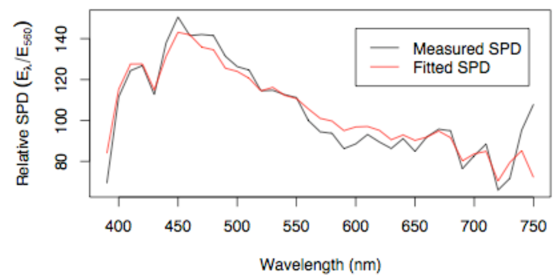

- Measured illumination (left) is similar to D_{65} although it is slightly bluer
- Model (by Judd et. al.) fits well (right top), with a slight temporal component to the error
- Even the worse error has minimal error and correct characteristic shape (right bottom)

Evaluation: Estimating Incident Illumination

- Accuracy of the Color by Correlation algorithm becomes reasonable (top) once enough training examples are used
- With 12 training examples, we find that error (bottom) clusters near zero
- Interestingly, performance was comparable with and without JPEG compression

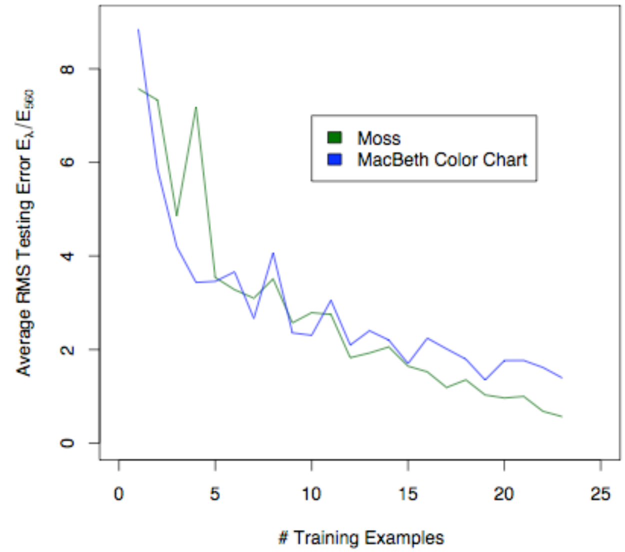

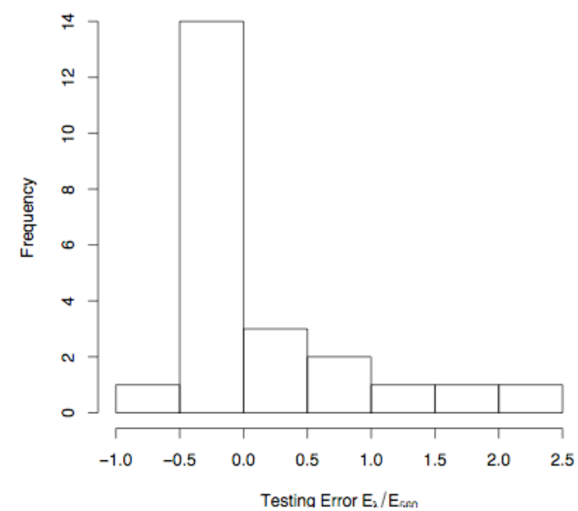

Evaluation: Lighting Transformation

- Images of a reference object (MacBeth color chart) shift over the course of the day (top)
- We visualized this change using the 2D Jenson-Shannon divergences of all pairs of images
  - Before the transform (middle), the image's divergences large
  - After the transform (bottom), the image's divergences were compressed towards zero

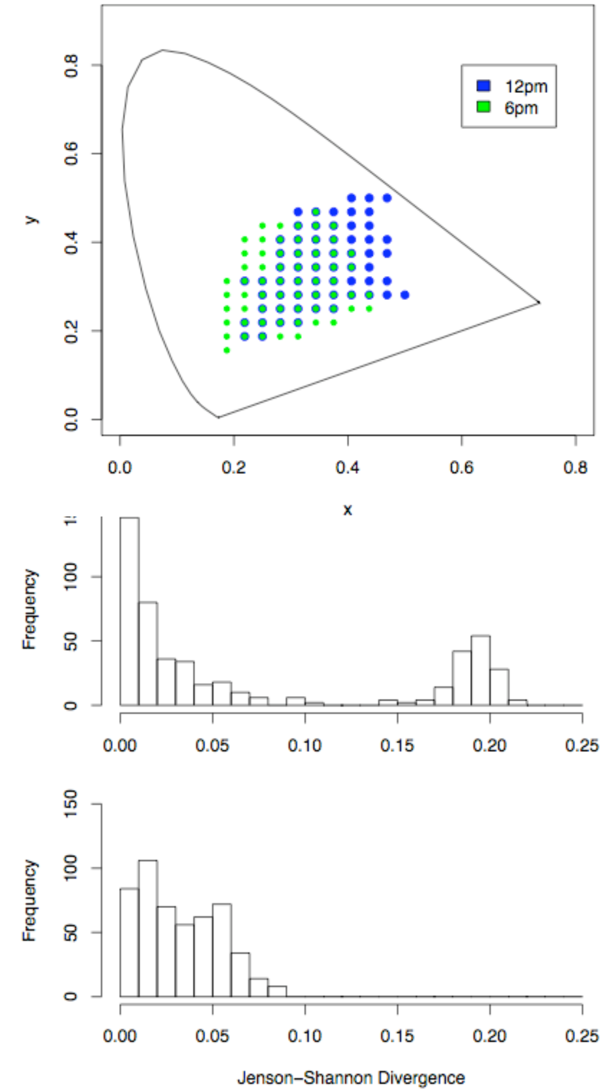

Evaluation: Estimating Spectral Reflectance

- We use only 3 basis functions: they contain 99.96% of the data
- The variation in the second and third basis functions (top) is expected:
  - variation low and high in the spectra caused by the sensor
  - variation in the middle caused by changes in the moss
- Predicted reflectance is quite good (middle)
- The fit with the maximum error (bottom) is very close below 700nm

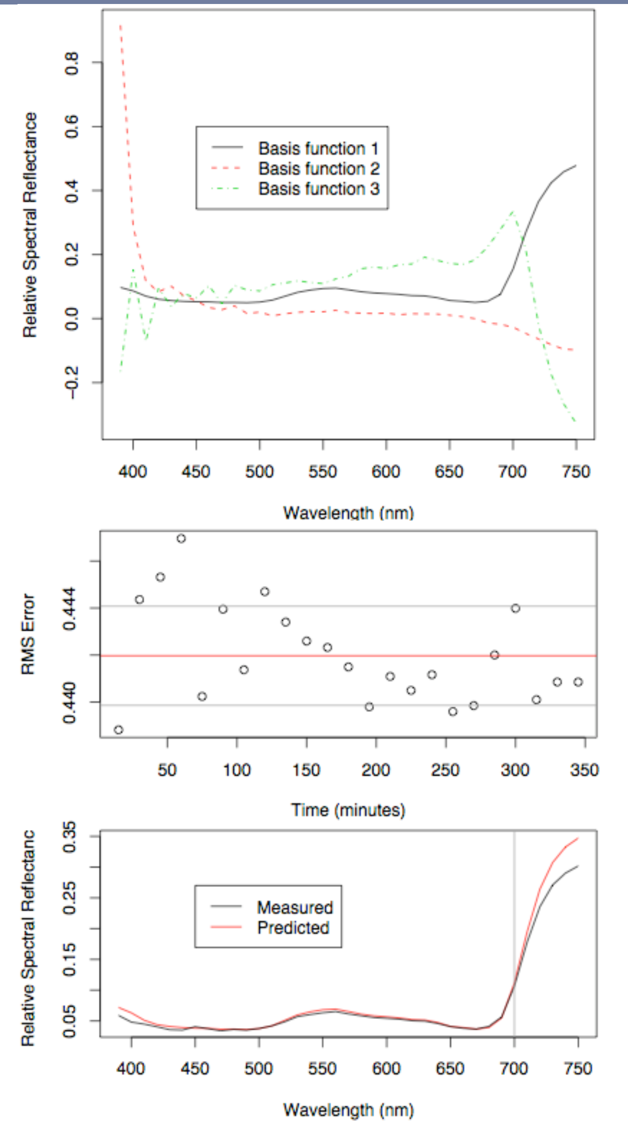

Incorporating Prior Work: Ecology Research

- Environmental Productivity Index (EPI)
  - Measures factors limiting maximum CO_2 uptake:
    - Temperature
    - Moisture
    - Light
  - Gathered CO_2 data in the laboratory
  - Tested in the field with Agave plants
  - Validated results using newly unfolding leaves
- Use this work as a foundation for model validation techniques

Incorporating Prior Work: Experimental Data

- Graham et. al. performed laboratory experiments on this particular moss plant
  - Measured CO_2 over a drydown with constant "day" and "night" temperatures and lighting
  - Measured CO_2 with varying lighting or temperature
- Use these accurate measurements to train our models

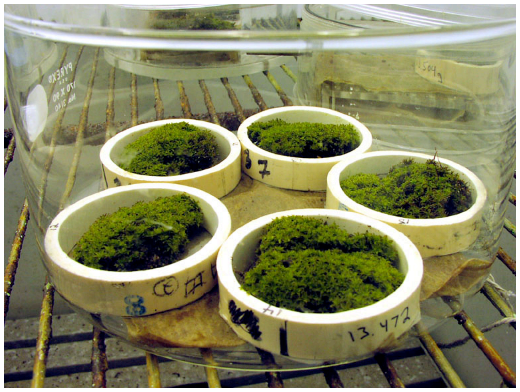

Moss in a temperature controlled chamber

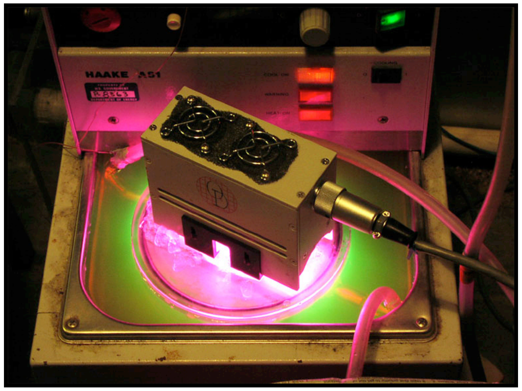

Infrared gas analyzer

Other Applications

- Measure CO_2 uptake in meadows using NEON flux towers
- Measure CO_2 release of soil below the forest canopy
- Measure particulate matter suspended in a river channel

Timeline

1. Reformulate and further test re-lighting (2 mo)
2. Generalize regression modifications (2 mo)
3. Complete moss photosynthesis evaluation (3 mo)
4. Build online system to predict moss photosynthesis (1 mo)
5. Apply to another application (6 mo)
6. Dissertation Writing (2-3 mo)

Planned Contributions

- **Application evaluated image-based sensor toolkit:** A procedure to correlate images to ecological signals of interest using a series best-of-breed computer vision, image processing, and statistical learning algorithms.
- **High-dimensional, highly structured data as regression inputs**
- **Field-robust algorithms and methodology**

Questions?

Related work

- Illumination Modeling: Judd et. al.
- Illumination Estimation: Finlayson et. al.
- Lighting Transform: Forsynth et. al.
- Regression modification: Johnson et. al.