Comparison of sampling strategies for 3D scene reconstruction from sparse multispectral lidar waveforms

Yoann ALTMANN

School of Engineering and Physical Sciences
Heriot-Watt University

New mathematical methods in computational imaging
June, 29th 2017

LEADERS IN **IDEAS** AND **SOLUTIONS**

Single-photon Lidar

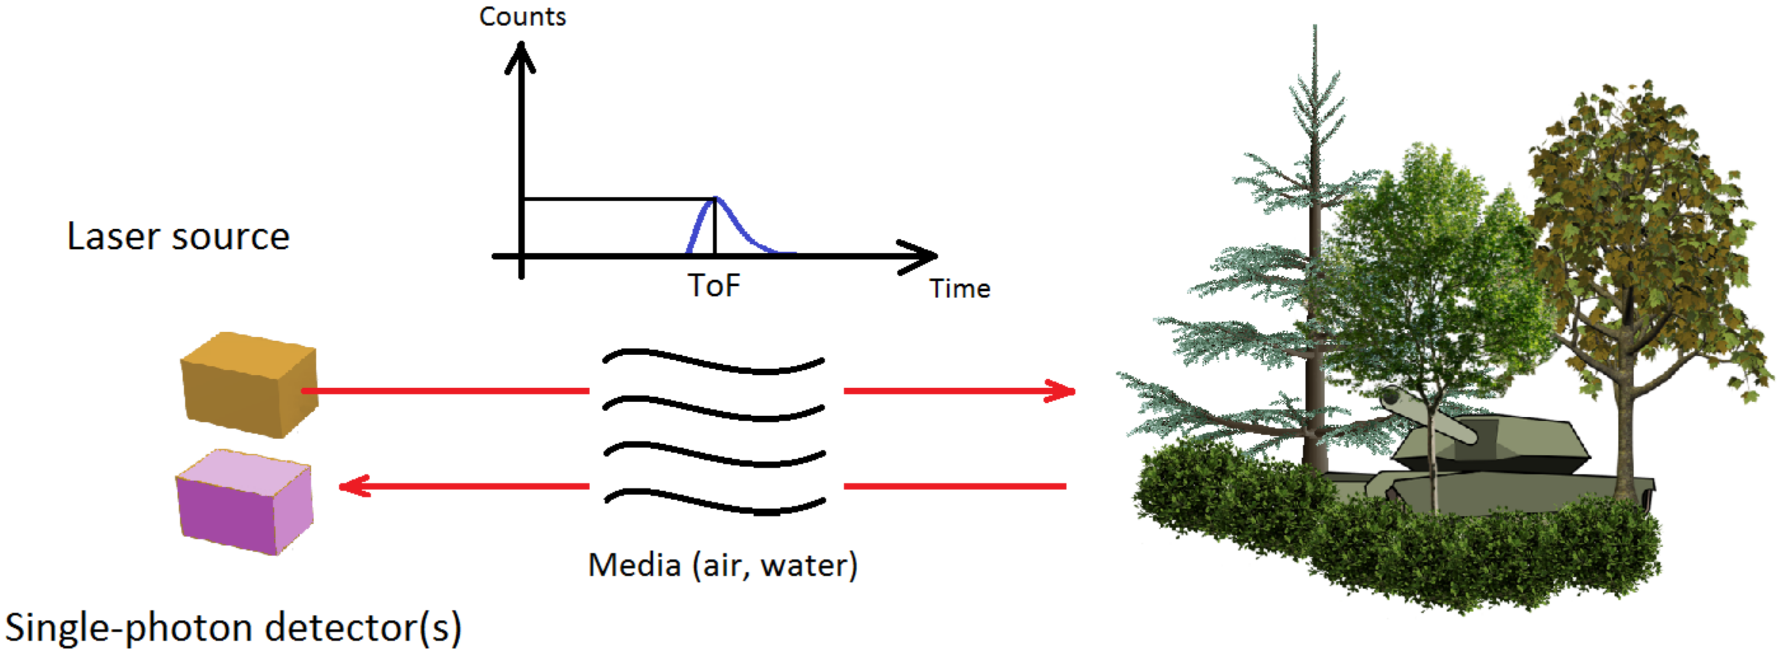

- Pulsed laser (20 MHz), low power (≈ μW)
- Detector: single-photon avalanche diode (SPAD)
- Time of flight: for each detected photon (precision ≈10^{-12} s)
  – Path length precision ≈ 600μm

Observation model

Ideal scenario

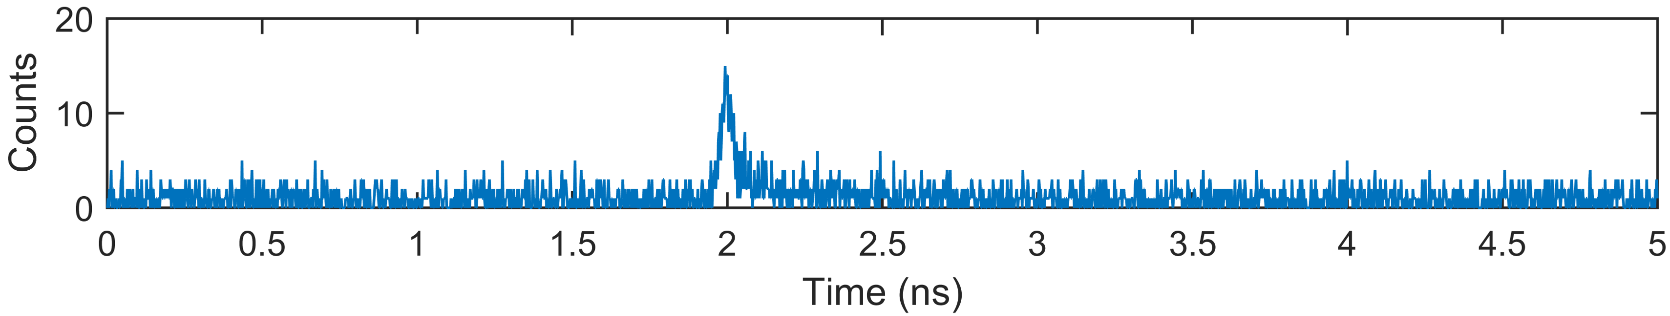

$$y_{n,t} \sim \mathcal{P}\left(r_n g_0(t - t_n) + b_n\right), \quad t \in \{1,\dots,T\}$$

- $y_{n,t}$: photon count in *t*th bin
- $g_0(\cdot)$: instrumental response
- T : Histogram length
- b_n : background level
- r_n : target reflectivity
- t_n : Time-of-flight (ToF)

Observation model

Low power/long range

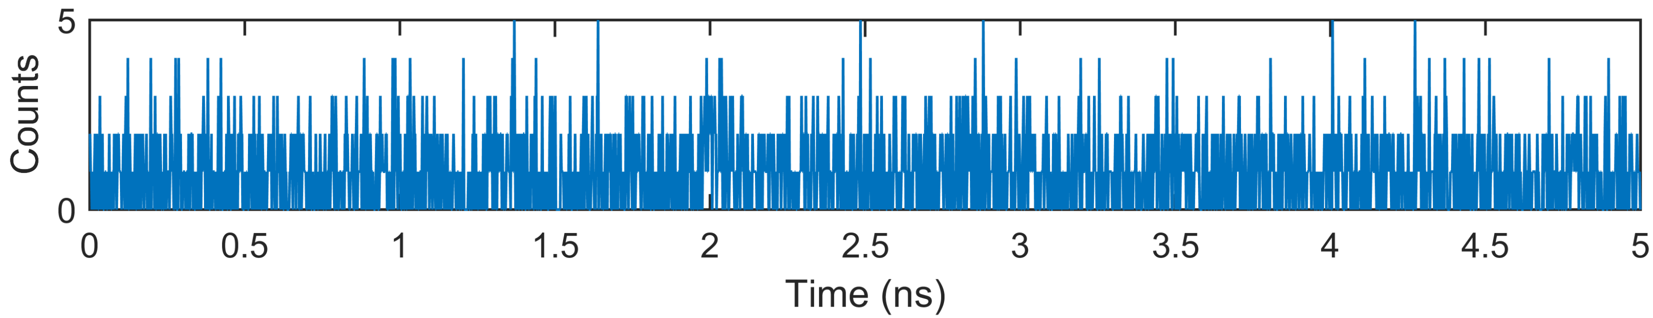

$$y_{n,t} \sim \mathcal{P}(r_n g_0(t - t_n) + b_n), \quad t \in \{1,\dots,T\}$$

- $y_{n,t}$: photon count in tth bin
- $g_0(\cdot)$: instrumental response
- T : Histogram length
- b_n : background level
- r_n : target reflectivity
- t_n : ToF

Observation model

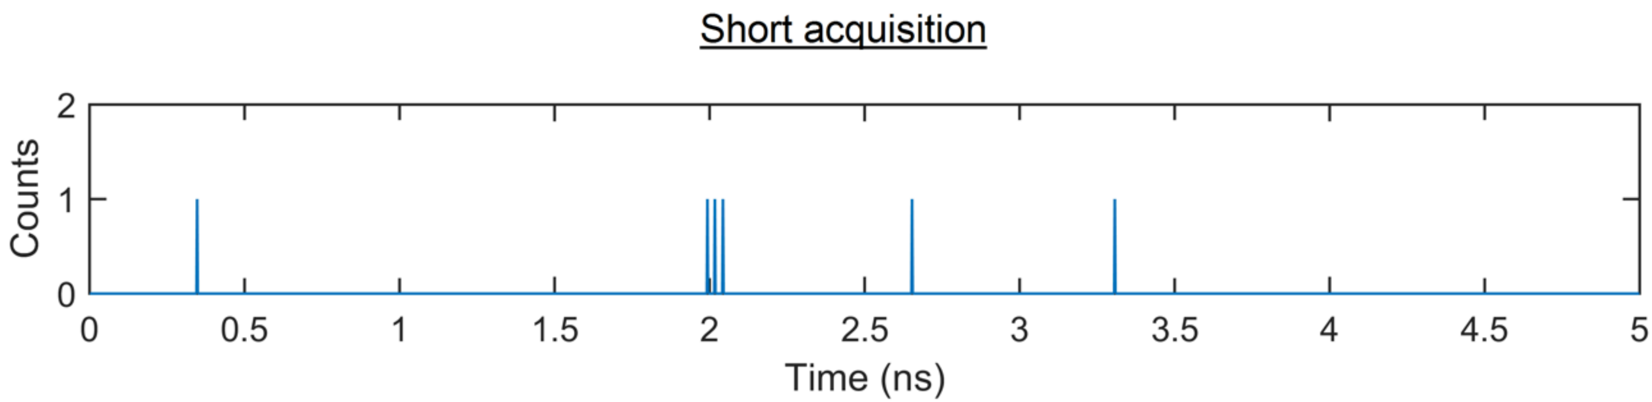

$$y_{n,t} \sim \mathcal{P}\left(r_n\, g_0(t - t_n) + b_n\right), \quad t \in \{1, \ldots, T\}$$

- $y_{n,t}$: photon count in tth bin
- $g_0(\cdot)$: instrumental response
- T : Histogram length
- b_n : background level
- r_n : target reflectivity
- t_n : ToF

Target identification from extremely sparse histograms

HERIOT WATT UNIVERSITY

Multispectral Lidar

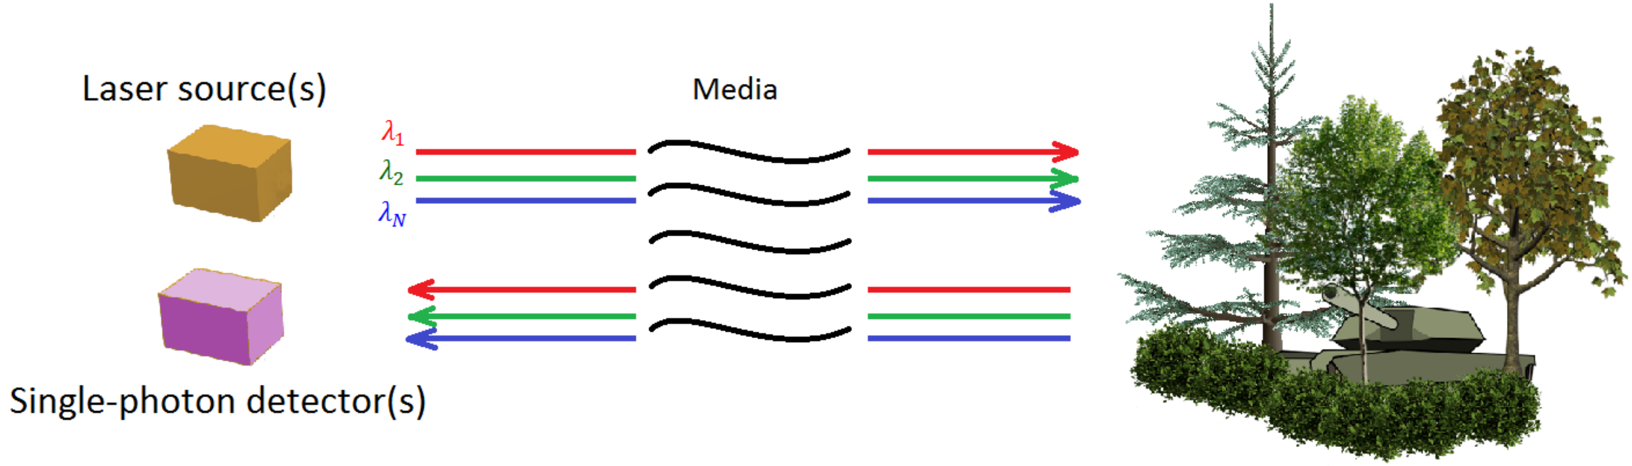

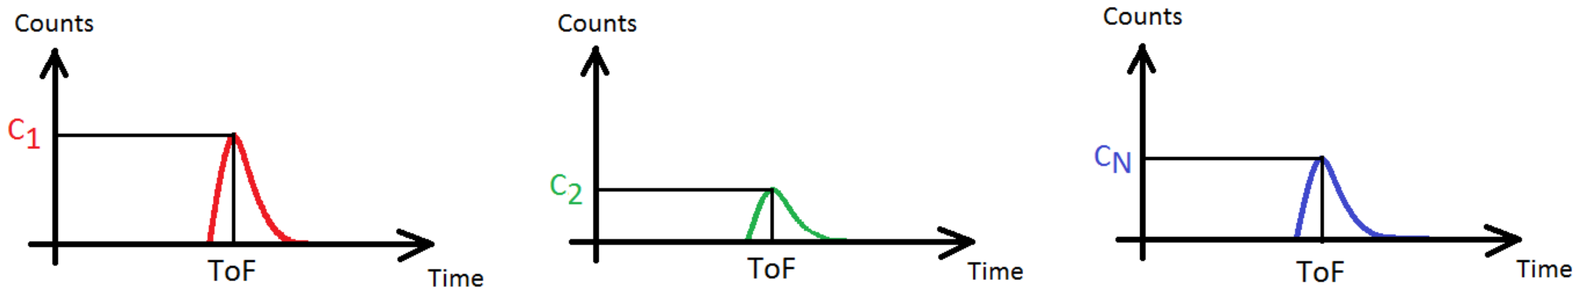

HERIOT WATT UNIVERSITY

Motivations

- Joint extraction of geometric and spectral information
  - Limited data registration issues (fusion Lidar/HSIs)

- Robustness
  - Energy spread across wavelengths (range estimation)
  - Illumination conditions (shadowing effects)

- Scene reconstruction with few photons
  - fast/long range imaging
  - < 10 of useful photons per pixel

Applications

- Defence

Long-range target identification

- Oil & gas, underwater imaging

Pipeline inspection

- Environmental sciences

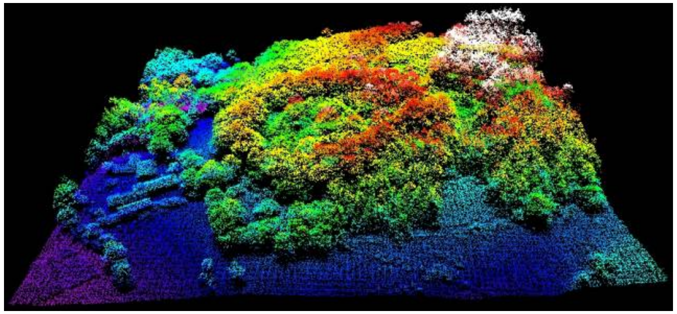

Forest canopy monitoring

Observation model

$$y_{n,\ell,t} \sim \mathcal{P}(r_{n,\ell}\, g_{0,\ell}(t - t_n) + b_{n,\ell})$$
$$t \in \{1,\dots,T\},\ \ell \in \{1,\dots,L\}$$

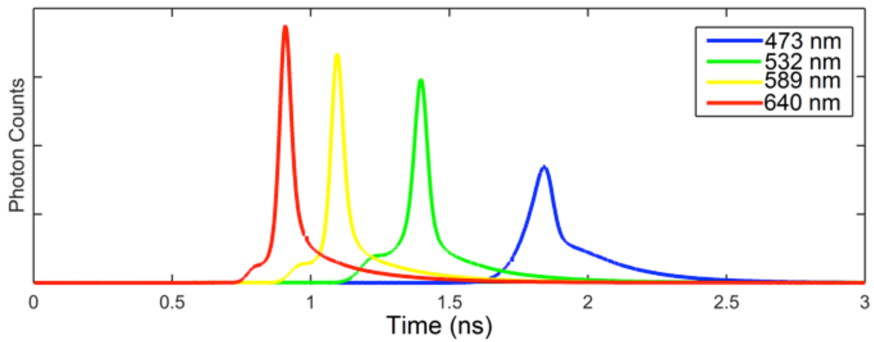

Examples of instrumental responses

- $y_{n,\ell,t}$: photon count in tth bin (ℓ th band)
- $g_{0,\ell}(\cdot)$: instrumental response
- T : Histogram length
- $b_{n,\ell}$: background level
- $r_{n,\ell}$: target reflectivity
- t_n : ToF

Estimation of t_n, $\{b_{n,\ell}\}$ and $\{r_{n,\ell}\}$ for each pixel

3D scene analysis

Scene of interest

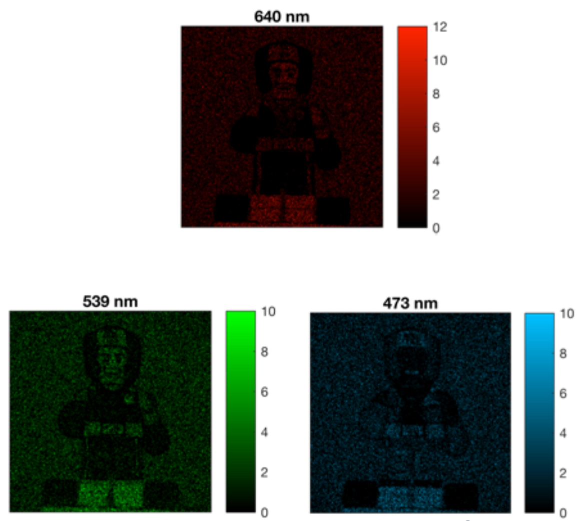

Observations (≈1 photon/pixel)

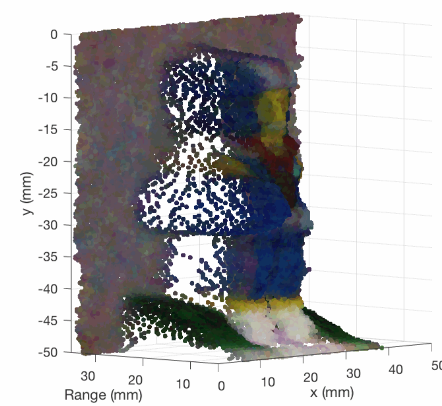

Estimated 3D profile

- Recovery of range and colour profiles from extremely low photon counts (denoising)

3D scene analysis

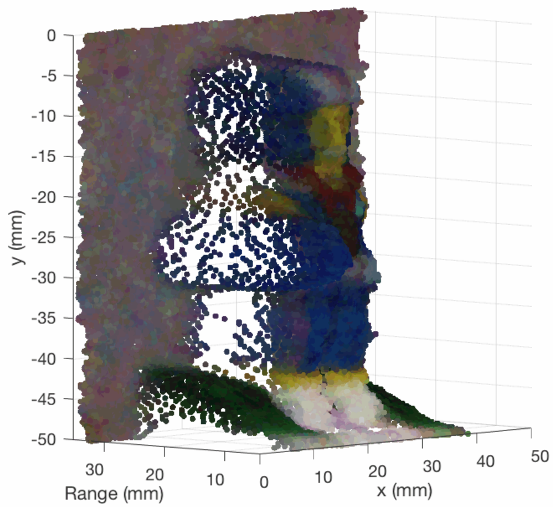

≈ 3 x 0.33 photons/pixel

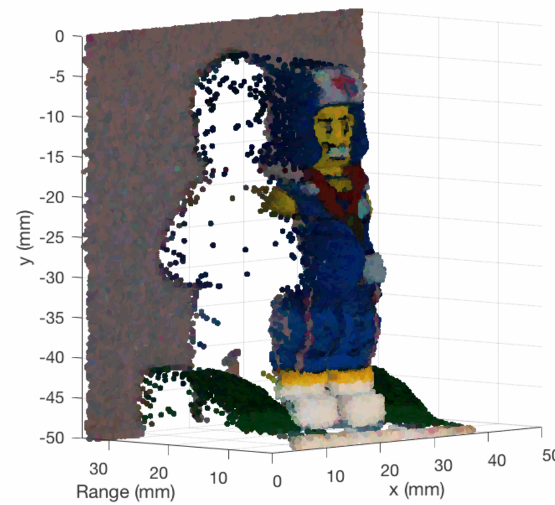

≈ 3 x 3.3 photons/pixel

HERIOT WATT UNIVERSITY

Spectral clustering/classification

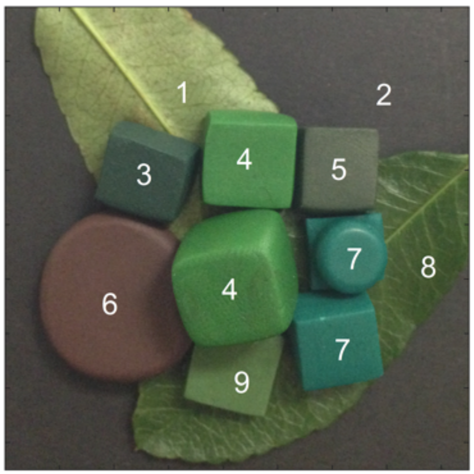

RGB image (5 x 5 cm)

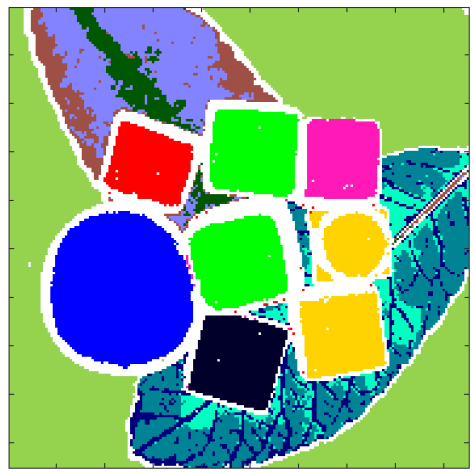

Unsupervised spectral clustering

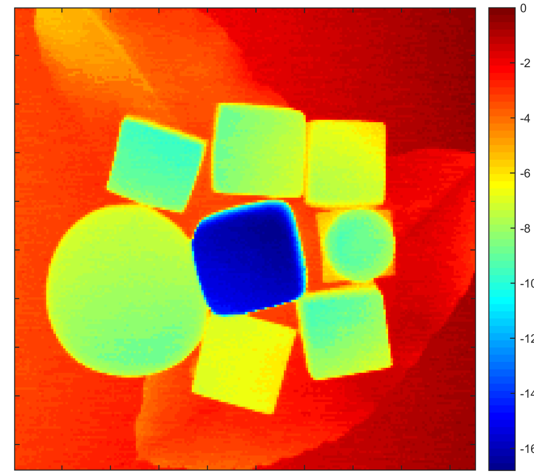

Estimated range profile (in mm)

- 33 spectral bands, 500– 820 nm, 200x200 pixels

First computational methods for spectral analysis from extremely photon-limited Multispectral Lidar data[1]

[1]Altmann et al., EUSIPCO 2016, IEEE SSP 2016, WHISPERS 2016

Spectral clustering/classification

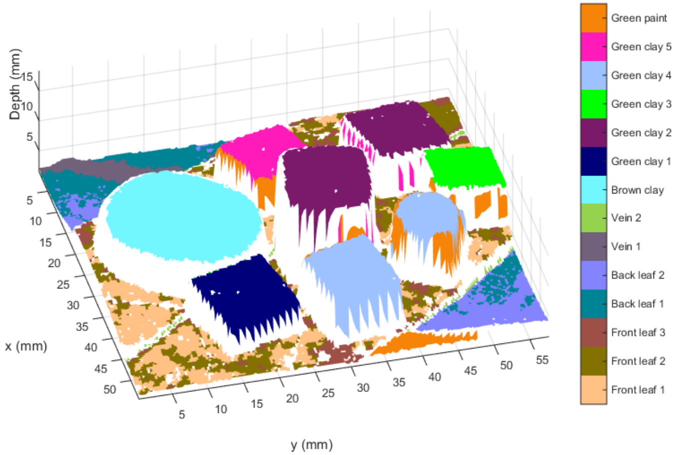

Depth/classification profile

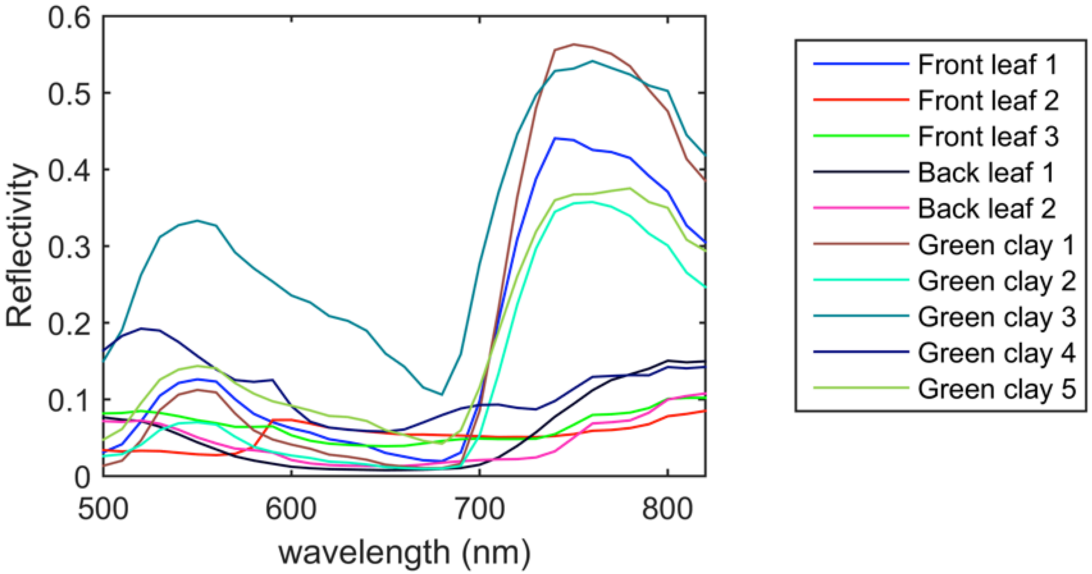

Main spectral signatures (mean of spectral classes) identified

Robust Spectral unmixing

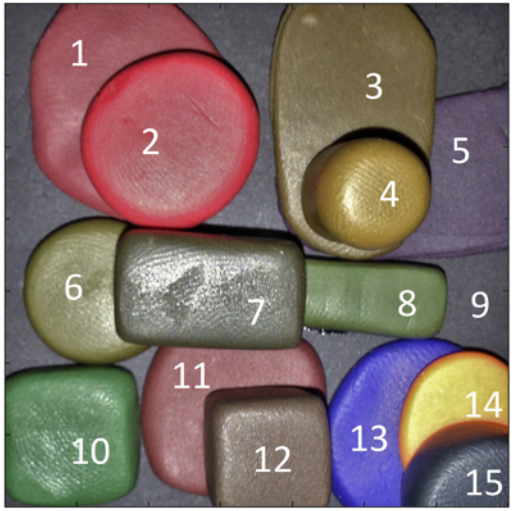

RGB image (5 x 5 cm)

- Material quantification, anomaly detection (range ≈ 1,80 m)
- Known spectral library
- Acquisition time per pixel: 1, 3 or 10 photons per band on average

Abundance estimation

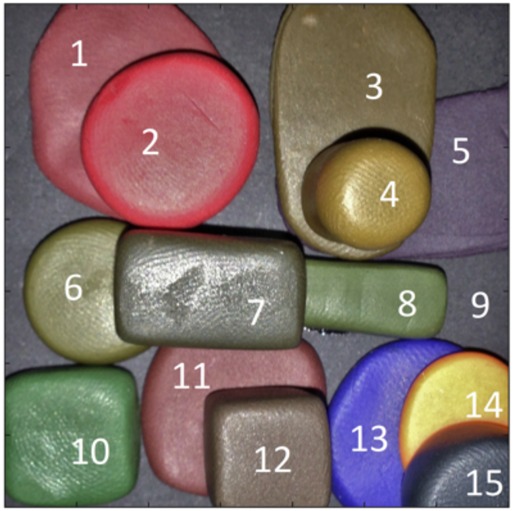

RGB image (5 x 5 cm)

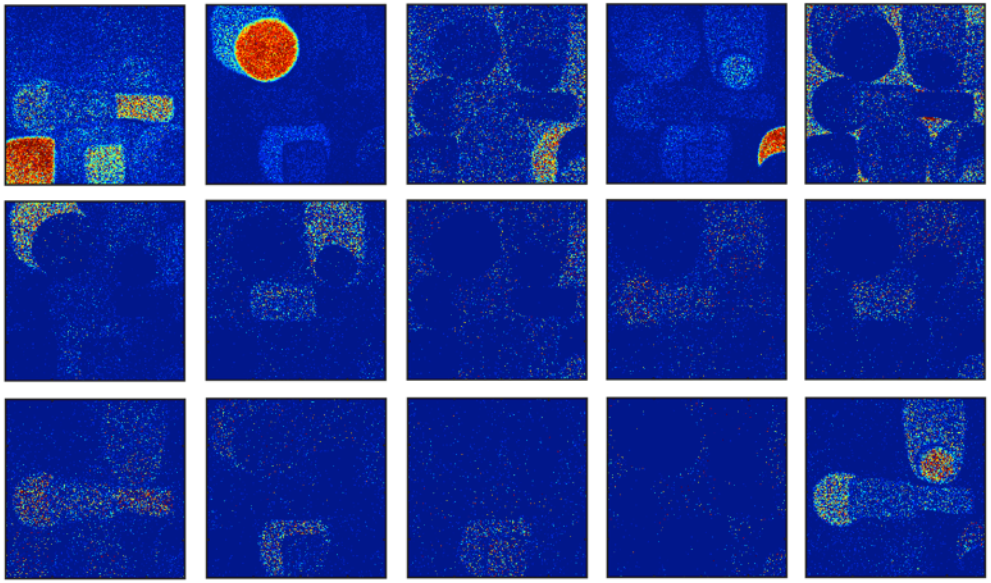

Estimated abundances: 1 x 33 photons per pixel (MLE).

Abundance estimation

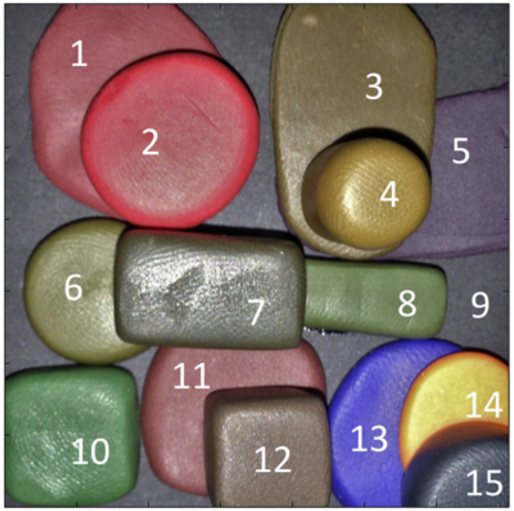

RGB image (5 x 5 cm)

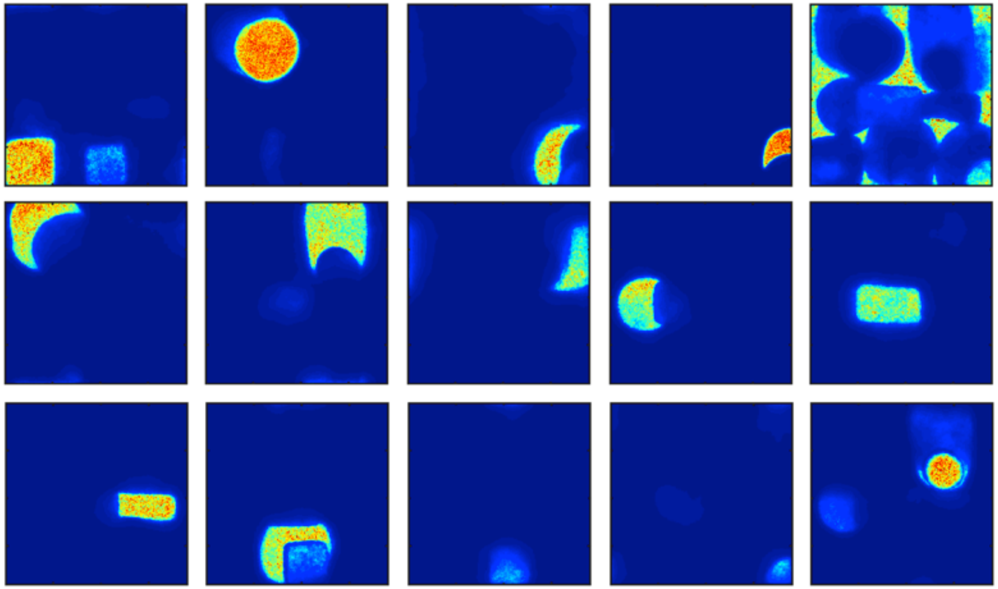

Estimated abundances: 1 x 33 photons per pixel (MMAP).

Outlier analysis

Glue
residue

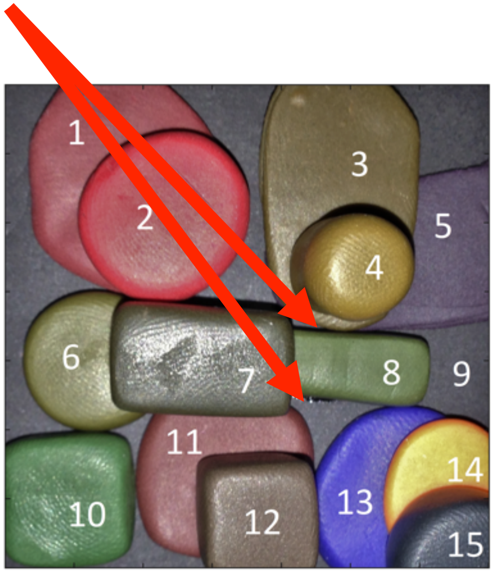

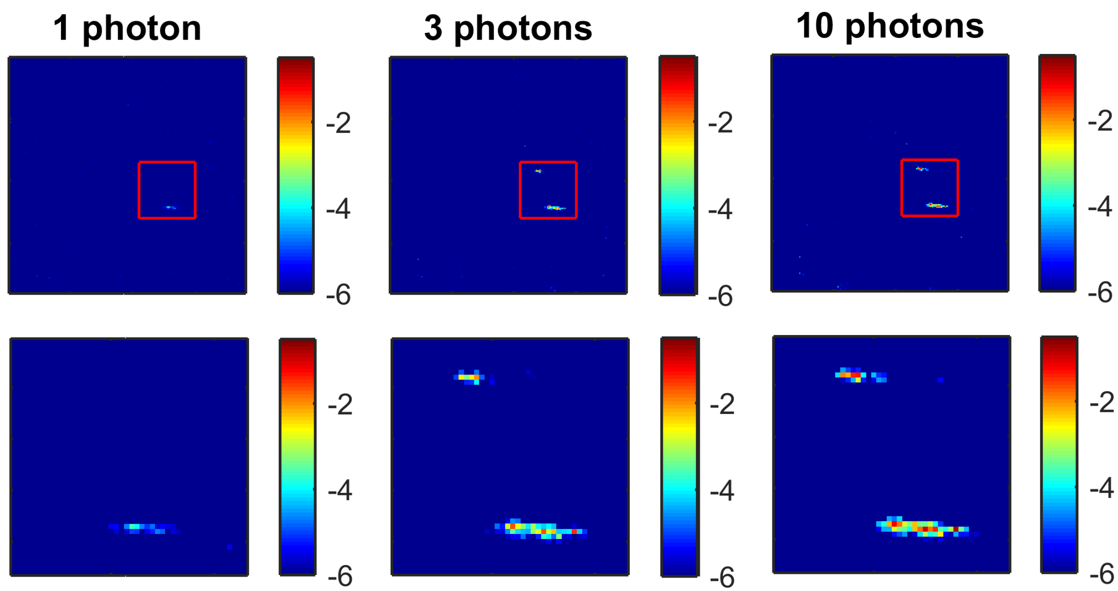

Anomaly maps ($\log(\|\mathbf{s}_n\|^2/L)$).

Range estimation

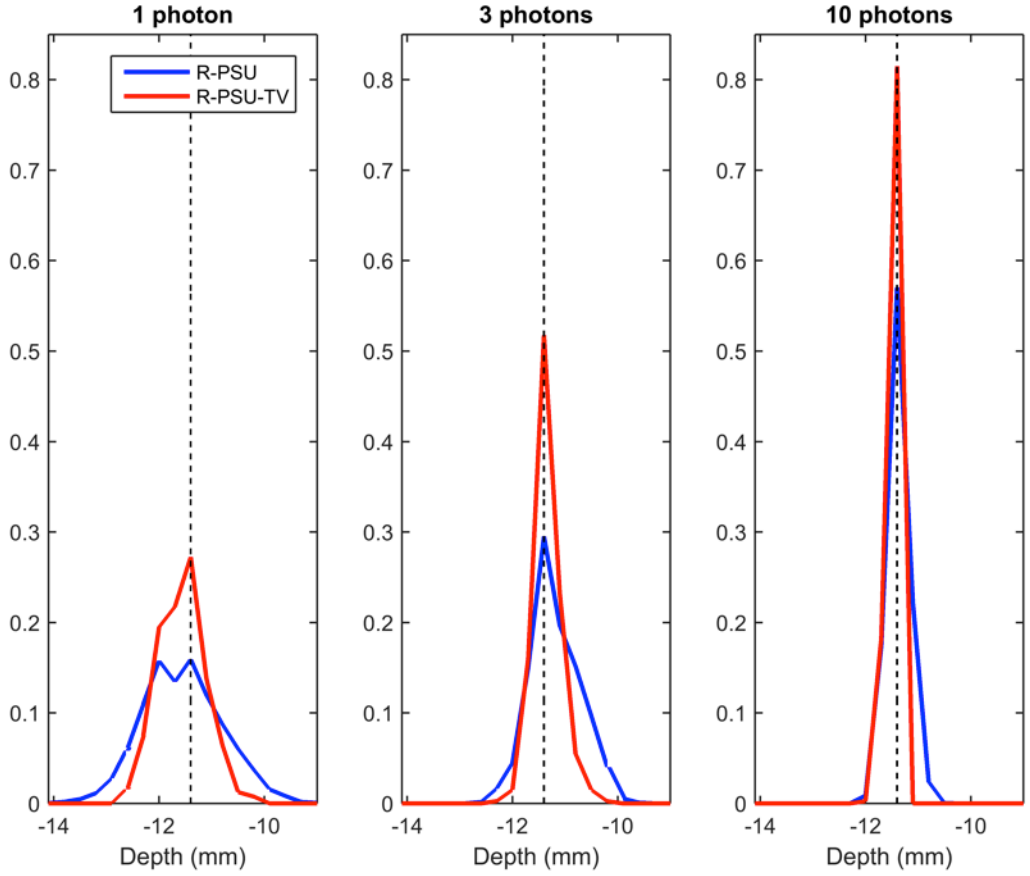

Marginal posterior distributions (range parameters).

Toward faster acquisition

- So far: all the pixels are observed
  - Reduced per-pixel acquisition time
- Different acquisition modes
  - Raster scans (sequential)
  - Detector array (parallel)
- Key question

More pixels, less photons vs less pixels, more photons

Toward faster acquisition

- Raster scan
  – Random/deterministic pixel selection?
- Detector array
  – One wavelength per pixel (mosaic filter)

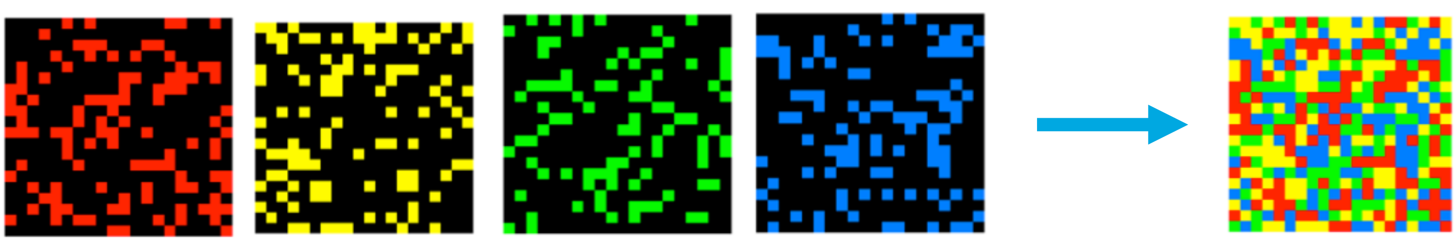

Different scenarios but same methodology for image inpainting/denoising

Proposed approach

- Bayesian modelling
  - Accurate observation model available
  - Regularization required
  - Uncertainty quantification
- Bayesian computation
  - Problem non-convex, multimodal
  - Monte Carlo simulation/hybrid methods

Observation model

$$y_{n,\ell,t} \sim \mathcal{P}\left(r_{n,\ell}\, g_{0,\ell}(t - t_n) + b_{n,\ell}\right)$$

$$t \in \{1,\ldots,T\},\ \ell \in \{1,\ldots,L\}$$

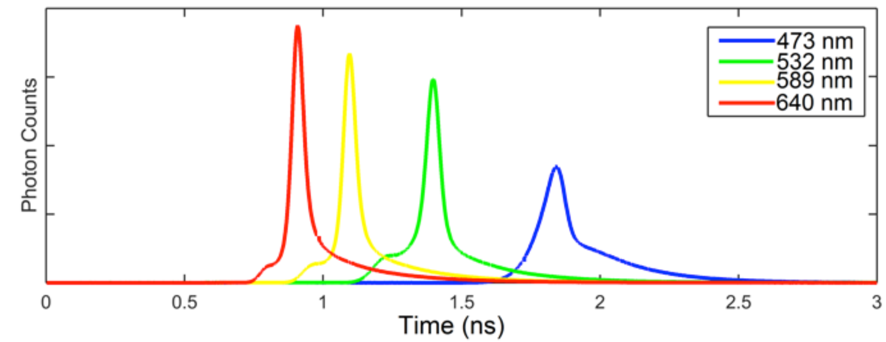

Examples of instrumental responses

- $y_{n,\ell,t}$: photon count in tth bin (ℓ th band)
- $g_{0,\ell}(\cdot)$: instrumental response
- T : Histogram length
- $b_{n,\ell}$: background level
- $r_{n,\ell}$: target reflectivity
- t_n : ToF

Estimation of t_n, $\{b_{n,\ell}\}$ and $\{r_{n,\ell}\}$ for each pixel

Prior models

- Markov random fields: spatial correlation
- Spectral correlation ignored
- Depth profile: Total variation (TV)-based

$$f(\mathbf{t}|\epsilon_t) \sim G(\epsilon_t)^{-1} \exp(\epsilon_t \phi(\mathbf{t}))$$

$$\phi(\mathbf{t}) = \sum_{n=1}^{N} \sum_{n' \in V(n)} |t_n - t_{n'}|$$

Depth parameters: discrete parameters

Prior models

- Background levels

$$f(\mathbf{b}_{\ell} \mid \epsilon_{b,\ell}) \sim G(\epsilon_{b,\ell})^{-1} \exp(\epsilon_{b,\ell}\ \phi(\mathbf{b}_{\ell})),\ \forall \ell$$

$$\phi(b_{\ell}) = \sum_{n=1}^{N} \sum_{n' \in V(n)} |b_{n,\ell} - b_{n',\ell}|$$

- Discrete background
  - Simple estimation of $\epsilon_{b,\ell}$
  - Relatively low cost

Prior models

- Target reflectivities

$$f(\boldsymbol{r}_\ell \mid \epsilon_{r,\ell}) \sim G(\epsilon_{r,\ell})^{-1} \exp(\epsilon_{r,\ell}\, \phi(\boldsymbol{r}_\ell)), \forall \ell$$

- More complex $\phi(\cdot)$ possible (discrete variables)

- Hyper-parameters estimated via maximum marginal likelihood estimation[1] (intractable normalisation constant)

[1]Pereyra et al., IEEE SSP 2014.

Estimation strategy

- Multimodal joint-posterior $f(\boldsymbol{B},\boldsymbol{R},\boldsymbol{t}|\boldsymbol{Y},\boldsymbol{E})$
  - $f(\boldsymbol{t}|\boldsymbol{Y},\boldsymbol{B},\boldsymbol{R},\boldsymbol{E})$ *multimodal*
  - $f(\boldsymbol{B},\boldsymbol{R}|\boldsymbol{Y},\boldsymbol{t},\boldsymbol{E})$ log-concave ($\boldsymbol{B},\boldsymbol{R}$ continuous)

- Sequential estimation of $\{\boldsymbol{b}_\ell\}_\ell,\{\boldsymbol{r}_\ell\}_\ell$ and $\boldsymbol{t}$
  - High-dimensional data
  - Fast/efficient algorithm required
    - Gibbs sampling here
  - Good performance in practice
    - Similar to MCMCs exploiting $f(\boldsymbol{B},\boldsymbol{R},\boldsymbol{t}|\boldsymbol{Y},\boldsymbol{E})$

1) Background estimation

- From calibration or subsets of the original data

$$y_{n,\ell,t} \sim \mathcal{P}(b_{n,\ell}), t \in \{1,\dots,T_b\}$$
$$s_{n,\ell} = \sum_{t=1}^{T_b} y_{n,\ell,t} \sim \mathcal{P}(T_b b_{n,\ell})$$

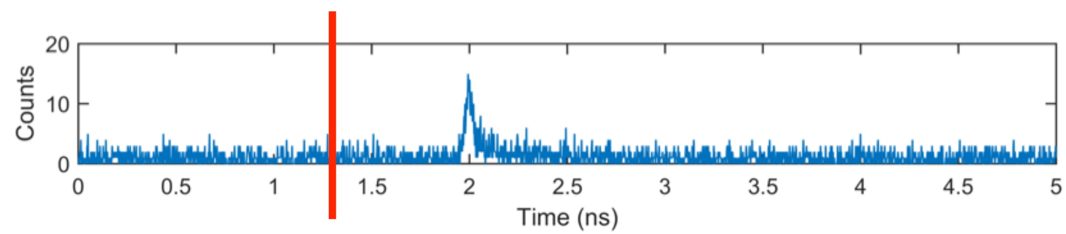

- Estimation of $\boldsymbol{b}_\ell, \forall \ell$ by exploitation of $f(\boldsymbol{b}_\ell | \boldsymbol{s}_\ell, \epsilon_{b,\ell})$

via Gibbs sampling ($\boldsymbol{s}_\ell = \{s_{n,\ell}\}_{n \in \mathrm{I}_{obs,\ell}}$)

- Inpainting of grey-scale images corrupted by Poisson noise

2) Reflectivity estimation

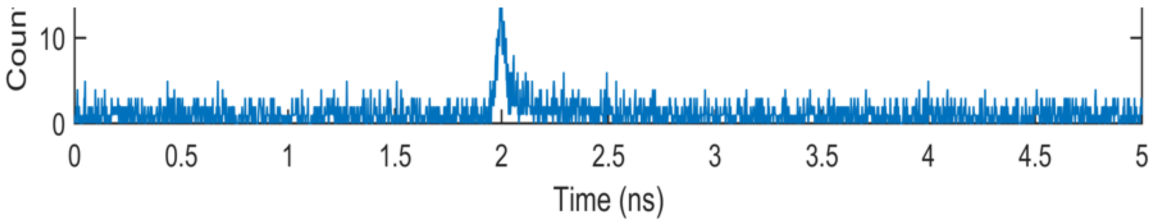

$y_{n,\ell,t} \sim \mathcal{P}(r_{n,\ell}\, g_{0,\ell}(t-t_n)+b_{n,\ell}), t\in\{1,\dots,T\}$

$$y_{n,\ell} = \sum_{t=1}^{T} y_{n,\ell,t} \sim \mathcal{P}(r_{n,\ell}\, G_{0,\ell} + T b_{n,\ell})$$

where $G_{0,\ell} = \sum_{t=1}^{T} g_{0,\ell}(t-t_n)$ (indep. of t_n)

- Estimation of $\mathbf{r}_\ell, \forall \ell$ by exploitation of $f(\mathbf{r}_\ell | \mathbf{y}_\ell, \mathbf{b}_\ell, \epsilon_{r,\ell})$ via Gibbs sampling ($\mathbf{y}_\ell = \{y_{n,\ell}\}_{n\in I_{obs,\ell}}$) (as for $\mathbf{b}_\ell$)
- Inpainting of grey-scale images corrupted by Poisson noise + known baseline

3) Depth estimation

$$y_{n,\ell,t} \sim \mathcal{P}\left(r_{n,\ell}\, g_{0,\ell}(t - t_n) + b_{n,\ell}\right), t \in \{1, \ldots, T_b\}$$

- Estimation of $\boldsymbol{t}$ by exploitation of $f(\boldsymbol{t} | \boldsymbol{Y}, \boldsymbol{B}, \boldsymbol{R}, \boldsymbol{E})$

via Gibbs sampling

- Each step: highly parallelizable

- Steps 1) and 2) can be performed faster (if hyper-parameters fixed): convex optimization

Results

- Two targets, 200x200 pixels (5x5 cm)
- Distance 1.8 m
- Four wavelengths

1) Scenario #1: full scans
2) Scenario #2: regular subsampling (25%), without overlap
3) Scenario #3: random subsampling (25%), without overlap
4) Scenario #4: random subsampling (25%), with overlap.

Reconstruction performance

$$DAE_n = |t_n - \hat{t}_n|, \quad RAE_n = \frac{1}{L}\sum_{\ell=1}^{L} |r_{n,\ell} - \hat{r}_{n,\ell}|$$

Reflectivity estimation

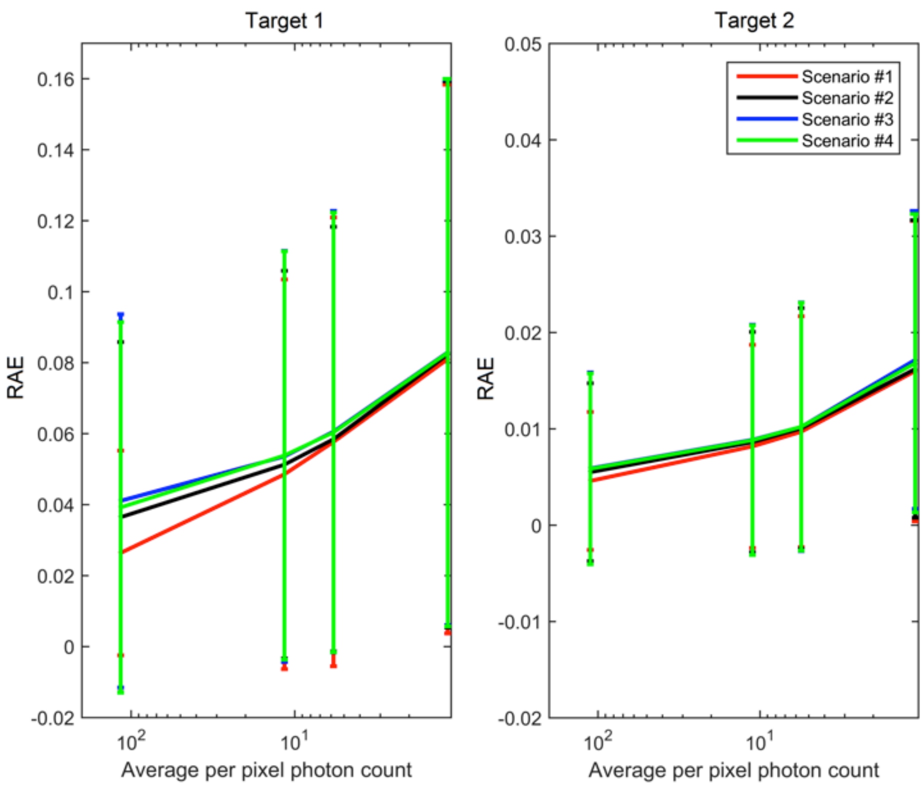

Without background illumination

Reflectivity estimation

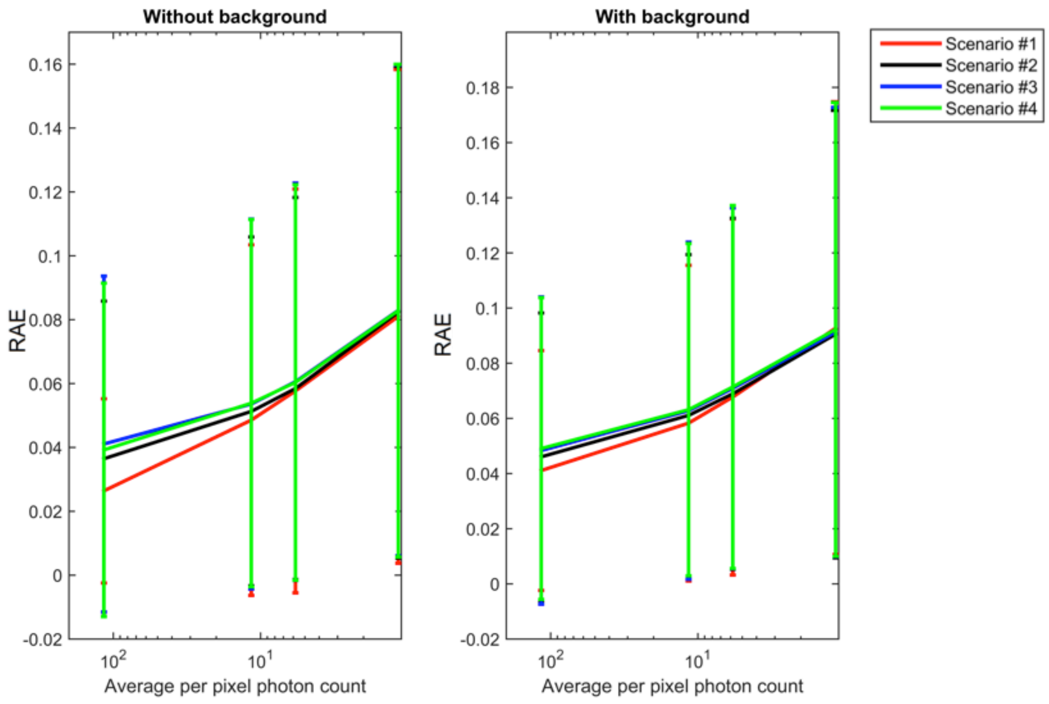

With/without background illumination (Target 1)

Depth estimation

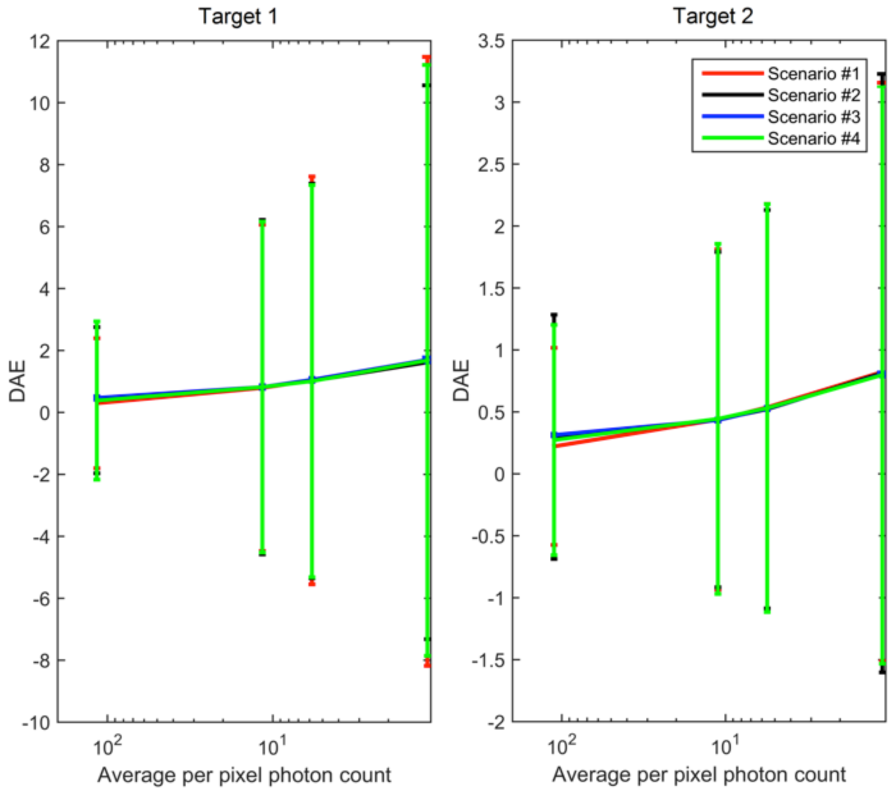

Partial conclusion

- Different sub-sampling strategies
  - Similar performance in low light flux
  - High photon counts: performance limited by the interpolation scheme
  - Possible solutions:
    - Joint prior models $f(\boldsymbol{B},\boldsymbol{R},\boldsymbol{t})$
    - Patch-based priors

- Data acquisition
  - Subsampling required?

Simultaneous acquisition

- Classical model

$$y_{n,\ell,t} \sim \mathcal{P}\left(r_{n,\ell}\, g_{0,\ell}(t - t_n) + b_{n,\ell}\right)$$

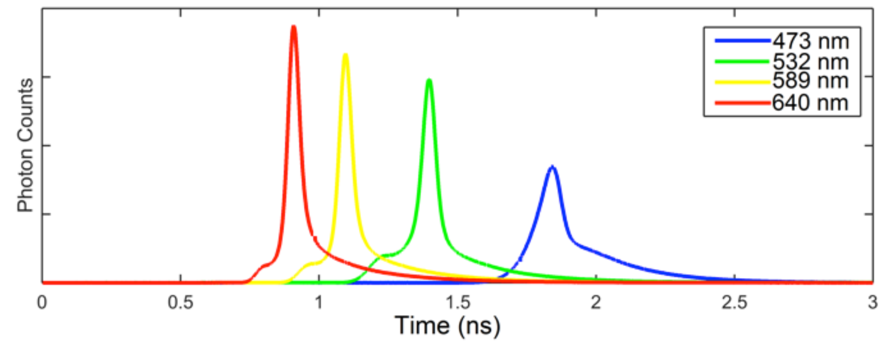

- New model

$$y_{n,t} \sim \mathcal{P}\left(\sum_{\ell=1}^{L} r_{n,\ell}\, g_{0,\ell}(t - t_n) + b_{n,\ell}\right)$$

  – “Lossless” acquisition but more challenging inverse problem

Parameter estimation

- Highly multimodal posterior
- Sequential estimation no longer possible
- Proposed solution: adaptive MH-WG sampler
  - Metropolis-Hastings updates
  - Multiple proposals (blocking schemes)
  - Discrete parameters (enhanced mixing properties)
  - Automated estimation of regularization parameters

Results

- 200x200 pixels (5x5 cm)
- Distance 1.8 m
- Four wavelengths

Scenario #1: 4 x 1 wavelength (X seconds)
Scenario #2: 2 x 2 wavelengths (X/2 seconds)
Scenario #3: 1 x 4 wavelengths (X/4 seconds)

Reflectivity estimation

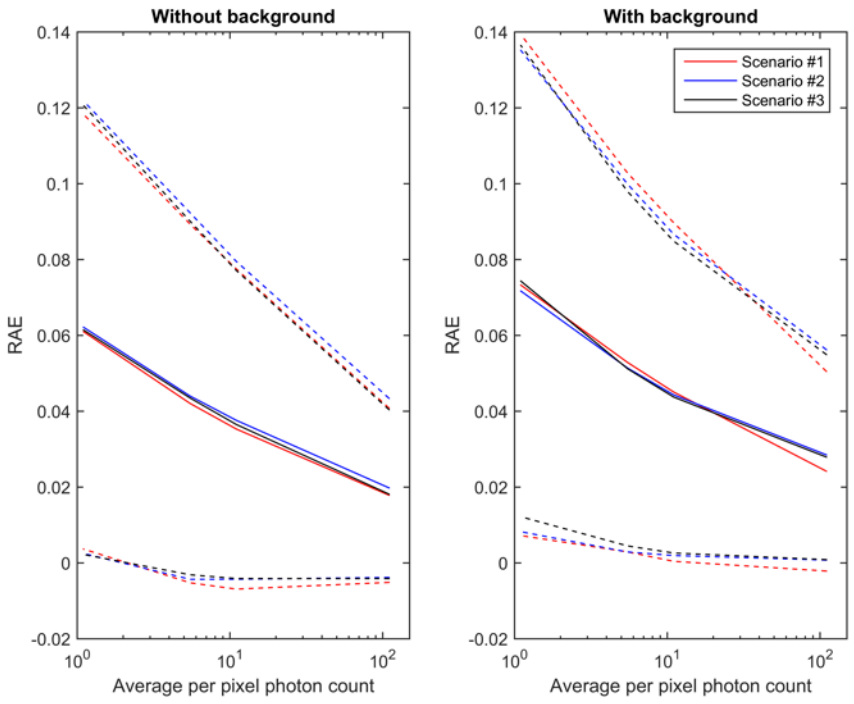

With/without background illumination

Range estimation

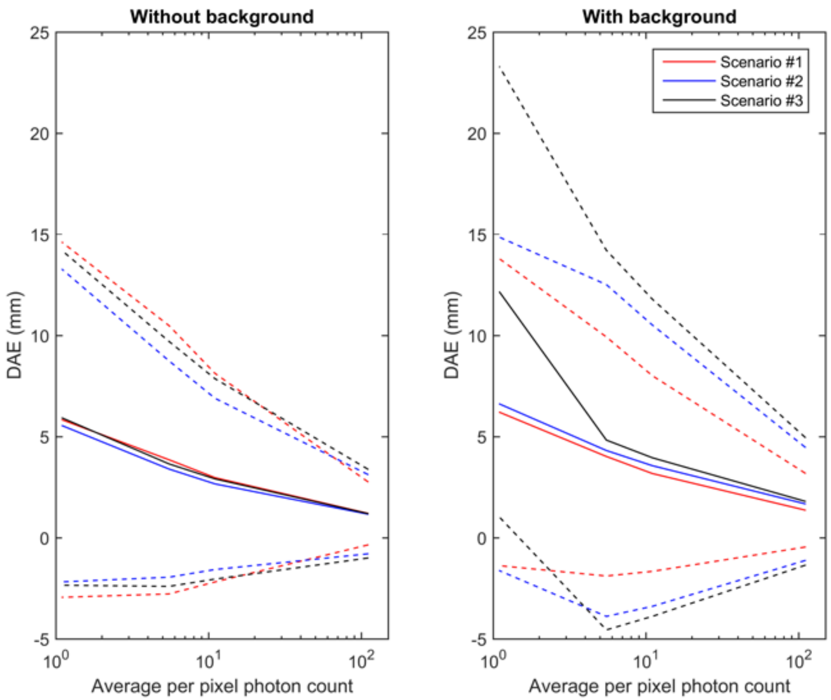

With/without background illumination

HERIOT WATT UNIVERSITY

Conclusion

- **Conclusions**
  - Several acquisition modes for
    - Scanning systems
    - Array-based systems
  - Similar estimation performance
    - Balance between fast acquisition and computational complexity

- **Future/ongoing work**
  - Hybrid methods for faster analysis
  - Generalization to actual 3D unmixing
    - multiple surfaces
  - Adaptive sampling strategies

Thanks for your attention !

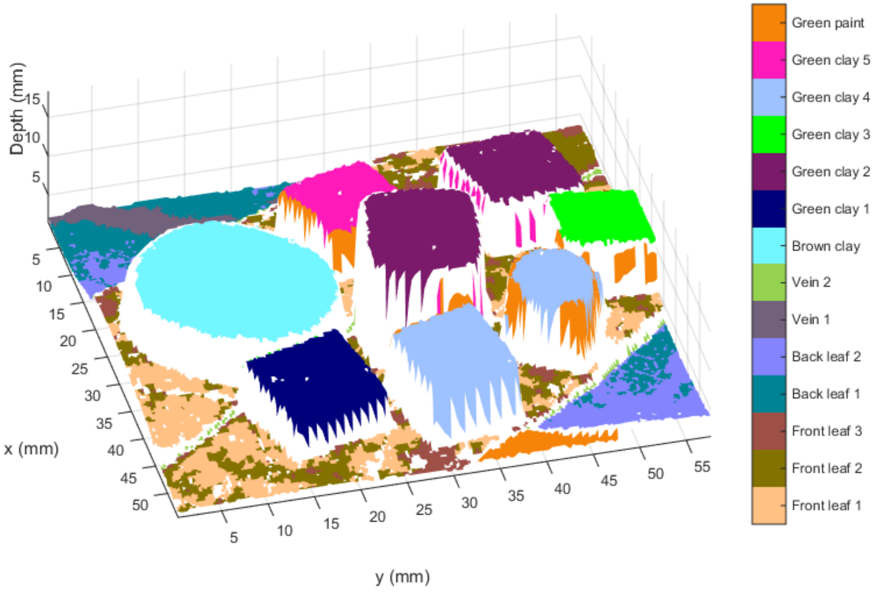

Y.Altmann@hw.ac.uk

HERIOT WATT UNIVERSITY

Spectral clustering/classification

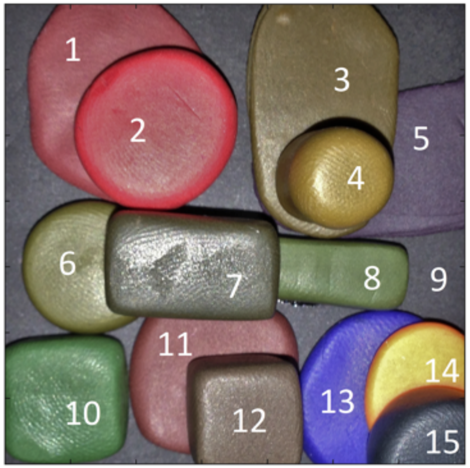

RGB image (5 x 5 cm)

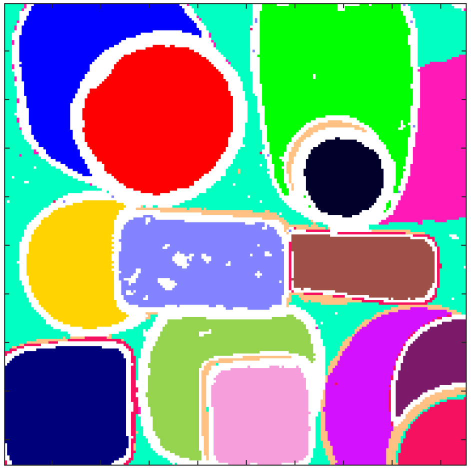

Unsupervised spectral classification

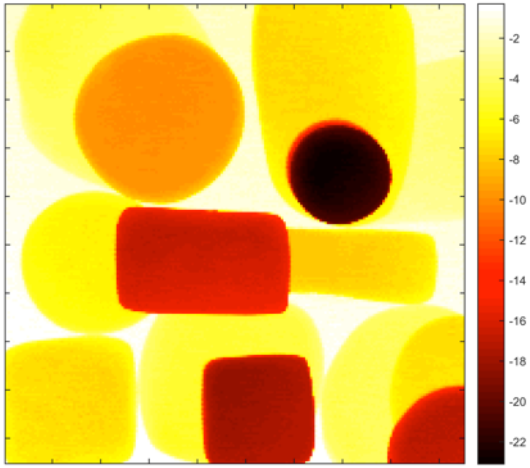

Estimated range profile (in mm)

- 33 spectral bands, 500– 820 nm, 200x200 pixels