HERIOT WATT UNIVERSITY

ROYAL ACADEMY OF ENGINEERING

LEADERS IN IDEAS AND SOLUTIONS

Online 3D imaging in high-background scenarios

Yoann Altmann

School of Engineering and Physical Sciences, Heriot-Watt University
Joint work with: Julian Tachella, Quentin Legros, Aurora Maccarone, Rachael Tobin, Aongus McCarthy, Gerald Buller, Stephen McLaughlin, Mike Davies

UDRC Themed Meeting: Imaging through obscure media
July, 22nd 2020

EPSRC Engineering and Physical Sciences Research Council

[dstl]

UDRC

Single-photon Lidar

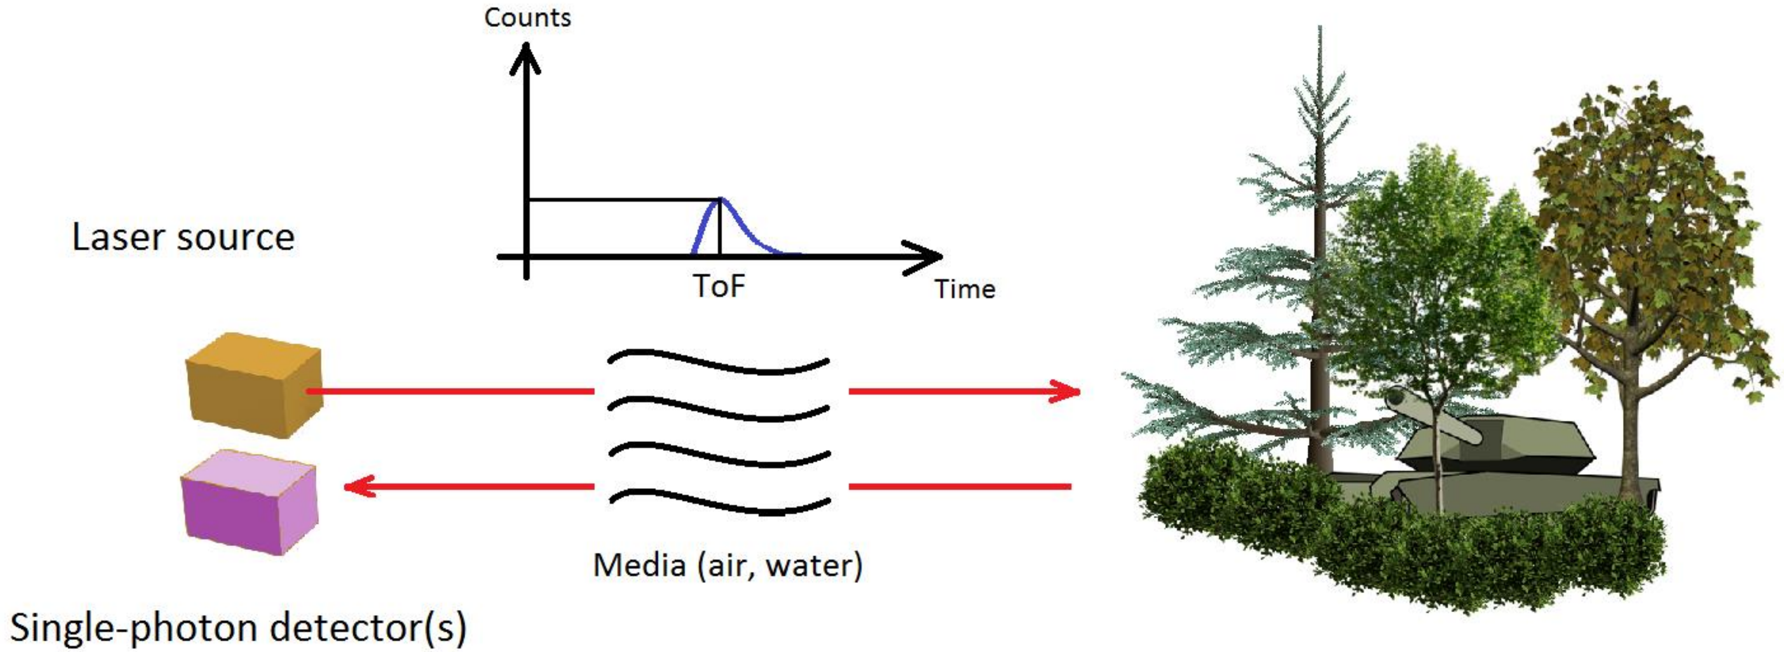

- Pulsed laser (20 MHz), low power (≈ μW)
- Detector(s): single-photon avalanche diode (SPAD)

High sensitivity and high temporal/ranging resolution

Challenges (I)

1. Few detected photons $s_t = r_0 g_0(t - t_0) \ll 1$

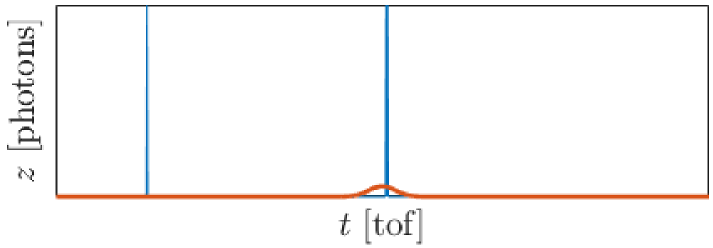

2. High background $b \gg s_t$

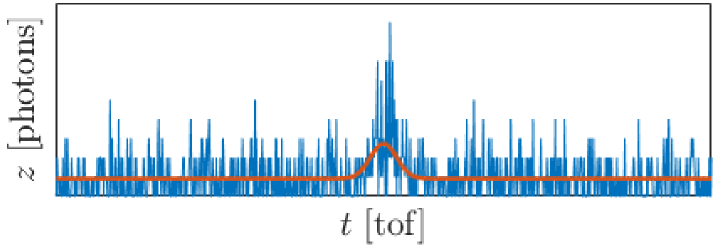

3. No target $s_t = 0$

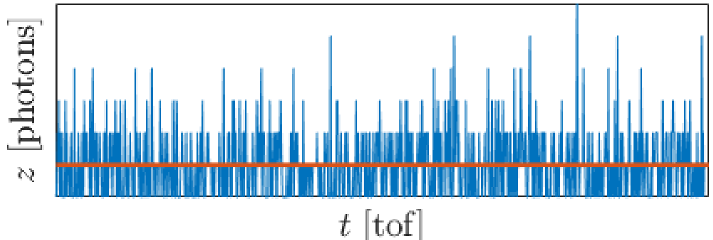

Challenges (II)

- Not only difficult inference problems (noise, convexity,…)
- Fast acquisition
- Large data volume/array size

- Optimization $\not\Rightarrow$ fast
  - Dimensionality of the data/unknowns
  - Convergence speed

- Need to redesign the inference process
  - Scalability
  - Robustness

Algorithmic scalability

- Graph-based representation

Physics-inspired
Sensor-dependent

Data
$f(\boldsymbol{y}|\boldsymbol{x})$

Prior information
$f(\boldsymbol{x})$

Sparsity, smoothness, bespoke denoisers, classifiers,…

$$f(\boldsymbol{x}|\boldsymbol{y}) \propto f(\boldsymbol{y}|\boldsymbol{x})f(\boldsymbol{x})$$

– Message passing structure
– Plug-and-play (PnP): computer graphics

RT3D algorithm

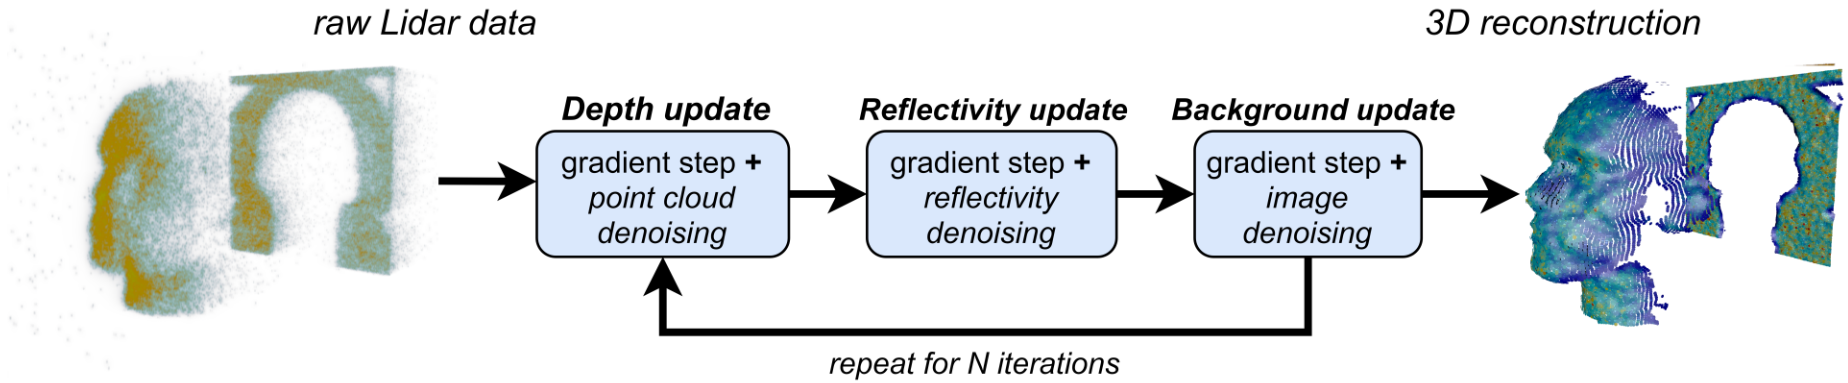

Significant gain using:

- Fast denoisers (e.g. parallel architectures)
- Application/variable-specific denoisers

J. Tachella et al. “Real-time 3D reconstruction from single-photon lidar data using plug-and-play point cloud denoisers”, Nature Comms., Nov 2019.

RT3D algorithm

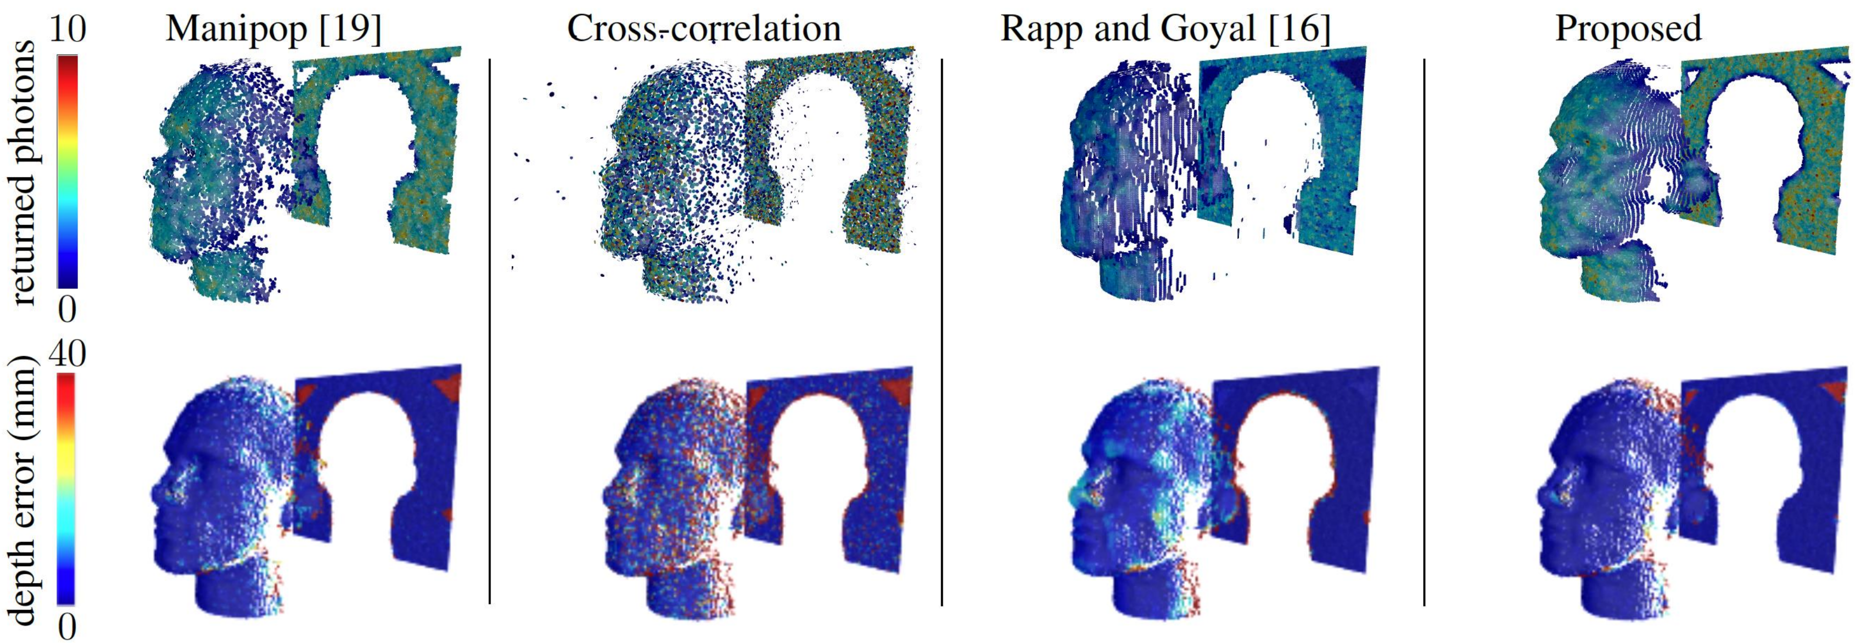

201 s | 1 ms | 201 s | 13 ms

J. Tachella et al. “Real-time 3D reconstruction from single-photon lidar data using plug-and-play point cloud denoisers”, Nature Comms., Nov 2019.

RT3D: distributed objects

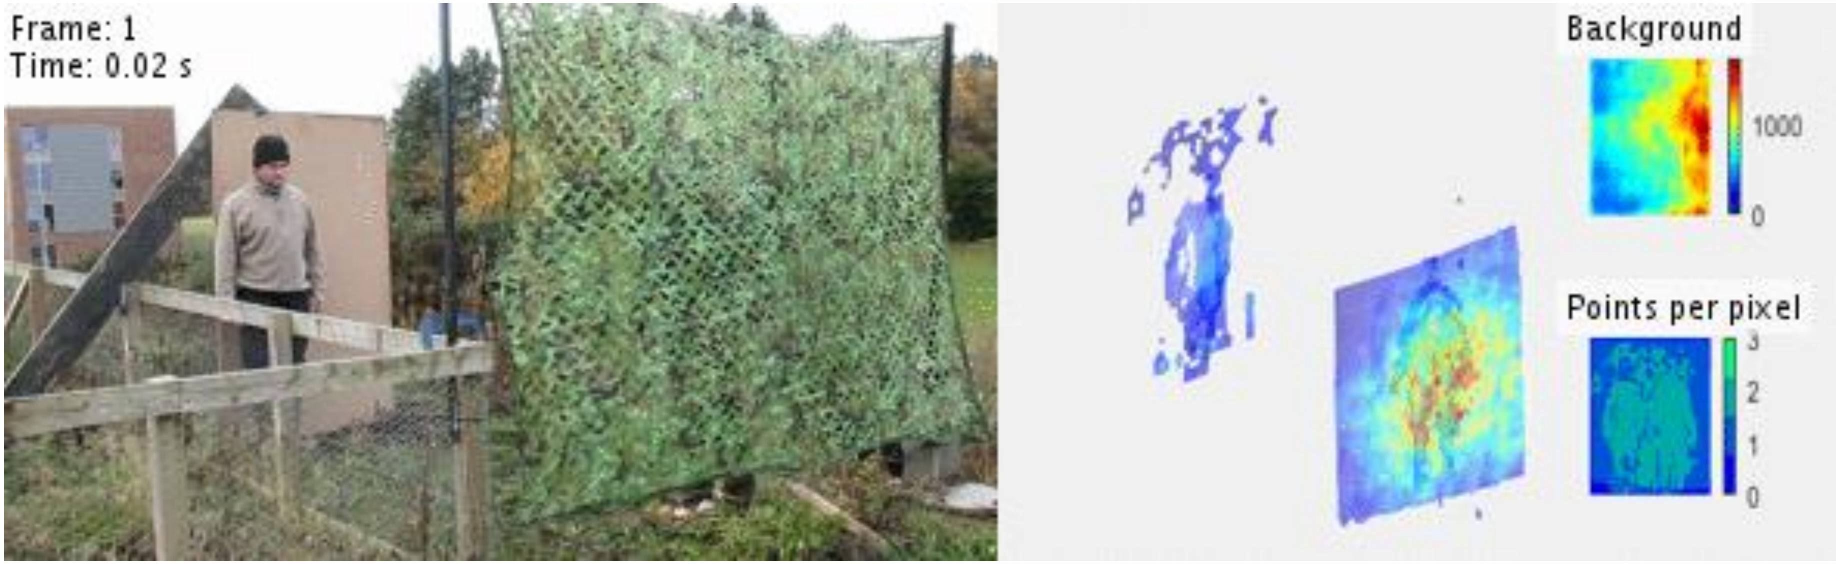

Real-time 3D reconstruction (50 fps) from 32x32 pixels (offline).
Princeton Lightwave camera, 1550 nm
Distance: 300 m

RT3D: real-time implementation

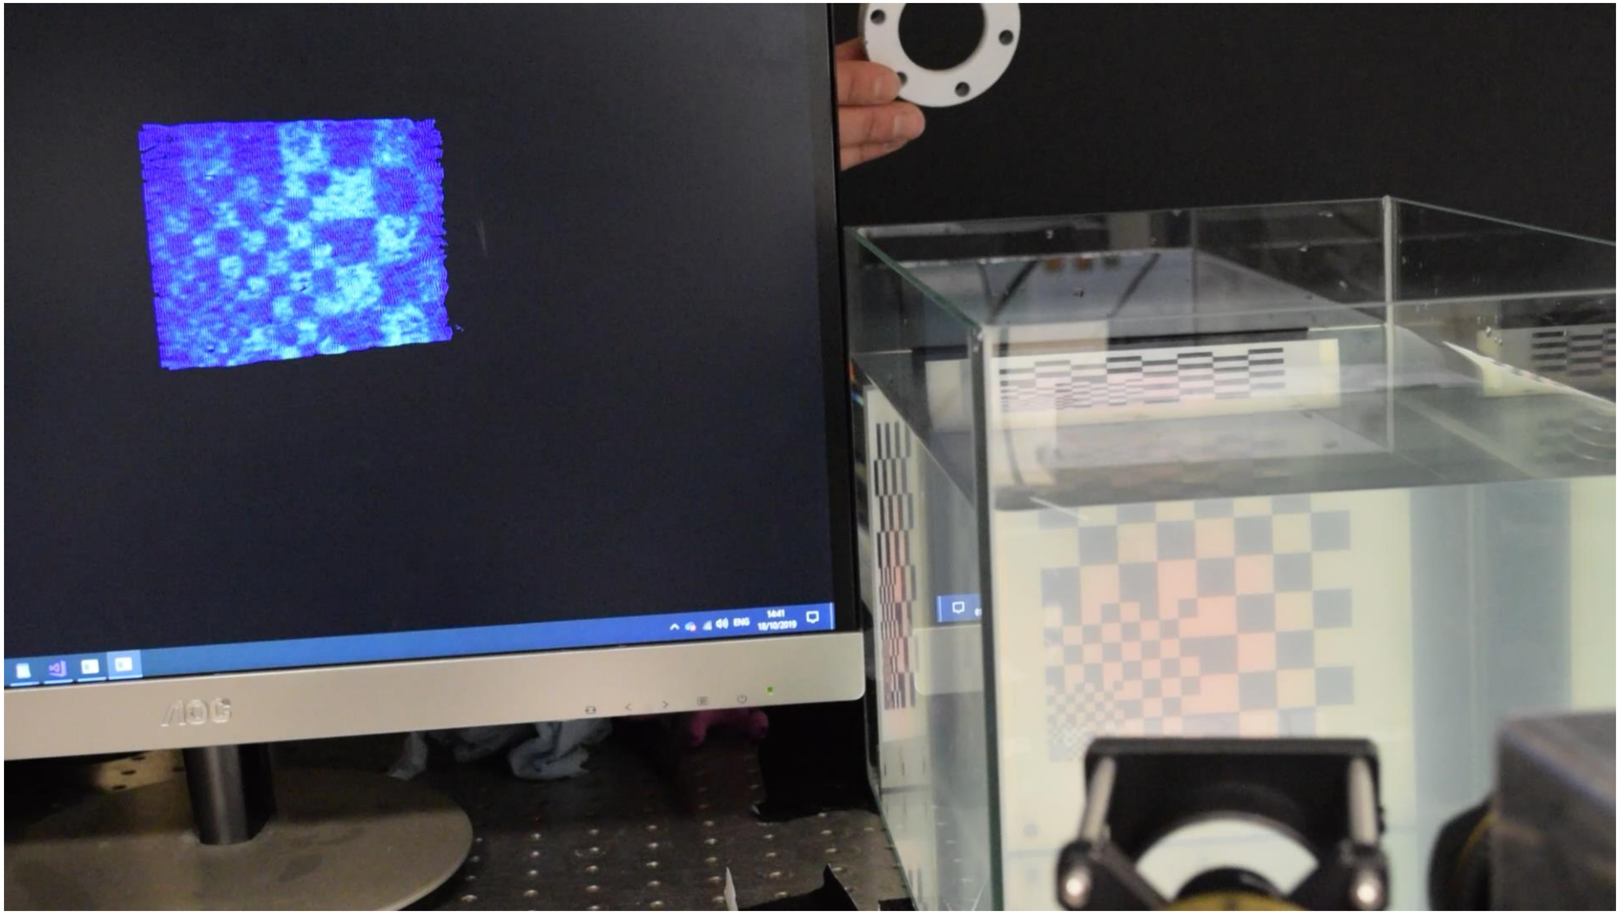

Real-time 3D reconstruction (20 fps) from 192x128 pixels (online).
QuantiCAM, imaging through murky water
Distance: 1 m

Scalable inference

- Algorithmic structure enabling video frame rates for complex scenes
  - New opportunities
    - Adaptive processing and sensing
    - Fusion
  - Remaining challenges
    - Scalability (still)
    - Robustness/reliability

Here: regularization via temporal information + robust statistics

Bayesian online inference

- Bayesian filtering

Estimation (t)	Prediction (t+1)	Detection (t+1)	Estimation (t+1)
Point cloud + UQ	Point cloud + UQ	Presence map	Point cloud + UQ

- Spatio-temporal model for prediction
- Detection for dimensionality reduction
- ~~Particle filtering~~ → variational approximation (ADF)
- Robustness?

J. Tachella et al., "Fast surface detection in single-photon Lidar waveforms", Proc. European Signal Processing Conf. (EUSIPCO), A coruna, Spain, Sept. 2019.
Y. Altmann et al., "Fast online 3D reconstruction of dynamic scenes from individual single-photon detection events", IEEE Trans. Image processing, vol. 29, 2019.

Pseudo-variational inference

Physics-inspired,
sensor-dependent
or surrogate models

Data
$f(\boldsymbol{y}|\boldsymbol{x})$

Prior information
$f(\boldsymbol{x})$

Bespoke denoisers,
classifiers,…

$$f(\boldsymbol{x}|\boldsymbol{y}) \propto f(\boldsymbol{y}|\boldsymbol{x}) f(\boldsymbol{x})$$

- Fast/efficient denoiser
- PnP approach possible in several blocks
  - PnP prior
  - PnP likelihood
- How to replace $f(\boldsymbol{y}|\boldsymbol{x})$?

Q. Legros et al., " Robust depth imaging in adverse scenarios using single-photon Lidar and beta-divergences", to appear in Proc. SSPD 2020.
Q. Legros et al., " Robust 3D reconstruction of dynamic scenes from single-photon lidar using Beta-divergence ", arxiv pre-print, 2020.

Robust inference

- Although not “optimal”, matched filtering (MF) works well in practice with large background
  - Fast, not iterative
  - Simple: does not require background estimation

- Here MF can be reinterpreted as a robust estimator

Robust inference

- Matched filter: $\{t_n\}_{n=1,\ldots N}$: set of photon ToAs

$$\max_{t_0} \left(\frac{1}{N} \sum_n g_0(t_n - t_0) \right)$$

- MLE (background-free) / LMF:

$$\max_{t_0} \left(\frac{1}{N} \sum_n \log(g_0(t_n - t_0)) \right) \Leftrightarrow \min_{t_0} \left(D_{KL}(\hat{f}||f_{t_0}) \right)$$

minimum KL-divergence estimator

- Robust estimator based on β-divergence

$$\max_{t_0} \frac{1}{N} \sum_n g_0(t_n - t_0)^{\beta} \Leftrightarrow \min_{t_0} \left(D_{\beta}(\hat{f}||f_{t_0}) \right)$$

Pseudo-likelihood → pseudo-posterior distribution

Comparison of estimators

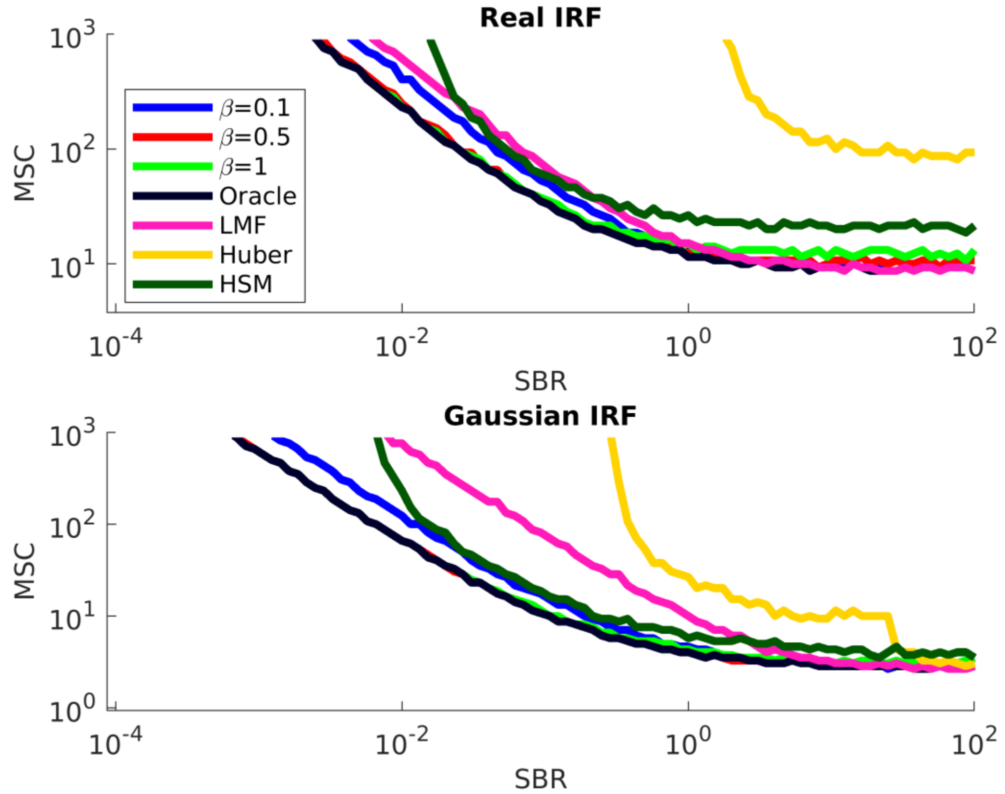

Detection threshold (> 85%) for different robust methods

Robust estimation close to oracle for $\beta \in [0.5,1]$

MSC: mean signal counts

SBR: signal-to-background ratio

Robust online 3D reconstruction

- 3D reconstruction (5000 fps) from 32x32 pixels.

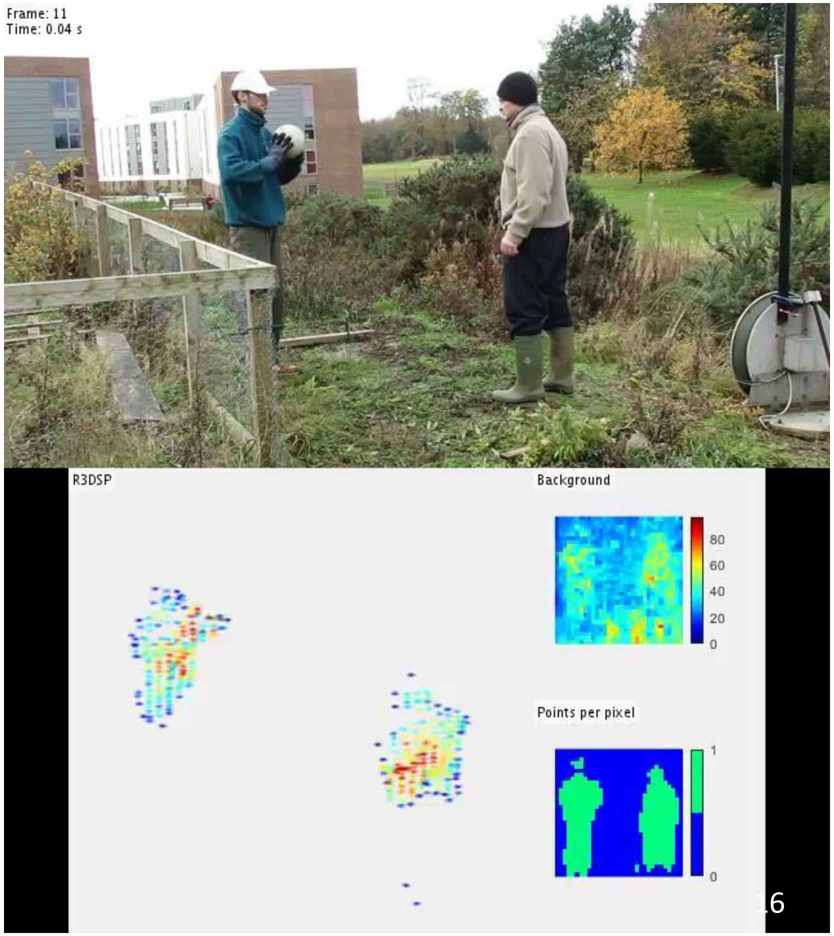

Q. Legros et al. "Robust 3D reconstruction of dynamic scenes from single-photon lidar using Beta-divergences", Arxiv pre-print, 2020

Robust online reconstruction

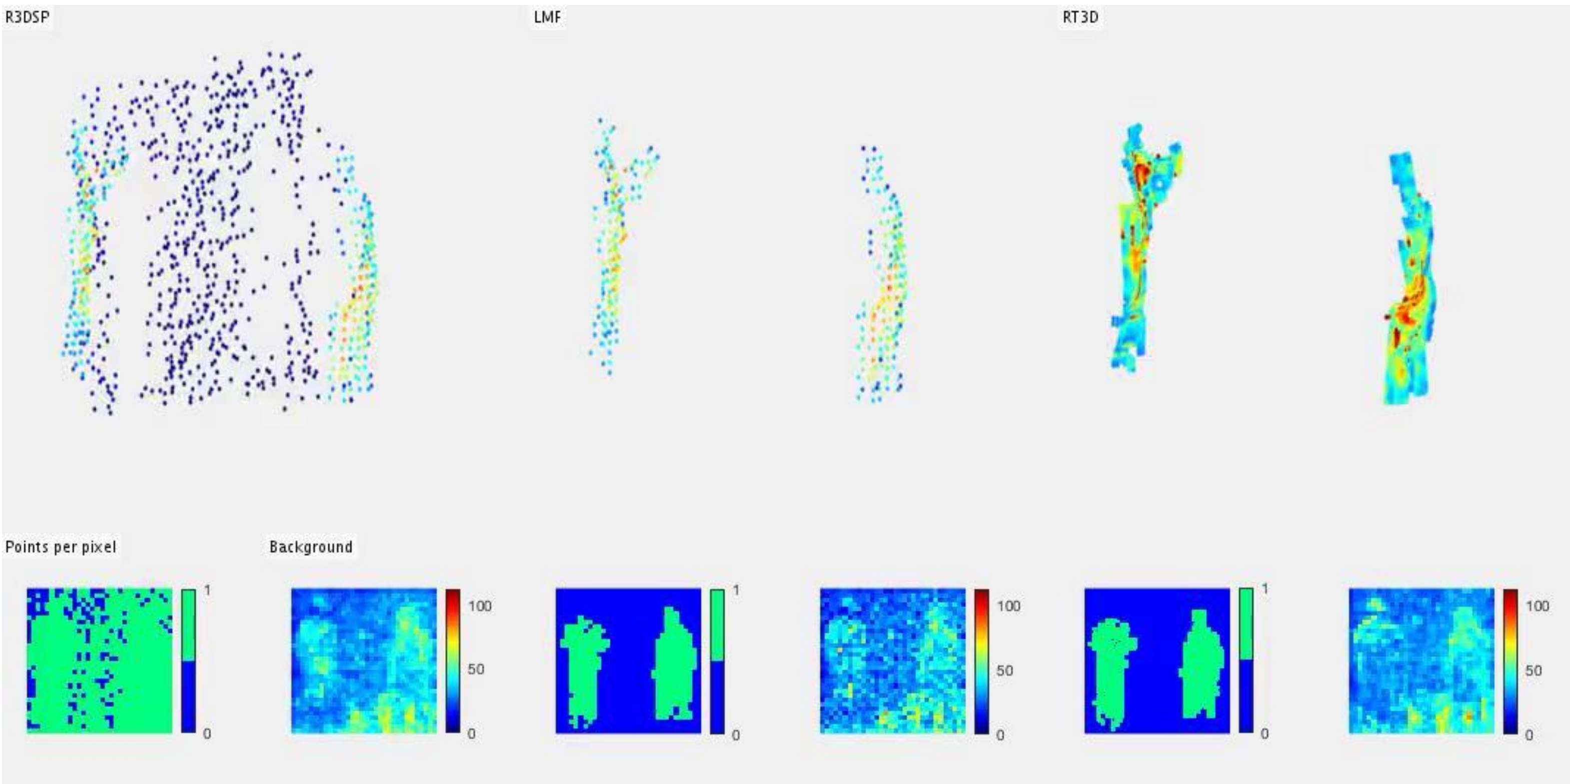

Online reconstruction with moderate solar illumination (MSC=55, SBR=1.6)

Robust online reconstruction

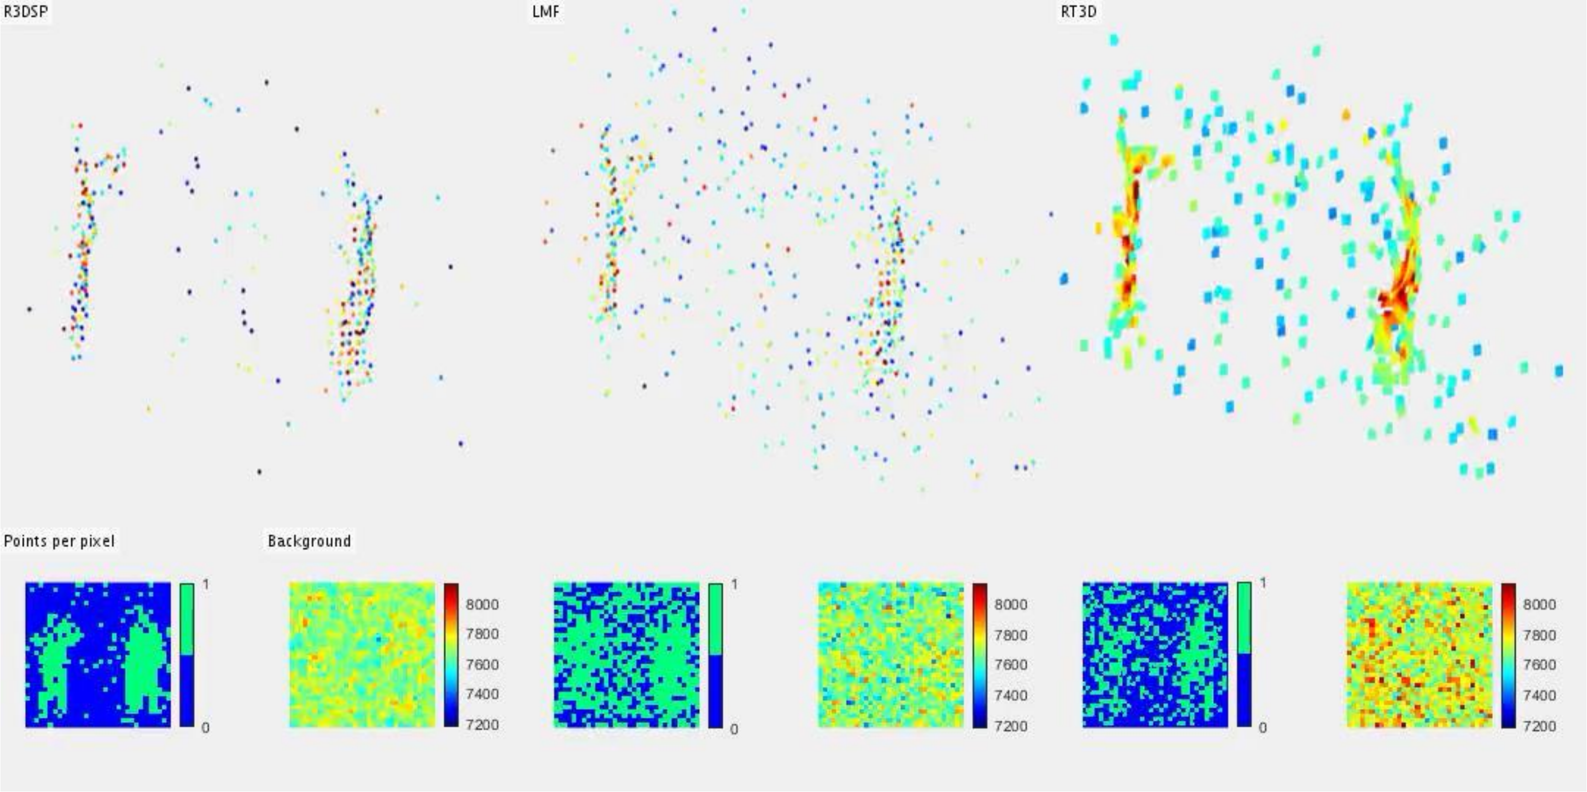

Online reconstruction with high solar illumination (MSC=55, SBR=7.10^{-3})

Robust online reconstruction

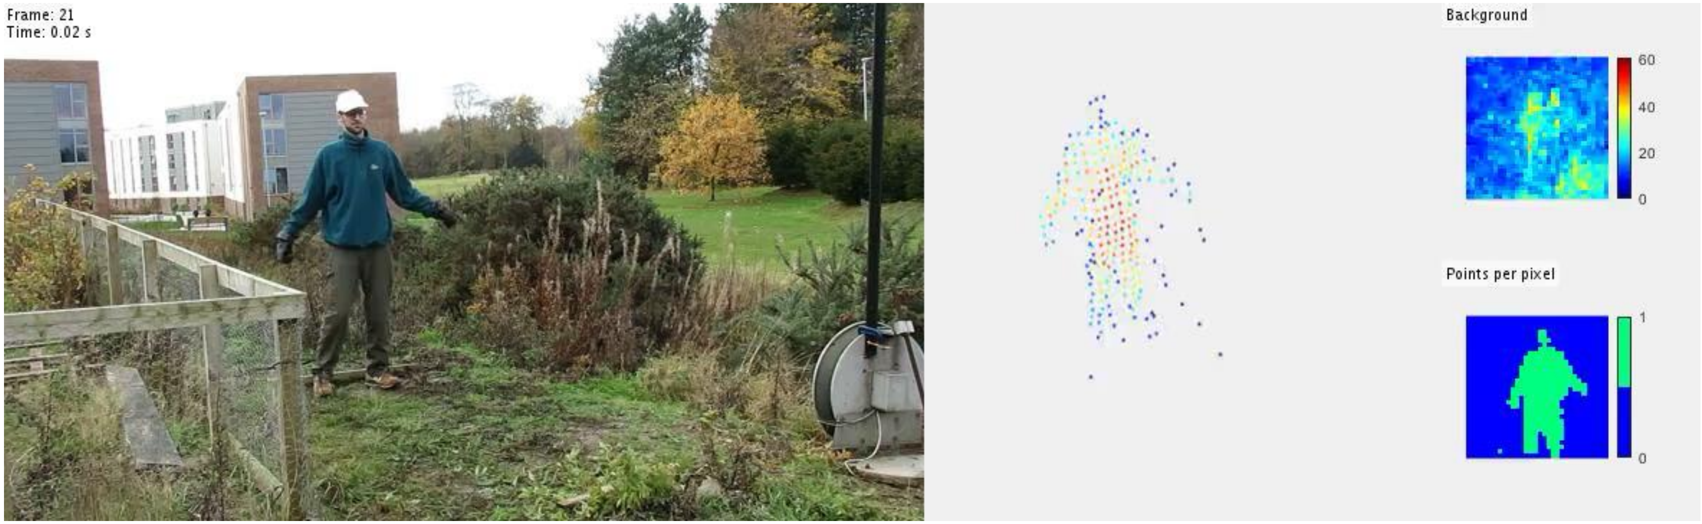

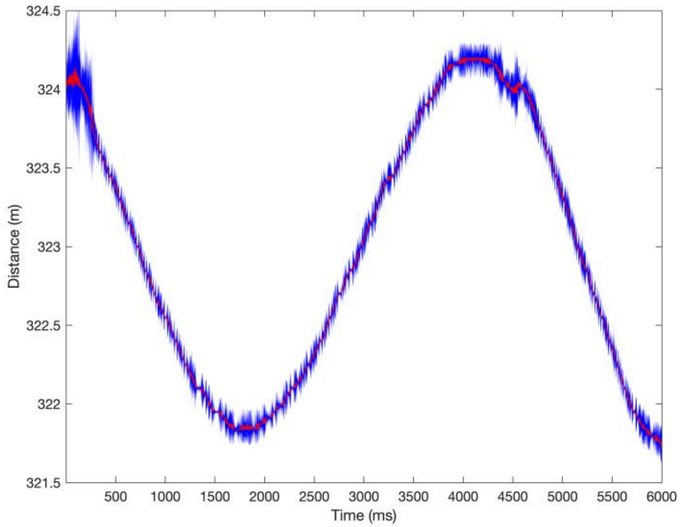

Red curve: depth posterior mean

Blue region: credible interval

Conclusion

Bayesian methods for online and robust 3D reconstruction

- Modular approach using approximate message passing for UQ
- PnP updates
  - Fast/scalable denoisers
  - Fast/robust “likelihood” terms
- Application-specific blocks
  - Peak-broadening, highly-scattering media

Q. Legros et al. “Robust depth imaging in adverse scenarios using single-photon Lidar and beta-divergences", Sensor Signal Processing for Defence (SSPD) Conference 2020, to appear.

Thanks for your attention!

Contact: Y.Altmann@hw.ac.uk