Cloud Detection Using Deep Learning on GPUs for Landsat 8

Alistair Francis[1,2], Panagiotis Sidiropoulos[1,2*,3], Andrea Mirabile[2],
Jan-Peter Muller[1]

[1]Mullard Space Science Laboratory, UCL
[2]Cortexica Vision Systems Ltd.
[3]Hummingbird Technologies Ltd.
*previously at…

Outline

- Context
  - Cloud masking
  - Fully convolutional networks
  - Landsat 8
- Design
- Performance
  - Accuracy
  - Speed
- Future Outlook
  - Large-scale processing of Landsat 8
  - Extension to other datasets
  - On-board processing

Cloud Masking

- Ubiquitous problem in EO community
- On average, two thirds of Earth is cloudy
- Opportunities for S/C operators:
  - On-board data reduction/real-time target acquire
  - Rich auxiliary dataset
- Opportunities for downstream users:
  - Automated candidate selection
  - Removal of non-surface pixels
- Direct applications:
  - Climate trends – cloud levels per year
  - Correlations with other variables – ecological or economic

Cloud Masking

- Since the 1980s, algorithms have been developed for masking clouds (Rossow et al., 1985)
- Most previous algorithms have been based on physical principles :
    - Derive physical quantities from bands (e.g. TOA radiance/derived temperature)
    - Thresholds applied to these physical values
    - Mostly over single pixels (or averaging over small number)
- Machine Learning techniques have recently come to prominence:
    - Single pixel Neural networks / Random Forest / SVM (etc.)
    - Some texture extractors (e.g. Gabor) or convolutional networks for spatial features
    - Shown to have **markedly higher performance** than physically-derived thresholds

State-of-the-art

- Many papers published with bespoke datasets/instruments: comparability is elusive
- CFMask (Confidence Function of Mask) used for Landsat 8. Decision tree style thresholding algorithm incorporating scene-wide statistics
- ACCA (Automated Cloud Coverage Assessment) used for Landsat 7
- Foga et al., 2017 gave comparison of 9 algorithms on same Landsat 8 dataset

Fully Convolutional Networks

- Convolutional layers only.
- Results in arbitrary input size : output map just grows with input
- Fully convolutional encoder-decoder ideal for segmentation tasks

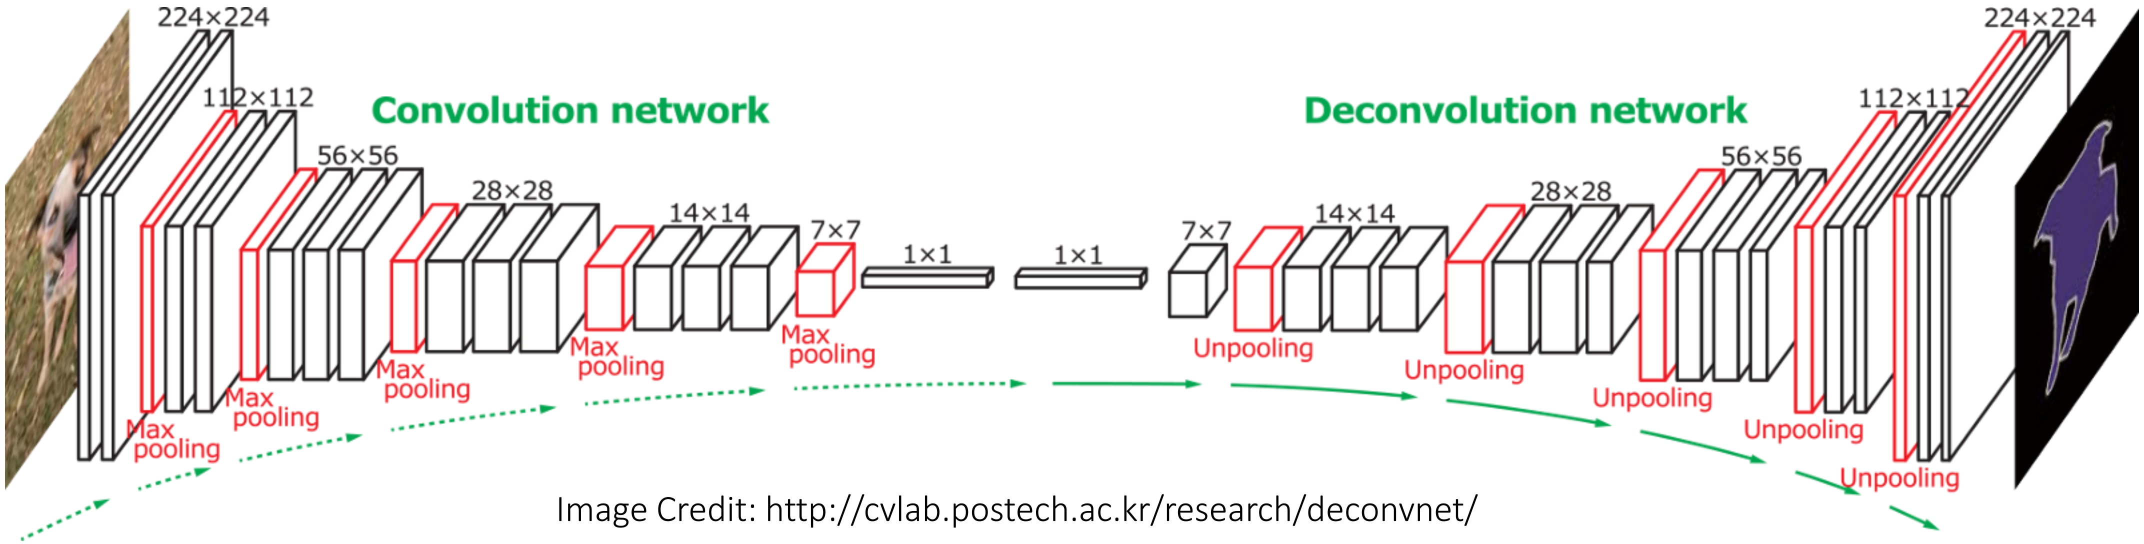

Image Credit: http://cvlab.postech.ac.kr/research/deconvnet/

CloudFCN: Design

- Original R&D as part of OVERPaSS project with Earth-i
  - Cloud detection in RGB satellite: Carbonite-2
- CloudFCN: redesigned for general purpose cloud detection for any EO imaging sensor
- Fully convolutional model architecture
  - Residual connections, inspired by U-net
  - Dimensionality reduction for increased speed, inspired by InceptionNet
  - Fuses pixel-level information with successively larger spatial features
  - Each output informed by individual pixel and areas around it

Layer	Stride	Filter Size	Jump	Receptive Field
Inception	1	5	1	1
Convolution	2	5	1	5
Convolution	3	7	2	9
Inception	1	5	6	21
Convolution	2	5	6	45
Inception	1	5	12	69
Convolution	2	5	12	117
Transpose	2	5	24	165
Convolution	1	1	12	261
Transpose	2	5	12	261
Transpose	3	7	6	309
Inception	1	5	2	345
Transpose	2	5	2	353
Inception	1	5	1	361
Convolution	1	1	1	365
Output	-	-	1	365

"PIXEL"
"TEXTURAL"
"SPATIAL"
"CONTEXTUAL"
"SPATIAL"
"TEXTURAL"
"PIXEL"

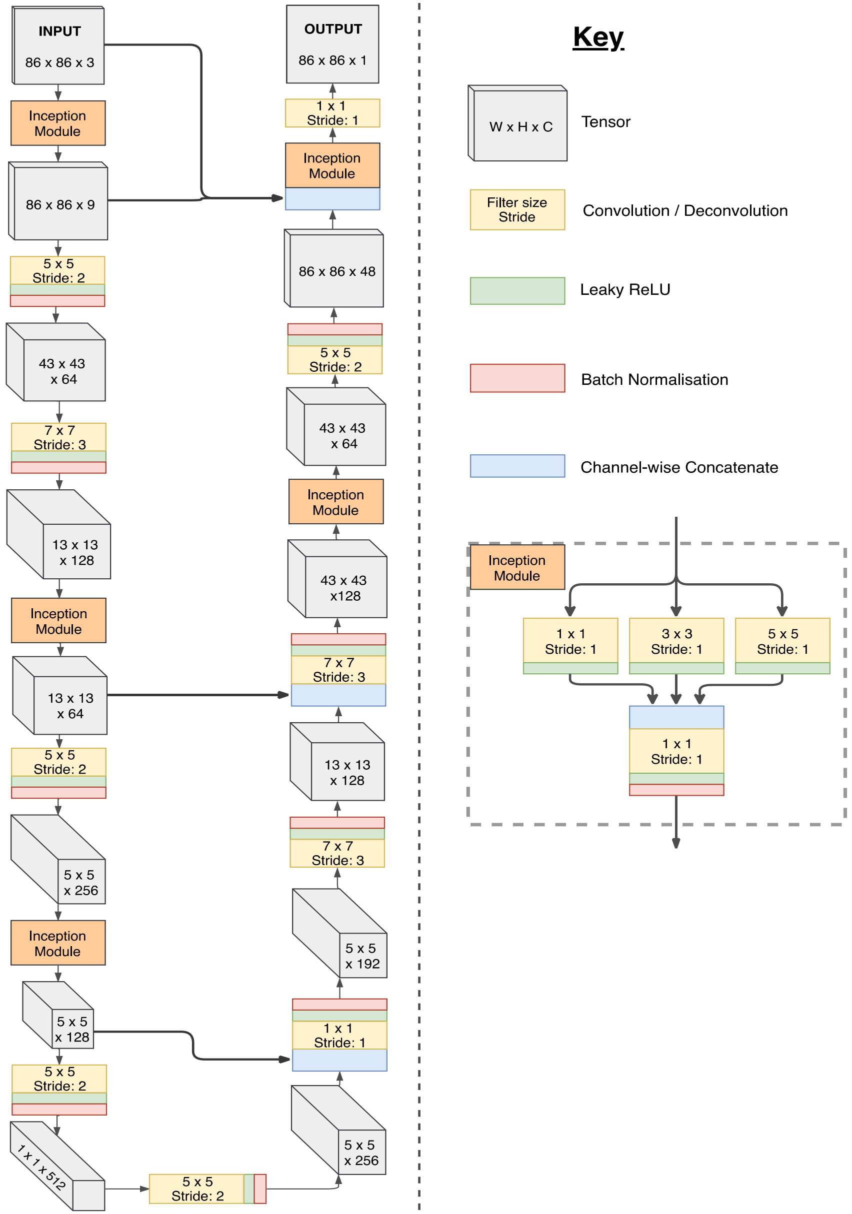

- Successive convolutions (LHS of diagram) produce higher-level features
- Deconvolutions (RHS of diagram) map features back to outputs
- Residual connections bypass deeper convolutions
- Concatenation followed by dimensionality reduction fuses information efficiently
- Receptive field gives idea of FOV of each feature layer

Dataset

- 96 tiles, each around 7000-by-7000 pixels
- 12 tiles for each of 8 terrains: Barren, Forest, Grass-Crops, Shrubland, Snow-Ice, Urban, Water, Wetlands
- Cloud masks hand-drawn by USGS
- Data taken from Foga et al., 2017

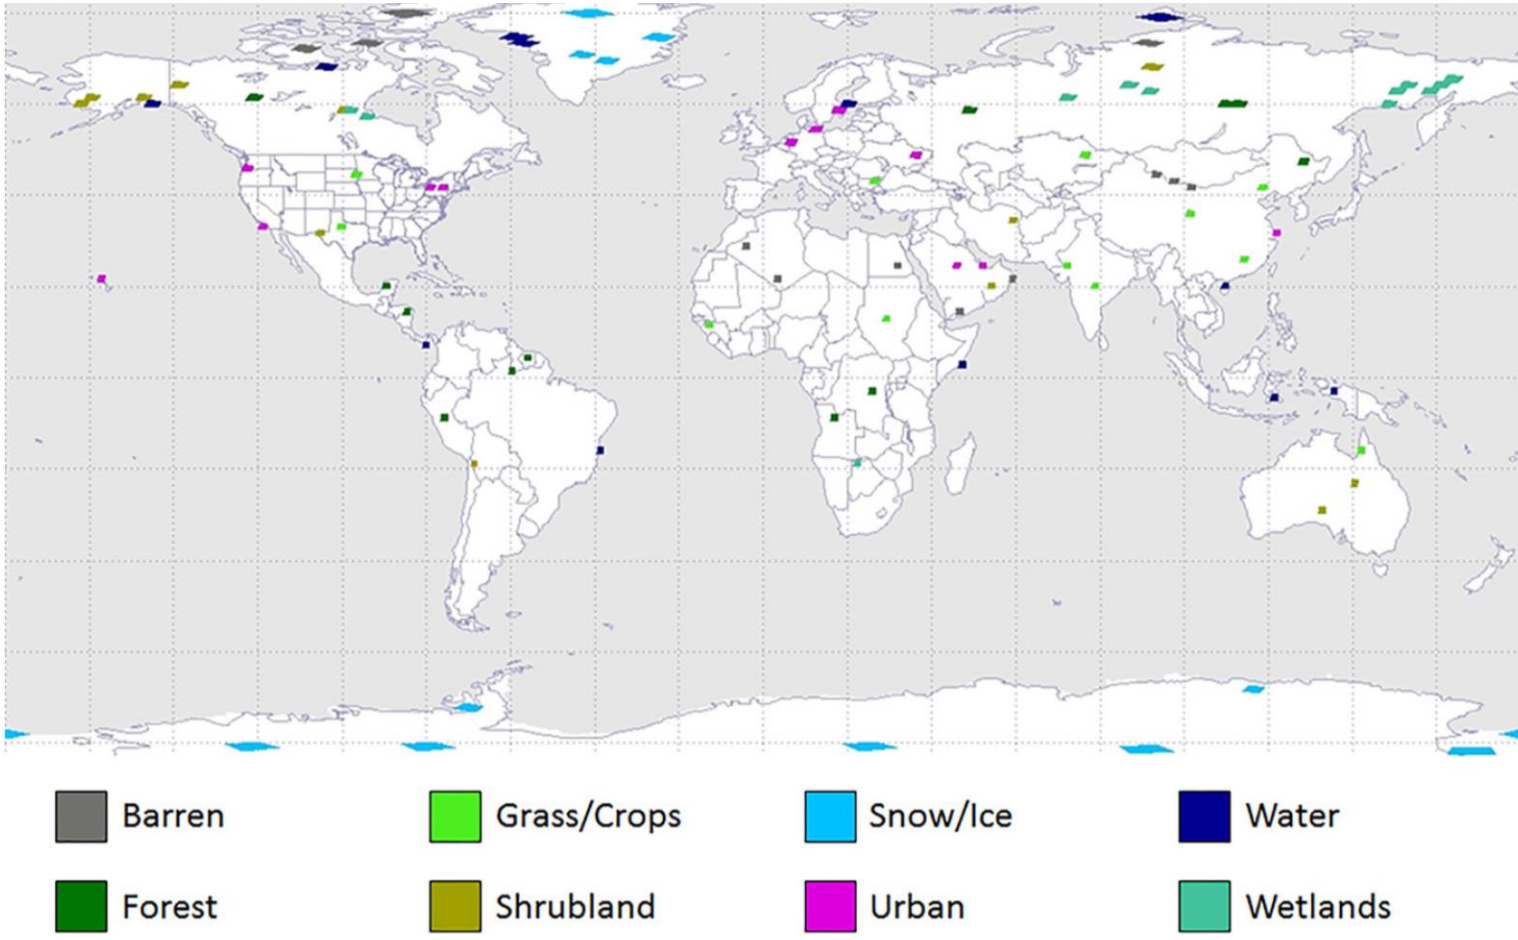

Courtesy of Foga, 2017

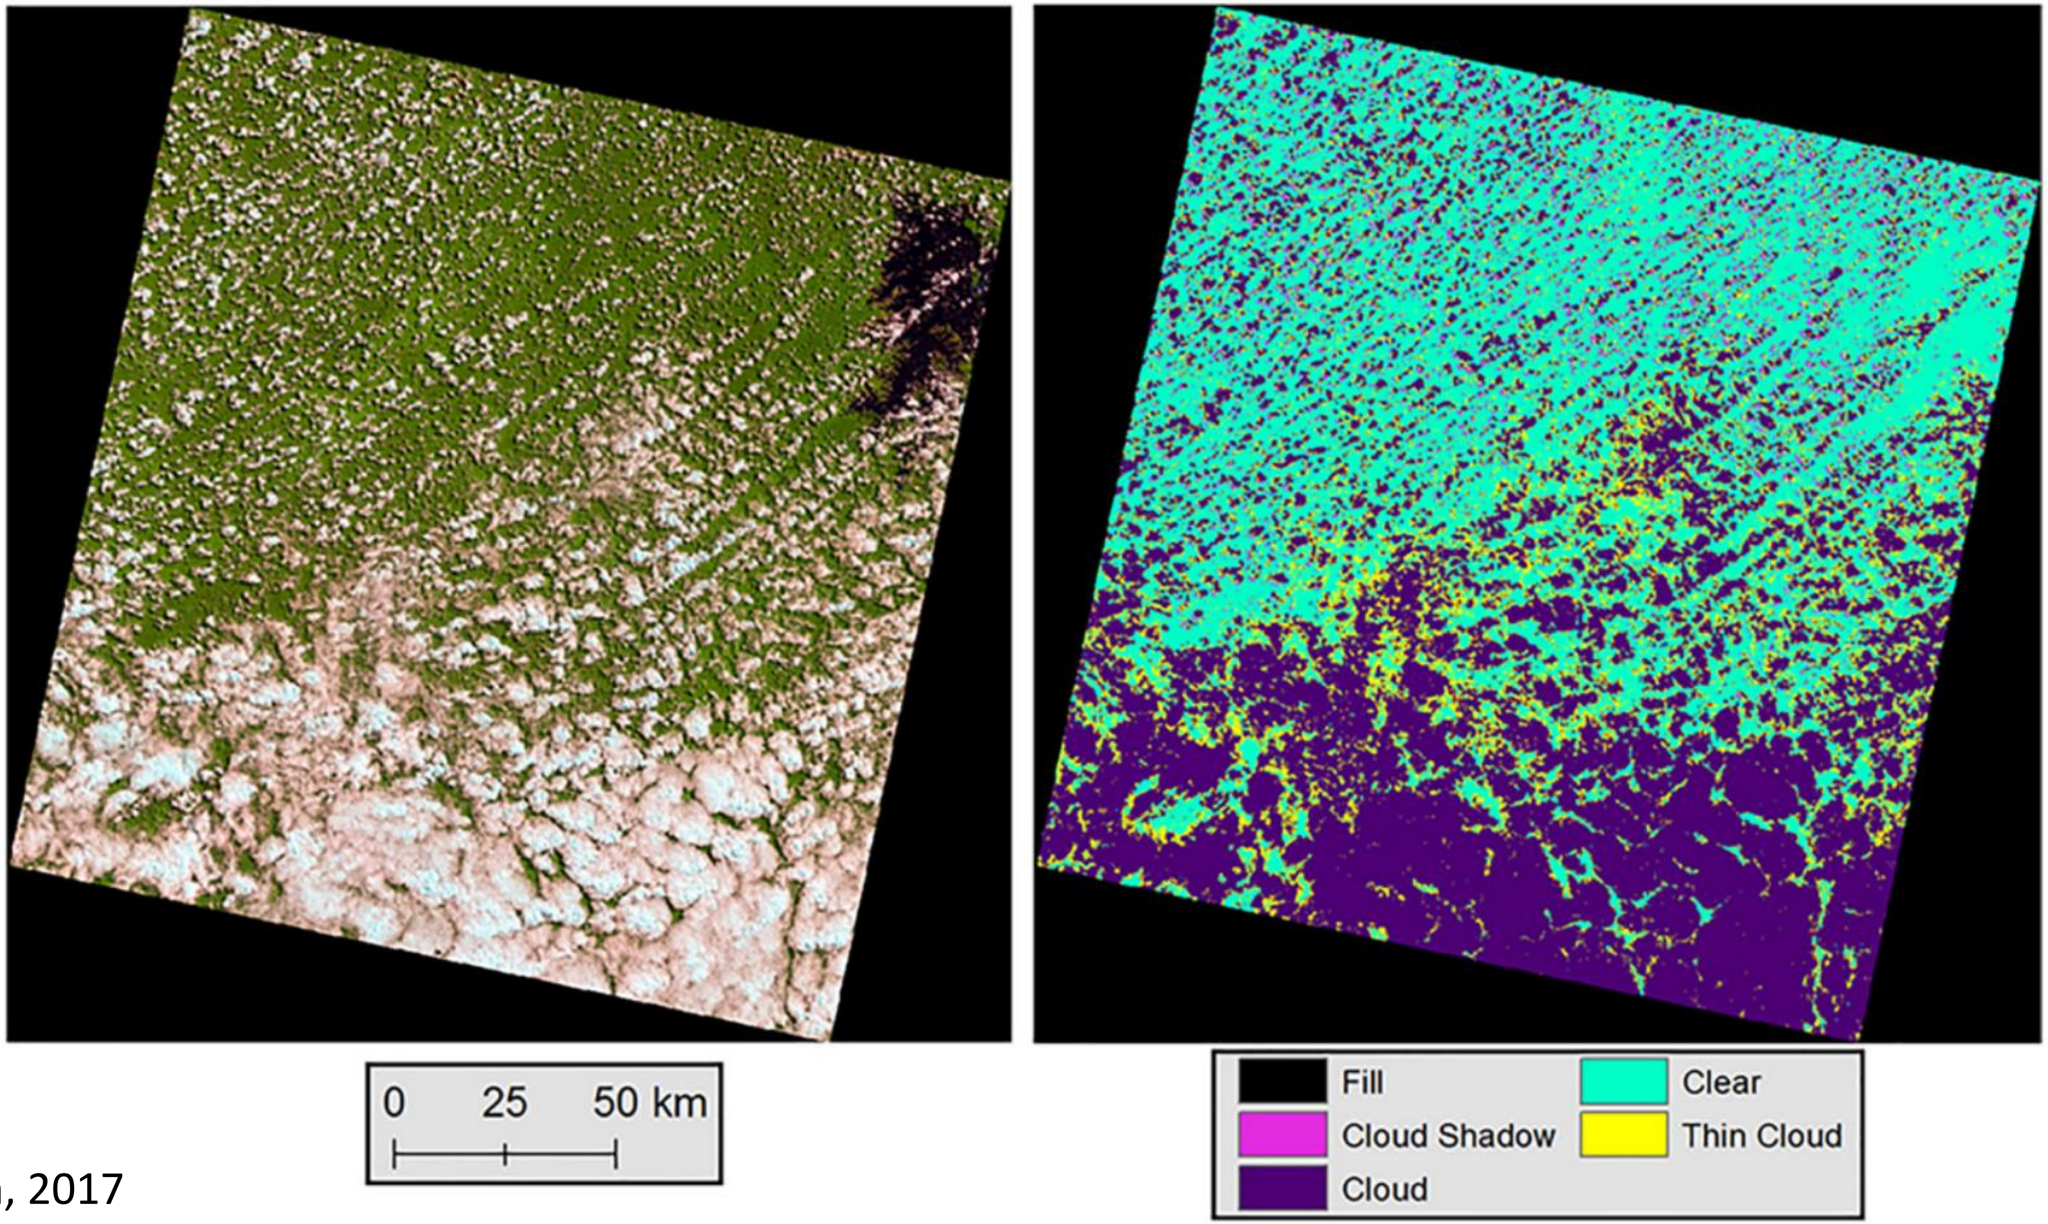

Courtesy of Foga, 2017

Performance Assessment

- Performance measured with 4 metrics:

1. %Correct = totalCorrect / total
2. %Omission = missedCloud / totalCloud
3. %Commission = falseCloud / totalNotCloud
4. %Quality = %Correct - %Omission - %Commission

- Quality provides balance between overall accuracy, missed detections and false detections

Performance Assessment

		Barren	Forest	Grass-Crops	Shrubland	Snow-Ice	Urban	Water	Wetlands	Average
CloudFCN (RGB)	%Correct	78.92	82.69	94.69	93.41	49.65	93.41	92.34	77.36	82.81
	%Omission	12.24	23.20	4.08	8.82	7.25	2.88	6.34	24.32	11.14
	%Commission	29.89	11.00	7.36	4.47	76.96	8.64	8.18	19.90	20.80
	%Quality	36.79	48.48	83.25	**80.11**	-34.56	81.89	77.82	33.14	50.87
CloudFCN (Multispectral)	%Correct	92.95	95.12	96.12	88.68	72.93	95.56	95.43	91.24	91.00
	%Omission	4.70	7.21	6.27	19.66	17.60	2.21	4.62	13.16	9.43
	%Commission	8.95	1.92	1.79	3.23	27.53	5.75	4.50	3.89	7.19
	%Quality	**79.30**	**85.99**	**88.06**	65.79	27.80	**87.61**	**86.31**	74.19	**74.38**
ACCA	%Quality	63.02	68.69	62	60.47	**36.25**	68.33	71.43	62.48	61.56
AT-ACCA	%Quality	66.67	73.83	74.09	70.65	35.86	74.06	70.51	**76.25**	67.72
cfmask	%Quality	77.1	67.27	85.74	75.53	26.37	74.72	50.98	65.97	65.69
cfmask_conf	%Quality	66.78	66.72	83.59	72.3	20.75	76.54	51.11	67.45	63.63
cfmask_nt_cirrus	%Quality	54.23	57.2	70.71	71.58	-15.87	74.37	50.23	47.16	51.62
cfmask_nt_cirrus_conf	%Quality	54.44	38.79	60.01	66.38	-43.68	73.2	49.04	35.14	41.66
cfmask_t_cirrus	%Quality	69.82	64.78	77.98	72.75	-24.1	72.42	57.21	53.27	49.01
cfmask_t_cirrus_conf	%Quality	69.37	43.99	77.76	72.34	-52.76	74.72	57.24	52.14	49.63
See5	%Quality	54.19	51.88	42.15	42.46	35.48	57.4	39.35	68.17	49.17

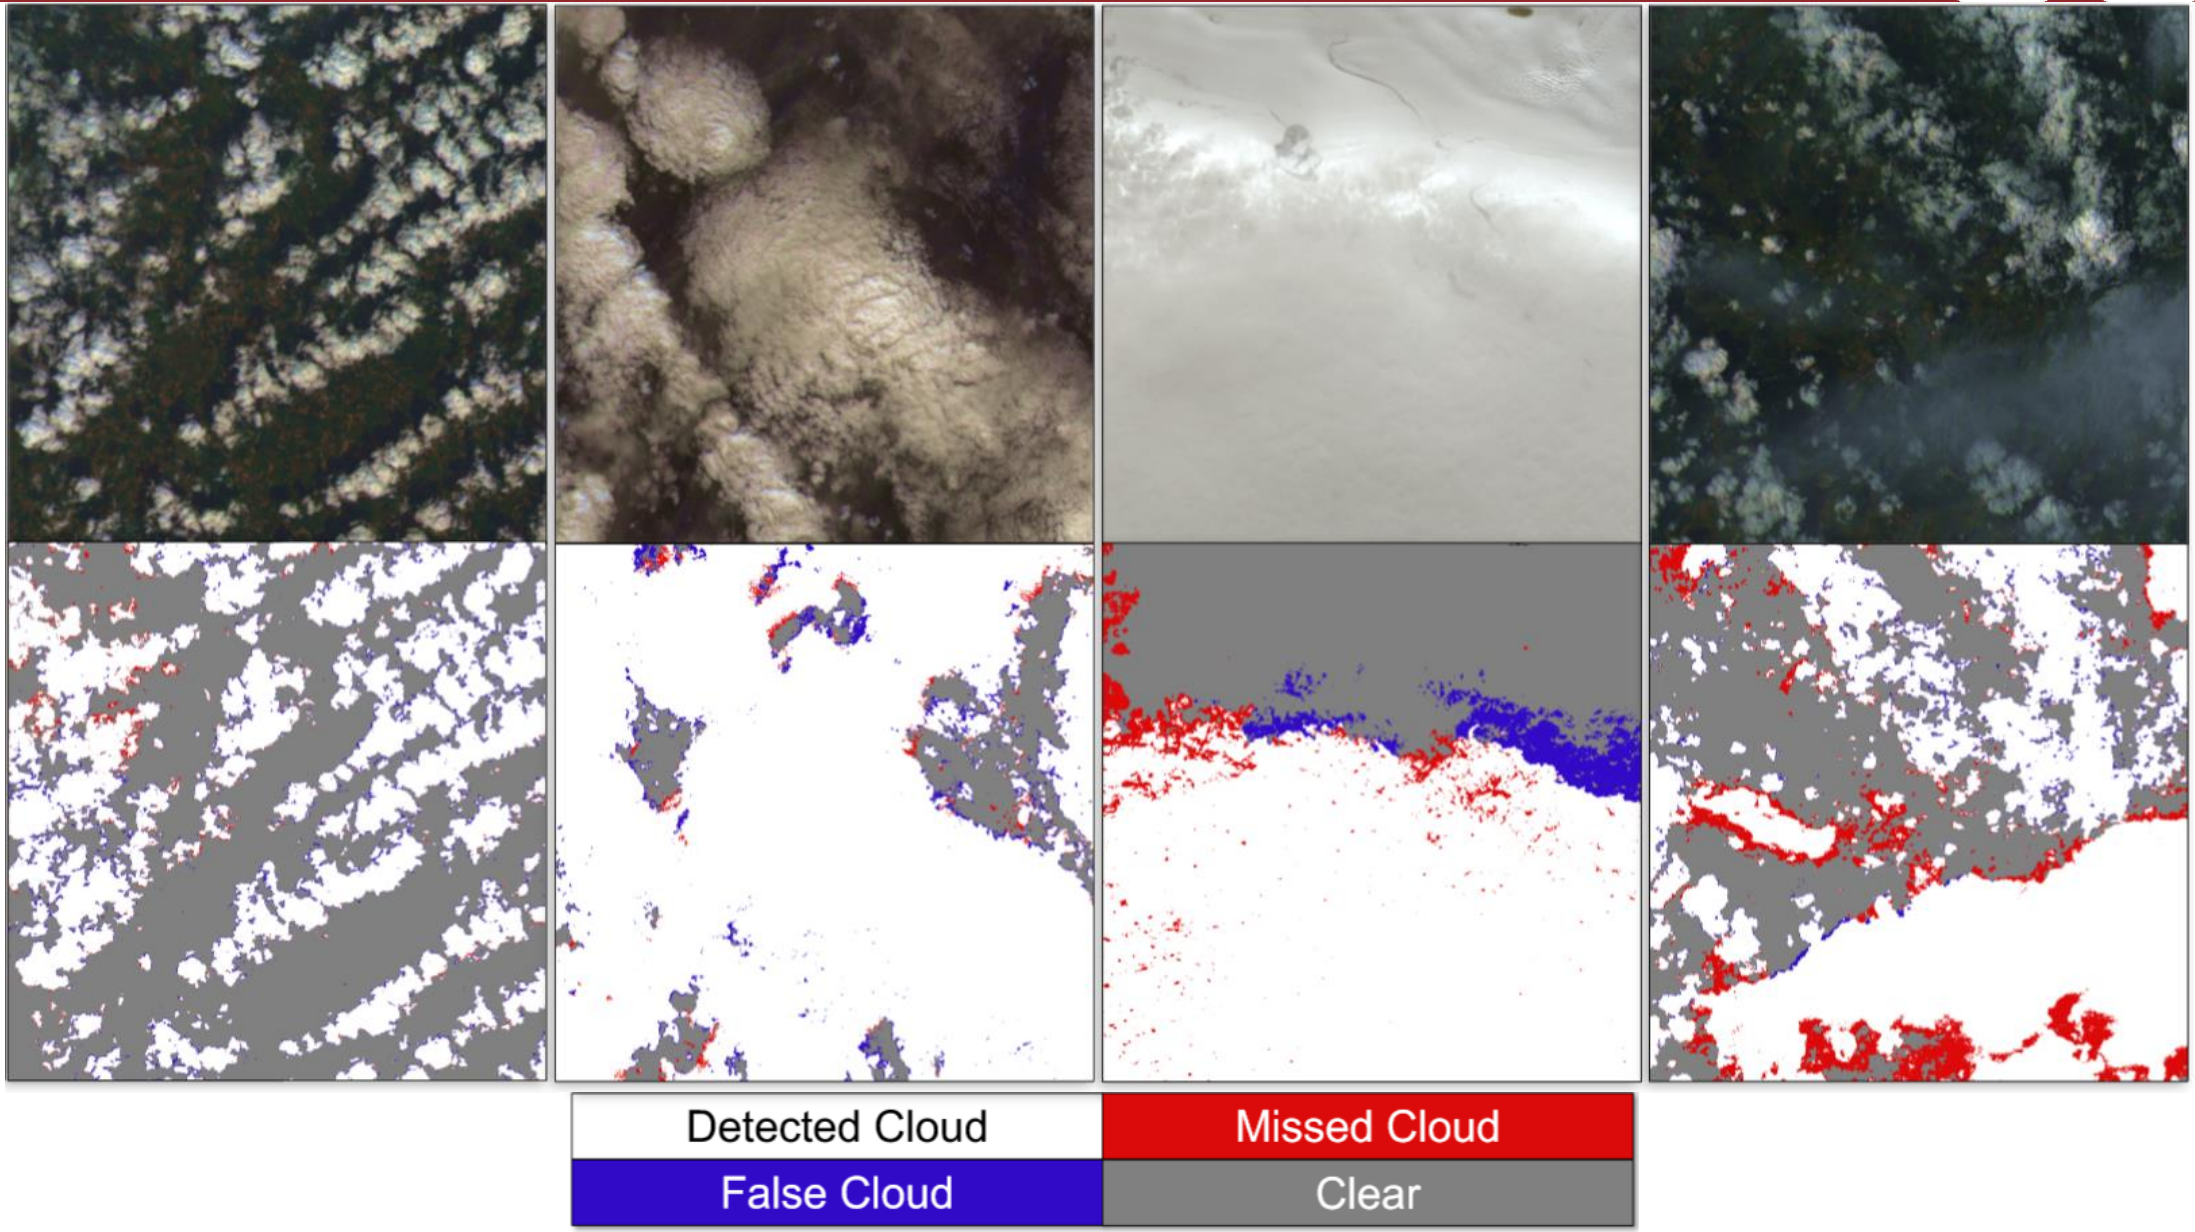
Detected Cloud
Missed Cloud
False Cloud
Clear

Performance

CPU: i7-8750H (2.2 GHz, 12 cores) – **0.1 MPix/sec**

GPU: GTX-1060 (6GB) – **1 MPix/sec**

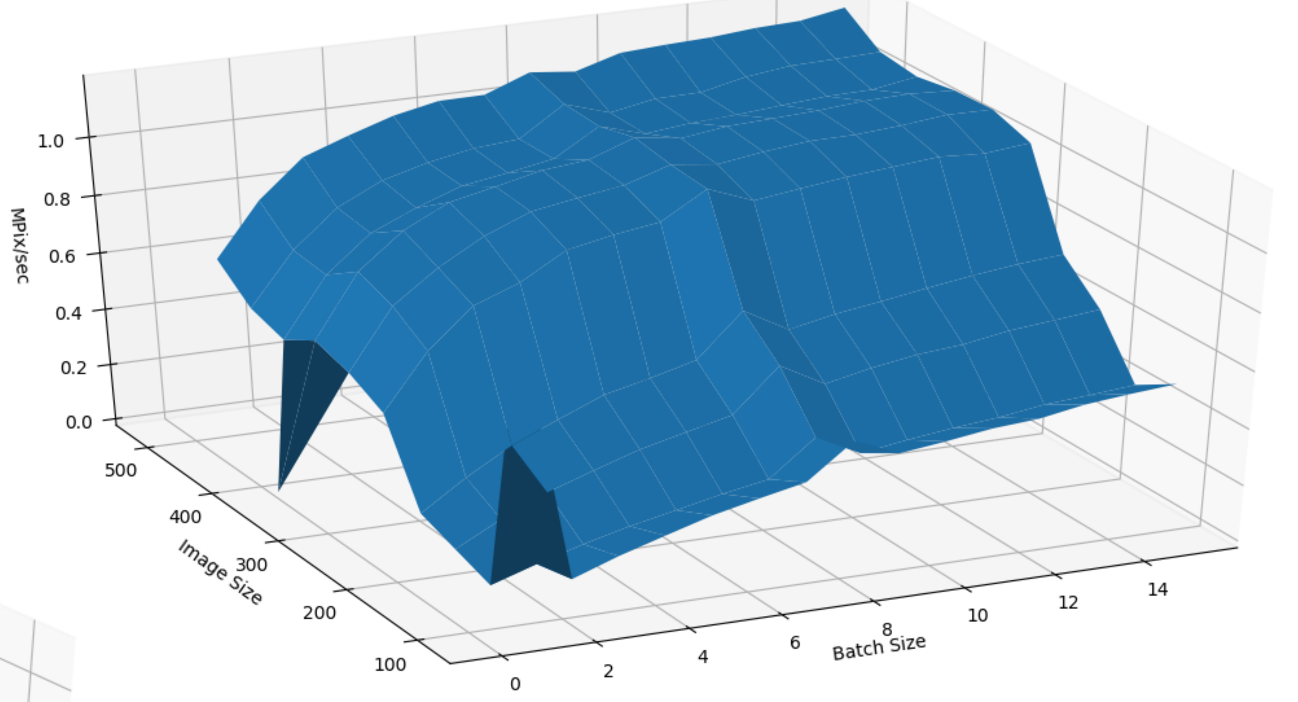

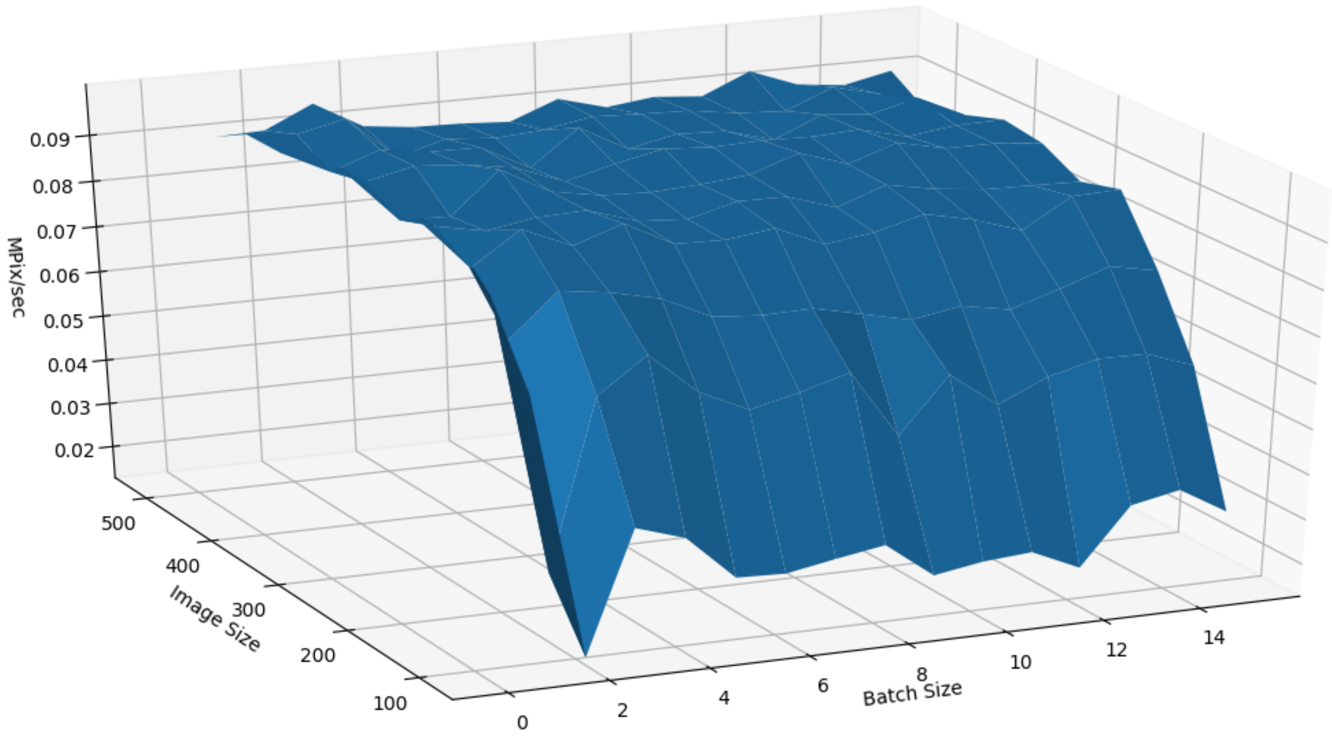

- High speed on consumer-grade GPU cards (£170)
- CPU performance unaffected by batch size
  - Calculated in series – no time save
- Down-sampling provides $O(n^2)$ speed-up

Outlook

- Processing of large amounts of Landsat 8 data
  - Provide community with alternative to fmask-derived quality assurance product
  - Further quantify performance on larger dataset
  - Single GPU @ 1MPix/sec = 4x data-rate of Landsat 8
- Other EO missions:
  - Sentinel 2/3 actively being worked on
  - Hyperspectral missions (e.g. Hyperscout-2, MODIS) being investigated
  - RGB satellites

Outlook

- On-board processing
  - OpenVINO, Movidius NCS, etc. provide platforms with ML capabilities
  - Data reduction
  - Dynamic data compression
  - Smart operational decision-making
- Wider applications
  - Generally applicable to segmentation problems
  - Land-use classification
  - Flood detection
  - Smoke/fire detection
  - And many others…

Predicted Impacts

- Affected missions
  - Short-term: All multispectral EO satellites
  - Long-term: With improved performance, panchromatic or otherwise limited spectral combinations
- New-mission concepts
  - Cloud mapping instruments without LIDAR/RADAR
  - Dynamically operated constellations: Prioritisation of repeat observations
- New science
  - Higher accuracy in existing science: fewer spurious datapoints
  - Big Data approaches to cloud cover trends for climate modelling
  - Other segmentation problems in EO: Sea ice detection, Land-use, etc.
- Commercialisation
  - Algorithms not directly protectable: these techniques are becoming standard in ML-sphere
  - Embedded systems
  - Pre-validated models and mission-phase training

Questions

Extra Slide

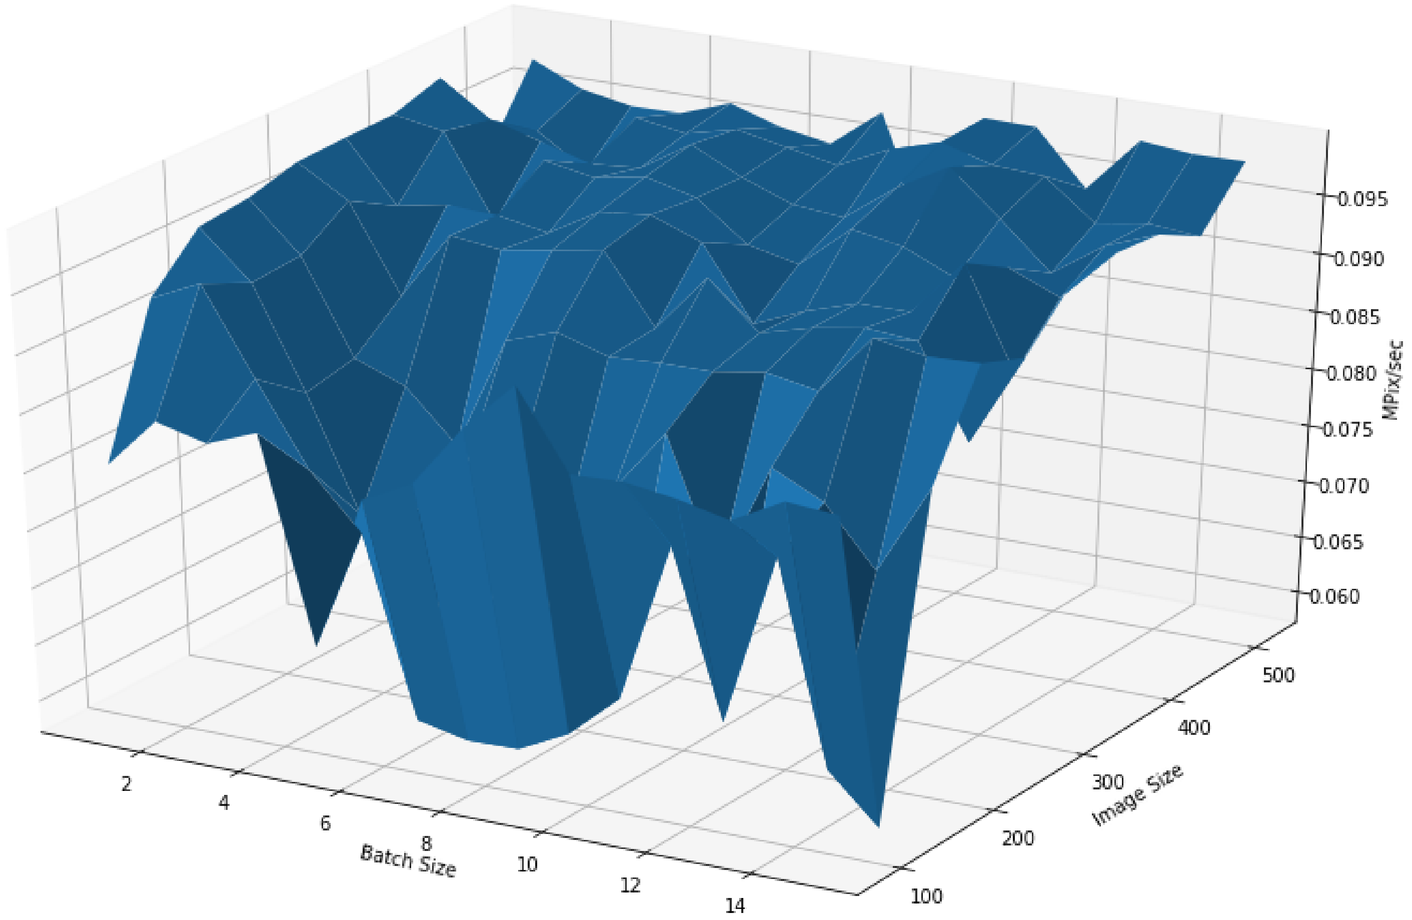

- Google TPU
- Suboptimal performance
- Low memory allowance is likely culprit