

Transfer Learning for Onboard Cloud Segmentation in Thermal EO

From Landsat to a CubeSat
Constellation

Niklas Wölki, Lukas Liesenhoff,
Christian Mollière, et al.

OroraTech




Goal and Challenges

Goal


Accurate onboard cloud segmentation.

Challenges

- Traditional Algorithms rely on multispectral input
 - FOREST-2: Single thermal band available for cloud masking at time of research
- Cloud Segmentation on thermal only largely underexplored
- Deep Learning Models emerging, but rely heavily on large, well-annotated datasets
 - minimal mission-specific data → transfer learning approach



Datasets




FOREST-2

Band LWIR-2: 10.5–12.6 μm

24 hand labelled images

→ 528 crops



Landsat-7

Band 6: 10.4 - 12.5 μm

200 pre-labelled images

→ 6000 crops



Transfer Learning Strategies

Baseline: f2-f2

- Train and test only on FOREST-2

Benchmark: l7-l7

- Train/test on Landsat-7

Generalisation: l7-f2

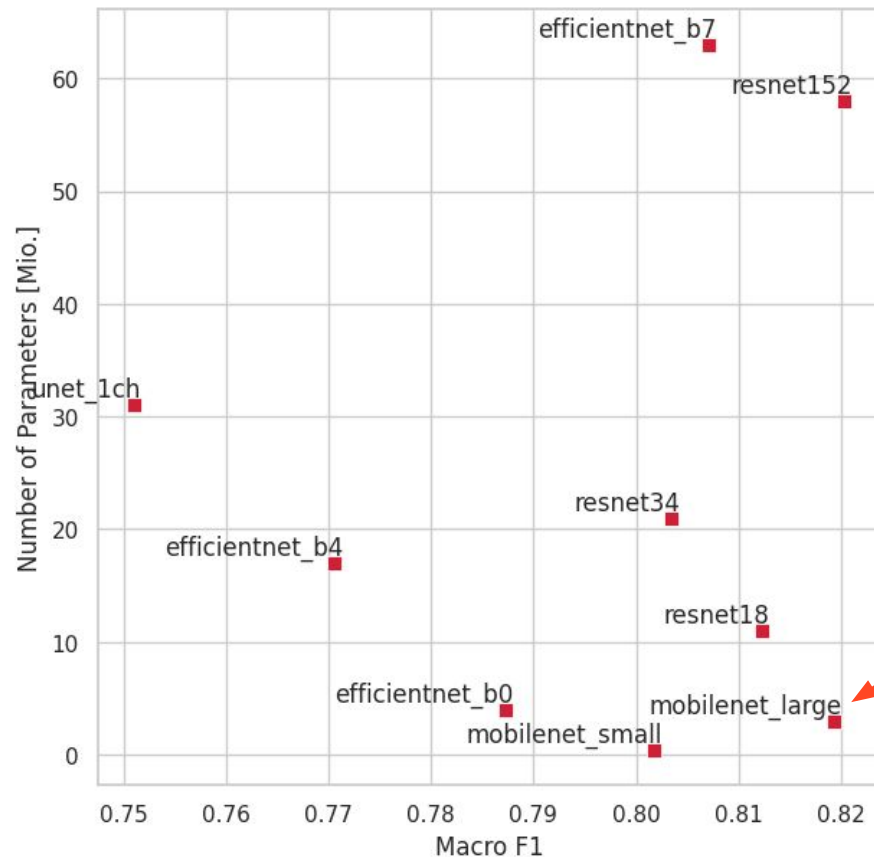
- Train on Landsat-7, test on FOREST-2

Transfer: joint-f2

- Train on Landsat-7 and FOREST-2, test on FOREST-2



Key Results




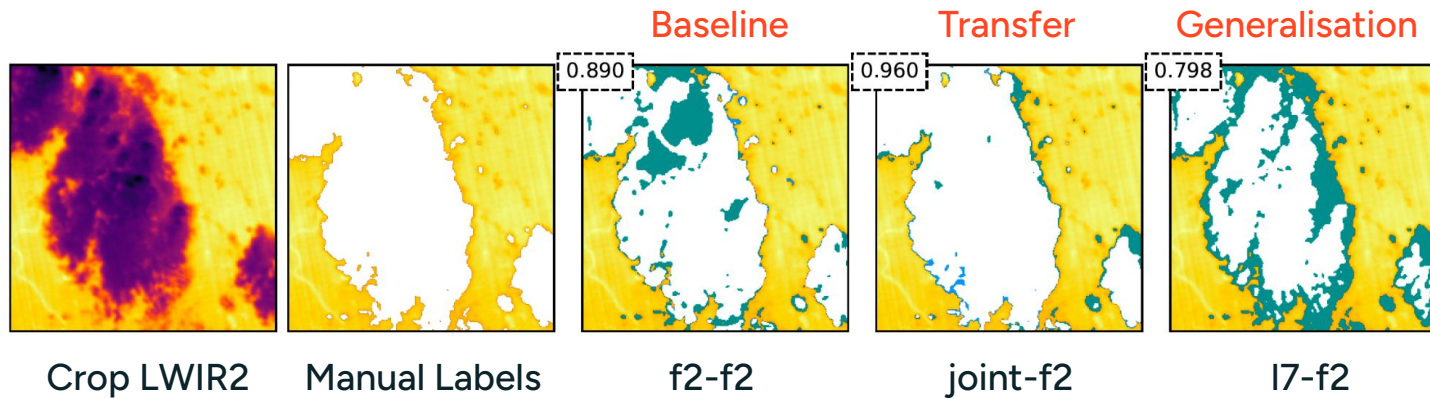
Key Results

Experiment Name		Macro F1	Accuracy
Baseline	f2-f2	0.850	0.889
Benchmark	l7-l7	0.903	0.917
Generalisation	l7-f2	0.796	0.830
Transfer	joint-f2	0.877	0.910

Onboard inference time: ~5s (2691x1762px)



Key Results



Takeaways

- Accurate cloud masking on thermal-only data is feasible
- Combining public and mission-specific data enables strong performance with minimal labelling effort
- Efficient model design + tiling strategy + GPU optimizations makes cloud masking feasible onboard
- Improve validation efforts and implement more sophisticated domain adaptation approaches in the future





OroraTech

Thank You!

Feel free to reach out

niklas.woelki@ororatech.com