ICEYE Lumi usecase

Tapio Friberg

ICEYE Oy

tapio.friberg@iceye.fi

An artists rendition of an ICEYE satellite.

ICEYE - Synthetic Aperture Radar (SAR)

GLOBAL COVERAGE

SHORT REVISIT TIMES

ANY PLACE ON EARTH
4-8 TIMES A DAY

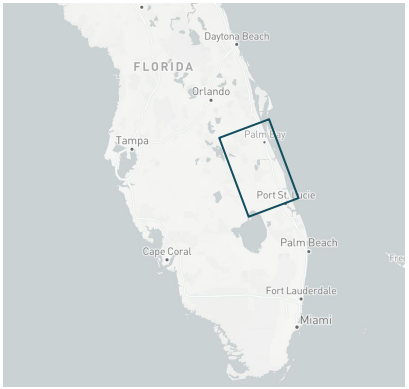

LARGE FOOTPRINT HIGH RESOLUTION

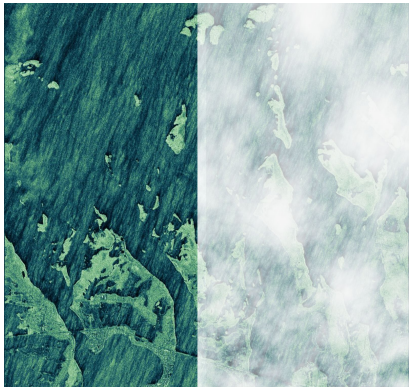

ALL WEATHER INFORMATION

THROUGH CLOUDS,
SMOKE AND DARKNESS

ICEYE

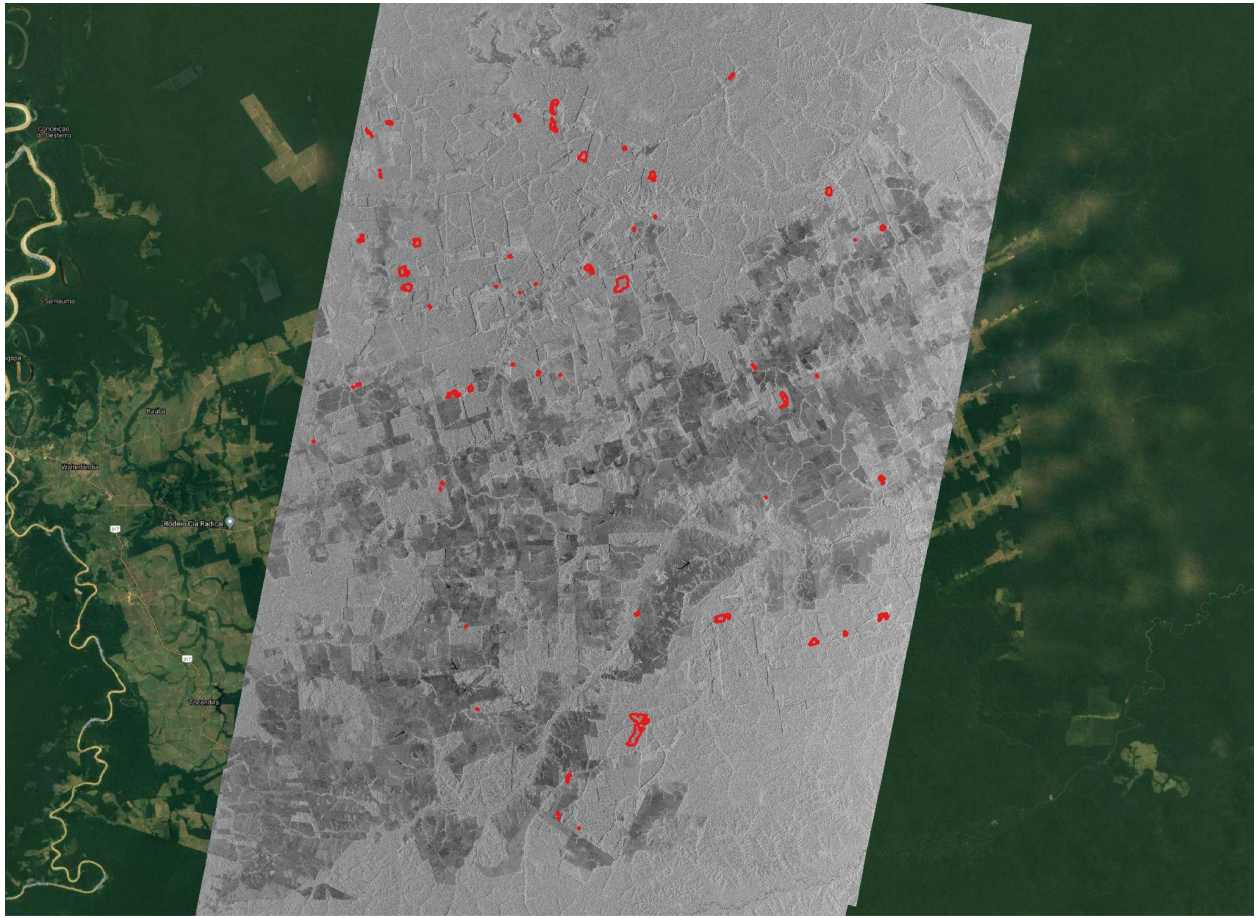

The ICEYE stripmap product might be ~45000 x 15000 pixels in full resolution.

A single frame is ~70km x ~30km.

Example datasets at :
https://www.iceye.com/downloads/datasets

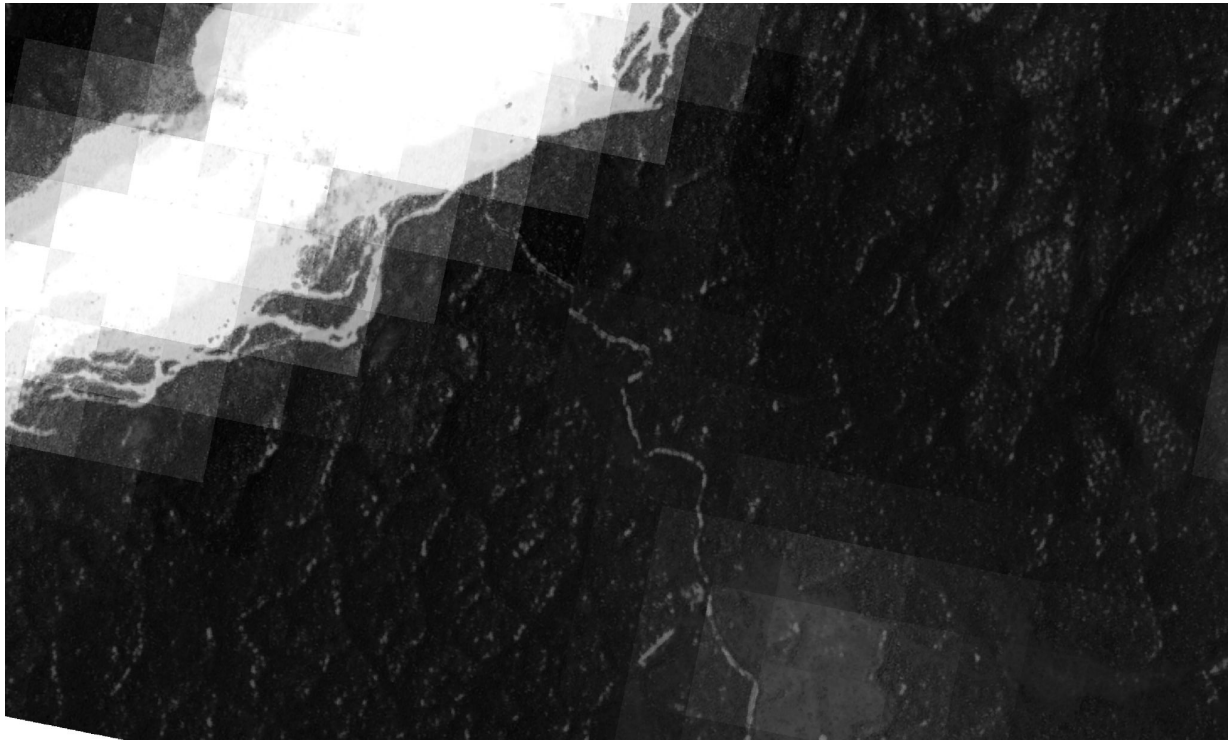

Water segmentation probability map.

“just “ 512 x 512 pixel patches are too parochial. Need more context.

Too much context -> problem dimensionality and memory costs explode.

Approach - masked autoencoders

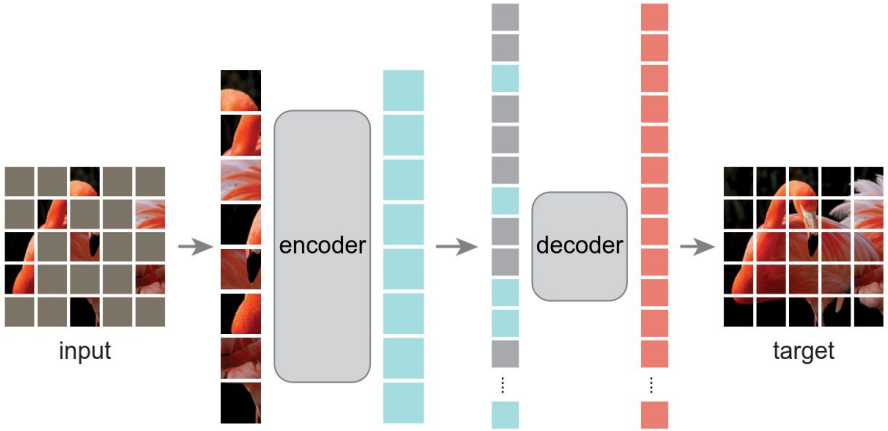

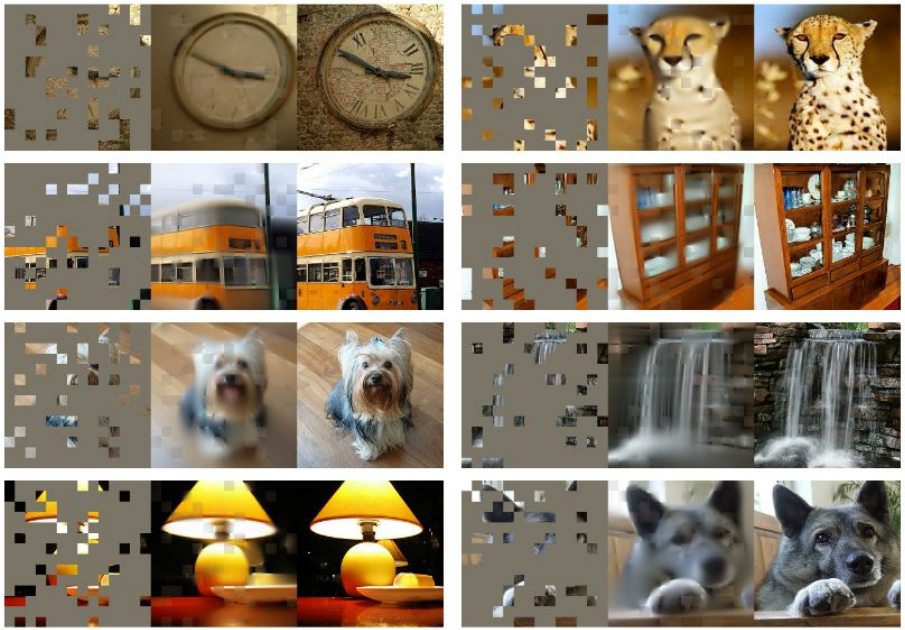

He, Kaiming, et al. "Masked autoencoders are scalable vision learners." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2022.

Approach - masked autoencoders

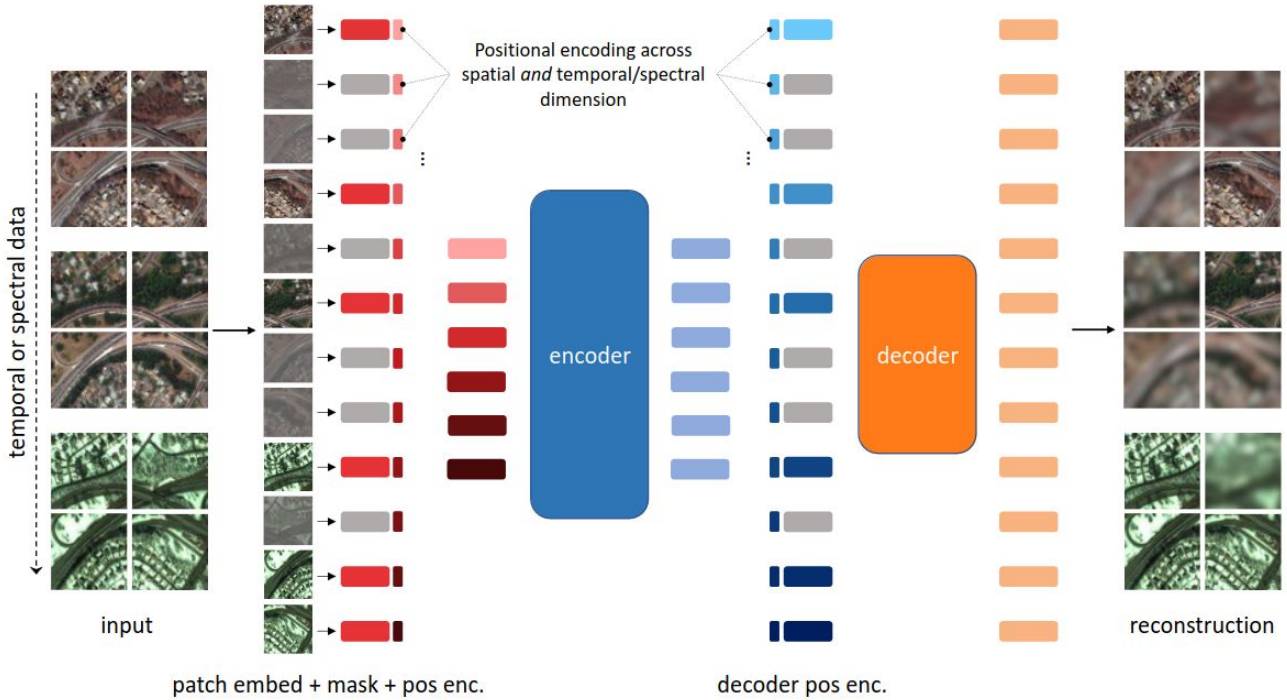

Cong, Yezhen, et al. "SatMAE: Pre-training transformers for temporal and multi-spectral satellite imagery." *Advances in Neural Information Processing Systems* 35 (2022): 197-211.

Varför Lumi vi har ju AWS

We rely heavily on cloud compute and storage.

Most of our workloads are bursty, unpredictable, need stability.

Lumi fills an important R&D niche for us:

- Single upfront IO cost, downloading 5TB+ of data can take days.
- Large scale flash storage good for IO bottlenecked operations.
- Self-supervised applications such as the MAE take days to compute.
- We want all the GPU memory, nothing is enough. The Lumi MI250X has a good GB/$.
- Multi-node training for ultra-large scale things (Monte Carlo simulation of Monte Carlo simulations)