Convolutional Neural Processes for Inpainting Satellite Images

Application to Water Body Segmentation

Alexander Pondaven*, Märt Bakler*, Donghu Guo, Hamzah Hashim, Martin Ignatov,
Samir Bhatt, Seth Flaxman, Swapnil Mishra, Elie Alhajjar, Harrison Zhu
Imperial College London
NeurIPS 2022 Workshop - Tackling Climate Change with Machine Learning

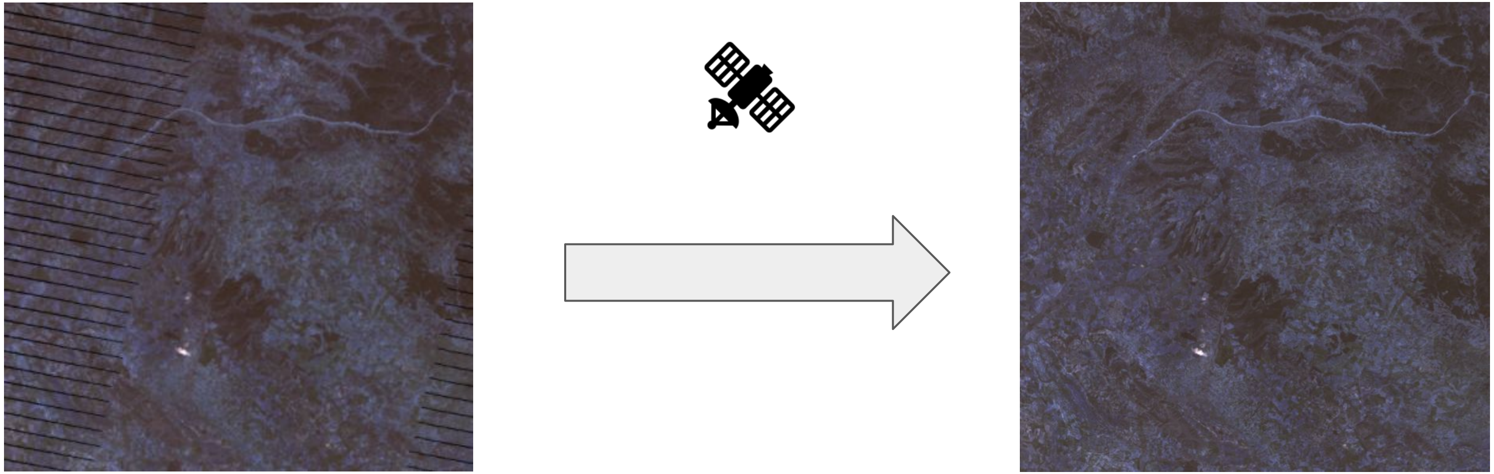

Our Contributions

- ☑ Repair LANDSAT 7 imagery with **Convolutional Neural Processes**
- ☑ State-of-the-art inpainting performance on in-distribution and **especially** out-of-distribution (OOD) inpainting
- ☑ Strong performance in climate **downstream** tasks

Satellite Imagery: LANDSAT 7

- LANDSAT 7 - images collected by NASA/USGS via the LANDSAT programme
- High-resolution (30m) images publicly available (**massive, terabytes**!)
- Scanline corrector (SLC) failure on 31st May 2003 → missing values at scanlines

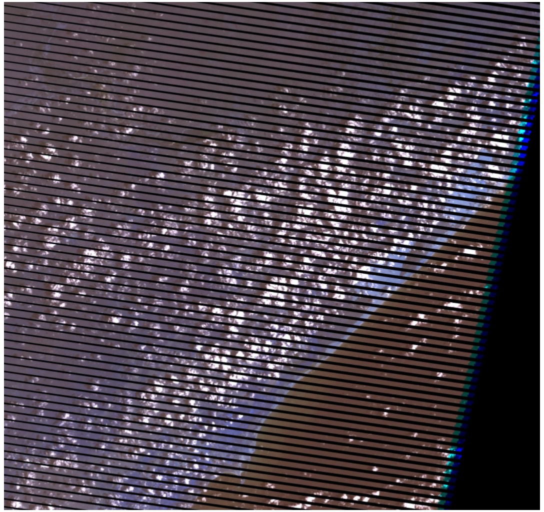

Figure 1: Snapshot in Kenya. Taken on 3rd January, 2005, after the SLC failure

Data from Google Earth Engine

- LANDSAT 7 Satellite images extracted using Google Earth Engine API (Gorelick et al. 2017)
- RGB channels/bands
- 256x256 images downloaded
- Cropped to 128x128 and 64x64 for training

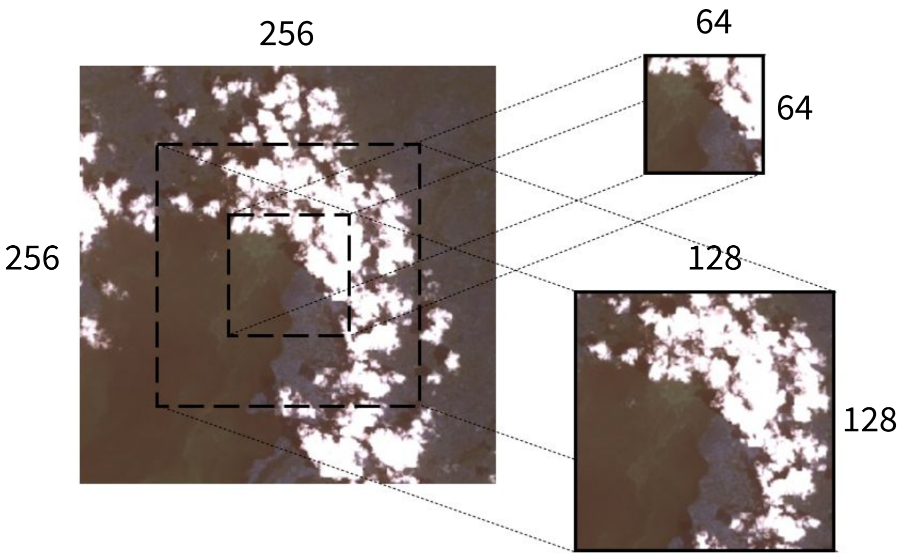

- In-distribution country
  - Kenya
- Out-of-distribution countries
  - UK
  - Brazil
  - Nepal
  - Norway

Data Processing for Training

- Post-2003 images used to extract scanline bitmasks
- Pre-2003 uncorrupted images used for training

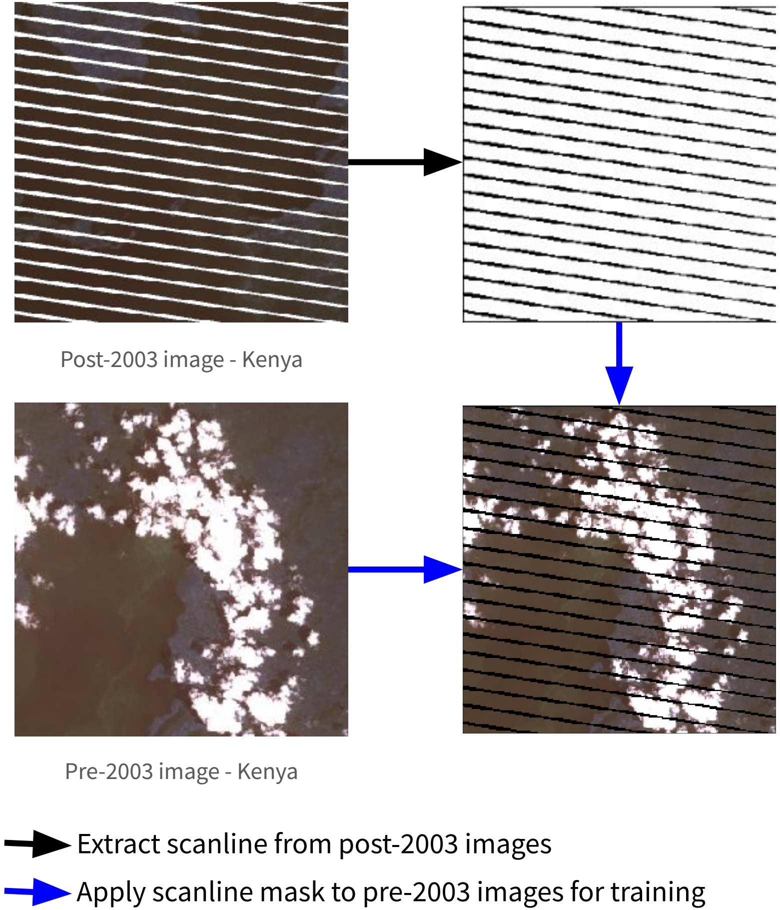

Baselines

Navier-Stokes

- ☑ Fast
- ☒ No information sharing between images

U-Net

- ☑ Expressive and works quite well for a lot of problems
- ☒ OOD requires large datasets and data augmentation

PartialConv

- ☑ Convolution takes into account of masks/missing pixels
- ☒ Requires large datasets and long training times

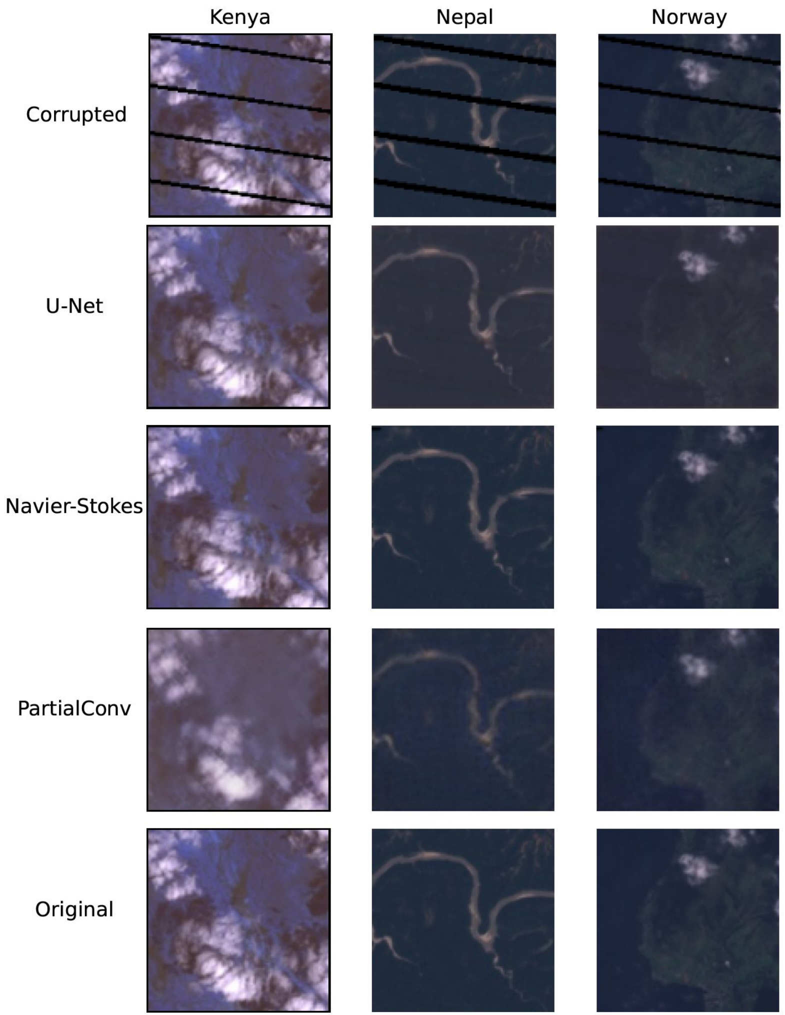

Neural Processes for Inpainting

- Satellite images are different regression problems
  - Different location and time
- Small dataset for each task

Context points are non-scanline pixels

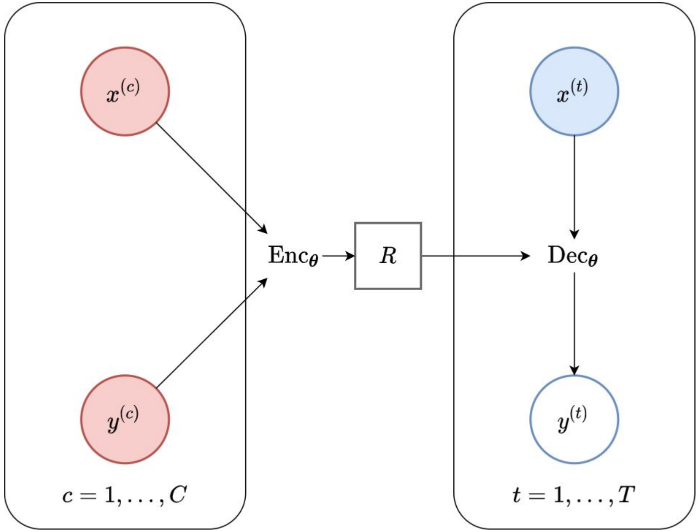

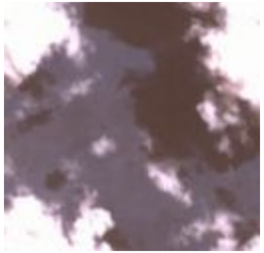

Target points are entire image (for continuity)

Convolutional Neural Processes

- Translational equivariance

- Convolutional Conditional Neural Processes
- Convolutional Latent Neural Processes

- Trained using Maximum Likelihood
- Multi-Scale Structural Similarity (MS-SSIM) Loss (Wang et al. 2003) generates sharper images

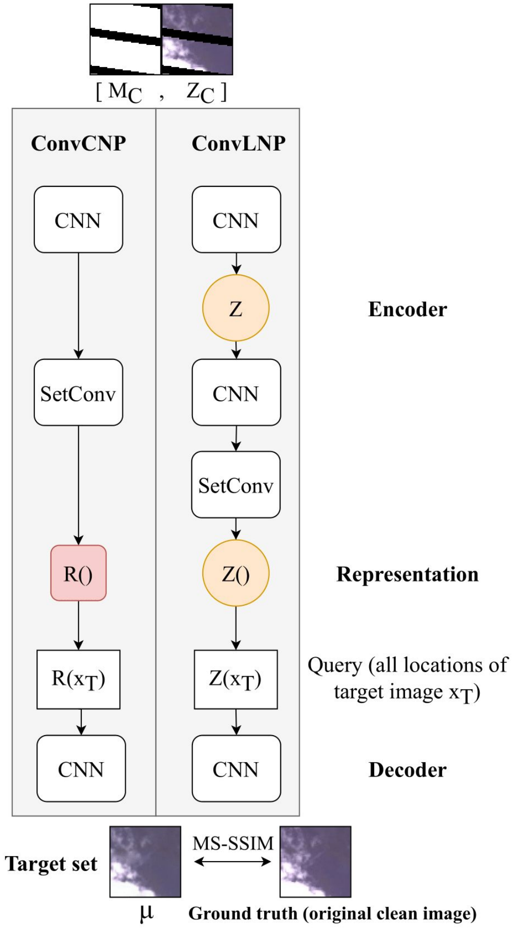

Inpainting results

- **ConvNPs** perform well both for in and out-of-distribution images and outperform baselines
- **ConvLNP** performs the best on average when also compared to ConvCNP
- Good generalisability of Meta-Learning is a result of treating input images as different tasks

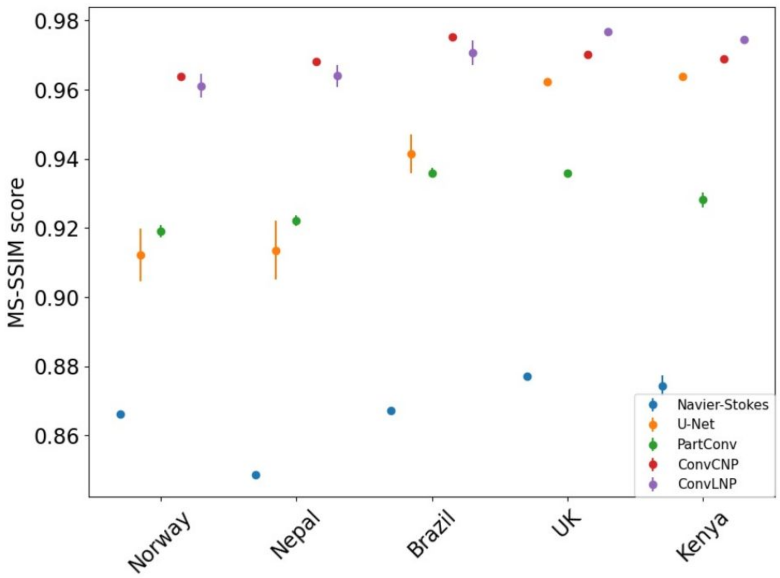

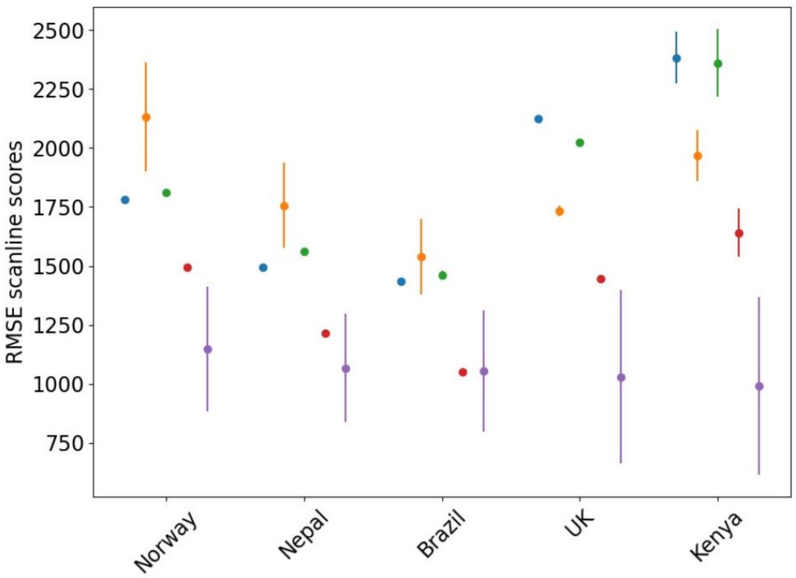

Experiment 2:
Water Body Segmentation Downstream Task

- Image segmentation of seasonality of water in Canada
- Classify pixels into 3 classes based on 3 months of imputed satellite images
- UNet with 3D convolutions and masked cross entropy loss

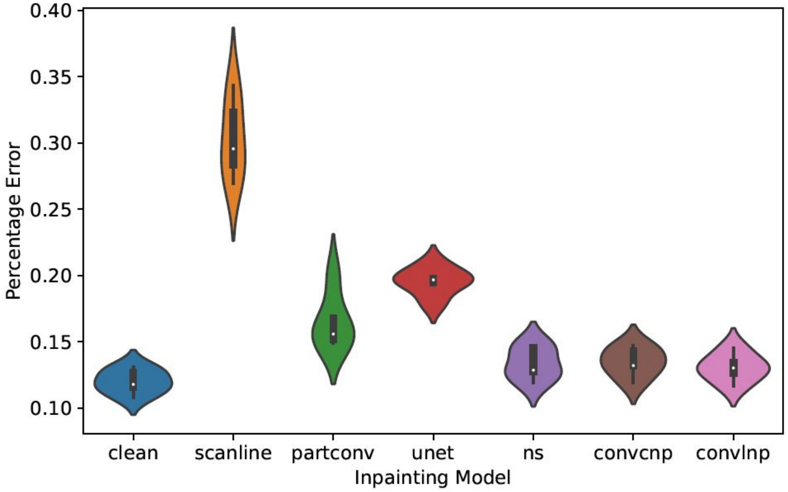

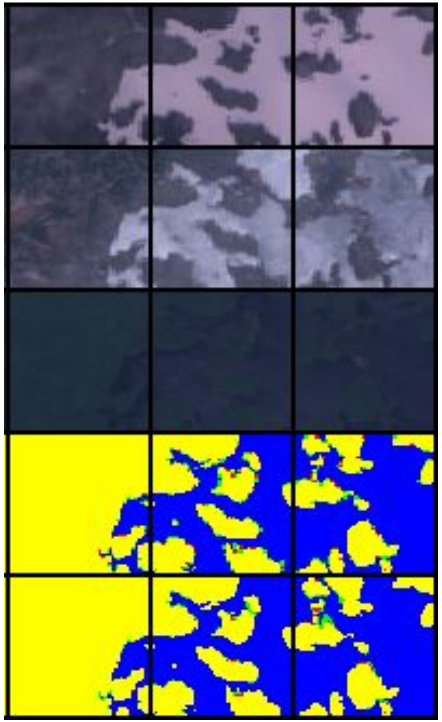

Conclusion and Discussion

- ☑ ConvNPs successful at inpainting in-distribution and out-of-distribution
  - ☑ Take advantage of different spatiotemporal structure of satellite images
  - ☑ Global inpainter for LANDSAT 7 by only training small subset of locations
- Bigger scanlines
- Cloud removal
- A wider array of downstream tasks

References

- Bertalmio et al. (2001). Navier-Stokes, Fluid Dynamics and Image and Video Inpainting. CVPR
- Dubois et al. (2020). Neural Process Family. https://yanndubs.github.io/Neural-Process-Family
- Foong et al. (2020). Meta-Learning Stationary Stochastic Process Prediction with convolutional Neural Processes. NeurIPS
- Garnelo et al. (2018). Conditional neural processes. ICML
- Garnelo et al. (2018). Neural Processes. *ICML Workshop on Theoretical Foundations and Applications of Deep Generative Models*
- Gordon et al. (2020). Convolutional conditional neural processes. *ICLR*
- Gorelick et al. (2017). Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*
- Liu et al. (2018). Image inpainting for irregular holes using partial convolutions. ECCV
- Lugmayr et al. (2022). Repaint: Inpainting using denoising diffusion probabilistic models
- Markou et al. (2022). Practical conditional neural processes via tractable dependent predictions. *ICLR*
- Nazabal et al. (2020). Handling incomplete heterogeneous data using vaes. *Pattern Recognition*
- Richard and Chang (2001). Fast digit image inpainting. VIIP
- Ronneberger et al. (2015). U-net: Convolutional networks for biomedical image segmentation. *International Conference on Medical image computing and computer-assisted intervention*
- Scaramuzza and Barsi(2005). Landsat 7 scan line corrector-off gap-filled product development
- Telea (2004). An image inpainting technique based on the fast marching method. *Journal of graphics tools*
- USGS. Landsat 7 Courtesy of the U.S. Geological Survey.
- Wang et al. (2003). Multiscale structural similarity for image quality assessment. *Asilomar Conference on Signals, Systems & Computers*