MICHAEL ZOTOV \ JEVGENIJ GAMPER

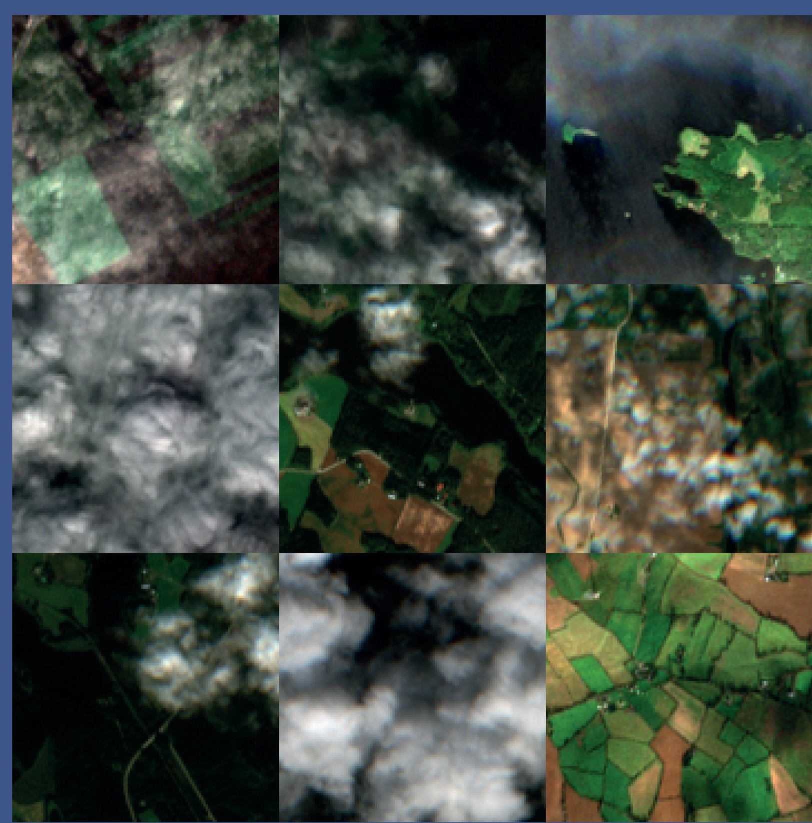

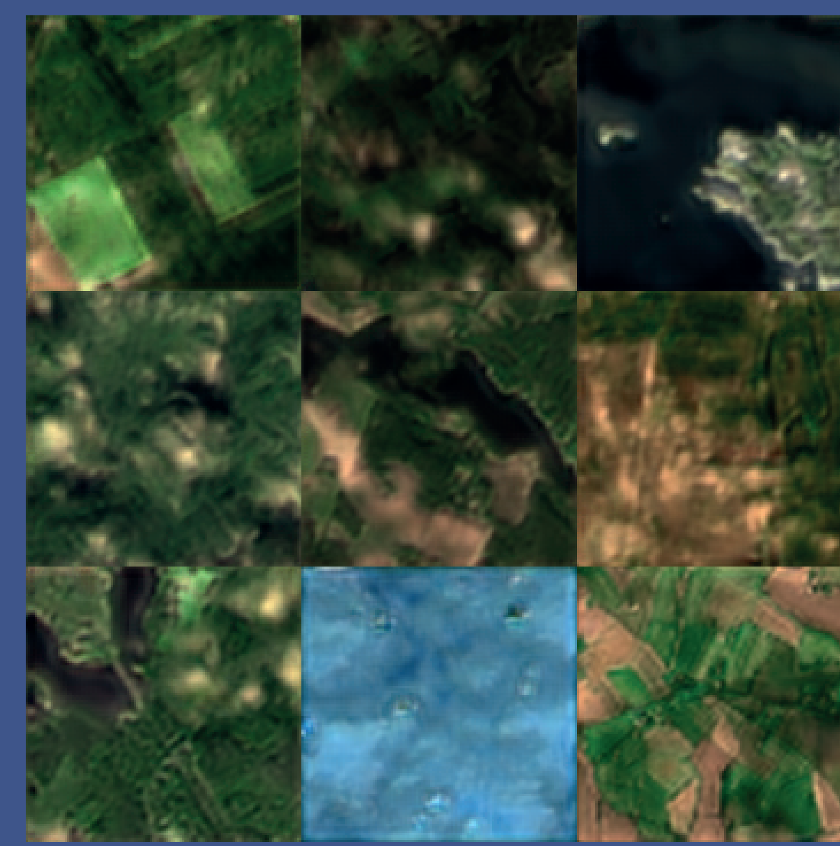

Cycle-Consistent Generative Adversarial Networks (Cycle GANs) can help obtain cleaner satellite imagery to improve precision and effectiveness of Earth science applications that use such data.

Examples of both good and poor performance of the generator that maps from the space of cloudy to the space of clear images.

Conditional Denoising of Remote Sensing Imagery Using Cycle-Consistent Deep Generative Models

The Problem

Satellite Imagery contains a lot of noise due to clouds and their shadows. This noise greatly hinders the effectiveness of many applications that use remote sensing data fo r Earth sciences. For instance, there is only a 25% chance that an arbitrary active field will appear in daily optical satellite imagery.

The Data

BigEarthNet dataset provides over 500000 labelled Sentinel-2 images in 12 bands, split into unpaired sets of cloudy and clear images

Proposed Solution

Train a Cycle GAN on unpaired sets of examples to produce images without clouds, conditioned on information contained in cloud penetrating satellite bands. We also introduce latent variables in training in hopes to encode the possible distribution of labels under the cloud, or the position of the cloud, depending on the direction of training.

The Good

The proposed solution provides very promising results.

The Bad

There is a lack of precision in available tools to evaluate GANs.

The Ugly

Capability of GANs to learn an approximate underlying data distribution is questionable. Convergence during training is also not established

Future Directions

Resolving the highlighted issues with Deep Generative Models like GANs will be fundamental for successfully applying adversarial inference to climate, remote sensing, or Earth science tasks. Additionally, using additional features from radar-based satellites is expected to yield better results.

Example of a full cycle during training. One of the issues we encountered is using L1 distance to compare the original and the restored image, which forces the generator to restore clouds in original locations.

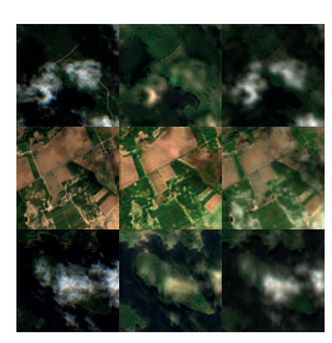

We use a Fréchet Distance-based metric to evaluate performance during training. It is, however, not ideal lower FD does not always correspond to high fidelity samples in this particular problem. A model that has suffered mode collapse may have lower FD on validation

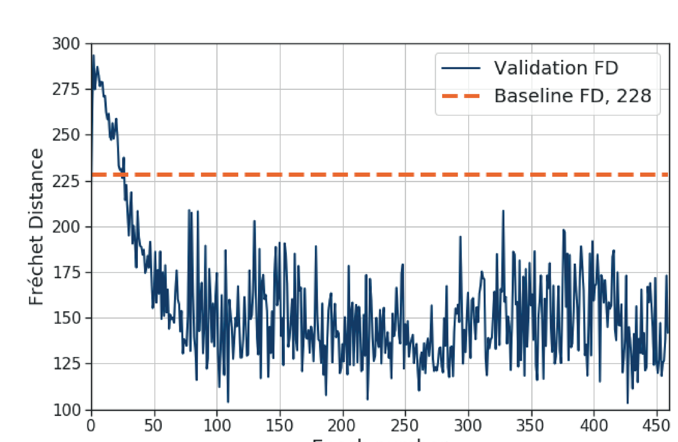

We demonstrate state-of-the-art results in multilabel land cover classification, and use the classifier to evaluate the performance of our model.

Model	$P(\%)$	$R(\%)$	F_1	F_2
S-CNN-All - used by authors of BigEarthNet	70	77	0.70	0.74
Regularized ResNet - our model	85	77	0.79	0.77

jev@cervest.earth

@ItsKoobik