Predicting Surface Elevation from a single RGB satellite image

Emmanouil Panagiotou
Georgios Chochlakis
Eleni Charou

The Idea

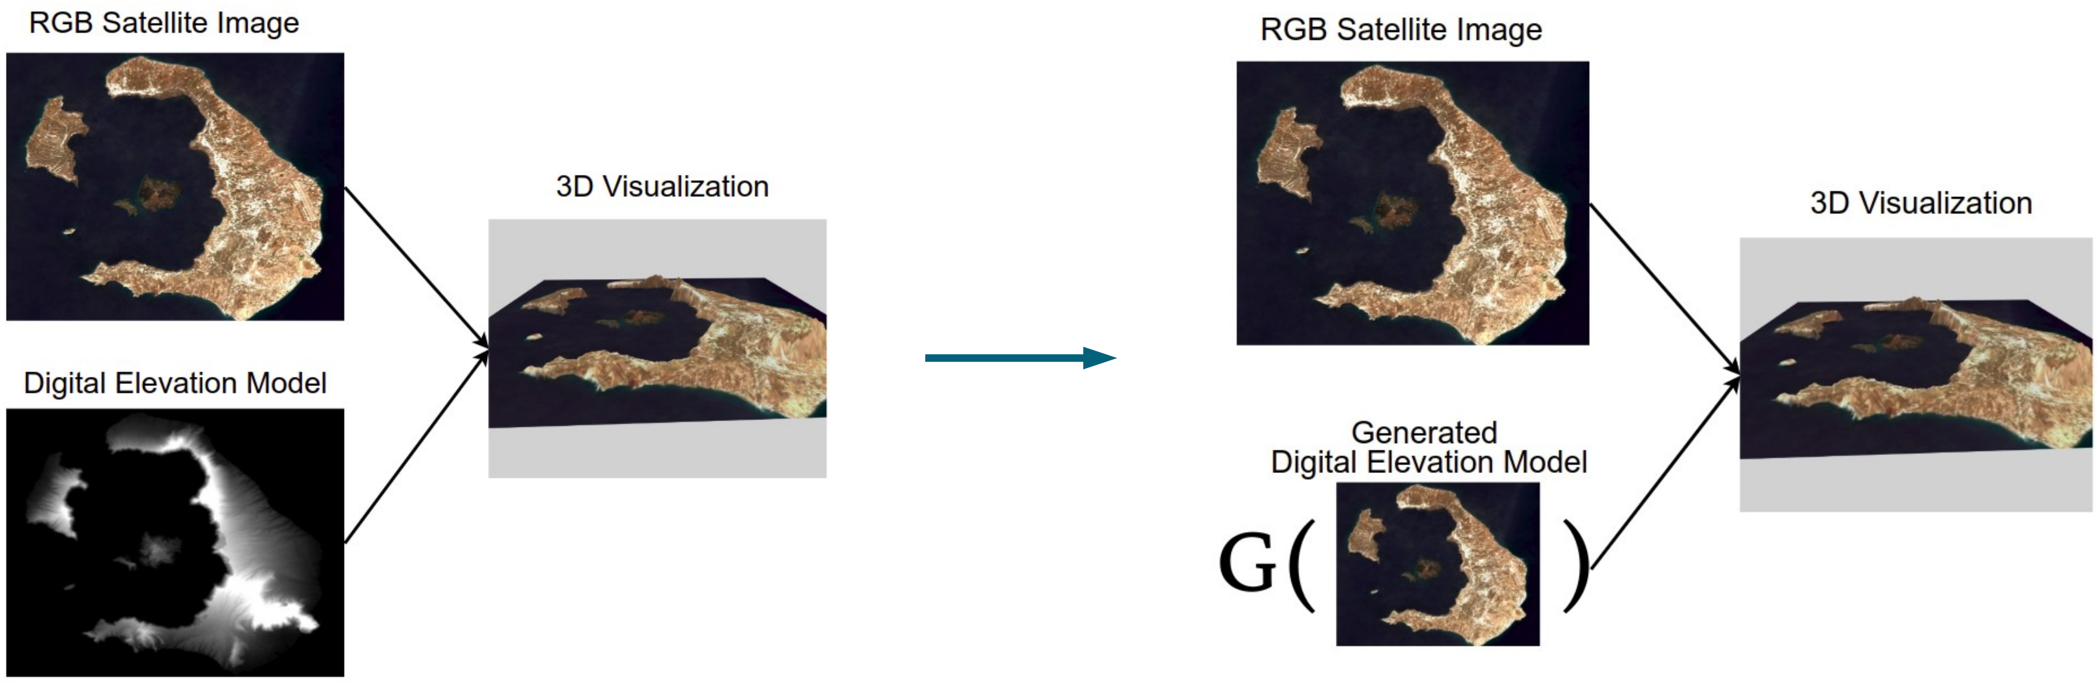

Use a Conditional Generative model, G, to predict the Digital Elevation Model corresponding to a satellite image

The Problem

Given a satellite image, x, we want to model the conditional probability of the DEM, Y, namely:

$$P(Y \mid X = x)$$

Background

- **Machine Learning**

  Algorithms that build a mathematical model based on training data in order to make predictions without being explicitly programmed for the task.

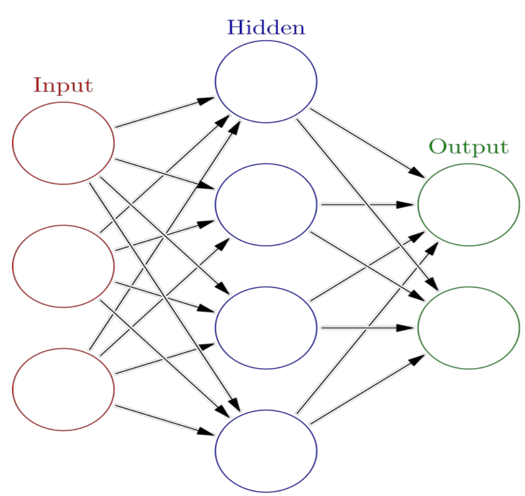

- **Discrete 2D Convolution**

$$[x*h][n_1,n_2] = \sum_{k_1}\sum_{k_2} x[k_1,k_2]h[n_1-k_1,n_2-k_2]$$

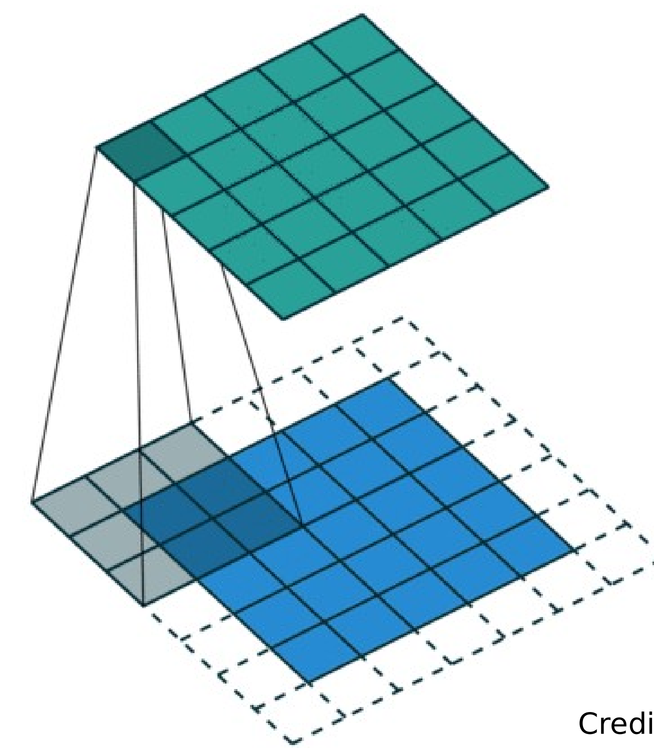

Credits: Sumit Saha

(Mildly) Related Work

IM2HEIGHT: Lichao Mou and Xiao Xiang Zhu

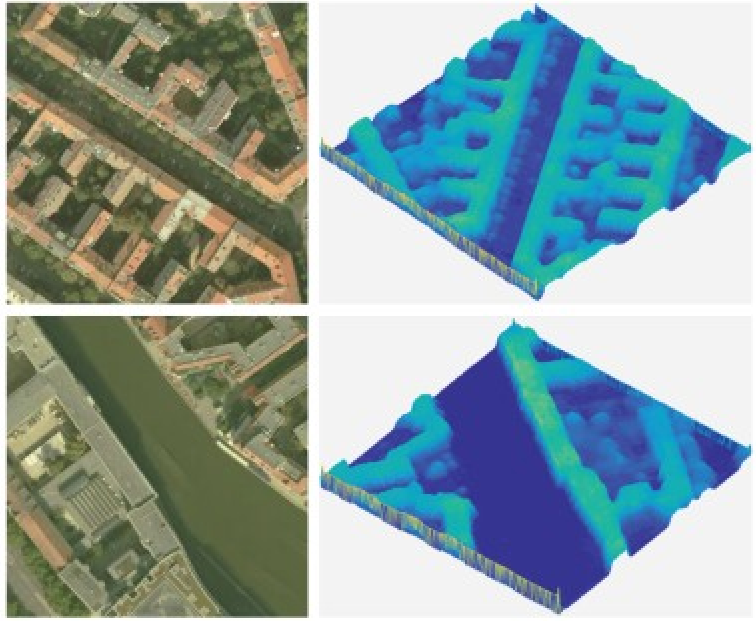

Point cloud visualization of the height estimates for two selected examples.

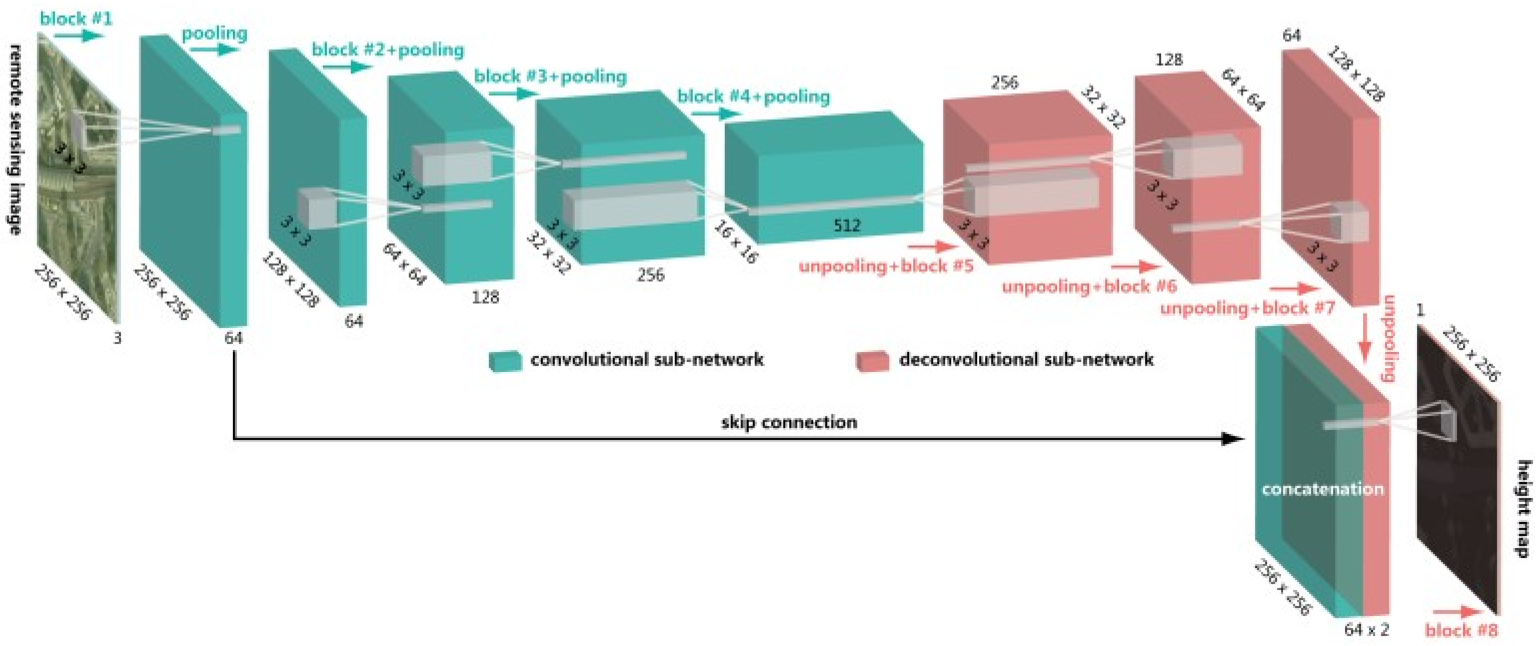

Final Network Architecture

(Mildly) Related Work

Adjusted IM2HEIGHT architecture for our dataset

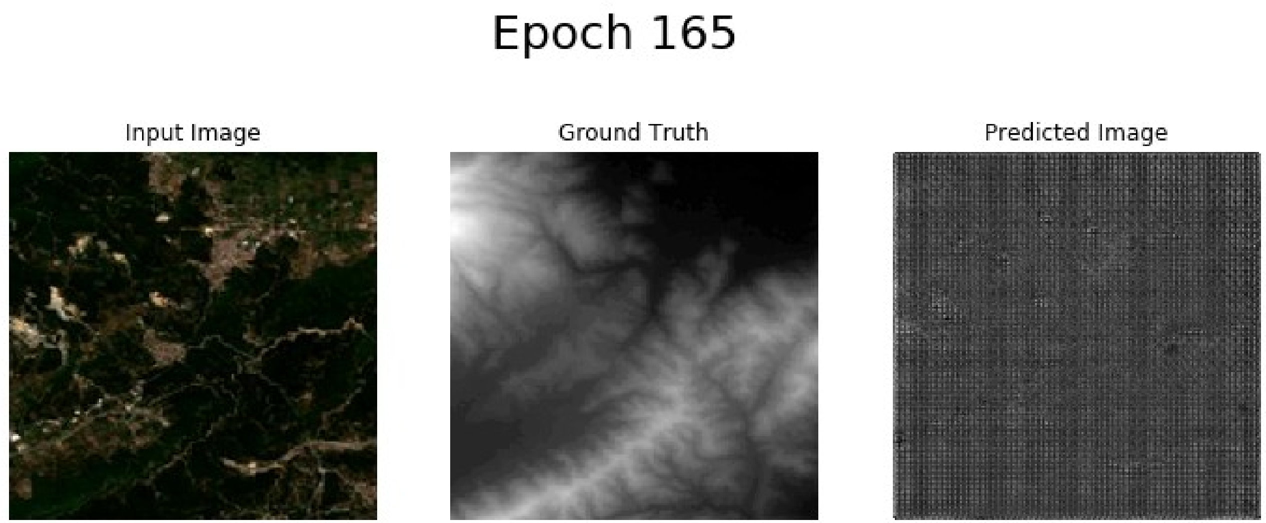

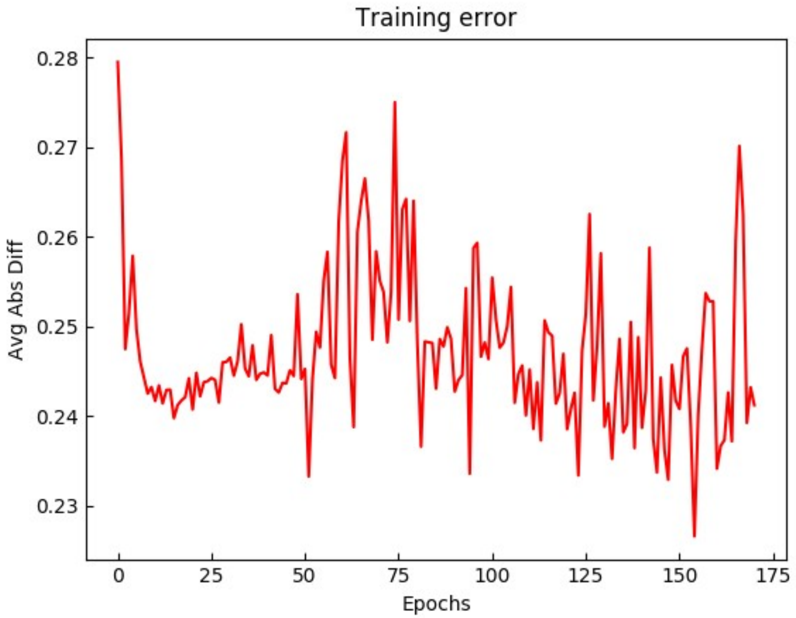

Our Solution

Conditional Generative Adversarial Networks (cGANs)

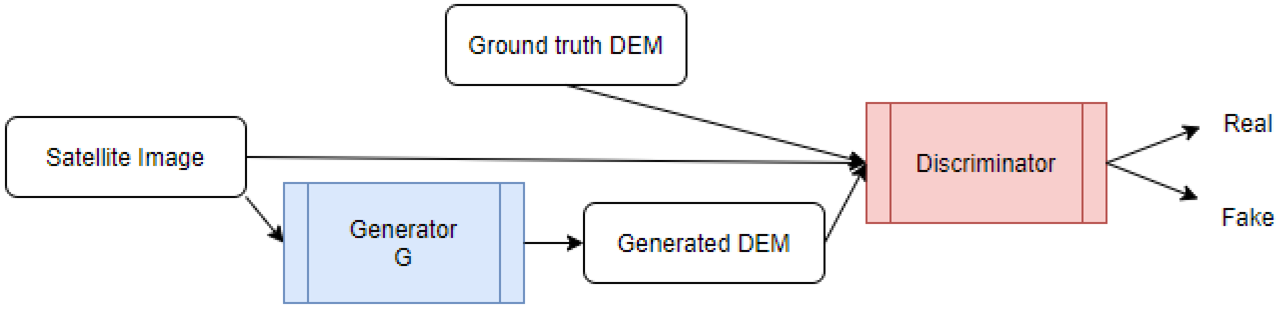

- Generator: generate DEMs, try to fool the Discriminator
- Discriminator: try to distinguish between real and generated DEMs
- Train them against each other

The Implementation (pix2pix)

- Generator G

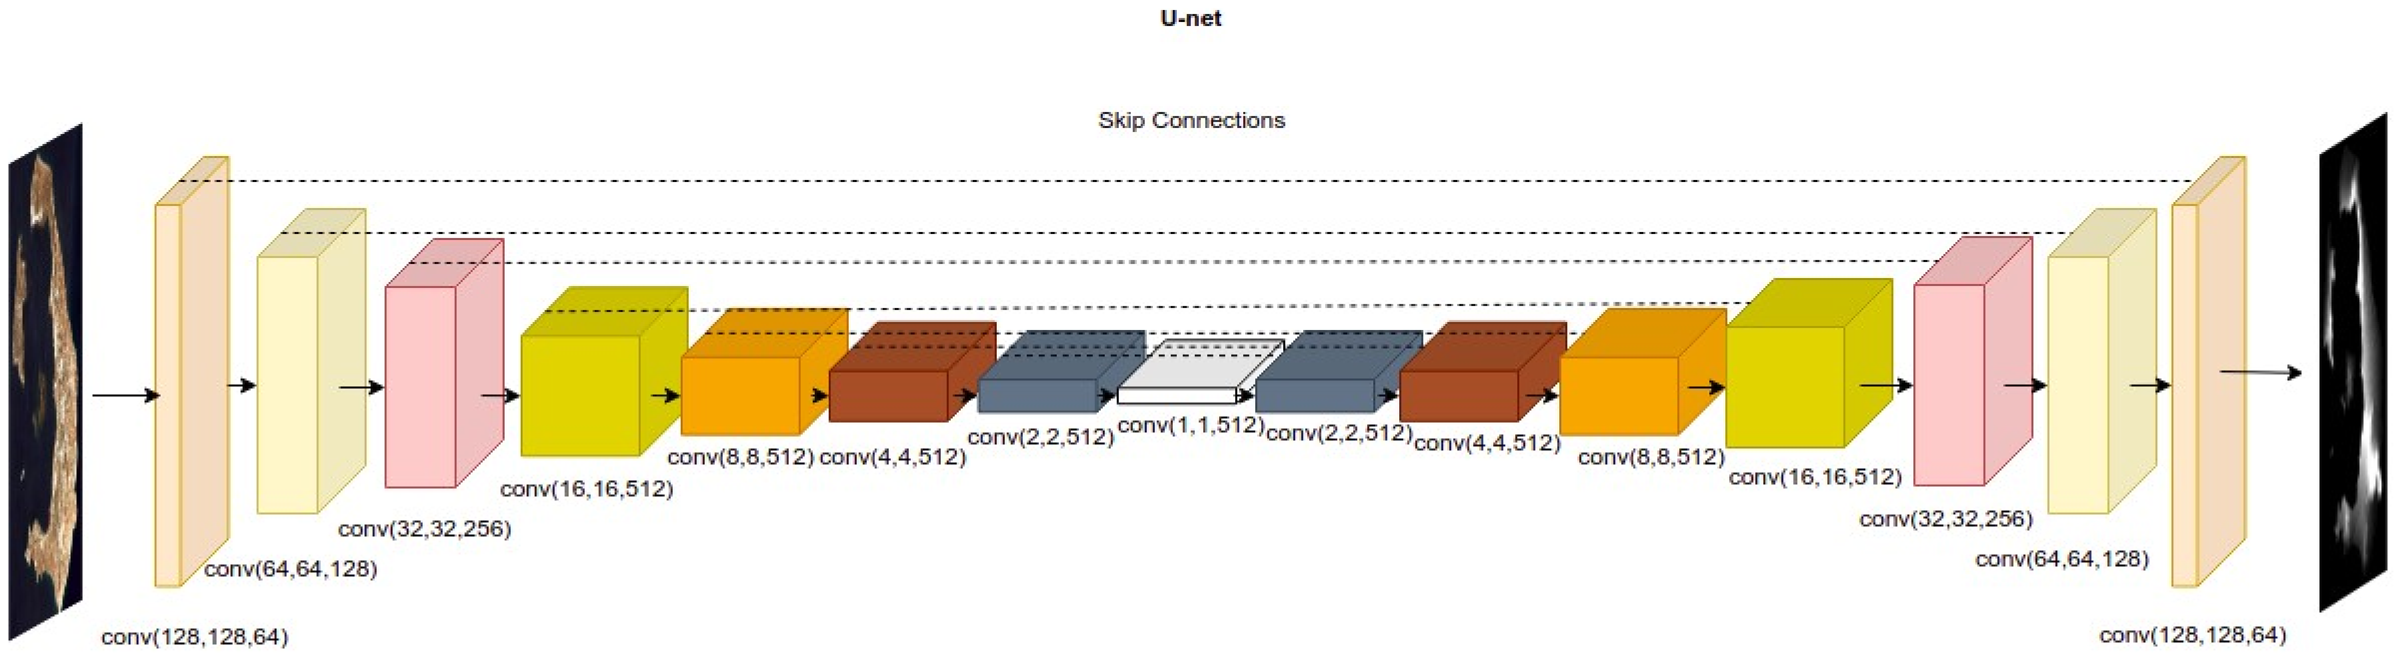

Skip connections preserve global structure (same in satellite image and DEM)

The Implementation (pix2pix)

- Discriminator

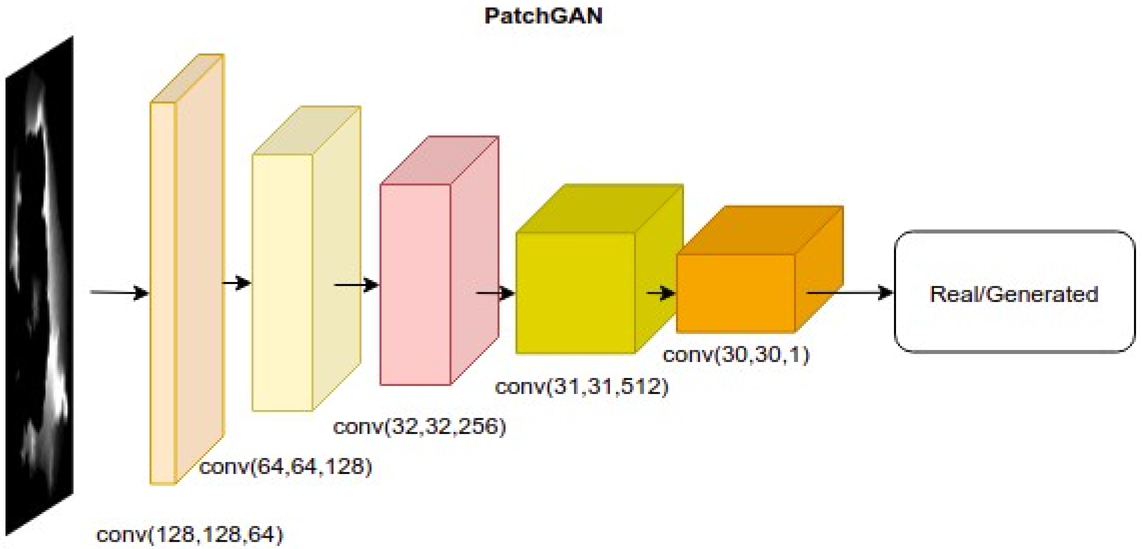

Assesses DEM’s plausibility patch-by-patch

Dataset Construction

- The satellite image is comprised of the bands [B4, B3, B2] (Copernicus Sentinel-2)
- The difference in resolution is handled by the API

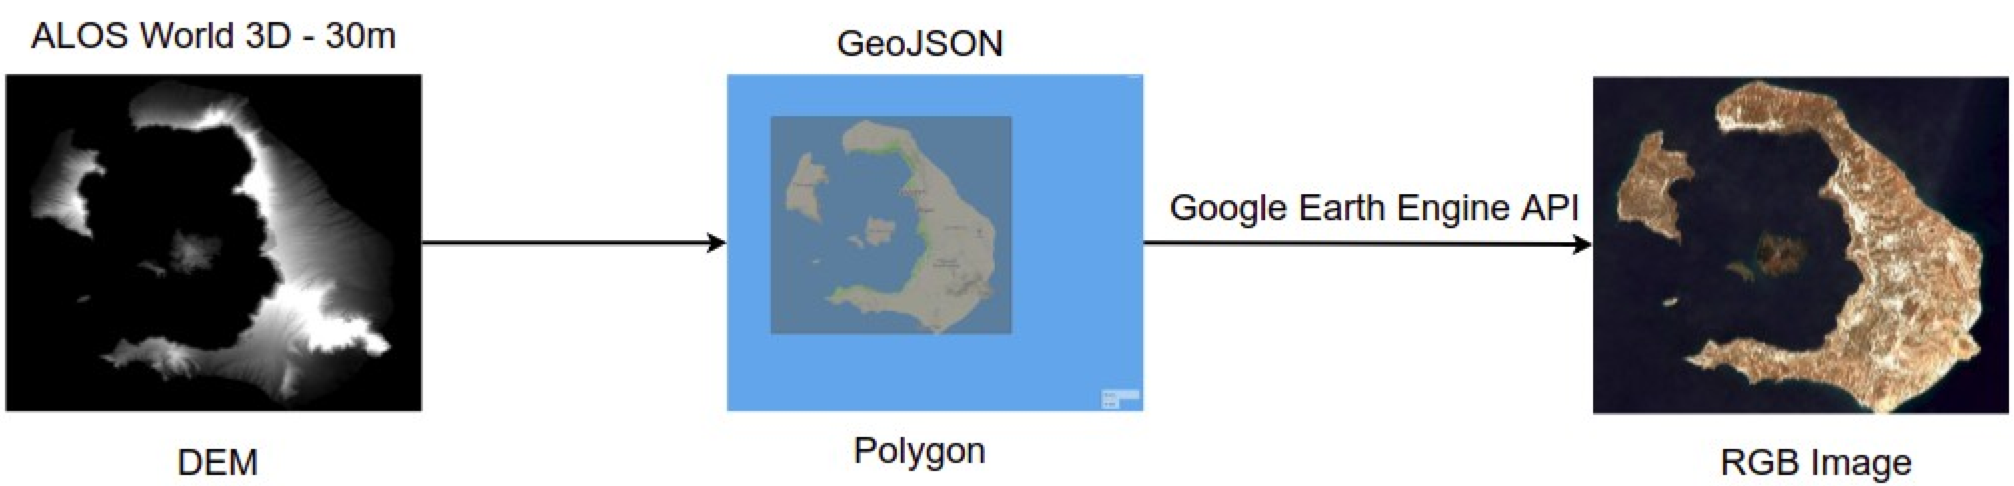

Our Results

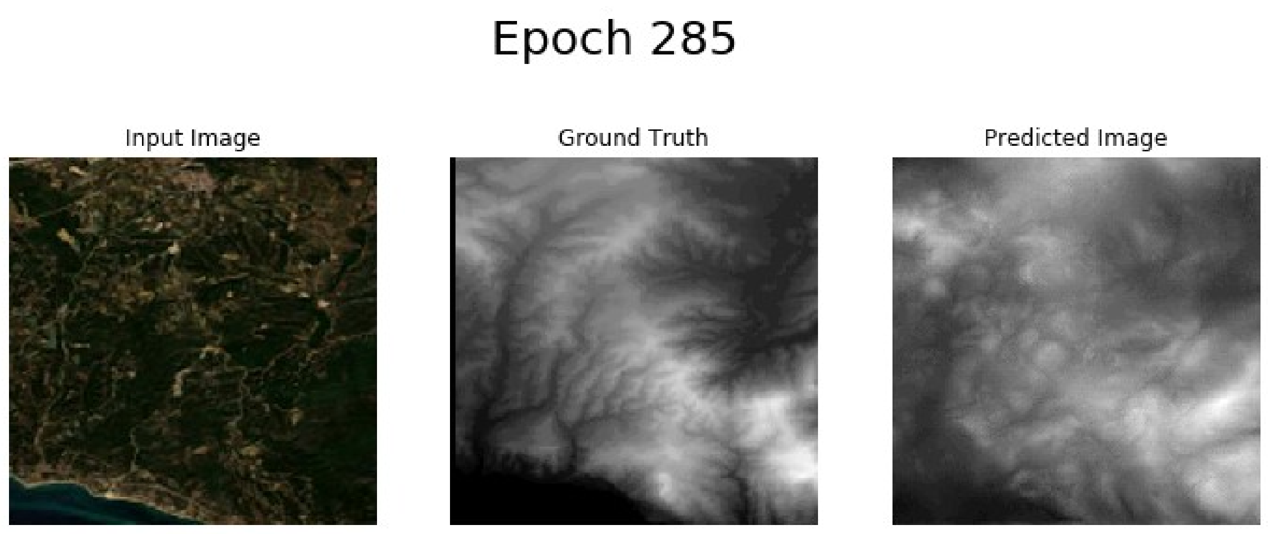
Epoch 285
Input Image
Ground Truth
Predicted Image

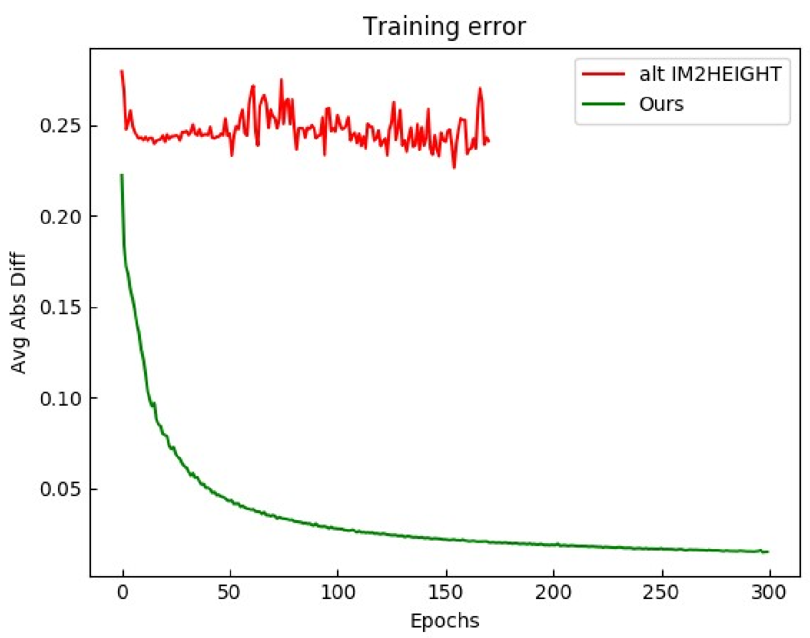
Training error
alt IM2HEIGHT
Ours
Avg Abs Diff
0.25
0.20
0.15
0.10
0.05
0
50
100
150
200
250
300
Epochs

Our Results

3D Visualization

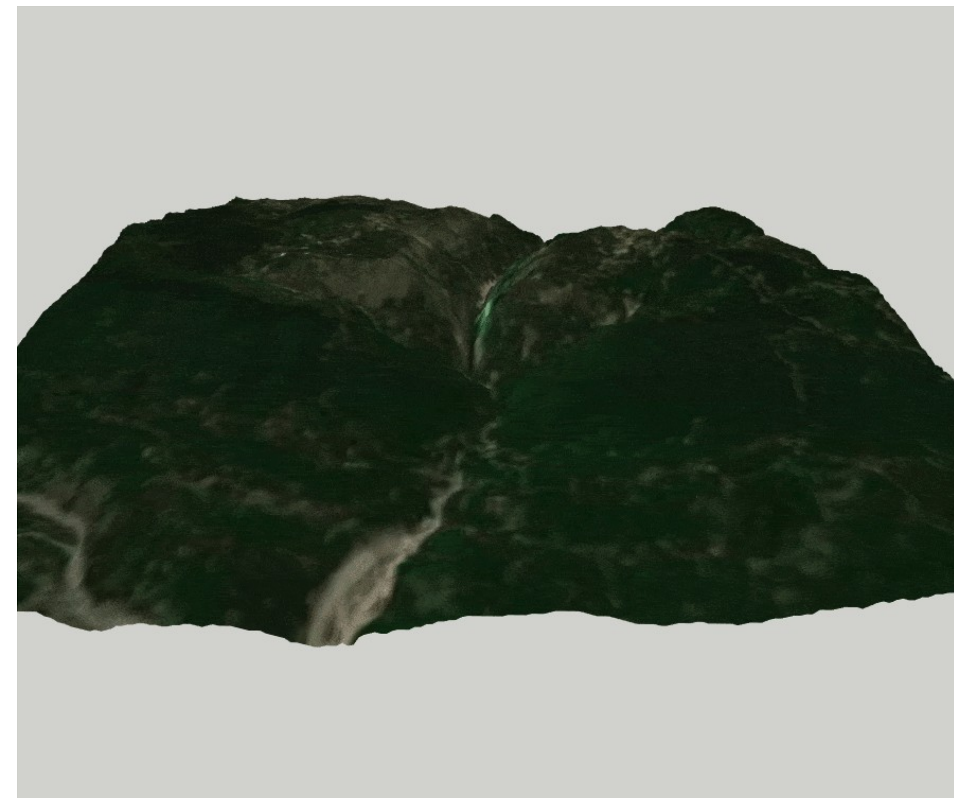

Input Image

Predicted Image

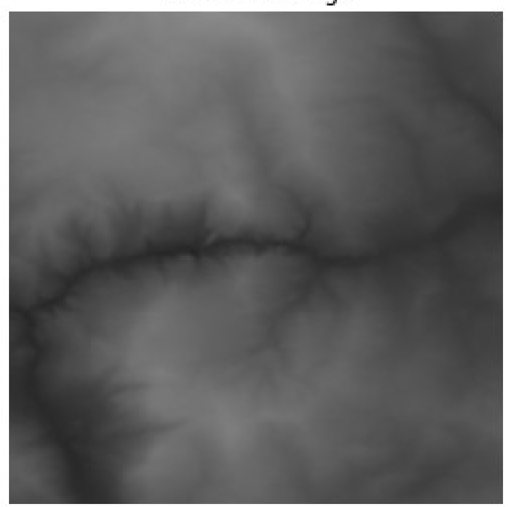

Our Results

3D
Visualization

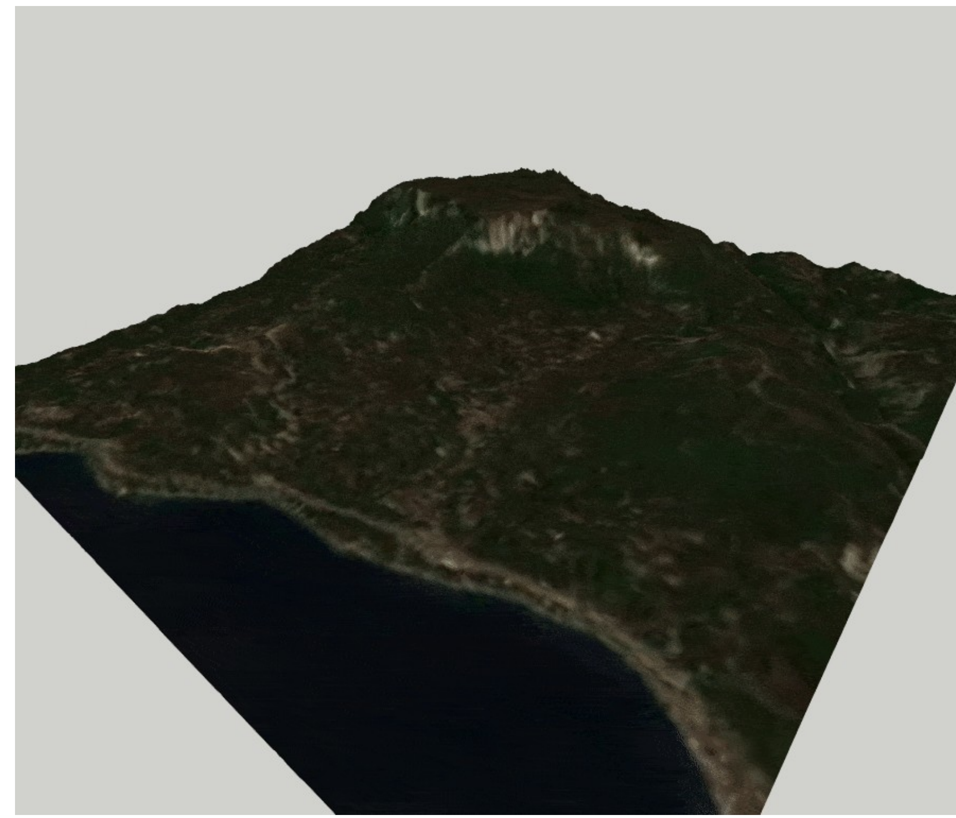

Input Image

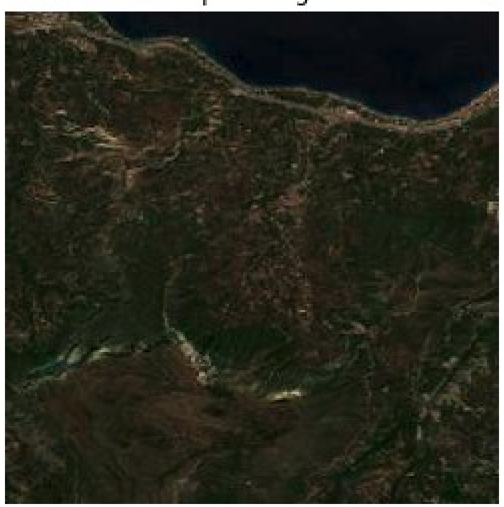

Predicted Image

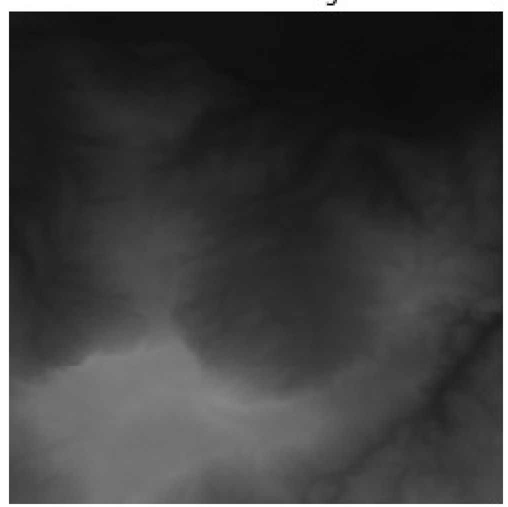

Possible Applications

- On-line Mapping
- Old aerial imagery reconstruction
- Applications in virtual environment rendering

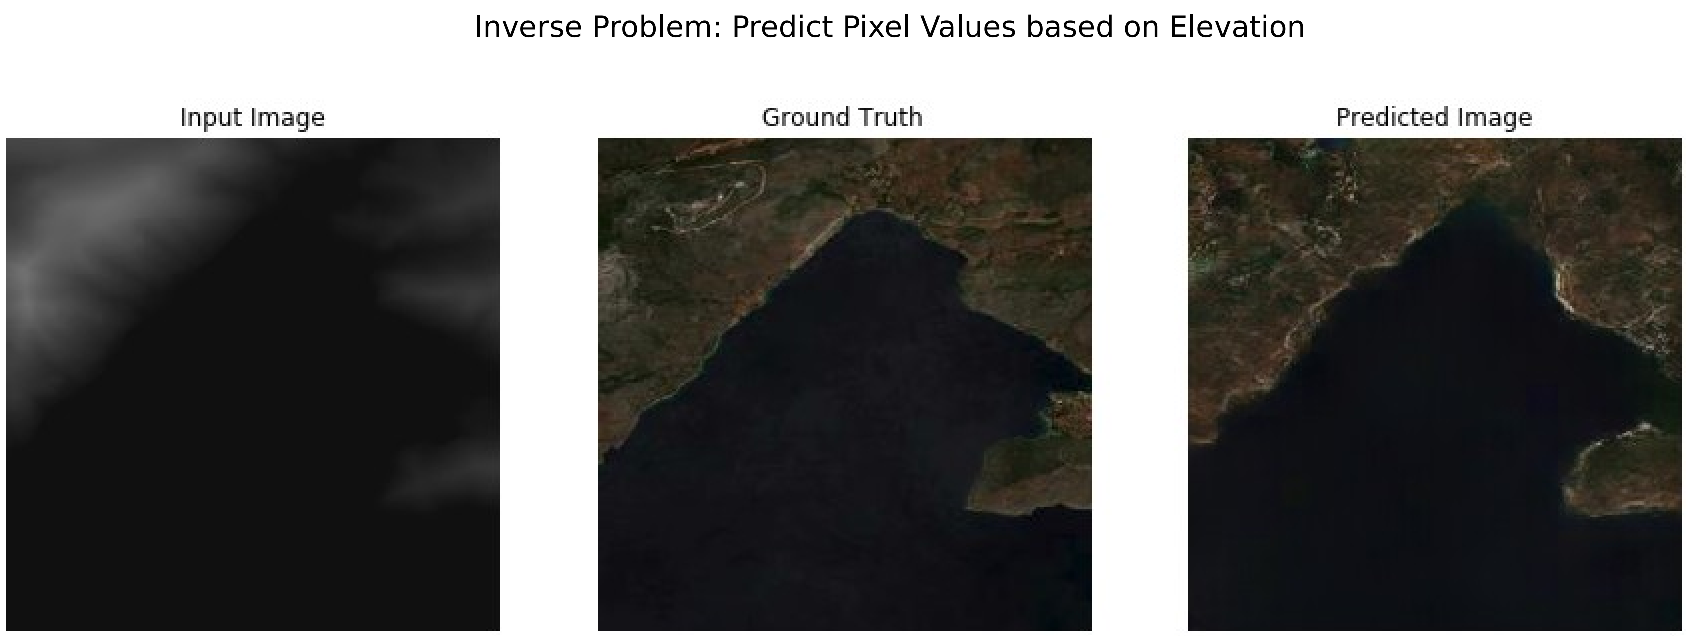

Possible Applications

Random Terrain Generation using the inverse model

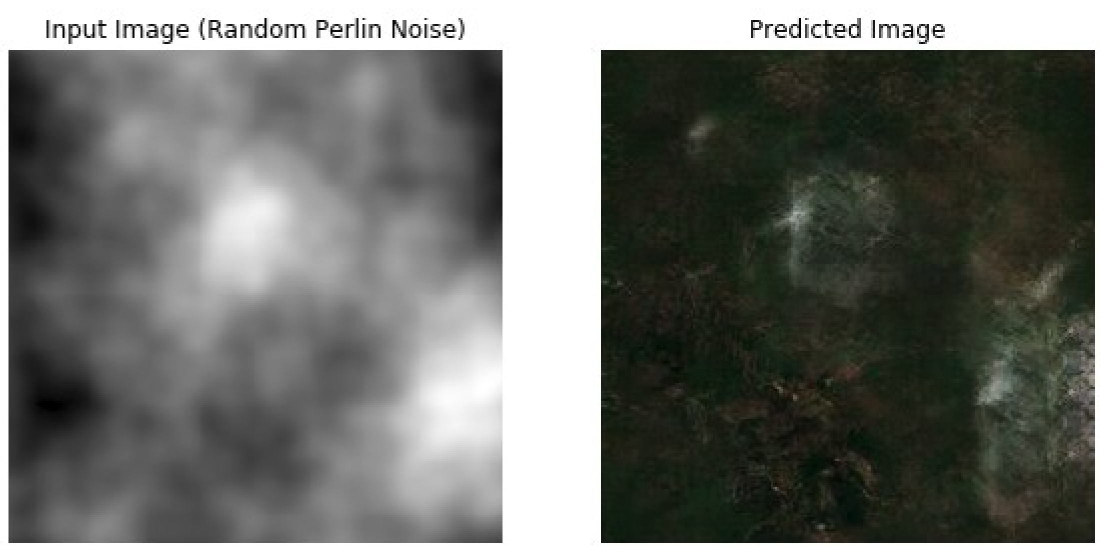

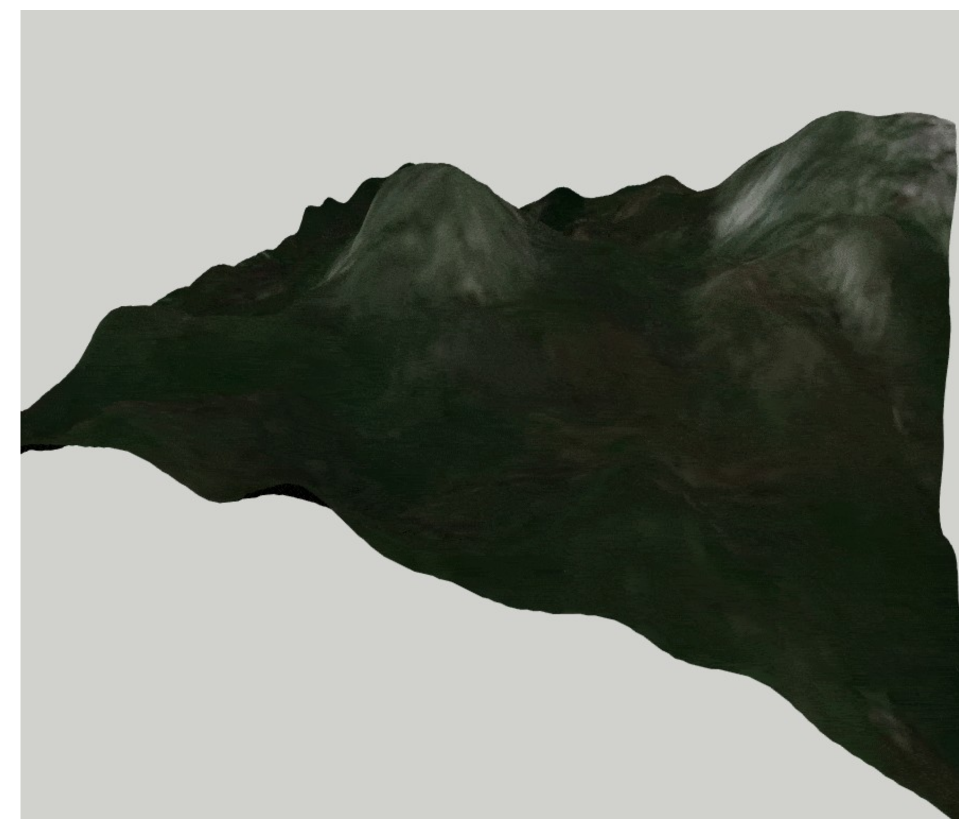

Limitations

- GANs are still an emerging and active area of research
- Predictions on never-before-seen data present a lot of variance
- Predictions of elevation are relative within a satellite image, not absolute.
- Hardware restrictions :(

Difficult Examples

Difficult Examples

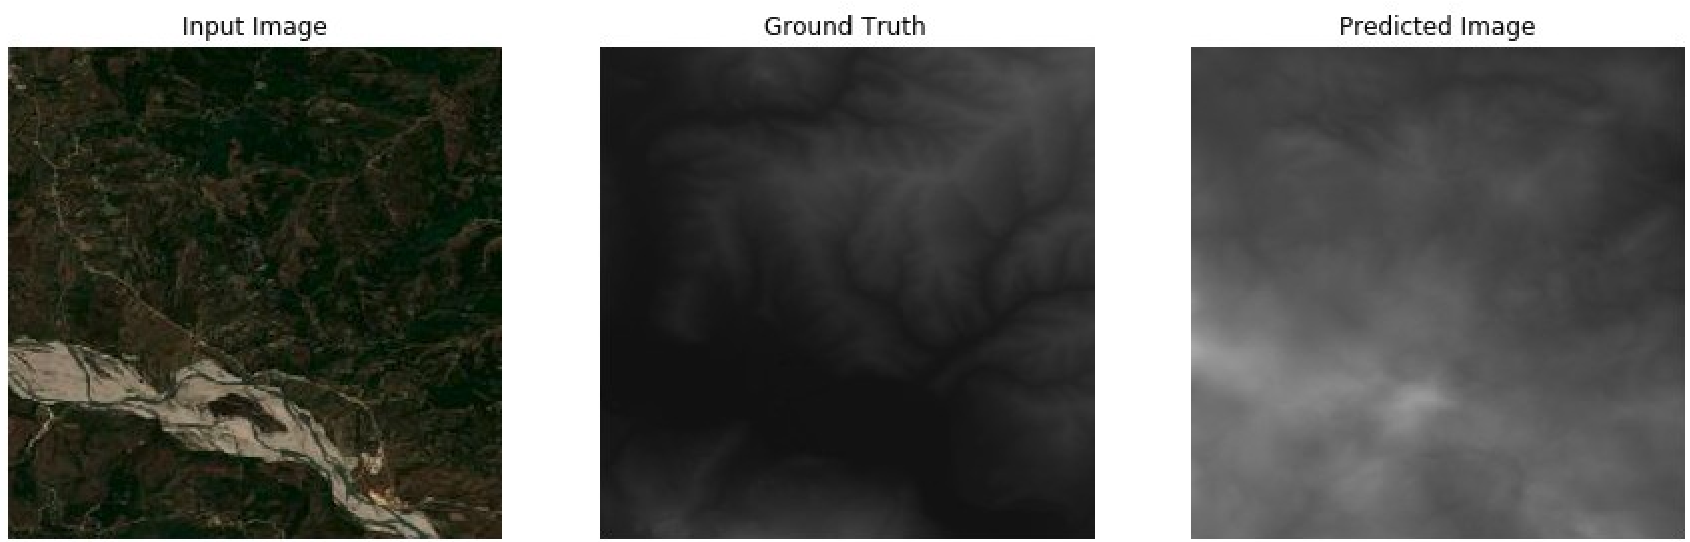

Solution?

More Data !

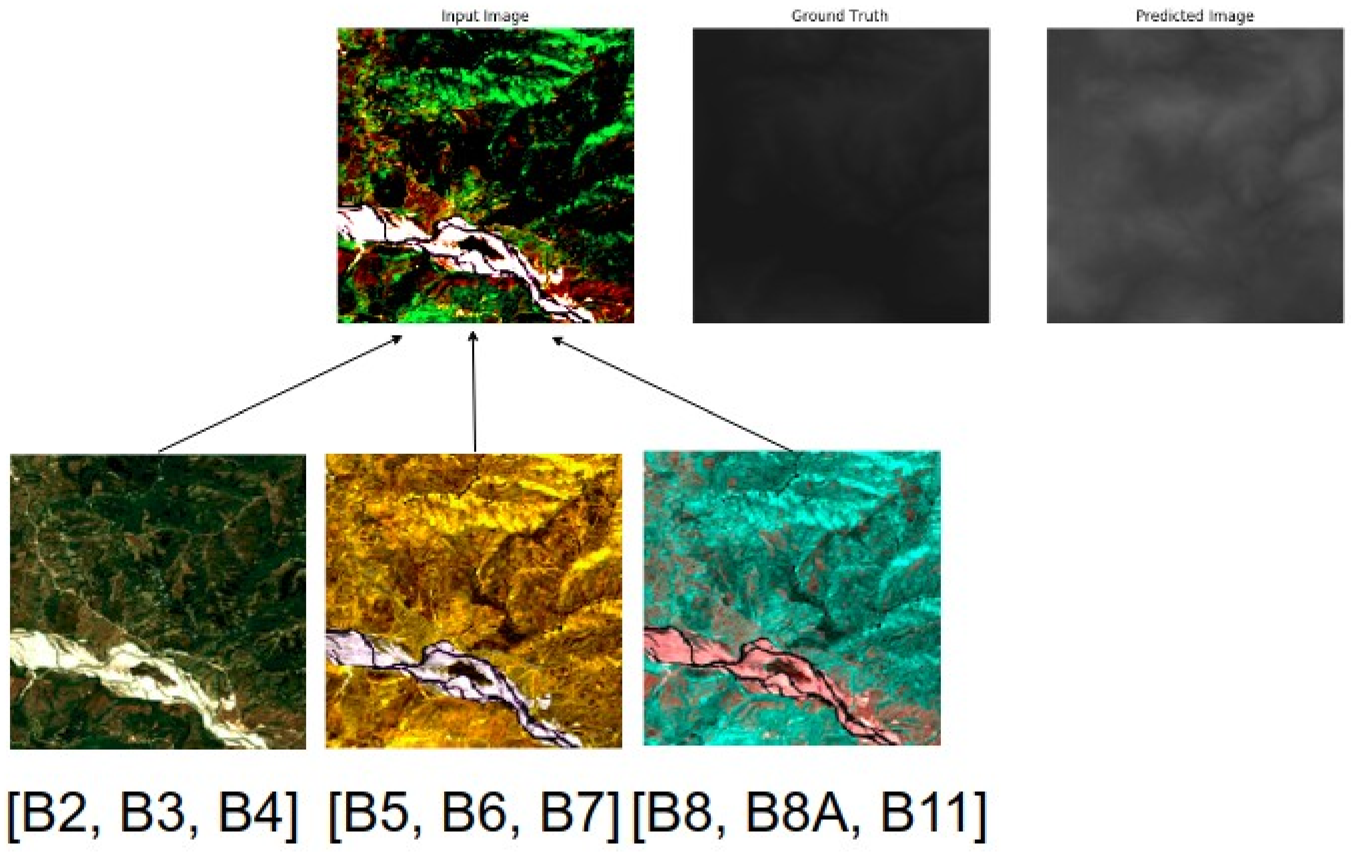

Future Work

- The process of synthesizing the data of our dataset itself contains many sources of error (precision, correspondence, lighting, …)
- We can mitigate their effects by working directly with the exact data captured by satellites, drones, …
- Depth Maps

References

- Isola, Phillip and Zhu, Jun-Yan and Zhou, Tinghui and Efros, Alexei A. Image-To-Image Translation With Conditional Adversarial Networks, in CVPR, 2017.
- Goodfellow, Ian and Pouget-Abadie, Jean and Mirza, Mehdi and Xu, Bing and Warde-Farley, David and Ozair, Sherjil and Courville, Aaron and Bengio, Yoshua. Generative Adversarial Nets, in NIPS, 2014.
- Lichao Mou and Xiao Xiang Zhu. **IM2HEIGHT:** Height Estimation from Single Monocular Imagery via Fully Residual Convolutional-Deconvolutional Network, arXiv:1802.10249, 2017.
- Salimans, Tim and Goodfellow, Ian and Zaremba, Wojciech and Cheung, Vicki and Radford, Alec and Chen, Xi and Chen, Xi. Improved Techniques for Training GANs, in NIPS, 2016.
- Sanjeev Arora and Yi Zhang. Do GANs actually learn the distribution? An empirical study, arXiv:1706.08224, 2018.
- Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D. & Moore, R. Google Earth Engine: Planetary-scale geospatial analysis for everyone. Remote Sensing Of Environment, 2017.