# Specialized Machine Learning Models for Satellite Imagery

Hamed Alemohammad[1], Milan Cvitkovic[2]

[1]) Radiant Earth Foundation
hamed@radiant.earth
[2]) California Institute of Technology
mcvitkov@caltech.edu

Satellite imagery is one of the only abundant and consistent sources of data about low–income parts of the world. While not a panacea [5], much useful work has been done using satellite imagery for projects relevant to the United Nations Sustainable Development Goals (SDGs), e.g. poverty prediction [6] or infrastructure assessment [9].

However, satellite data have many unique properties, and little work has been done to develop machine learning models to exploit/compensate for them. We argue that designing specialized machine learning models for satellite imagery is an academically interesting research problem with great potential impact.

Of course, the machine learning research community is unlikely to focus on any topic without easy access to data and concrete performance metrics. Thanks to data engineering efforts like those of the Radiant Earth Foundation, however, it has never been easier for practitioners to train machine learning models on satellite imagery. The data have been gathered, the data pipelines built, and the benchmarks set; it's time to build the specialized models.

## The Unique Properties of Satellite Imagery

There are 5 major properties of satellite imagery that distinguish them from general image data:

- The presence of reliable image metadata, such as precise geolocation of images on the Earth's surface, camera position and orientation, image acquisition time, and pixel resolution.
- The inherent time–series nature of satellite imagery, which consists of repeated images of the same location across time.
- Imagery that is captured at different wavelengths and modes, each with unique properties, e.g. optical vs. near–infrared or passive vs. active/radar.
- The presence of cloud occlusion (particularly in passive measurements), cloud shadows, and the ensuing illumination variability these both cause.
- Multiple resolutions of imagery for the same ground truth, and the frequent need to transfer models trained on one resolution to another.

The first three of these properties are more–or–less opportunities to be exploited. In the same ways that convolutional networks benefit from their inbuilt equivariance to translations and scaling, so too should models for satellite imagery benefit from being designed to be equivariant (in some sense) to different camera positions, image aquisition times, wavelengths, etc. Whether this equivariance is best accomplished by some form of regularization during training or by model architecture choices, as in e.g. models for multi–view image data [11], is an open question. Some work has begun to exploit the image properties in the third bullet above, including [3], [10], [4] and [9]. [8] and [1] give accessible explanations of these properties. [7] includes a good overview of different types of remote sensing metadata that could be combined with satellite imagery.

The issue of cloud occlusion is a major one, and is more of a problematic property of satellite imagery than an opportunity. Nevertheless, it is a more circumscribed problem than the general problem of occlusion in computer vision, and thus is potentially more addressable.

Finally, the issue of (sometimes vastly) differing resolutions between and within datasets is a mixed blessing. The reader will be unsurprised to learn that the wealthier the region, the higher the resolution of its satellite imagery. So where there are regions with both high– and low–resolution imagery, there is an interesting opportunity for transfer learning. However, the domain shift between satellite imagery of wealthy and poorer regions is substantial, not just in terms of image resolution but also image content. This presents a major modeling challenge, as has been thoroughly explored in [5].

As an aside, we should also mention the perennial problem of a lack of labeled data for satellite imagery of poorer regions. This is not an unique property of satellite imagery, but it is one that must be overcome, with e.g. transfer learning approaches, for satellite imagery models to be relevant to SDGs.

## Concrete Problems and Datasets

The Radiant Earth Foundation will be publishing a suite of open source benchmark datasets for assessing the performance of satellite imagery models at `mlhub.earth/` in the near future. Until then, useful machine learning tasks on satellite imagery for which there are open source training data include:

- Crop suitability, weather, air quality, and population prediction (data at Radiant Earth API below)
- Cloud detection, segmentation, and in–painting with generative models (data at any API below)
- Transfer learning from high– to low–resolution data and between data of differing wavelengths and modes (data at any API below)
- Building and road detection (data at `spacenetchallenge.github.io`)

Training imagery for these tasks can be accessed via a number of well–engineered APIs, including those of Radiant Earth (`doc.radiant.earth/` for imagery, `doc-api.radiant.earth/` for labels, `github.com/radiantearth/radiantearth-python-client` for a python client, and `help.radiant.earth/` for tutorials), Google Earth Engine ( `earthengine.google.com/`), and Descartes Labs (`www.descarteslabs.com/`). There is also an ongoing community effort to standardize data catalogues for satellite imagery and target labels [2].

# References

[1] State of Satellite Imagery. URL https://landscape.satsummit.io/.

[2] MLHubEarth, August 2018. URL https://github.com/MLHubEarth/training-data-stac-spec.

[3] Patrick Doupe, Emilie Bruzelius, James Faghmous, and Samuel G. Ruchman. Equitable Development Through Deep Learning: The Case of Subnational Population Density Estimation. In *Proceedings of the 7th Annual Symposium on Computing for Development*, ACM DEV '16, pages 6:1–6:10, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4649-8. doi: 10.1145/3001913.3001921. URL http://doi.acm.org/10.1145/3001913.3001921.

[4] P. Gentine and S. H. Alemohammad. RSIF (Reconstructed Solar Induced Fluorescence): a machine-learning vegetation product based on MODIS surface reflectance to reproduce GOME-2 solar induced fluorescence. *Geophysical Research Letters*, mar 2018. ISSN 00948276. doi: 10.1002/2017GL076294. URL http://doi.wiley.com/10.1002/2017GL076294.

[5] Andrew Head, Mlanie Manguin, Nhat Tran, and Joshua E. Blumenstock. Can Human Development Be Measured with Satellite Imagery? In *Proceedings of the Ninth International Conference on Information and Communication Technologies and Development*, ICTD '17, pages 8:1–8:11, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-5277-2. doi: 10.1145/3136560.3136576. URL http://doi.acm.org/10.1145/3136560.3136576.

[6] Neal Jean, Marshall Burke, Michael Xie, W. Matthew Davis, David B. Lobell, and Stefano Ermon. Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301):790–794, August 2016. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.aaf7894. URL http://science.sciencemag.org/content/353/6301/790.

[7] Ron Mahabir, Arie Croitoru, Andrew T. Crooks, Peggy Agouris, and Anthony Stefanidis. A critical review of high and very high-resolution remote sensing approaches for detecting and mapping slums: Trends, challenges and emerging opportunities. *Urban Science*, 2(1), 2018. ISSN 2413-8851. doi: 10.3390/urbansci2010008. URL http://www.mdpi.com/2413-8851/2/1/8.

[8] Anne Hale Miglarese. Observing the Earth, Fueling Global Development Solutions, April 2018. URL https://medium.com/radiant-earth-insights/observing-the-earth-fueling-global-development-solutions-1c69fd5632bc.

[9] Barak Oshri, Annie Hu, Peter Adelson, Xiao Chen, Pascaline Dupas, Jeremy Weinstein, Marshall Burke, David Lobell, and Stefano Ermon. Infrastructure quality assessment in africa using satellite imagery and deep learning. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '18, pages 616–625, New York, NY, USA, 2018. ACM. ISBN 978-1-4503-5552-0. doi:

10.1145/3219819.3219924. URL `http://doi.acm.org/10.1145/3219819.3219924`.

[10] Anthony Perez, Christopher Yeh, George Azzari, Marshall Burke, David B. Lobell, and Stefano Ermon. Poverty Prediction with Public Landsat 7 Satellite Imagery and Machine Learning. abs/1711.03654, 2017. URL `https://arxiv.org/abs/1711.03654`.

[11] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik G. Learned-Miller. Multi-view Convolutional Neural Networks for 3d Shape Recognition. *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 945–953, 2015.