EGU22-8852
https://doi.org/10.5194/egusphere-egu22-8852
EGU General Assembly 2022

Time-dependent Hillshades: Dispelling the Shadow Curse of Machine Learning Applications in Earth Observation

Freddie Kalaitzis[1,5], Gonzalo Mateo-Garcia[2,5], Kevin Dobbs[5], Dolores Garcia[5], Jason Stoker[4], and Giovanni Marchisio[3]
[1]University of Oxford
[2]University of Valencia
[3]Planet
[4]USGS
[5]Frontier Development Lab

We show that machine learning models learn and perform better when they know where to expect shadows, through hillshades modeled to the time of imagery acquisition.

Shadows are detrimental to all machine learning applications on satellite imagery. Prediction tasks like semantic / instance segmentation, object detection, counting of rivers, roads, buildings, trees, all rely on crisp edges and colour gradients that are confounded by the presence of shadows in passive optical imagery, which rely on the sun's illumination for reflectance values.

Hillshading is a standard technique for enriching a mapped terrain with relief effects, which is done by emulating the shadow caused by steep terrain and/or tall vegetation. A hillshade that is modeled to the time of day and year can be easily derived through a basic form of ray tracing on a Digital Terrain Model (DTM) (also known as a bare-earth DEM) or Digital Surface Model (DSM) given the sun's altitude and azimuth angles. In this work, we use lidar-derived DSMs. A DSM-based hillshade conveys a lot more information on shadows than a bare-earth DEM alone, namely any non-terrain vertical features (e.g. vegetation, buildings) resolvable at a 1-m resolution. The use of this level of fidelity of DSM for hillshading and its input to a machine learning model is novel and the main contribution of our work. Any uncertainty over the angles can be captured through a composite multi-angle hillshade, which shows the range where shadows can appear throughout the day.

We show the utility of time-dependent hillshades in the daily mapping of rivers from Very High Resolution (VHR) passive optical and lidar-derived terrain data [1]. Specifically, we leverage the acquisition timestamps within a daily 3m PlanetScope product over a 2-year period. Given a datetime and geolocation, we model the sun's azimuth and elevation relative to that geolocation at that time of day and year. We can then generate a time-dependent hillshade and therefore locate shadows in any given time within that 2-year period. In our ablation study we show that, out of all the lidar-derived products, the time-dependent hillshades contribute a 8-9% accuracy improvement in the semantic segmentation of rivers. This indicates that a semantic segmentation

machine learning model is less prone to errors of commission (false positives), by better disambiguating shadows from dark water.

Time-dependent hillshades are not currently used in ML for EO use-cases, yet they can be useful. All that is needed to produce them is access to high-resolution bare-earth DEMs, like that of the US National 3D Elevation Program covering the entire continental U.S at 1-meter resolution, or creation of DSMs from the lidar point cloud data itself. As the coverage of DSM and/or DEM products expands to more parts of the world, time-dependent hillshades could become as commonplace as cloud masks in EO use cases.

[1] Dolores Garcia, Gonzalo Mateo-Garcia, Hannes Bernhardt, Ron Hagensieker, Ignacio G. Lopez-Francos, Jonathan Stock, Guy Schumann, Kevin Dobbs and Freddie Kalaitzis Pix2Streams: Dynamic Hydrology Maps from Satellite-LiDAR Fusion. *AI for Earth Sciences Workshop, NeurIPS 2020*