


HOW DEEP LEARNING ENABLES LARGE-SCALE ANALYSIS

Use of AI to detect rooftop solar potential



Global climate change and, in some areas, the lack of a cheap and reliable energy supply are contributing to the expansion of photovoltaics, both regionally and worldwide. In addition to ground-mounted and agrisolar systems, rooftop systems are considered a promising factor in the generation of solar energy. Besides being subject to less complex planning processes and legal frameworks, they are also affordable, reliable and make use of existing rooftop infrastructure.

Installation of a photovoltaic system is usually preceded by an estimate of the solar potential. In terms of rooftop systems, this analysis is mainly performed individually for each building. This requires the measurement of various parameters such as roof pitch, orientation

and any roof superstructures, either manually or as part of a Lidar project. For large target areas, the estimation work quickly becomes infeasible, either due to the huge manual effort involved or the high costs associated with high-resolution Lidar measurements. By applying AI to digital orthoimagery, the German geointelligence company Deepeer.technology has developed a method to overcome these bottlenecks and determine the net solar potential for vast regions and even entire countries.

Multiple models for a single dataset

By using proprietary AI methods in combination with open-source machine learning libraries, the method applies multiple neural networks in sequence, each of which makes its own contribution to the process. At the beginning, a building detection network evaluates orthoimagery with a resolution of 20cm and extracts the buildings as a binary raster (Figure 1). However, in order to determine not only the pure building area but also the actual usable area, two further steps must be taken: 1) the roof surface of the building must be broken down into its individual convex surfaces in order to consider each roof surface individually, and 2) areas that cannot be used due to obstructions need to be excluded from the calculation. In order to accomplish this, two further neural networks are applied, one handling the decomposition into linear surfaces and the other extracting roof superstructures as well as other non-usable roof areas from the orthophotos.



Figure 1: Multiple buildings extracted from orthoimagery as a binary raster. This intermediate result shows vectorized polygons with the detected buildings crosshatched in blue.

After the building detection network, the linear decomposition network and the superstructure detection network have been applied, knowledge is gathered about the total roof area and the non-usable surface. In order to obtain the net useful area of the roofs, in the next step only the difference between these quantities needs to be calculated. Thus, the intermediate result after applying the three above-mentioned networks is a layer containing the net useful area for each building in the target area.

Not just a matter of ‘where’

At this point in the process, knowledge about the location of the buildings as well as the net usable area has been generated. However, it is not yet clear how useful these roof areas will be in reality. To calculate the actual solar potential, it is necessary to include additional information about elevation. This is estimated by another deep learning model in a process called monocular depth estimation. By processing digital orthophotos together with a digital surface model as the input data, this model learns to estimate the height and 3D shape of each building. As a next step, the elevation estimating network is applied to the target area. This generates crucial information about the height and slope of individual roof surfaces (Figure 2) as well as about potential shading by other objects such as neighbouring buildings or trees. If this data is then combined with the information about the roof orientation towards the sun, the solar potential can be determined fairly accurately based on the angle of inclination and azimuth of the roof surfaces.

In detail, this is achieved by calculating for each roof surface the percentage yield deviation from the optimum orientation (south) and ideal tilt angle (0 degrees). Then, based on the size of the percentage deviation, the degree of shading and the nominal power of a standard solar module, the potential net useful power of each roof surface is determined (Figure 3). Therefore, the final result shows not only the theoretically usable area, but also the effective output for each roof surface.



Figure 2: Output of the elevation estimating network showing the digital elevation model obtained, with a scale from very high (red) to very low (blue).

Benefits and limitations

The great advantage of this method is that it can be applied to an area of any size. Furthermore, the input data is readily available at large scales, which might not be the case for high-resolution Lidar measurements. This allows larger areas such as provinces and even entire countries to be computed in very little time (Figure 4). Depending on the computing capacity and the technology used, the solution can analyse several square kilometres per minute. This creates a considerable spatial as well as temporal advantage compared to traditional methods such as manual annotations and calculations, and a cost several orders of magnitude cheaper than a Lidar project. On the other hand, of course, this method is not completely error-free. Due to the dependence on no less than four neural networks, it cannot be ruled out that objects will not be recognized correctly. This applies to individual buildings as well as single roof surfaces, superstructures and shading. However, since the error rate is comparable to that of a human GIS annotator, the percentage of errors can be regarded as negligible, especially considering the scale and the advantages mentioned above.

This method offers benefits for a wide variety of organizations, ranging from companies in the energy industry to public administrative bodies. At a time when there is often little to no inventory data available on solar potential, the main advantage is that this method generates key metrics that provide a quick situational overview of the buildings in a particular area as the basis for decision-making. Moreover, the intersection with commercial register data also allows energy suppliers to address building owners more specifically. In addition to that, there are opportunities for more targeted grid-expansion planning to further drive the expansion of renewable energies.



Figure 3: The same buildings analysed for their solar potential, with an efficiency scale from yellow (high potential) to brown (low potential). This shows that most of the really useful areas are oriented to the south. As expected, areas that are more inclined to the north have poorer yields.

Conclusion

Large-scale geospatial analysis is often associated with considerable effort and high costs. AI can help limit both the effort and the cost involved. Especially in applications that are of strong urgency due to their political or societal necessity, intelligent algorithms can be a decisive accelerator. In the case of the expansion of solar energy systems, deep learning algorithms show that the rooftop solar potential in entire countries can be analysed in just a few days. The geointelligence method presented above is not impeccable. However, as a complement to existing processes, it offers a way in which the geospatial industry can make greater use of the acceleration opportunities provided by AI, machine learning and automation to play its part in helping to meet climate-related and energy-related policy goals.



Figure 4: The computed solar potential for each building in the German district of Aachen, a district with an area of around 550km².