

Effective, low-cost methods of applying computer vision to public Earth observation data

Michael Evans¹ and Taylor Minich²

¹Center for Conservation Innovation, Defenders of Wildlife, 1130 17th St. NW, Washington, DC 20036

²Nicholas School of the Environment, Duke University, Grainger Hall 9 Circuit Drive, Durham, NC 27708

Abstract

We have an unprecedented ability to analyze and map the Earth's surface, as deep learning technologies are applied to an abundance of Earth observation systems collecting images of the planet daily. In order to realize the potential of these data to improve conservation outcomes, simple, free, and effective methods are needed to enable a wide variety of stakeholders to derive actionable insights from these tools. In this paper we demonstrate simple methods and workflows using free, open computing resources to train well-studied convolutional neural networks and use these to delineate objects of interest in publicly available Earth observation images. With limited training datasets (<1000 observations), we used Google Earth Engine and Tensorflow to process Sentinel-2 and National Agricultural Imaging Program data and use these to train U-Net and DeepLab models that delineate ground mounted solar arrays and parking lots in satellite imagery. The trained models achieved 81.5% intersection over union between predictions and ground-truth observations in validation images. These images were generated at different times and from different places from those upon which they were trained, indicating the ability of models to generalize outside of data on which they were trained. The two case studies we present illustrate how these methods can be used to inform and improve the development of renewable energy in a manner that is consistent with wildlife conservation.

Introduction

The proliferation of Earth observation data has made possible the mapping of features on the Earth's surface at an extraordinary frequency and level of detail. Publicly available images of the entire Earth are being collected as often as every 5 days and global images with ≤ 1 m resolution are collected daily by private companies. This availability of data has created an unprecedented ability to map the distribution and state of forest loss (Song et al. 2018), urban growth (Watson & Venter 2019), water resources (Pekel et al. 2016) and other changes to the Earth's surface. These endeavors are critical to better understand and conserve natural resources. More recent advances in mapping capabilities are being made as researchers combine Earth observation data with powerful deep learning approaches (Lecun et al. 2015; Mahdianpari et al. 2018; Li et al. 2019). Yet the ability to apply this work to conservation challenges remains out of reach for many conservationists, as early progress has largely focused on high-resolution imagery, large training datasets, and required substantial computing resources (Li et al. 2019). A transition is needed to unlock the insights and impacts possible from this synergy for the purposes of conservation.

The application of computer vision methods to Earth observation has transformed the types of features and precision with which researchers are able to map the Earth's surface. Computer vision describes the automated identification of objects in images, traditionally using deep learning models to recognize, locate, and delineate objects like cats, cars, and faces in photographs (Simonyan & Zisserman 2015). More recently, these same models have been applied to satellite and aerial imagery to identify and classify land cover (Nogueira et al. 2017). The promising development is the application of convolutional neural networks (CNNs) – deep learning architectures that can take advantage of the shape and context of an object – marking a transition away from traditional pixel-based classification approaches like those used to create land cover maps (Jin et al. 2019). Incorporating information about spatial context can be critical for distinguishing features with similar spectral characteristics. For instance, an oil drilling pad may be spectrally like a large surface mine, but the two features are easily distinguishable based on shape, configuration, and context.

A variety of CNN models have been applied to various Earth observation data sources to successfully delineate various features (Mahdianpari et al. 2018; Wiratama et al. 2020). Thus far, many of the advances have focused on model development, training feature engineering, often in specific use cases taken at a snapshot in time. In this paper, we focus on the advancement of the breadth of applicability to conservation challenges. To be of maximal use to conservation, CNN models using Earth observation data need to:

1. Use publicly available data
2. Use data that is updated regularly
3. Be free or of low cost to implement
4. Work with limited training datasets
5. Be generalizable

Methods and models must meet these five criteria in order to enable conservationists to regularly use them to create actionable outputs. To address these needs, we sought to develop simple, replicable computer vision models that delineate features of interest using free, publicly available earth observation data. The launch of Google Earth Engine (Gorelick et al. 2017) and its recent integration with the Tensorflow deep learning library (Abadi et al. 2015) now make this possible. As a demonstration of concept, we focus on mapping two types of infrastructure related to renewable energy.

As renewable energy proliferates, there is interest from conservationists in promoting the development of renewable energy facilities in a manner that is consistent with the goals of wildlife conservation. A wildlife-friendly approach to renewable energy development involves both understanding the current extent of renewable energy and identifying sites for future facilities that minimize disturbance to wildlife and wildlife habitat. Solar energy production has grown rapidly in the United States over the past decade, experiencing a 35-fold increase from 2008 to 2018 (Margolis et al. 2018). Conservationists are interested in mapping the spatial distribution of these sites to understand their potential impacts to habitat availability and fragmentation, and help determine the potential for these sites to be simultaneously used for habitat restoration benefiting native species (Beatty et al. 2017; Sinha et al. 2018).

Simultaneously, conservationists want to prioritize the development of future solar energy facilities at sites that minimize impacts to wildlife. This requires accurate, up-to-date maps of potential low-conflict sites in order to assess generation potential and prioritize siting in conjunction with other factors (e.g., distribution capacity, demand, zoning, etc.)

In this paper, we apply two well-studied CNNs to delineate the footprints of large (> 2 ac) parking lots in Long Island, NY and ground-mounted solar arrays across North Carolina. These case studies are both part of ongoing conservation initiatives. The solar roadmap (<http://solarroadmap.org/>) seeks to help Long Island meet the goal of 70% renewable energy by 2030 by identifying low-conflict sites for solar development. North Carolina produces the second most solar energy in the United States, generating more than 4% of its annual energy from solar arrays following substantial growth over the past decade (Margolis et al. 2018) – and there is interest from The Nature Conservancy in understanding the current configuration of arrays relative to wildlife habitat. Our goal was to demonstrate the ability to delineate features of interest at landscape scales using free resources by combining CNN models with publicly available earth observation data.

Methods

Training Data

Our analyses focused on mapping parking lots greater than 2 ac in the town of Huntington, NY (located on Long Island), and ground mounted solar arrays in the state of North Carolina. We hereafter refer to parking lots and solar arrays as target features. The Nature Conservancy's North Carolina chapter provided us with a shapefile containing footprints of ground mounted solar arrays in North Carolina as of 2016. The Nature Conservancy's Long Island chapter provided a shapefile containing 645 hand-digitized boundaries of parking lots > 2ac in size in Suffolk County, NY. After ingesting these shapefiles into Google Earth Engine, we refined and updated the solar data to create 663 solar array polygons. Finally, we converted target feature polygons into single band label images with pixel values equal to 1 in areas covered by a target feature and 0 elsewhere.

Earth Observation Data

To delineate ground mounted solar arrays in North Carolina, we used top of atmosphere reflectance data collected by the Sentinel-2 imaging system (Drusch et al. 2012). Sentinel-2 data contain 13 multispectral bands collected at 10, 20, and 60 m resolution. We accessed all Sentinel 2 images intersecting the boundary of North Carolina collected between 2016-01-01 to 2016-12-31 in which less than 20 percent of pixels were labeled as cloudy, as recorded in the image metadata. We then masked clouds from these images using the included quality assurance band, which flags cloudy pixels. Finally, we selected values from the three visible red, green, and blue (RGB), near infrared (NIR), and two short-wave infrared bands (SWIR1 & SWIR2) as training features for the CNN model used to delineate solar arrays.

To make the model robust to phenological variability, we separated the 2016 collection of Sentinel-2 images into four seasonal collections: Spring (01Mar16 - 31May16); Summer (01Jun16 – 31Aug16); Fall (01Sep16 – 30Nov16); and Winter (01Dec16 – 28Feb16). Each of these four collections images were then reduced to a single image per season by taking the

median value at each pixel among all images in the collection. Thus, we generated four six-band images containing Sentinel-2 reflectance data covering North Carolina.

We used 1 m² resolution National Agricultural Imaging Program (NAIP) imagery to precisely map parking lots across Long Island. The (sometimes) small size of parking lots necessitated higher resolution imagery to train a useful image segmentation model. NAIP imagery is collected approximately once every two years per state, and the most recent images covering Long Island were from 2016. These images contained blue, green, red, and near-infrared reflectance values, and we used RGB bands as training features. NAIP images are collected from fixed-wing aircraft and are cloud free. Therefore, we did not apply any preprocessing to these data. We used Google Earth Engine to access and process all Earth observation data.

We overlaid the respective solar array and parking lot label images with Sentinel-2 and NAIP images to create images with 7 and 4 bands, respectively (Fig. 1). We sampled image chips from each of these training images for model training and evaluation. Because solar arrays and parking lots can be relatively sparse features on a landscape, we took two steps to ensure our model had enough positive examples from which to learn to recognize these features. First, we sampled image chips at the centroids of each digitized training feature. We then generated a random sample of 1000 points within 5km of these features and sampled image chips at each point. The resulting sets of image chips were then divided into 70% training and 30% evaluation sets.

Model Training & Evaluation

To delineate solar arrays in Sentinel-2 imagery, we trained a U-Net model (Ronneberger et al. 2015) taking 256 x 256 x 6 pixel image chips as input. To map parking lots, we trained a DeepLab v3 model with a Resnet backbone (Chen et al. 2017) taking 512 x 512 x 3 pixel image chips as input. The spectral information contained in NAIP data is like that in photographs, for which pre-trained Resnet weights exist. We used weights previously trained the ImageNet collection (Russakovsky et al. 2015).

During training we rescaled image chips to standardize model input. For Sentinel-2 data, we calculated per-band means and variances for each incoming chip and normalized each to have mean = 0 and standard deviation = 1. For NAIP data, we centered each image chip using pre-calculated means per band from the ImageNet collection. ImageNet photographs contain red, green, and blue values on a 0-256 scale, which is the same as NAIP imagery, and this centering created input data consistent with that expected by the pretrained ResNet backbone. We additionally implemented morphological and spectral image augmentation to artificially increase the variability of training data. At training time, we randomly applied a rotation of 0, 90, 180, or 270 degrees as well as random horizontal and vertical flips to image chips including multispectral training and label data (e.g., Peng et al. 2019). We then separated labels from multispectral bands and augmented the color of the latter.

We trained the U-Net solar array model using Keras with Tensorflow as the backend using batches of 16 images per step for 50 epochs, optimizing a weighted binary cross entropy loss function using the Adam optimizer with initial learning rate of 1e-4 and a decay rate of beta1 =

0.9, $\beta_2 = 0.999$. We evaluated model performance on the evaluation data set at each epoch in terms of intersection over union (IoU) between predictions and labels. We saved the weights maximizing the IoU metric during training using Keras callbacks. We trained the DeepLab v3 parking lot model with Tensorflow 2.0 using batches of 16 images per step for 50 epochs. We optimized weighted binary cross entropy loss using a momentum optimizer with initial learning rate $7e^{-3}$, final learning rate $1e^{-6}$, and a polynomial decay. We retained model weights from checkpoints at which IoU was maximized. For both models, we used Tensorboard to track model performance during training and evaluation.

We tested the predictive performance of each model on test images that were used neither during training nor evaluation. Performance was measured in terms of commission rates and area (i.e., false positives), omission rates and area (i.e., false negatives), and intersection over union relative to ground truth features in validation images. Solar array ground truth features included 126 arrays in four 1800 km² images from North Carolina in 2018. These four validation areas were selected to capture a sample of solar arrays falling under each of the four Sentinel orbit paths covering the state. Parking lot ground truth features included 135 hand digitized parking lots in two 15 km² images from Monroe county, NY in 2016. We calculated each metric using polygons generated by applying varying thresholds to output prediction probabilities and identify the optimal threshold as the probability that produced the greatest IoU. For parking lots, we provide a comparison of the number of parking lot polygons generated using the optimal threshold that overlap parking lot features in OpenStreetMap data (Contributors 2017).

All data preparation and sampling, as well as model training and prediction were performed using Google Collaboratory notebooks with a Python 3 runtime. During Keras and Tensorflow model training and prediction we used the available graphical processing unit hardware accelerator.

Results

At the end of training, the U-Net solar array model had a mean IoU of 85.9% on evaluation data. Predictions made by the model improved in all metrics as the probability threshold delineating solar arrays from background was increased. The threshold resulting in polygons with the greatest IoU was 0.99 (Table 1). Using output polygons created with this threshold, we identified 159 confirmed new solar arrays covering 119.25 km² in North Carolina from between 2016 and 2019.

The DeepLab v3 parking lot model had a mean IoU of 72.3% on evaluation data at the end of training. The threshold resulting in polygons with the greatest IoU was 0.6 (Table 1). Using output polygons created using this threshold, we identified 148 parking lots > 2 ac in the town of Huntington, covering 2.81 km² in 2016. Of these, 123 (81%) were absent from the Open Street Maps dataset. The output of both models can be interactively visualized through a Google Earth Engine app (Evans 2019).

Discussion

Free, simple, effective workflows are needed to leverage the available capacity to combine Earth observation imagery with deep learning methods to solve conservation challenges. In this paper

we demonstrate the ability to delineate features of interest using public data, free computing resources, and limited training data sets. Our results indicate that the methods presented can achieve useful predictions and produce outputs that can be used in conservation research and planning. Importantly, the validation data on which we evaluate performance were generated from images that were spatially or temporally distinct from model training and evaluation data. Good performance on out-of-sample data suggest our models generalize reasonably well. Thus, this study addresses two of the major limitations in applying deep learning to Earth observation data – generalizability and limited training data (Li et al. 2019). The features we delineate can be used to better understand the current and potential impacts of solar energy development to wildlife habitat, as well as identify and prioritize sites for future development that are low impact. Furthermore, the trained models may be useful for delineating these features in other geographies or at future times.

Our trained CNNs were effective at accurately delineating solar arrays and parking lots and were able to identify instances of each feature that had not previously been mapped. The U-Net model delineating solar arrays was more accurate and precise than the DeepLab v3 model delineating parking lots. Both models had low rates of omission, but the DeepLab v3 model had relatively high commission errors. These were due mostly to small false positives, as indicated by relatively small areas of commission. It is possible that solar arrays are inherently more easily distinguished from the surrounding landscape than parking lots. Solar arrays have both distinct reflectance characteristics and spatial configurations, whereas parking lots are spectrally similar to other impervious surfaces and can take a variety of forms. The U-Net model trained to identify solar panels also used more training features due to the greater spectral resolution of Sentinel-2 data relative to NAIP data. It is also possible that the use of pre-trained weights with DeepLab v3 affected model performance. Theoretically, transfer learning should improve model performance and training time by leveraging low-level features learned during training on similar, larger datasets (Huang et al. 2017). However, other research has found this strategy does not necessarily improve the final target task accuracy (He et al. 2019). Our results indicate that parking lots might be more effectively delineated by incorporating additional infrared reflectance data available in NAIP imagery, and training either a DeepLab v3 or U-Net model from scratch.

It is critical that CNN models perform sufficiently well with relatively small training datasets to be of use to conservation. Limited training data has been a primary bottleneck in the advancement of applications of computer vision using Earth observation imagery. As such, conservationists often do not have access to large training datasets for a feature of interest (although more are becoming available e.g., (Dunnett et al. 2020; Rand et al. 2020)) and collecting these data can be cost-prohibitive. We trained models on less than 1,000 examples, and data augmentation methods were likely important to the success of models trained on such small datasets. Image augmentation, and label overloading (Robinson et al. 2019) - pairing label images with earth observation data taken at different times – increased the variability of examples seen by models within realistic bounds. We were unable to apply image overloading in training our DeepLab v3 model, due to the coarse temporal resolution of NAIP data, which may also have contributed to the decreased performance of this model relative to the U-Net model. The image augmentation procedures we used are commonly applied in the training of CNNs, and

should be implemented as standard practice when working with limited training data to help improve model generalizability.

A complimentary solution to limited training data is a stable of pre-trained models that can be repurposed for other tasks. Because we use native Sentinel-2 and NAIP bands as training features, our models may be a useful starting point in training CNNs to perform similar tasks using either of these free sources of Earth observation data. The ability to delineate a variety of target features using Sentinel-2 imagery is particularly appealing because the system's high revisit rate makes rapid updating of maps possible. Global datasets like OpenStreetMaps (OSM) are not always complete and rely on updates from a user community. For instance, a recent inventory of OSM solar panel data revealed a 50% match with locations derived from renewable energy specific datasets (Dunnett et al. 2020). Most large parking lots delineated by our final DeepLab v3 model were not included in the OSM dataset. Thus, a relatively lightweight model that can be produce predictions at regular intervals could represent a substantial advance in the ability to track ever-changing conditions. Trained models that do not rely on engineered training features are a necessary first step in this direction.

In application, practitioners may benefit from using ancillary datasets to perform simple, common sense corrections to CNN model outputs. Such adjustments can help improve the usefulness of outputs by eliminating unreasonable predictions. Both models had higher rates of commission errors than omission errors and were more likely to misidentify a non-target feature than miss a true instance of a solar array or parking lot. For example, large patches of clouds reflected in coastal waters were assigned a high probability of being solar arrays by our U-Net model. Occasionally, patches of forest were also identified as solar arrays. In a post-hoc analysis, we eliminated false positives over water using annual surface water data (Pekel et al. 2016), and those corresponding to forest patches by calculating the normalized difference vegetation index in the Sentinel-2 image on which predictions were made and masking pixels with an index over 0.2. These two adjustments increased IoU to 0.97. Similarly, rectangular patches of bare ground were mis-identified as parking lots by the DeepLab v3 model, and a similar approach using a bare soil index might improve the final output.

Finally, we provide annotation datasets that can be used for other training tasks as georeferenced polygons, with temporal metadata, and recommend that future training datasets be provided in a similar format. This format is important for three reasons. First, polygons can be used for both image segmentation and localization tasks, whereas labeled images or object locations can only be used for the latter (e.g. UC Merced dataset (Yang & Newsam 2010)). Second, polygons that are georeferenced enable the introduction of temporal variability through label data overloading. Third, as opposed to a label image, polygons can be easily combined with any available imagery allowing for the training of models using different, or multiple Earth observation systems (Robinson et al. 2019).

Data Availability

Ground truth data and trained model weights are available in an Open Science Framework repository: osf.io/g463z

Code used to generate images, sample training data, train models, and run predictions is available in the GitHub repo: https://github.com/mjevans26/Satellite_ComputerVision/

Tables

Table 1. Performance metrics for convolutional neural networks used to delineate solar arrays from Sentinel-2 data and parking lots from National Agricultural Imagery Program data. Metrics were calculated between ground truth polygons and polygons produced using different probability thresholds on model predictions.

Model	Threshold	Commission		Omission		IoU
		Area (%)	Rate	Area (%)	Rate	
DeepLab v3 parking lots	0.50	0.142	0.592	0.007	0.025	0.523
	0.55	0.125	0.564	0.007	0.025	0.530
	0.60	0.112	0.535	0.010	0.030	0.530
	0.65	0.103	0.515	0.013	0.040	0.522
	0.70	0.088	0.450	0.021	0.061	0.503
	0.75	0.071	0.390	0.044	0.091	0.470
	0.80	0.050	0.325	0.049	0.112	0.419
	0.85	0.031	0.281	0.113	0.238	0.347
	0.90	0.024	0.184	0.182	0.355	0.254
	0.95	0.012	0.051	0.340	0.604	0.125
U-Net solar arrays	0.50	0.110	0.484	0.014	0.023	0.793
	0.55	0.109	0.495	0.014	0.023	0.794
	0.60	0.108	0.489	0.014	0.023	0.795
	0.65	0.106	0.489	0.014	0.023	0.797
	0.70	0.105	0.473	0.014	0.023	0.798
	0.75	0.104	0.461	0.014	0.023	0.799
	0.80	0.102	0.478	0.014	0.023	0.800
	0.85	0.100	0.442	0.014	0.023	0.802
	0.90	0.097	0.415	0.014	0.023	0.805
	0.95	0.091	0.417	0.014	0.023	0.808
	0.99	0.078	0.366	0.014	0.023	0.815

Figures

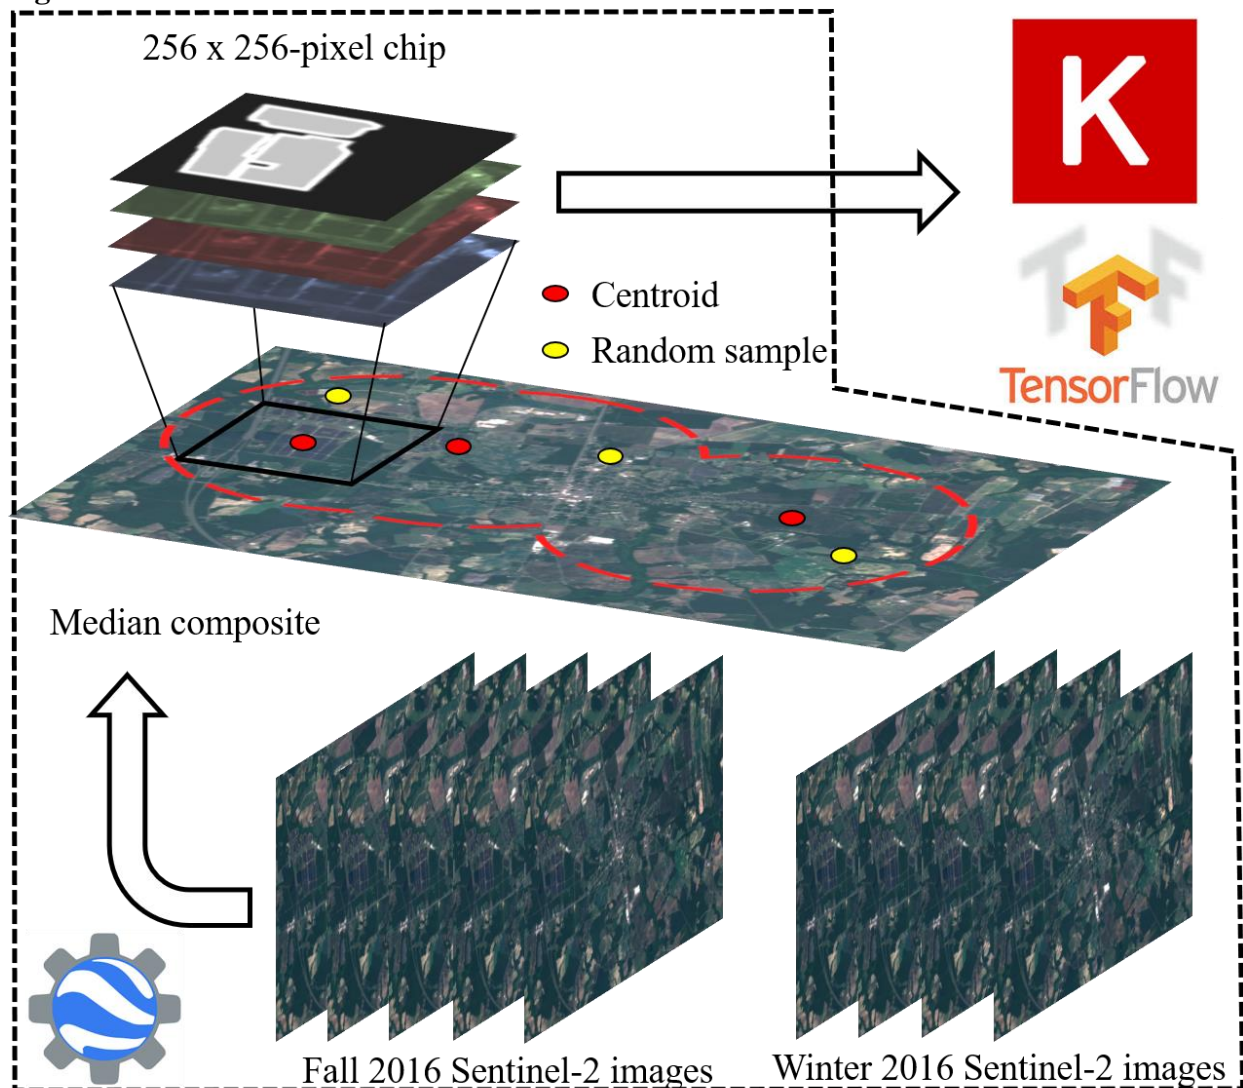


Figure 1. Flow-chart showing data collection process used to train deep learning models delineating solar arrays and parking lots. Steps contained within the dashed border are performed using Google Earth Engine, and consist of dividing imagery into seasonal collections (when available), creating median composites per collection, and sampling these composites at target feature centroids and random locations.

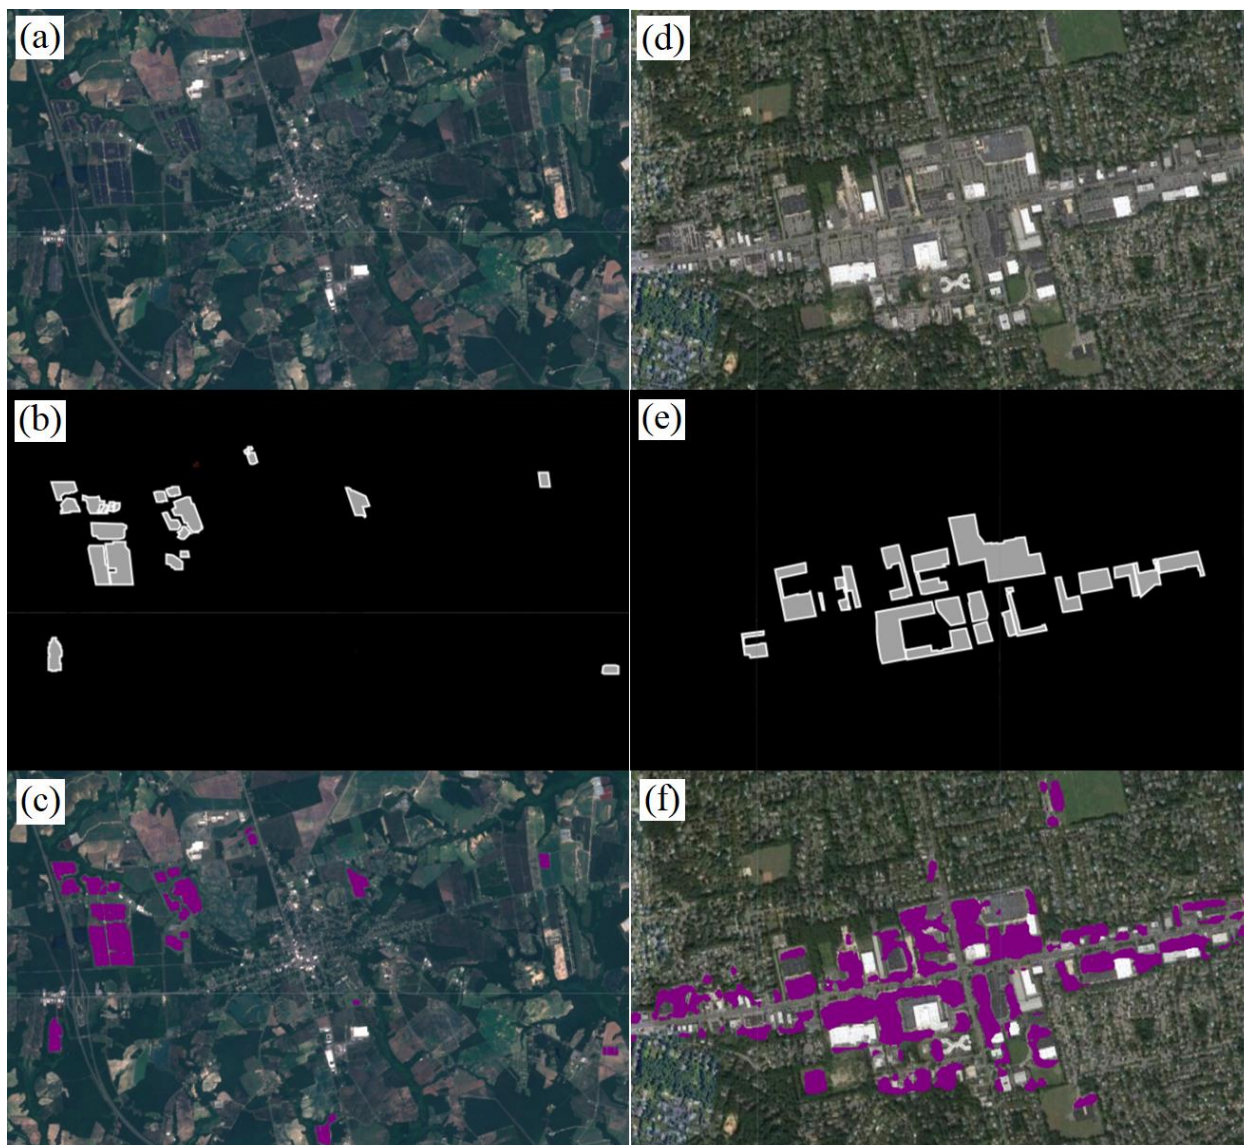


Figure 2. Predictions made by convolutional neural networks trained to delineating solar arrays from Sentinel-2 data (a, b, c) and parking lots using National Agricultural Imagery Program data (d, e, f). Panels flow from top to bottom showing the raw imagery used to generate training features (a, d), ground truth labels (b, e), and model output predictions (c, d).

References

- Abadi M et al. 2015. Tensorflow: Large-scale machine learning on heterogeneous systems. <http://tensorflow.org>.
- Beatty B, Macknick J, Mccall J, Braus G, Buckner D. 2017. Native vegetation performance under a solar PV array at the National Wind Technology Center. National Renewable Energy Laboratory. Available at: www.nrel.gov/publications
- Chen L-C, Papandreou G, Schroff F, Adam H. 2017. Rethinking Atrous Convolution for Semantic Image Segmentation. arXiv:1706.05587v3.
- Drusch M et al. 2012. Sentinel-2: ESA's optical high-resolution mission for GMES operational services. *Remote Sensing of Environment* **120**: 25–36.
- Dunnett S, Sorichetta A, Taylor G, Eigenbrod F. 2020. Harmonised global datasets of wind and solar farm locations and power. *Scientific Data* **7**:130.
- Evans MJ. 2019. CompViz web application. Available from <https://defendersofwildlifegis.users.earthengine.app/view/compviz>.
- Gorelick N, Hancher M, Dixon M, Ilyushchenko S, Thau D, Moore R. 2017. Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment* **202**:18–27.
- He K, Girshick R, Dollar P. 2019. Rethinking ImageNet pre-training. *Proceedings of the IEEE International Conference on Computer Vision*: 4917–4926.
- Huang Z, Pan Z, Lei B. 2017. Transfer learning with deep convolutional neural network for SAR target classification with limited labeled data. *Remote Sensing* **9**: 907
doi:10.3390/rs9090907.
- Jin S, Homer C, Yang L, Danielson P, Dewitz J, Li C, Zhu Z, Xian G, Howard D. 2019. Overall methodology design for the United States national land cover database 2016 products. *Remote Sensing* **11**: doi:10.3390/rs11242971.
- Lecun Y, Bengio Y, Hinton G. 2015, May 27. Deep learning. *Nature* **521**: 436-444.
- Li J, Huang X, Gong J. 2019. Deep neural network for remote-sensing image interpretation: Status and perspectives. *National Science Review* **6**:1082–1086.
- Mahdianpari M, Salehi B, Rezaee M, Mohammadimanesh F, Zhang Y. 2018. Very deep convolutional neural networks for complex land cover mapping using multispectral remote sensing imagery. *Remote Sensing* **10**: 1119; doi:10.3390/rs10071119.
- Margolis R, Feldman D, Boff D. 2018. Solar Industry Update. National Renewable Energy Laboratory:1–83.
- Nogueira K, Penatti OAB, Santos JA. 2017. better exploiting convolutional neural networks for remote sensing scene classification. *Pattern Recognition* **61**: 539-556.
- OpenStreetMap Contributors. 2017. OpenStreetMap.

- Pekel JF, Cottam A, Gorelick N, Belward AS. 2016. High-resolution mapping of global surface water and its long-term changes. *Nature* **540**:418–422.
- Peng D, Zhang Y, Guan H. 2019. End-to-end change detection for high resolution satellite images using improved UNet++. *Remote Sensing* **11**: 1382.
- Rand JT, Kramer LA, Garrity CP, Hoen BD, Diffendorfer JE, Hunt HE, Spears M. 2020. A continuously updated, geospatially rectified database of utility-scale wind turbines in the United States. *Scientific Data* **7**:1–12.
- Robinson C, Hou L, Malkin K, Soobitsky R, Czawlytko J, Dilkina B, Jojic N. 2019. Large scale high-resolution land cover mapping with multi-resolution data. Pages 12718–12727 *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*.
- Ronneberger O, Fischer P, Brox T. 2015. U-Net: Convolutional networks for biomedical image segmentation. arXiv: doi.1505.04597.
- Russakovsky O et al. 2015. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision* **115**:211–252.
- Simonyan K, Zisserman A. 2015. Very deep convolutional networks for large-scale image recognition. Page 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings.
- Sinha P, Hoffman B, Sakers J, Althouse L. 2018. Best Practices in Responsible Land Use for Improving Biodiversity at a Utility-Scale Solar Facility. *Case Studies in the Environment* **2**:16-12.
- Song X-P, Hansen MC, Stehman SV., Potapov PV., Tyukavina A, Vermote EF, Townshend JR. 2018. Global land change from 1982 to 2016. *Nature* **560**:639–643.
- Watson JEM, Venter O. 2019. Mapping the Continuum of Humanity’s Footprint on Land. *One Earth* **1**:175–180.
- Wiratama W, Lee J, Sim D. 2020. Change Detection on Multi-Spectral Images Based on Feature-level U-Net. *IEEE Access* **8**:12279–12289. Institute of Electrical and Electronics Engineers (IEEE).
- Yang Y, Newsam S. 2010. Bag-of-visual-words and spatial extensions for land-use classification. *GIS: Proceedings of the ACM International Symposium on Advances in Geographic Information Systems*:270–279.