
OGNet: Towards a Global Oil and Gas Infrastructure Database using Deep Learning on Remotely Sensed Imagery

Hao Sheng^{*1}, Jeremy Irvin^{*1}
Sasankh Munukutla¹, Shawn Zhang¹, Christopher Cross¹
Kyle Story², Rose Rustowicz², Cooper Elsworth², Zutao Yang¹
Mark Omara³, Ritesh Gautam³, Robert B. Jackson¹, Andrew Y. Ng¹
¹Stanford University, ²Descartes Labs, ³Environmental Defense Fund
{haosheng, jirvin16, sasankh, szhang22, chrisglc}@cs.stanford.edu
{kyle, rose, cooper}@descarteslabs.com, {momara, rgautam}@edf.org
rob.jackson@stanford.edu, ang@cs.stanford.edu

Abstract

At least a quarter of the warming that the Earth is experiencing today is due to anthropogenic methane emissions. There are multiple satellites in orbit and planned for launch in the next few years which can detect and quantify these emissions; however, to attribute methane emissions to their sources on the ground, a comprehensive database of the locations and characteristics of emission sources worldwide is essential. In this work, we develop deep learning algorithms that leverage freely available high-resolution aerial imagery to automatically detect oil and gas infrastructure, one of the largest contributors to global methane emissions. We use the best algorithm, which we call OGNet, together with expert review to identify the locations of oil refineries and petroleum terminals in the U.S. We show that OGNet detects many facilities which are not present in four standard public datasets of oil and gas infrastructure. All detected facilities are associated with characteristics known to contribute to methane emissions, including the infrastructure type and the number of storage tanks. The data curated and produced in this study is freely available at <http://stanfordmlgroup.github.io/projects/ognet>.

1 Introduction

Methane is the second-largest contributor to climate warming after carbon dioxide and accounts for at least one-quarter of present-day warming [1]. Around a third of current global anthropogenic methane emissions arise from the fossil fuel sector, and reducing these emissions is critical to mitigate further warming in the near-term [2, 3, 4, 5]. In the last decade, quantifying methane emissions from the oil and gas (O&G) sector has become a significant area of research and interest to both governments and the O&G industry [6, 7, 8, 9]. Recent progress in quantifying methane emissions from the O&G sector has been largely due to advancements in satellite-based observations [10, 11, 12, 8, 13]. Currently, multiple global mapping satellite missions (SCIAMACHY, GOSAT, TROPOMI) provide atmospheric methane concentrations data [14] and limited coverage satellites (GHGSat) can detect high emission rates from specific facilities [15]. Furthermore, a new satellite (MethaneSAT) planned for launch in 2022 will provide publicly-available methane emissions data made up of diffuse emission fields and point sources to help drive action towards reducing emissions and for companies and countries to track their emission reduction targets [16, 17].

^{*}Equal contribution.



Figure 1: Example NAIP imagery of positive examples (top row) and negative examples (bottom row) in the dataset used for model development. The positive images capture oil refineries, characterized by a large footprint and consistent features like large round storage tanks. The negative images consist of randomly sampled locations as well as objects and landscapes which look similar to oil refineries.

To attribute methane emissions detected by these satellites to facilities on the ground, a comprehensive database of the locations and characteristics of O&G infrastructure worldwide is critically required. There have been prior efforts to construct public datasets of O&G facilities. The global O&G infrastructure database (GOGI) was curated by developing machine learning models to parse the web for publicly-listed instances of infrastructure, but this process leads to data gaps due to limitations with searching for records published online [18]. Other national programs which host O&G infrastructure data include the Homeland Infrastructure Foundation-Level Data (HIFLD) by the Department of Homeland Security [19], the U.S Energy Mapping System from the Energy Information Administration (EIA) [20], and the Greenhouse Gas Reporting Program (GHGRP) by the EPA [21]. The data sources for these programs include required reporting by facility operators and publicly available data, which may be outdated or have incomplete coverage [22].

The recent unprecedented availability of high-resolution satellite and aerial imagery, together with advancements in deep learning methods, offer a powerful opportunity to create high-quality, granular, and large-scale databases of global infrastructure [23, 24]. Deep learning techniques have been increasingly utilized for automatically detecting objects in remotely sensed imagery, from building footprint detection [25] and urban land use classification [26] to energy infrastructure classification, including solar photovoltaics [27] and wind turbines [28]. Only recently have these models been deployed on high-resolution imagery to construct large-scale databases [27, 29].

In this work, we develop deep learning models to automatically detect O&G infrastructure in freely available aerial imagery. We deploy the best models in the full continental U.S. to map oil refineries and petroleum terminals and attribute each facility with characteristics that are known to contribute to methane emissions. To our knowledge, this is the first study to leverage high-resolution imagery to map O&G facilities in the continental U.S. All data curated and generated in this study is available at <http://stanfordmlgroup.github.io/projects/ognet>.

2 Methods

Task We train deep learning models to perform a canonical binary classification task, namely identifying whether an aerial image contains an oil refinery. We focus on oil refineries as they have large footprints and consistent features, which make them feasible to detect in aerial imagery, but we demonstrate that the deep learning models can also detect other O&G facilities.

Dataset Instances To develop and validate the models, we use 149 locations of oil refineries in the U.S. obtained from the Enverus Drillinginfo database, a commercial database containing point locations of various types of O&G infrastructure [30]. We additionally curate a large number of negative examples to include in the dataset. We first include a random sample of locations in the U.S. that capture a variety of landscapes. However, this set of negatives does not include facilities that may appear similar to oil refineries or landscapes that may commonly appear near oil refineries. In

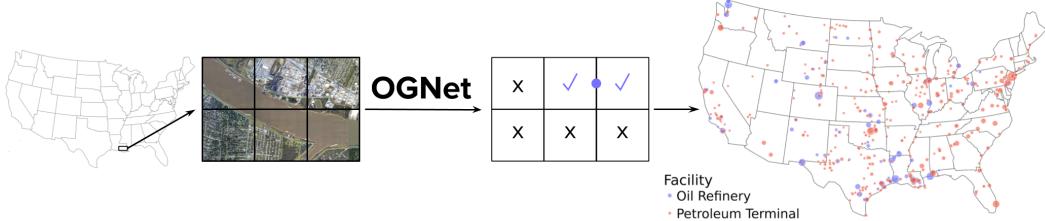


Figure 2: OGNet classifies NAIP imagery for the presence of oil and gas infrastructure. We deploy OGNet on the entire U.S. to identify the locations of oil refineries and petroleum terminals. The radius of the circles is proportional to the number of storage tanks at each facility.

order to address this, we sample locations containing visually similar objects and landscapes (difficult negatives) using an open-source GeoVisual search tool [31]. The query locations used in the search were obtained by finding objects and landscapes near the locations of the oil refineries in the dataset (see Appendix for details). The full dataset contains 7,066 examples in total and was split into a training set to learn model parameters (127 positive examples, 5,525 negative examples), a validation set to select model hyperparameters (13 positive examples, 693 negative examples), and a test set to evaluate the performance of the best model (9 positive examples, 697 negative examples).

Aerial Imagery We use aerial imagery between 2015 and 2019 from the National Agriculture Imagery Program (NAIP) [32], which captures the continental U.S. at a minimum of 1m resolution. We construct images of size 500 x 500 pixels downsampled to a resolution of 2.5m. We found that this area was sufficient to capture a large portion of the facility in most of the cases while also balancing memory usage and resolution for accurate detection. The images are mosaics of the most recent captures of the location and do not suffer from cloud cover or haze because images were acquired aerially on days with low cloud cover. The images of positive examples are centered around the coordinates associated with the oil refineries. Example images are shown in Figure 1.

OGNet Development and Deployment The deep learning model training procedure and architectures that were explored are provided in the Appendix. We call our best model OGNet. We convert predicted probabilities to binary predictions using the threshold that achieves the highest precision (0.81) subject to a recall of 1.0 on the validation set. To identify infrastructure in the continental U.S. using OGNet, we partition the region into equal 500 x 500 pixel tiles at 2.5m resolution leading to 5,082,722 tiles in total and run each of the tiles through OGNet. Any tile which was assigned a positive prediction is greedily merged with any positively classified adjacent tiles, and the mean of the centroids of each of the tiles within the merged group is used as the detected location (see Figure 2).

Human Review OGNet detections in the continental U.S. were manually reviewed, and all false positive detections were removed. We discovered that OGNet was able to detect O&G facilities other than oil refineries, so the remaining facilities were classified as oil refineries or petroleum terminals (including liquefied natural gas terminals and crude oil terminals). During this process, we also identified the number of storage tanks at each facility. An expert interpreter designed the annotation procedure (see Appendix) and verified the detected facilities and their attributed characteristics.

3 Results

OGNet Test Set Performance The best performing model, which we call OGNet, was a 121-layer DenseNet model. At its operating point, it achieved an accuracy of 0.996, a precision of 0.75, a recall of 1.0, and an F1 of 0.857 on the test set. The false positives all contained characteristic features of oil refineries but were other facilities, including a wastewater treatment plant (see Appendix).

OGNet Deployment Results OGNet detected a total of 1,902 instances of potential facilities in the U.S. after filtering two regions that contained many false positive detections. After manual review, 114 of the detected facilities were classified as oil refineries, and 336 detected facilities were classified as petroleum terminals. The remaining detections were false positives, and typical examples that

	Oil Refinery	Petroleum Terminal
Total Detections	114	336
Coverage of Benchmark Datasets	73.5% (108/147)	23.9% (292/1222)
New Detections	6	142
Example Image		

Table 1: Statistics of the OGNet detections and comparison to benchmark datasets. We compare detections to the union of four public databases of O&G infrastructure, including GOGI [18], GHGRP [21], HIFLD [19], and EIA [20]. New detections do not include any instances in the training set.

OGNet assigned high probability were grain elevators and water treatment and chemical plants. To interpret model predictions, we use class activation maps (see Appendix for details and examples).

Comparison with Public Datasets We compare the list of manually verified facilities detected by OGNet with four publicly available datasets, namely GOGI, GHGRP, HIFLD, and EIA. We combine the datasets and remove duplicate records by combining coordinates within 2 km of each other. To match the coverage of NAIP imagery, we additionally remove any samples outside the continental U.S. The final combined dataset has 147 oil refineries and 1,222 petroleum terminals in total.

For each facility in the combined dataset, we count the facility as covered if there is a facility detected by OGNet within 3 km, chosen to account for any perturbations in coordinates due to how centroid coordinates were determined in the detections. OGNet detected 73.5% of the oil refineries and 23.9% of the petroleum terminals in the combined dataset. Close to half of the “missed” oil refineries were due to inaccurate locations reported in the public datasets, with reported coordinates deviating more than 5 km away from the actual site. We also count the number of detected facilities that neither occur in the public datasets nor the training set. OGNet detected 6 oil refineries (including one abandoned facility) and 142 petroleum terminals that were not present in the public datasets (Table 1).

4 Discussion

Our results demonstrate the capability of deep learning models to detect large O&G facilities in remotely sensed imagery. We believe that the development and deployment of these models can fill existing gaps in publicly reported data. While we demonstrate the feasibility of detecting oil refinery and petroleum terminals in the U.S., our approach can be extended to other O&G production regions around the world using high-resolution imagery and thus has the potential to address critical data gaps, particularly in regions with a lack of data. We anticipate that the methods presented in this work can be extended to identify other types of O&G infrastructure and combined with other sources of data to create a comprehensive, granular, and up-to-date global database of O&G infrastructure.

We believe the development of this database will be important to the site-level quantification and attribution of greenhouse gas emissions, especially when combined with coarse measurements of emissions from other sources [16, 17]. Our work is a first step towards building a comprehensive infrastructure inventory to support multiple methane-measuring satellite missions aiming to attribute emissions and monitor progress toward emissions abatement. Focusing on the O&G sector is crucial as it is among the largest contributors to global anthropogenic methane emissions, and furthermore, there is evidence that methane emissions from fossil fuels have been significantly underestimated [3].

This study has important limitations. First, high-resolution imagery like NAIP is not widely publicly available worldwide. Future work should investigate the use of imagery with global coverage like Sentinel-2 10m optical imagery [33]. Second, while the model dramatically reduces the necessary amount of manual review, manual review is still necessary, and the current approach does not leverage the new annotations during training. Human-in-the-loop machine learning techniques could be used to incorporate human feedback [34]. Third, a canonical classification approach may not be effective for O&G infrastructure with smaller footprints like well pads and compressor stations. Object detection and semantic segmentation models may be more effective for localizing other O&G infrastructure.

References

- [1] G. Myhre, D. Shindell, F. Bréon, W. Collins, J. Fuglestvedt, J. Huang, D. Koch, J. Lamarque, D. Lee, B. Mendoza, *et al.*, “Anthropogenic and natural radiative forcing. climate change 2013: The physical science basis. contribution of working group i to the fifth assessment report of the intergovernmental panel on climate change, 659–740,” 2013.
- [2] J. D. Maasakkers, D. J. Jacob, M. P. Sulprizio, T. R. Scarpelli, H. Nesser, J.-X. Sheng, Y. Zhang, M. Hersher, A. A. Bloom, K. W. Bowman, *et al.*, “Global distribution of methane emissions, emission trends, and oh concentrations and trends inferred from an inversion of gosat satellite data for 2010–2015.,” *Atmospheric Chemistry & Physics*, vol. 19, no. 11, 2019.
- [3] B. Hmiel, V. Petrenko, M. Dyonisius, C. Buizert, A. Smith, P. Place, C. Harth, R. Beaudette, Q. Hua, B. Yang, *et al.*, “Preindustrial 14 ch 4 indicates greater anthropogenic fossil ch 4 emissions,” *Nature*, vol. 578, no. 7795, pp. 409–412, 2020.
- [4] J. K. Shoemaker, D. P. Schrag, M. J. Molina, and V. Ramanathan, “What role for short-lived climate pollutants in mitigation policy?,” *Science*, vol. 342, no. 6164, pp. 1323–1324, 2013.
- [5] R. B. Jackson, M. Saunois, P. Bousquet, J. G. Canadell, B. Poulter, A. R. Stavert, P. Bergamaschi, Y. Niwa, A. Segers, and A. Tsuruta, “Increasing anthropogenic methane emissions arise equally from agricultural and fossil fuel sources,” *Environmental Research Letters*, vol. 15, no. 7, p. 071002, 2020.
- [6] R. A. Alvarez, D. Zavala-Araiza, D. R. Lyon, D. T. Allen, Z. R. Barkley, A. R. Brandt, K. J. Davis, S. C. Herndon, D. J. Jacob, A. Karion, *et al.*, “Assessment of methane emissions from the us oil and gas supply chain,” *Science*, vol. 361, no. 6398, pp. 186–188, 2018.
- [7] “United states environmental protection agency, global methane initiative.” <https://www.epa.gov/gmi>. Accessed: 2020-09-15.
- [8] Y. Zhang, R. Gautam, S. Pandey, M. Omara, J. D. Maasakkers, P. Sadavarte, D. Lyon, H. Nesser, M. P. Sulprizio, D. J. Varon, *et al.*, “Quantifying methane emissions from the largest oil-producing basin in the united states from space,” *Science Advances*, vol. 6, no. 17, p. eaaz5120, 2020.
- [9] T. R. Scarpelli, D. J. Jacob, J. D. Maasakkers, M. P. Sulprizio, J.-X. Sheng, K. Rose, L. Romeo, J. R. Worden, and G. Janssens-Maenhout, “A global gridded (0.1×0.1) inventory of methane emissions from oil, gas, and coal exploitation based on national reports to the united nations framework convention on climate change.,” *Earth System Science Data*, vol. 12, no. 1, 2020.
- [10] E. A. Kort, C. Frankenberg, K. R. Costigan, R. Lindenmaier, M. K. Dubey, and D. Wunch, “Four corners: The largest us methane anomaly viewed from space,” *Geophysical Research Letters*, vol. 41, no. 19, pp. 6898–6903, 2014.
- [11] S. Pandey, R. Gautam, S. Houweling, H. D. Van Der Gon, P. Sadavarte, T. Borsdorff, O. Hasekamp, J. Landgraf, P. Tol, T. Van Kempen, *et al.*, “Satellite observations reveal extreme methane leakage from a natural gas well blowout,” *Proceedings of the National Academy of Sciences*, vol. 116, no. 52, pp. 26376–26381, 2019.
- [12] J. A. de Gouw, J. P. Veefkind, E. Roosenbrand, B. Dix, J. C. Lin, J. Landgraf, and P. F. Levelt, “Daily satellite observations of methane from oil and gas production regions in the united states,” *Scientific reports*, vol. 10, no. 1, pp. 1–10, 2020.
- [13] O. Schneising, M. Buchwitz, M. Reuter, S. Vanselow, H. Bovensmann, and J. P. Burrows, “Remote sensing of methane leakage from natural gas and petroleum systems revisited,” *Atmospheric Chemistry and Physics*, vol. 20, no. 15, pp. 9169–9182, 2020.
- [14] D. J. Jacob, A. J. Turner, J. D. Maasakkers, J. Sheng, K. Sun, X. Liu, K. Chance, I. Aben, J. McKeever, and C. Frankenberg, “Satellite observations of atmospheric methane and their value for quantifying methane emissions,” *Atmospheric Chemistry and Physics*, 2016.
- [15] D. Varon, J. McKeever, D. Jervis, J. Maasakkers, S. Pandey, S. Houweling, I. Aben, T. Scarpelli, and D. Jacob, “Satellite discovery of anomalously large methane point sources from oil/gas production,” *Geophysical Research Letters*, vol. 46, no. 22, pp. 13507–13516, 2019.
- [16] S. C. Wofsy and S. Hamburg, “Methanesat-a new observing platform for high resolution measurements of methane and carbon dioxide,” *AGUFM*, vol. 2019, pp. A53F–02, 2019.
- [17] “Methanesat.” <https://www.methanesat.org/>. Accessed: 2020-09-30.
- [18] K. Rose, J. Bauer, V. Baker, A. Barkhurst, A. Bean, J. DiGiulio, K. Jones, T. Jones, D. Justman, M. Sabatino, *et al.*, “Global oil & gas features database,” tech. rep., National Energy Technology Laboratory-Energy Data eXchange; NETL, 2018.
- [19] U. D. of Homeland Security, “Homeland infrastructure foundation-level data (hifld).” <https://hifld-geoplatform.opendata.arcgis.com/>. Accessed: 2020-10-1.

[20] U. E. I. Administration, “U.s energy mapping system.” <https://www.eia.gov/state/maps.php>. Accessed: 2020-10-3.

[21] “Greenhouse gas reporting program (ghgrp).” <https://www.epa.gov/ghgreporting>. Accessed: 2020-09-30.

[22] National Academies of Sciences, Engineering, and Medicine and others, *Improving characterization of anthropogenic methane emissions in the United States*. National Academies Press, 2018.

[23] A. Karpatne, I. Ebert-Uphoff, S. Ravela, H. A. Babaie, and V. Kumar, “Machine learning for the geosciences: Challenges and opportunities,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 31, no. 8, pp. 1544–1554, 2018.

[24] K. Janowicz, S. Gao, G. McKenzie, Y. Hu, and B. Bhaduri, “Geoai: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond,” 2020.

[25] W. Li, C. He, J. Fang, J. Zheng, H. Fu, and L. Yu, “Semantic segmentation-based building footprint extraction using very high-resolution satellite images and multi-source gis data,” *Remote Sensing*, vol. 11, no. 4, p. 403, 2019.

[26] P. Zhang, Y. Ke, Z. Zhang, M. Wang, P. Li, and S. Zhang, “Urban land use and land cover classification using novel deep learning models based on high spatial resolution satellite imagery,” *Sensors*, vol. 18, no. 11, p. 3717, 2018.

[27] J. Yu, Z. Wang, A. Majumdar, and R. Rajagopal, “Deepsolar: A machine learning framework to efficiently construct a solar deployment database in the united states,” *Joule*, vol. 2, no. 12, pp. 2605–2617, 2018.

[28] S. Zhou, J. Irvin, Z. Wang, E. Zhang, J. Aljubran, W. Deadrick, R. Rajagopal, and A. Ng, “Deepwind: Weakly supervised localization of wind turbines in satellite imagery,” in *NeurIPS*, Workshop on Tackling Climate Change with Machine Learning, 2019.

[29] X. Hou, B. Wang, W. Hu, L. Yin, and H. Wu, “Solarnet: A deep learning framework to map solar power plants in china from satellite imagery,” *arXiv preprint arXiv:1912.03685*, 2019.

[30] “Enverus drillinginfo.” <https://www.enverus.com/industry/midstream/>. Accessed: 2020-09-16.

[31] R. Keisler, S. W. Skillman, S. Gonnabathula, J. Poehnelt, X. Rudelis, and M. S. Warren, “Visual search over billions of aerial and satellite images,” *Computer Vision and Image Understanding*, 2019.

[32] U.-F.-A. A. P. F. Office, “National geospatial data asset (ngda) naip imagery.” <https://gis.apfo.usda.gov/arcgis/rest/services/NAIP>, 2015. Accessed: 2020-10-1.

[33] M. Drusch, U. Del Bello, S. Carlier, O. Colin, V. Fernandez, F. Gascon, B. Hoersch, C. Isola, P. Laberinti, P. Martimort, *et al.*, “Sentinel-2: Esa’s optical high-resolution mission for gmes operational services,” *Remote sensing of Environment*, vol. 120, pp. 25–36, 2012.

[34] D. Xin, L. Ma, J. Liu, S. Macke, S. Song, and A. Parameswaran, “Accelerating human-in-the-loop machine learning: Challenges and opportunities,” in *Proceedings of the Second Workshop on Data Management for End-To-End Machine Learning*, pp. 1–4, 2018.

[35] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

[36] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.

[37] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255, Ieee, 2009.

[38] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.

[39] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, “Learning deep features for discriminative localization,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2921–2929, 2016.

Appendix

Examples



Figure A1: Examples of features (in bounding boxes (A)(B)(C)) associated with oil and gas infrastructure. (A) *Storage containers and tank farms* tend to be common indicators for all types of infrastructure (both oil refineries and petroleum terminals). (B) *Jetties and piers* tend to be exclusive to LNG terminals and occasionally coastal refineries. (C) *Distillation units* are typically unique to refineries. The high resolution images shown here are obtained from Google Earth.



Figure A2: Examples of features (in bounding boxes (A-F)) associated with common facilities which mimic oil and gas infrastructure. (Left) Chemical plants are usually indicated by (A) *clarifiers*, (B) *small processing tanks*, and overall complex footprints similar to refineries. (Middle) Grain processing facilities are usually indicated by (C) *grain bins* and (D) *storage warehouses*. (Right) Wastewater treatment facilities are usually indicated by (E) *sedimentation tanks* and (F) *wastewater clarifiers*. The high resolution images shown here are obtained from Google Earth.

Difficult Negative Sampling

We sampled difficult negatives to include in the training, validation, and test sets in order to capture infrastructure and landscapes which may be misclassified by the model. We sampled these examples in multiple stages based on the performance of the model on the deployment set. First, we included locations with common detailed features similar to those of oil refineries including highly-populated urban landscapes, suburban landscapes, and townships with sparse shrubbery. Second, we included more challenging examples that led to frequent false positives including well pads, arable land/crop fields, and dense forestry. Third, we included snowy regions that the model tended to conflate with the features of oil and gas infrastructure (e.g. tops of white storage tanks). We found that the snowy regions reduced sensitivity to refineries, so the final model did not use these images in the dataset.

Architecture and Training Procedure

We experimented with ResNet [35] and DenseNet [36] architectures with different numbers of layers. We applied a variety of augmentations during training including random vertical and horizontal flips, random affine transformations and random color jitter. We initialized the networks with weights from a model pre-trained on ImageNet [37] and normalized the images based on the mean and standard deviation of images in the ImageNet training set. The networks were trained end-to-end using a unweighted binary cross entropy function, Adam



Figure A3: Examples of Class Activation Maps (CAMs) on OGNet detections from deployment. CAMs are produced by OGNet and overlaid over the NAIP images to highlight features of the image which contribute most to its prediction. The CAMs shown here highlight indicative features like storage tanks and distillation units mentioned in Figure A1.

with standard parameters [38], and a learning rate of 0.0001. During training, we evaluated the network on the validation set after each epoch and save the checkpoint with the lowest validation loss. We call the model which achieved the best performance on the validation set OGNet.

Deployment Run-time

Constructing and downloading the mosaics used in deployment took 20 days on a single machine with multiprocessing, and generating the predictions on the images using a single NVIDIA TITAN-Xp GPU took 40 hours.

Annotation Procedure

An expert on the oil and gas sector designed the manual annotation procedure for verifying and characterizing facilities detected by OGNet. Each detection was classified as negative, oil refinery, crude oil terminal, or liquefied natural gas terminals. The oil and gas facilities were identified using characteristic features visible from sub-meter resolution imagery provided by Google Maps and Google Earth Pro (Figure A1). Other facilities with similar features were distinguished by identifying the presence of other objects that are not present in the oil and gas facilities (Figure A2) and using place names in Google Maps. Finally, the number of big (internal and external) floating roof tanks was counted at each facility, verified using Google Earth's 3D building view.

Class Activation Maps

We used class activation maps (CAMs) to interpret OGNet predictions [39]. CAMs highlight the regions of the image which contribute to a positive classification. The CAM for an image was computed by taking a weighted average between the feature maps produced by OGNet for that image and the weights of the fully connected layer, followed by setting all negative values to zero. For an input image of size 500 x 500 pixels, the CAM was of size 15 x 15. We upsampled the CAM to size 500 x 500 and overlaid the image to highlight salient regions. Examples of class activation maps produced by OGNet are shown in Figure A3.