High Performance Geocomputation for Assessing Human Dynamics at Planet Scale

B. Bhaduri*, D. Lunga, H. Yang, J. McKee, M. Laverdiere, K. Kurte, and J. Sanyal

Oak Ridge National Laboratory, Oak Ridge, TN 37831 USA
*Email: bhaduribl@ornl.gov

Abstract

Settlements are key indicators of human presence on the landscape. Large scale mapping of human settlements and their morphology from very high-resolution satellite images is a critical step towards developing an interpretative understanding of population distribution and the sociocultural attributes of the built environment they live in. Convolutional neural network (CNN) based Deep Learning experiments indicates that such computations can be scaled to some of the largest high-performance computing (HPC) architectures. While early results are encouraging for developing settlement and corresponding population maps at unprecedented speed and spatial resolutions, characterizing human dynamics at planet scale with high temporal resolution will require the community to develop novel geocomputational infrastructures and ecosystems.

Keywords: Population data, human settlements, satellite imagery, HPC.

1. Introduction

High resolution (sub-Census local level) population data is instrumental for decision support systems in various applications including good governance, humanitarian missions, poverty reduction strategies, human security, and prosperity in social, economic, and environmental health. Geospatial data driven automated systems offer novel approaches to decompose aggregated Census data into finer spatial and temporal units enabling decision support to be executed at an individual unit e.g. at individual building level as compared to communal level. At Oak Ridge National Laboratory (ORNL), our focus through the LandScan population modelling program (Dobson et al., 2000; Bhaduri et al., 2002) has been to develop high resolution population distribution data from local to global scales while taking advantage of data driven automated systems. LandScan Global database has a spatial resolution of 30 arc seconds resolution (approximately 1 km cells). Since its inception in the late 1990s, it has emerged as one of the leading community resources for population data with globally expansive user-base, including key applications focusing on sustainable development including many social science applications. There have been two prominent desires from the user community: first being the increasing resolution of the data in spatial, temporal, and sociocultural dimensions and secondly, the uncertainty quantification, verification and validation of the models and database. While research is ongoing on the later, here we discuss progress thus far on how to increase the data resolution in spatial and sociocultural dimensions while leveraging high performance computing resources available on supercomputers e.g. Titan and Summit (OLCF 2019).

2. Enhancing Spatial Resolution

Nucleus to the LandScan population data, is the spatial data-driven model that provides foundational information for estimating populations. The model benefits mostly from moderate to high resolution land use land cover (LULC) data, as derived from NASA MODIS (250-500m) or Landsat TM (30m), to facilitate the development and continued enhancement of population distribution data. However, the model is sensitive to the quality and resolution of the ancillary input variables that are generated from LULC data. Furthermore, in order to assess the true magnitude and extent of the human footprint in population datasets, it is critical to understand the distribution and relationships of the small and medium-sized human settlements. These structures remain mostly undetectable from medium resolution satellite derived LULC data. This is particularly true in suburban and rural areas, where the population is dispersed to a greater degree than in urban areas. Our greater efforts have been to seek advanced machine learning methods to extract high-resolution settlement data as a proxy to improved population distribution models. We have demonstrated through large scale experiments (Lunga et al. 2018; Yang et al. 2018) that extracting settlement information from very high-resolution, earth observation imagery provides a key pathway for rapid estimation and revision of settlement and population distribution data.

3. Enhancing Sociocultural Resolution

Population distribution data is key for humanitarian missions, where the vulnerable, such as those living in refugee camps, informal settlements and slums need to be efficiently and precisely located during, for example, disease eradication campaigns, in wake of natural disasters. It is for these reasons that the exploitation of very high resolution (sub-meter resolution) imagery should go beyond generating accurate human settlement maps but also for studying the characterization of population (social and vulnerability. For example, the work of Weeks et al, has long highlighted the potential of mapping poverty from remote sensing imagery data. From our recent efforts, we have not only validated the above hypothesis, but extended data-driven methods to the delineation of a new paradigm product, from rapid ingestion and analysis of high-resolution imagery to yield the “neighborhood” segmentation of human built structures. Such an exploitation of large volumes of imagery as enabled by capable geocomputation methods allows the community to move toward the goal of creating a global foundation level database for impervious surfaces and “neighborhoods”. Apart from the tremendous promise to advance the state of currency, this could open avenues for timely flow of critical information to benefit numerous sustainable development programs.

4. High Performance Geocomputation Architecture

The creation of any imagery based global foundation level database with such enhanced resolution requires advanced algorithms capable of extracting, representing, modelling, and interpreting image scene features. The key tasks being to characterize the spatial, structural, and semantic attributes from peta-scale volumes of imagery. Although higher resolution imagery continues to enable the characterization of such attributes at finer spatial units, the scalability of algorithms remains a challenge. At half-meter pixel resolution, the earth's land surface has roughly 100 Trillion pixels. The requirement to process at this scale at repeated intervals demands highly scalable solutions. We have prototyped, over several years, multiple workflows that are GPU-based to provide efficient and scalable computation frameworks. For example, our current advanced workflow is designed to identify damaged critical infrastructure from large-scale satellite to assess vulnerable population during the 2018 California Fire.

Feature learning and feature engineering from imagery data are known to be computationally expensive for automated image scene analysis. The computational process requires dozens of floating-point computations per pixel, which can result in slow runtime even for the fastest of CPUs. The slow speed of a CPU is a serious hindrance to productivity for time critical missions. Our GPU-accelerated computing solution provides an order of magnitude or more in performance by offloading compute-intensive portions of the application to the GPU, while the remainder of the code still runs on the CPU. The implementation further scales linearly with the available nodes thus enabling the processing of large-scale data on high end GPU based cluster-computers.

As early as 2007 (Vjayaraj et al., 2007; Cheriyadat et al. 2008), automated feature extraction algorithms implemented on available CPU-based architectures demonstrated radical improvement in image analysis efficiency; although such efforts relied heavily on human effort. However, with the introduction of automated methods, manual settlement identification for a 100 km^2 area was reduced from 10 hours to 30 minutes. However, this scaled inefficiently beyond 10 nodes and at that rate processing 57 million km^2 habitable area would take decades.

Persisting to seek efficient workflows, most recently, we have leveraged new advances from deep learning, deep convolutional neural networks (CNN), to provide a scalable object detector which has led to the detection and mapping of human settlements at individual building level. We have demonstrated that buildings for the contiguous United States could be mapped utilizing 1m NAIP imagery data (Yang et al., 2018), and a single CNN-model at an average processing time less than ~56 km^2 per minute with 8 NVIDIA-Tesla K80 GPUs. We further adapted the workflow in an attempt to test scalability of mapping buildings with very high-resolution imagery. Using 30cm imagery, we scaled the CNN-model to 4758 Nvidia Tesla K20X GPU nodes on Titan Supercomputer at ORNL. Satellite images covering the entire country of Yemen, a total of pixels covering ~1.9M (1,925,180.66) km^2 were processed in less than 2 hours to map all buildings with a hybrid workflow (Figure 1) where large amount of GPU nodes of Titan, and higher computing capacity offered by DGX-1 were exploited. The process area is significantly larger than actual area of Yemen (527,970 km^2) because of the overlap among all the image strips that cover the entire country. The workflow in Figure 1 is also ready to be adapted to Summit where we expect to have a unified workflow to derive similar country-scale feature mapping.

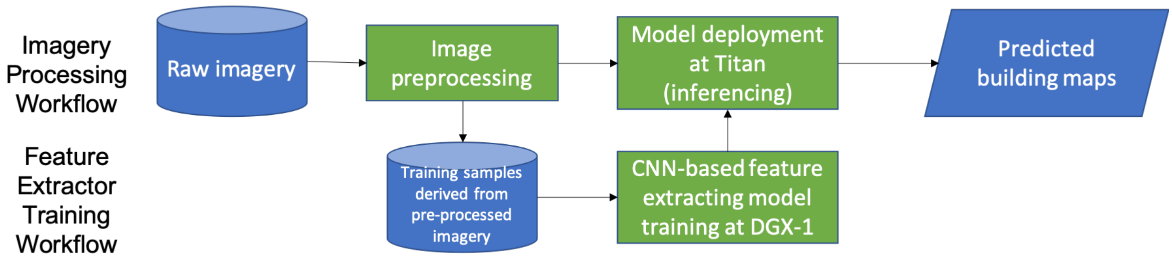

Figure 1: A hybrid high performance CNN-based feature extractor and deployment workflow

5. Conclusion

It is encouraging to realize that combined with the CNN-based algorithmic advancements and emerging computational architectures such as GPU based processing, our recent experimental benchmarks indicate that increasing the spatial resolution of the global LandScan database from 30 to 3 arc seconds can be achieved within a few years (and not a few decades as it appeared a few years back). Moreover, this approach was successfully utilized in a first ever attempt to develop a bottom up, high resolution (90m) population estimate for Nigeria (Weber et al., 2018), a country still awaiting an official census since 2006.

6. Acknowledgements

Support for this research was provided by the U.S. Government including the Department of Energy's (DOE) Advanced Scientific Computing Research (ASCR) Leadership Computing Challenge (ALCC). This paper has been authored by employees of UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the U.S. Department of Energy. Accordingly, the United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

7. References

Bhaduri, B., Bright, E., Coleman, P., & Dobson, J. 2002. LandScan: Locating People is What Matters. Geoinformatics, 5, pp. 34-37.

Dobson, J., Bright, E., Coleman, P., Durfee, R., & Worley, B. 2000. A Global Population database for Estimating Populations at Risk. Photogrammetric Engineering & Remote Sensing, 66(7), pp. 849-857.

Cheriyadat, A., Bright, E., Bhaduri, B, and Potere, D., 2007. Mapping of Settlements in High Resolution Satellite Imagery Using High Performance Computing. GeoJournal, 69(1-2), pp.119-129.

Lunga, D., Huang, H. L., Reith, A., Weaver, J., Yuan, J., and Bhaduri, B. 2018. Domain-Adapted Convolutional Networks for Satellite Image Classification: A Large-Scale Interactive Learning Workflow. IEEE Journal of Selected Topics in Earth Observations and Remote Sensing, 11(3), pp. 962-977.

OLCF (Oak Ridge Leadership Computing Facility). 2019. Compute Systems [Online]. Accessed 14 August 2019 Available from https://www.olcf.ornl.gov/olcf-resources/compute-systems/. Last.

Vijayaraj, V., Bright E., and Bhaduri, B., 2007. High Resolution Urban Feature Extraction for Global Population Mapping using High Performance Computing, Proceedings of the IEEE International geosciences and *remote sensing symposium (IGARSS),* July 23-27 2007, Barcelona, Spain.

Weber, E., Seaman, V., Stewart, R., Bird, T., Tatem, A., McKee, J., Bhaduri, B., Moehl, J. and Reith, A., 2018. Census-Independent Population Mapping in Northern Nigeria, *Remote Sensing of Environment*, 204, pp. 786-798,

Yang, H. L., Yuan, J., Lunga, D., Laverdiere, M., Rose, A., and Bhaduri, B. 2018. Building Extraction at Scale Using Convolutional Neural Network: Mapping of the United States. IEEE Journal of Selected Topics in Earth Observations and Remote Sensing, 11(8), pp. 2600-2614.