GeoAI at ACM SIGSPATIAL: Progress, Challenges, and Future Directions

Yingjie Hu[1], Song Gao[2], Dalton Lunga[3], Wenwen Li[4], Shawn Newsam[5], Budhendra Bhaduri[3]
[1]University at Buffalo, USA
[2]University of Wisconsin, Madison, USA
[3]Oak Ridge National Laboratory, USA
[4]Arizona State University, USA
[5]University of California, Merced, USA

Abstract

Geospatial artificial intelligence (GeoAI) is an interdisciplinary field that has received tremendous attention from both academia and industry in recent years. This article reviews the series of GeoAI workshops held at the Association for Computing Machinery (ACM) International Conference on Advances in Geographic Information Systems (SIGSPATIAL) since 2017. These workshops have provided researchers a forum to present GeoAI advances covering a wide range of topics, such as geospatial image processing, transportation modeling, public health, and digital humanities. We provide a summary of these topics and the research articles presented at the 2017, 2018, and 2019 GeoAI workshops. We conclude with a list of open research directions for this rapidly advancing field.

1 Introduction

While the field of artificial intelligence (AI) was born in the 1950s, it has recently experienced a tremendous resurgence and is currently drawing significant attention from academia, industry, news media, and the general public. Through machine learning and in particular deep learning, AI has provided novel solutions to a variety of challenging problems ranging from computer vision to natural language processing, achieving near human-level performance. The impact of deep learning has reached many application domains, and geography is no exception. Remote sensing and geospatial image processing [57] is one area in geography and geographic information science (GIScience) that has been quick to embrace artificial intelligence techniques. For example, deep learning techniques have been adopted and further improved for tasks such as hyperspectral image analysis [2] and high-resolution satellite image interpretation [56]. Beyond remote sensing, researchers have also utilized deep learning techniques to extract information from other sources of geospatial images, such as Google Street View [35] and scanned historical maps [8]. In natural language processing (NLP), deep learning models, such as deep recurrent neural networks, have been employed to improve the accuracy of place name extraction from textual data [23, 43]. Other neural network based NLP techniques, such as word embeddings, have been employed to help quantify changes in stereotypes and attitudes toward women and ethnic minorities over a 100 year study period in the United States [10]. And, there are many other examples of research that integrates geography and AI, such as extracting building footprints using convolutional neural networks (CNNs) [50, 54], deep semantic segmentation for automated driving [32], vehicle trajectory prediction [48], indoor navigation [17], gazetteer conflation [31], and spatial epidemics [44]. The integration of geography and AI has given rise to the new and exciting interdisciplinary field of geospatial artificial intelligence (GeoAI).

Three factors in particular are contributing to the emergence of GeoAI as a field. First, the growing availability of large amounts of geospatial data continue to enable the training of increasingly complex AI models. Large datasets have long existed in the field of geography, with examples including remote sensing images, national-level census data, road networks, and land use and land cover (LULC) data. However, a variety of newer geospatial datasets, such as location-based social media and high resolution GPS trajectory data, have become available in the past two decades. In addition, companies such as Yelp and Uber have started to share their data (such as Yelp points of interest (POI) and Uber vehicle trajectories). These large and diverse geospatial datasets capture diverse aspects of natural and social environments, and enable the training of AI models to address a rich variety of problems.

The second factor contributing to GeoAI is that novel AI models and other computational methods are starting to be developed specifically for geographic problems. For example, Li et al. [18] extend a Faster-RCNN (Faster Region-based CNN) model to support the automatic detection of natural terrain features from remote sensing imagery. The model addresses a number of unique geospatial challenges, such as the ambiguous boundary of natural features in comparison to man-made features such as buildings and roads. Marcos et al. [27] propose a neural network model called RotEqNet for land cover mapping. This model encodes rotation equivariance in a CNN, and can address the challenge of recognizing rotated versions of the same object in remote sensing images. It is worth noting that geographic location and geographic information systems (GIS) play critical roles in developing GeoAI models since these models often integrate heterogeneous data from different sources. Geographic location and GIS are essential to establishing links between the data (e.g., linking streets to nearby residential areas and green spaces).

The third factor is that the increasing availability of high-performance computing hardware makes it possible to efficiently train GeoAI models with big geo data. GIS, as a special type of computer and information system, has already been integrated with supercomputing infrastructure (e.g., CyberGIS) [42]. Training GeoAI models on supercomputing infrastructure can significantly reduce the often lengthy training times, and enable broader hyperparameter and architecture search to identify the optimal models. Major companies such as ESRI and Microsoft have begun offering new computing resources (e.g., GeoAI Data Science Virtual Machine) with the goal of bringing AI, geospatial analysis, and high-performance computing together. Accelerated inference for deployment with AI models is now achievable via production ready system including the use of the Apache Spark computing framework [22].

It is in this context that a series of GeoAI workshops have been organized at ACM SIGSPATIAL, the premier conference at the intersection of geospatial data analysis and computer science. GeoAI 2017 and GeoAI 2018 took place in Los Angeles and Seattle respectively [26, 11], and, at time of writing, GeoAI 2019 has confirmed the list of accepted papers and will take place in Chicago. These GeoAI workshops bring together GIScientists, computer scientists, engineers, entrepreneurs, and decision makers from academia, industry, and government to discuss the latest trends, successes, challenges, and opportunities in this interdisciplinary field of GeoAI. The 2017 and 2018 GeoAI workshops were some of the more popular workshops at the ACM SIGSPATIAL conference based on conference reported statistics, and we expect the 2019 workshop to be similarly well attended. The papers accepted and presented at the workshops have covered a wide range of GeoAI research topics. In this article, we systematically review and summarize these papers. We then propose future GeoAI research directions.

2 Research Topics at the GeoAI Workshops

In this section, we review the research topics and problems covered at the three GeoAI workshops. We first review the set of papers presented at the 2017, 2018, and 2019 workshops. We then provide a summary of the topics covered.

2.1 GeoAI'17

ACM SIGSPATIAL 2017 marked the first gathering of the GeoAI workshop series and took place in Los Angeles. This workshop received a total of 14 submissions, and after a rigorous review process, 8 papers were accepted and presented at the workshop. This workshop also featured two keynotes: one was an academic keynote by Dr. Shawn Newsam (University of California, Merced) on "Geographic knowledge discovery using ground-level images and videos", and the other was an industry keynote from Dr. Saikat Basu (Facebook) on "Using AI to help generate roads for OpenStreetMap". Among the accepted 8 papers, there were two major research topics: geospatial image processing and transportation modeling. On the topic of geospatial image processing, Li et al. [19] used the Faster-RCNN model to support the automatic detection of terrain features, such as craters, from remote sensing images. Law et al. [16] applied a CNN to classify street frontages using both Google Street View images and 3D models produced by ESRI's City Engine. Collins et al. [5] developed a deep CNN model to enhance low resolution multispectral satellite imagery where no corresponding high resolution images are available. Duan et al. [8] proposed an algorithm to automatically align vector data to scanned historical maps to generate large datasets for training a CNN to recognize geographic features in such maps.

On the topic of transportation modeling and analysis, Kulkarni and Garbinato [14] explored the effectiveness of RNNs for generating synthetic traffic trajectory datasets. Murphy et al. [28] proposed an image-based classification method that uses CNNs to classify the noise level in GPS trajectories with the goal of improving the accuracy of distance estimation for ride-sharing applications. Within the general scope of transportation and navigation, Li et al. [17] developed a method that combines deep learning based object recognition and second-order hidden Markov models to detect indoor landmarks. In addition to the papers which focused on the two main themes, Majic et al. [25] developed an unsupervised approach to identify equivalent OpenStreetMap keys based on their co-occurring patterns and the geometries of the features annotated by the keys.

A special issue on "GeoAI: Artificial Intelligence Techniques for Geographic Knowledge Discovery" was organized in the *International Journal of Geographical Information Science* [12], which complemented the GeoAI'17 Workshop. This special issue attracted full paper submissions from not only the GeoAI 2017 workshop but also interested scholars from the community.

2.2 GeoAI'18

Due to the popularity and success of the first GeoAI workshop, the second workshop, GeoAI 2018, was held at ACM SIGSPATIAL in Seattle. This workshop received 19 submissions, and 10 papers were accepted after review. Dr. Rangan Sukumar (Cray Inc.) gave an industry keynote on "The AI Journey in Geospatial Discovery: Navigating Shapes, Sizes and Spaces of Data", and Dr. Bruno Martins (University of Lisbon) gave an academic keynote on "GeoAI Applications in the Spatial Humanities". Many of the accepted papers continued the two major research topics from GeoAI 2017, namely geospatial image processing and transportation modeling. On the topic of geospatial image processing, Xu et al. [47] reviewed the use of computer vision and CNN models for locating aerial images collected by unmanned aerial vehicles (UAVs). Sun et al. [38] used a CNN model, a modified U-Net, to combine satellite imagery and GPS data for road extraction. And, Srivastava et al. [36] trained a CNN model for classifying the multiple functions of buildings from Google Street View images in the city of Amsterdam.

On the topic of transportation modeling and analysis, Van Hinsbergh et al. [41] combined GPS trajectory with in-car signal data, and used a location extraction technique, Gradient-based Visit Extractor, to extract interesting locations of a trajectory, such as drop-off, parked, and pick-up. Pourebrahim et al. [30] examined the potential of adding Twitter data to neural network and gravity models to enhance commuter trip prediction.

Two papers at GeoAI 2018 specifically focused on methodology. Swan et al. [39] explored the negative impact of noise in training data on trained deep neural network models, and found that models trained with small amounts of noise had little change in precision but considerable increases in recall; however, as noise levels continued increasing, both precision and recall decreased. Aydin et al. [1] proposed an unsupervised and consensus-based regionalization algorithm, SKATER-CON, to create spatially contiguous regions, which addresses the chaining problem of computing minimum spanning trees by using random spanning trees and outperformed two state-of-the-art regionalization methods, SKATER and ARISEL.

A number of other research topics were also covered at GeoAI 2018. Xi et al. [44] proposed a deep residual network for infuenza prediction by integrating the spatial-temporal properties of infuenza at an intra-urban scale, and applied their model to a dataset in Shenzhen, China. Chow [4] discussed the possibility and challenges of integrating AI and agent-based models (ABMs) to simulate and estimate the head count of a moving crowd. Elgarroussi et al. [9] developed Aconcagua which is a spatio-temporal framework that can perform emotion analysis on tweets and monitor the change of positive and negative emotions over time and space.

2.3 GeoAI'19

The third GeoAI workshop will be held at ACM SIGSPATIAL in Chicago in November 2019. This workshop received 25 submissions, and 17 papers were accepted after review. Dr. Xin Chen (HERE Technologies) will give an industry keynote on "HD Live Map for Automated Driving: An AI Approach", and Dr. Raju Vatsavai (North Carolina State University) will give an academic keynote on "Geospatial AI for Monitoring Crops to Nuclear Proliferation Using Global Earth Observations". Many papers continue the two major research topics from GeoAI 2017 and GeoAI 2018, geospatial image processing and transportation modeling. On the topic of geospatial image processing, Chen et al. [3] propose a ChangeNet to identify relevant pixelwise changes in time-varying images taken at the same location and utilize conditional generative adversarial networks (GANs) to improve classification results. Dorji et al. [7] present a machine learning approach to estimate the median income levels of sub-districts in Thailand using nighttime satellite images, population density, road density, and distance to major metropolitan areas. Peng et al. [29] propose a residual patch similarity CNN (ResPSNet) to map urban flood hazard zones using bi-temporal high resolution (3 meters) pre- and post-flooding multispectral surface reflectance satellite imagery. Law and Neira [15] propose a convolutional principle component analysis (ConvPCA), which is applied to both street level and street network images to find a set of uncorrelated and ordered visual latent components to predict urban characteristics such as street quality classification and street enclosure regression tasks. Liang and Newsam [21] explore deep learning regression approaches to estimate the spatial resolution of very high-resolution overhead imagery and show that a stacked auto-encoder (SAE) frontend outperforms a standard CNN feature extractor. Xin and Adler [45] use a convolutional long short-term memory (Convolutional-LSTM) neural network model for grass identification in multi-temporal Sentinel-2 images. Tavakkol et al. [40] introduce the Kartta Labs open-data project aimed at unrendering and organizing the world's historical maps, and support user queries for the corresponding vector-based geospatial content as well as recreating the historical maps in various cartographic styles.

On the topic of transportation modeling, Yin et al. [53] propose a web-based analytic platform that leverages deep learning techniques in computer vision for vehicle plate number identification and illegal parking detection. Xing et al. [46] from Didi utilize street images captured by driving vehicle recorders (DVR) and develop a traffic sign discovery driven system for various types of traffic rule automatic update such as no left/right/U turn, no parking, speed limit, etc. Several works use GPS trajectories to gain insights. J. Krumm and K. Krumm [13] show how to use mobility traces from 1.8 million users to infer the number of businesses and residences in each 250m * 250m grid cell of the Seattle metropolitan area. Yin et al. [52] introduce a multi-scale graph CNN for road intersection detection from ride-hailing service GPS trajectories in Singapore. Mai et al. [24] learn the probability of pick-up and drop-off locations and times from taxi GPS trajectories and develop a spatial-temporal intelligent recommendation system to improve the expected net revenue of the electric-taxi.

Beyond geospatial image processing and transportation modeling, Soliman and Terstriep [34] develop Keras Spatial, a Python package for preprocessing and augmenting geospatial data. Yuan and Crooks [55] use a CNN to extract user opinions in natural language from more than 3 million user-contributed Yelp restaurant reviews.

They discover homogeneity among cities regarding the average proportions of aspects (i.e., terms and categories of semantics) in restaurant reviews. Snyder et al. [33] combine a deep neural network (Deepgeo) with a social media analytics and reporting toolkit (SMART) to infer the city-level geolocation of tweets and create real-time visual analytics. Finally, there are two vision papers focusing on the spatiotemporal intelligence of human behaviors and decision making. Li and Huang [20] outline a novel spatial-temporal imitation learning (STIL) framework that defines, investigates, and addresses the emerging research challenges of analyzing and learning human decision-making strategies from human-generated spatial-temporal data. And, Yang and Jankowska [51] share thoughts on how health behaviors can be contextualized in space and time through the use of GeoAI just-in-time adaptive interventions models, and highlight challenges such as on-the-fly feature engineering and spatiotemporal contextual uncertainty that need to be addressed in future work.

2.4 Summary

Table 1 groups the papers from the three GeoAI workshops based on their research topics. As can be seen, the majority of papers focus on geospatial image processing and transportation modeling. The popularity of these two topics can be attributed to the many advancements made by deep learning in image processing and the widespread interest of the ACM SIGSPATIAL community in transportation modeling. Other topics addressed include digital humanities, cartography, public health, disaster response, and social media analysis. Several papers specifically focus on methodological research, while almost all papers adapt and improve existing AI methods to target geographic problems. Interesting applications and visions, such as integrating AI and agent-based models (ABMs) for crowd analysis, have also been discussed.

3 Future Research Directions for GeoAI

Many global challenges are yet to benefit from GeoAI. With the ever growing collections of geospatial data we anticipate new GeoAI technological advances to emerge. However, near-human performance with AI seems to hinge on the availability ground-truth data for training. More opportunities could be opened up by taking advantage of the vast amount of unlabeled data for training unsupervised AI based methods. Furthermore, domain problems need to be fully understood and problem owners and AI experts from academia and industry need to be engaged during the design of the GeoAI solutions. Below, we list and organize possible future research directions for GeoAI into four categories.

Research directions and questions motivated by the uniqueness of geospatial problems and phenomena:

- Building spatiotemporally explicit models. In order to better understand the complex geospatial contexts and geographical process on the ground, it is crucial to employ spatiotemporally explicit models and evaluate the results by integrating both human intelligence and machine intelligence evaluations [49].
- Enhancing model generalizability in the context of geography. Geographical data sets are always collected from certain regions. How can we ensure that the GeoAI models trained using data from one geographic area can generalize to other geographic areas?
- Accommodating uncertainty in geospatial problems and datasets. Deep learning methods have traditionally not been designed with data uncertainty in mind yet uncertainty is a fundamental concept in geography. Can geography actually contribute to deep learning more broadly by developing methods to imbue the models with uncertainty analysis?

Research directions and questions motivated by the wealth of geospatial data:

- Developing new and open geospatial data infrastructures. ImageNet [6] played a key role in revolutionizing the field of computer vision. Future GeoAI applications could benefit from similar platforms by

Table 1: Papers presented in the three GeoAI workshops.

Research Topic	Related Papers in GeoAI Workshops
Geospatial image processing	- Li, W. et al. (2017) [19] - Law, S. et al. (2017) [16] - Collins, C.B. et al. (2017) [5] - Duan, W. et al. (2017) [8] - Xu, Y. et al. (2018) [47] - Sun, T. et al. (2018) [38] - Srivastava, S. et al. (2018) [35] - Chen et al. (2019) [3] - Dorji et al. (2019) [7] - Law and Neira (2019) [15] - Liang and Newsam (2019) [21] - Xin and Adler (2019) [45]
Transportation modeling and analysis	- Kulkarni, V. and Garbinato, B., (2017) [14] - Murphy, J. et al. (2017) [28] - Li, Q. et al. (2017) [17] - Sun, T. et al. (2018) [38] - Van Hinsbergh, J. et al (2018) [41] - Pourebrahim, N. et al (2018) [30] - Yin et al. (2019) [53] - J. Krumm and K. Krumm (2019) [13] - Xing et al. (2019) [46] - Yin et al. (2019) [52] - Mai et al. (2019) [24]
Digital humanities	- Duan, W. et al. (2017) [8] - Tavakkol et al. (2019) [40]
Public health	- Xi, G. et al. (2018) [44] - Yang and Jankowska (2019) [51]
Disaster response	- Peng et al. (2019) [29]
Social media and geo-text analysis	- Pourebrahim, N. et al (2018) [30] - Elgarroussi, K., et al. (2018) [9] - Yuan and Crooks (2019) [55] - Snyder et al. (2019) [33]
Methods and techniques	- Swan, B. et al. (2018) [39] - Aydin, O. et al. (2018) [1] - Soliman and Terstriep (2019) [34]
Novel applications and visions	- Majic, I. et al. (2017) [25] - Chow, T. E. (2018) [4] - Li and Huang (2019) [20]

investigating open and indexable rich geospatial data archives. A great example is the BigEarthNet platform [37] which has been demonstrated to be significantly larger than the existing archives in remote sensing and has been used as a diverse training source in the context of deep learning.

- Fusing multi-source geospatial data for knowledge discovery. The fusion of diverse geospatial datasets at different spatiotemporal resolutions through feature engineering and deep learning can enable novel geographic knowledge. How can machine learning help automate, streamline, or assist geospatial data integration?

Research directions and questions motivated by the need of labeled training data:

- Building domain datasets via novel techniques. As most GeoAI models involve supervised learning, the availability of large, benchmark datasets becomes essential to promote research across geospatial communities and to validate the generalizability of the research results. Can we leverage certain novel techniques to facilitate the development of domain datasets? For example, is there a role for generative adversarial networks (GANs) to play in augmenting GeoAI training data?
- Reducing the need of labeled data through self-supervised learning. Self-supervised learning methods have demonstrated their capability in other domains with their potentially unlimited capacity to uncover patterns in unlabeled data. Could they offer better scalability and reduce the need for labeled training data in the geography domain?

Research directions and questions motivated by the need of reliable and explainable models:

- Building robust and reliable GeoAI models. How can we ensure that future GeoAI systems are robust and reliable, and how do we evaluate them at scale?
- Building explainable GeoAI models. Most AI learning systems remain a black box. While these systems have demonstrated good performance in object inspection and classification, it is important to understand their learning and decision making process when applied to a variety of geospatial problems in both the physical and social sciences domains.

The research directions and questions discussed above are not exhaustive, and there exist many other important directions to be explored. We look forward to seeing exciting new research to be shared and published in the ACM SIGSPATIAL GeoAI workshops and other venues in the coming years.

References

[1] O. Aydin, M. V. Janikas, R. Assunção, and T.-H. Lee. SKATER-CON: Unsupervised regionalization via stochastic tree partitioning within a consensus framework using random spanning trees. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 33–42. ACM, 2018.

[2] Y. Chen, Z. Lin, X. Zhao, G. Wang, and Y. Gu. Deep learning-based classification of hyperspectral data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 7(6):2094–2107, 2014.

[3] Y. Chen, X. Ouyang, and G. Agam. ChangeNet: Learning to detect changes in satellite images. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 24–31. ACM, 2019.

[4] T. E. Chow. When GeoAI meets the crowd. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, GeoAI'18, pages 52–53, New York, NY, USA, 2018. ACM.

[5] C. B. Collins, J. M. Beck, S. M. Bridges, J. A. Rushing, and S. J. Graves. Deep learning for multisensor image resolution enhancement. In *Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery*, pages 37–44. ACM, 2017.

[6] 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*, pages 248–255. Ieee, 2009.

[7] U. J. Dorji, A. Plangprasopchok, N. Surasvadi, and C. Siripanpornchana. A machine learning approach to estimate median income levels of sub-districts in Thailand using satellite and geospatial data. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 11–14. ACM, 2019.

[8] W. Duan, Y.-Y. Chiang, C. A. Knoblock, V. Jain, D. Feldman, J. H. Uhl, and S. Leyk. Automatic alignment of geographic features in contemporary vector data and historical maps. In *Proceedings of the 1st workshop on artificial intelligence and deep learning for geographic knowledge discovery*, pages 45–54. ACM, 2017.

[9] K. Elgarroussi, S. Wang, R. Banerjee, and C. F. Eick. Aconcagua: A novel spatiotemporal emotion change analysis framework. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 54–61. ACM, 2018.

[10] N. Garg, L. Schiebinger, D. Jurafsky, and J. Zou. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644, 2018.

[11] Y. Hu, S. Gao, S. Newsam, and D. Lunga. GeoAI 2018 workshop report the 2nd acm sigspatial international workshop on GeoAI: AI for geographic knowledge discovery seattle, wa, usa-november 6, 2018. *SIGSPATIAL special*, 10(3):16–16, 2019.

[12] K. Janowicz, S. Gao, G. McKenzie, Y. Hu, and B. Bhaduri. GeoAI: Spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond. *International Journal of Geographical Information Science*, pages 1–13, 2020.

[13] J. Krumm and K. Krumm. Land use inference from mobility traces. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 1–4. ACM, 2019.

[14] V. Kulkarni and B. Garbinato. Generating synthetic mobility traffic using rnns. In *Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery*, pages 1–4. ACM, 2017.

[15] S. Law and D. M. N. Alvarez. An unsupervised approach to geographical knowledge discovery using street level and street network images. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 56–65. ACM, 2019.

[16] S. Law, Y. Shen, and C. Seresinhe. An application of convolutional neural network in street image classification: The case study of london. In *Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery*, pages 5–9. ACM, 2017.

[17] Q. Li, J. Zhu, T. Liu, J. Garibaldi, Q. Li, and G. Qiu. Visual landmark sequence-based indoor localization. In *Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery*, pages 14–23. ACM, 2017.

[18] W. Li and C.-Y. Hsu. Automated terrain feature identification from remote sensing imagery: a deep learning approach. *International Journal of Geographical Information Science*, pages 1–24, 2018.

[19] W. Li, B. Zhou, C.-Y. Hsu, Y. Li, and F. Ren. Recognizing terrain features on terrestrial surface using a deep learning model: An example with crater detection. In *Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery*, pages 33–36. ACM, 2017.

[20] Y. Li and W. Huang. Imitation learning from human-generated spatial-temporal data. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 9–10. ACM, 2019.

[21] H. Liang and S. Newsam. Estimating the spatial resolution of very high-resolution overhead imagery. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 77–80. ACM, 2019.

[22] D. Lunga, J. Gerrand, H. L. Yang, C. Layton, and R. Stewart. Apache spark accelerated deep learning inference for large scale satellite image analytics, 2019.

[23] A. Magge, D. Weissenbacher, A. Sarker, M. Scotch, and G. Gonzalez-Hernandez. Deep neural networks and distant supervision for geographic location mention extraction. *Bioinformatics*, 34(13):i565–i573, 2018.

[24] K. Mai, W. Tu, Q. Li, H. Ye, T. Zhao, and Y. Zhang. STIETR: Spatial-temporal intelligent e-taxi recommendation system using gps trajectories. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 5–8. ACM, 2019.

[25] I. Majic, S. Winter, and M. Tomko. Finding equivalent keys in openstreetmap: semantic similarity computation based on extensional definitions. In *Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery*, pages 24–32. ACM, 2017.

[26] H. Mao, Y. Hu, B. Kar, S. Gao, and G. McKenzie. Geoai 2017 workshop report: the 1st acm sigspatial international workshop on geoai:@ ai and deep learning for geographic knowledge discovery: Redondo beach, ca, usa-november 7, 2016. *SIGSPATIAL Special*, 9(3):25–25, 2018.

[27] D. Marcos, M. Volpi, B. Kellenberger, and D. Tuia. Land cover mapping at very high resolution with rotation equivariant cnns: Towards small yet accurate models. *ISPRS journal of photogrammetry and remote sensing*, 145:96–107, 2018.

[28] J. Murphy, Y. Pao, and A. Haque. Image-based classification of gps noise level using convolutional neural networks for accurate distance estimation. In *Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery*, pages 10–13. ACM, 2017.

[29] B. Peng, X. Liu, Z. Meng, and Q. Huang. Urban flood mapping with residual patch similarity learning. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 40–47. ACM, 2019.

[30] N. Pourebrahim, S. Sultana, J.-C. Thill, and S. Mohanty. Enhancing trip distribution prediction with twitter data: comparison of neural network and gravity models. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 5–8. ACM, 2018.

[31] R. Santos, P. Murrieta-Flores, P. Calado, and B. Martins. Toponym matching through deep neural networks. *International Journal of Geographical Information Science*, 32(2):324–348, 2018.

[32] M. Siam, S. Elkerdawy, M. Jagersand, and S. Yogamani. Deep semantic segmentation for automated driving: Taxonomy, roadmap and challenges. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–8. IEEE, 2017.

[33] L. Snyder, M. Karimzadeh, R. Chen, and D. Ebert. City-level geolocation of tweets for real-time visual analytics. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 85–88. ACM, 2019.

[34] A. Soliman and J. Terstriep. Keras Spatial: Extending deep learning frameworks for preprocessing and on-the-fly augmentation of geospatial data. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 69–76. ACM, 2019.

[35] S. Srivastava, J. E. Vargas Muñoz, S. Lobry, and D. Tuia. Fine-grained landuse characterization using ground-based pictures: a deep learning solution based on globally available data. *International Journal of Geographical Information Science*, pages 1–20, 2018.

[36] S. Srivastava, J. E. Vargas-Muñoz, D. Swinkels, and D. Tuia. Multilabel building functions classification from ground pictures using convolutional neural networks. In *Proceedings of the 2nd ACM SIGSPATIAL international workshop on AI for geographic knowledge discovery*, pages 43–46. ACM, 2018.

[37] G. Sumbul, M. Charfuelan, B. Demir, and V. Markl. Bigearthnet: A large-scale benchmark archive for remote sensing image understanding. *CoRR*, abs/1902.06148, 2019.

[38] T. Sun, Z. Di, and Y. Wang. Combining satellite imagery and gps data for road extraction. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 29–32. ACM, 2018.

[39] B. Swan, M. Laverdiere, and H. L. Yang. How good is good enough?: Quantifying the effects of training set quality. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 47–51. ACM, 2018.

[40] S. Tavakkol, Y.-Y. Chiang, T. Waters, F. Han, K. Prasad, and R. Kiveris. Kartta labs: Unrendering historical maps. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 48–51. ACM, 2019.

[41] J. Van Hinsbergh, N. Griffiths, P. Taylor, A. Thomason, Z. Xu, and A. Mouzakitis. Vehicle point of interest detection using in-car data. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 1–4. ACM, 2018.

[42] S. Wang. A cybergis framework for the synthesis of cyberinfrastructure, gis, and spatial analysis. *Annals of the Association of American Geographers*, 100(3):535–557, 2010.

[43] X. Wang, C. Ma, H. Zheng, C. Liu, P. Xie, L. Li, and L. Si. Dm_nlp at semeval-2018 task 12: A pipeline system for toponym resolution. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 917–923, Stroudsburg, PA, USA, 2019. ACL.

[44] G. Xi, L. Yin, Y. Li, and S. Mei. A deep residual network integrating spatial-temporal properties to predict influenza trends at an intra-urban scale. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 19–28. ACM, 2018.

[45] Y. Xin and P. R. Adler. Mapping miscanthus using multi-temporal convolutional neural network and google earth engine. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 81–84. ACM, 2019.

[46] T. Xing, Y. Gu, Z. Song, Z. Wang, Y. Meng, N. Ma, P. Xu, R. Hu, and H. Chai. A traffic sign discovery driven system for traffic rule updating. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 52–55. ACM, 2019.

[47] Y. Xu, L. Pan, C. Du, J. Li, N. Jing, and J. Wu. Vision-based uavs aerial image localization: A survey. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 9–18. ACM, 2018.

[48] Y. Xu, Z. Piao, and S. Gao. Encoding crowd interaction with deep neural network for pedestrian trajectory prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5275–5284, 2018.

[49] B. Yan, K. Janowicz, G. Mai, and S. Gao. From itdl to place2vec: Reasoning about place type similarity and relatedness by learning embeddings from augmented spatial contexts. In *Proceedings of the 25th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, page 35. ACM, 2017.

[50] H. L. Yang, J. Yuan, D. Lunga, M. Laverdiere, A. Rose, and B. Bhaduri. Building extraction at scale using convolutional neural network: Mapping of the united states. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 11(8):2600–2614, Aug 2018.

[51] J.-A. Yang and M. Jankowska. Contextualizing space and time for geoai jitais (just-in-time adaptive interventions). In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 66–68. ACM, 2019.

[52] Y. Yin, A. Sunderrajan, X. Huang, J. Varadarajan, G. Wang, D. Sahrawat, Y. Zhang, R. Zimmermann, and S.-K. Ng. Multi-scale graph convolutional network for intersection detection from gps trajectories. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 36–39. ACM, 2019.

[53] Z. Yin, H. Xiong, X. Zhou, D. Goldberg, D. Bennett, and C. Zhang. A deep learning based illegal parking detection platform. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 32–35. ACM, 2019.

[54] J. Yuan. Learning building extraction in aerial scenes with convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 40(11):2793–2798, 2017.

[55] X. Yuan and A. Crooks. Assessing the placeness of locations through user-contributed content. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 15–23. ACM, 2019.

[56] F. Zhang, B. Du, and L. Zhang. Scene classification via a gradient boosting random convolutional network framework. *IEEE Transactions on Geoscience and Remote Sensing*, 54(3):1793–1802, 2015.

[57] X. X. Zhu, D. Tuia, L. Mou, G.-S. Xia, L. Zhang, F. Xu, and F. Fraundorfer. Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geoscience and Remote Sensing Magazine*, 5(4):8–36, 2017.