[AM-08-093] Artificial Intelligence Approaches

Abstract

Artificial Intelligence (AI) has received tremendous attention from academia, industry, and the general public in recent years. The integration of geography and AI, or GeoAI, provides novel approaches for addressing a variety of problems in the natural environment and our human society. This entry briefly reviews the recent development of AI with a focus on machine learning and deep learning approaches. We discuss the integration of AI with geography and particularly geographic information science, and present a number of GeoAI applications and possible future directions.

Keywords: AI, artificial intelligence (AI), deep learning

Author & citation

Hu, Y., Li, W., Wright, D., Aydin, O., Wilson, D., Maher, O, and Raad, M. (2019). Artificial Intelligence Approaches. The Geographic Information Science & Technology Body of Knowledge (3rd Quarter 2019 Edition), John P. Wilson (ed.).
DOI: 10.22224/gistbok/2019.3.4.

Explanation

1. Definitions
2. AI and Geography
3. Main AI Approaches
4. Applications of AI in Geography
5. Possible Future Directions

1. Definitions

artificial intelligence: the study and design of machines or computational methods that can perform tasks that normally require human intelligence

GeoAI: an interdisciplinary field of geography and artificial intelligence

machine learning: a sub field in AI that relies on statistical methods or numerical optimization techniques to derive models from data without explicitly programming every model parameter or computing step

deep learning: a special type of machine learning that leverages multiple layers of nonlinear processing units, or neurons, to learn representations from raw data to achieve the goal of automatic learning for completing various AI tasks

2. AI and Geography

Artificial Intelligence (AI) has received tremendous attention in recent years from academia,

industry, and the general public. Despite its recent popularity, the field was born back in 1956 at a workshop at Dartmouth College (McCarthy, 1956). AI is a broad field from its beginning, and has many different definitions (Russell et al., 2003). Some definitions focus on designing intelligent machines that can act like humans. For example, the famous Turing Test was designed to see if the responses of a machine can be indistinguishable from those of a real person (Turing, 1950). Some other definitions focus on designing and developing computational methods to complete tasks that typically require human intelligence, such as recognizing objects from images or understanding the meaning of natural language sentences. This entry is primarily based on the second type of definitions.

The development of AI has experienced falls and rises. Following its early optimism in 1960s and 70s, AI research went through the “AI winter” due to the failures of AI methods in addressing real-world problems. The following decades witnessed several other waves of optimism and disappointment. Since the 21st century, and especially after 2010, there has been significant progress in AI research. Three major factors have contributed to this fast advancement of AI: big data, novel algorithms, and immense computational power. The emergence of ubiquitous sensors and user-generated content on the Web allows large amounts of data to be generated and collected at a rapid pace. Big data enables computers to “observe” many different aspects of the world, to learn the ways in which the world functions, and to predict the future based on existing observations. Meanwhile, novel algorithms and models have been developed, and the AI community has embraced various ideas and theories from other fields, such as statistics, economics, biology, and cognitive science, in addition to its tradition of logics. Third, high performance computing (HPC) provides the essential power for linking big data and new computational models, and allows the training of sophisticated models on large datasets to be completed within hours or days rather than weeks or months. These three major factors, namely big data, novel algorithms, and immense computational power, greatly fueled the remarkable development of AI in recent years.

In this context, it is probably unsurprising to see the integration of AI and geography, particularly GIScience and GI Systems (GIS). A high volume of data nowadays contains georeferenced information, that is, information about locations on or near the surface of the Earth. Examples include GPS trajectories, remote sensing images, location-based social media, spatial footprints of buildings, roads, and parcels, global elevation data, land use and land cover data, population distribution, and so forth. These georeferenced data are critical inputs for many models that address a wide range of problems related to our human society and the natural environment, and GIS is essential for effectively processing, managing, and visualizing these big geo data. In addition, many AI models need to synthesize heterogeneous data from different sources, while geographic location is often the only factor that can link such heterogeneous datasets. Finally, existing research has already demonstrated successful integrations between GIS and HPC (Wang, 2010), and integrated systems of GIS and AI also have the capability of leveraging high computational power.

Many efforts have already been devoted by both academia and industry to facilitating the integration of geography and AI, and the outcome is an exciting new and interdisciplinary area--GeoAI. In academia, the first GeoAI workshop was held in November 2017 in a major GIS conference, ACM SIGSPATIAL, which attracted more than 100 participants (Mao et al., 2018a). In April 2018, the first AI and Deep Learning Symposium was held in the largest geography conference, Annual Meeting of American Association of Geographers (AAG). In

industry, one of the major GIS companies, Esri Inc., has been collaborating with Microsoft to offer the GeoAI Data Science Virtual Machine (DSVM) with the goal of bringing AI, cloud computing, geospatial analytics, and geo-visualizations together. It is also worth noting that there are already studies reported in the literature that leverage machine learning for geographical problem-solving before the GeoAI workshop, and many of which, such as hyperspectral image analysis (Chen et al., 2014), high-resolution satellite image interpretation (Zhang et al., 2015), and 3D reconstruction (Blaha et al., 2016), are from the remote sensing community (Tuia et al., 2009; Zhu et al., 2017). Given this continuing interest, it is likely that there will be more interactions between geography and AI in the foreseeable future. This entry aims to provide interested readers, especially students who will become future GIS leaders, with some background of GeoAI. This field is rapidly developing, and this entry, completed between middle 2018 and early 2019, reflects the progress so far. New advancements are likely to happen in the following years.

3. Main AI Approaches

AI, as a broad field, encompasses many different approaches ranging from top-down knowledge representation to bottom-up machine learning. There are three related concepts that have been frequently used in recent years: AI, machine learning, and deep learning. In general, AI is the broadest concept, machine learning is a sub field in AI, and deep learning is a special type of machine learning. Figure 1 illustrates the relations among these three concepts. While the broad field of AI includes many approaches, its recent popularity is largely due to the outstanding performances of machine learning, especially deep learning. Therefore, this entry focuses on discussing these two types of AI approaches.

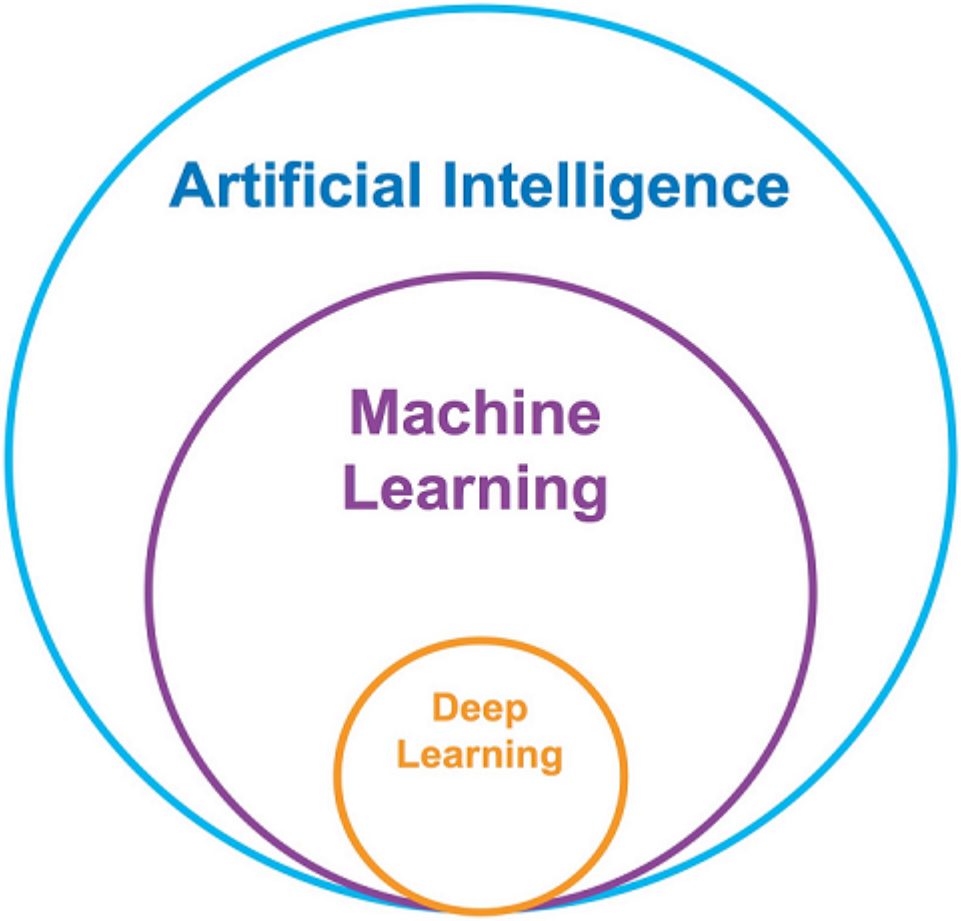


Figure 1. Relations among AI, machine learning, and deep learning (Bennett 2018).

3.1 Machine learning

Machine learning is a sub field in AI (see Figure 1) that usually relies on statistical methods or numerical optimization techniques to derive models from data without explicitly programming every model parameter or computing step (Valiant, 1984). One important characteristic shared by many machine learning models is the use of probability to represent the uncertainty that widely exists in real-world problems. There are three main types of learning: supervised learning, unsupervised learning, and reinforcement learning. Supervised learning requires labeled data for training a computational model, while unsupervised learning examines unlabeled data to discover patterns. Reinforcement learning does not require labeled data but needs action-based feedbacks, such as rewards or punishments, to help a computational model to learn. Machine learning tasks can be categorized in various ways. Based on their goals, we can identify tasks such as classification, clustering, and prediction (Bennett, 2018). In classification, the goal is to classify a target into a category, e.g., classifying a land parcel into a category such as Commercial or Agricultural. In clustering, the goal is to detect clusters from data, e.g., finding the clusters of vehicles based on their locations to detect traffic congestions. In prediction, the goal is to predict unknown values, e.g., predicting the average temperatures of several locations in the near future based on their historical temperatures and other variables using a regression model. There are also other tasks that can fall under the umbrella of machine learning, such as anomaly/novelty detection, data generation,

visualization, feature learning, and others. A variety of machine learning models have been developed, such as regression, decision tree, random forest, support vector machine (SVM), naïve Bayesian classifier, density-based clustering, hidden Markov model (HMM), artificial neural network (ANN), and numerous others. These methods are discussed with details in machine learning textbooks, such as Flach (2012). While most machine learning methods can be directly applied to geographic data, they typically do not take into account the uniqueness of geographic phenomena, such as spatial autocorrelation and spatial nonstationarity. There exist some methods, such as Empirical Bayesian Kriging (EBK) regression (Krivoruchko, 2012) and spatial Principal Component Analysis (sPCA) (Jombart et al., 2008), that explicitly model the spatial aspect of geographic problems by, e.g., including spatial weights. Some classic spatial models, such as geographically weighted regression (GWR) (Fotheringham et al., 2003), can also be used in a machine learning manner, i.e., they can be firstly trained on one dataset and then tested on other datasets.

3.2 Deep learning

Deep learning is a special type of machine learning that focuses on developing and using deep neural networks (DNN) for machine learning tasks. DNN is a special type of artificial neural network which has multiple layers (also called hidden layers) between the input and the output layers. Each layer consists of a collection of computing units, called neurons, which take the input from the previous layer and generate a non-linear output to the next layer. Deep learning has gained a huge amount of interest in recent years due to its outstanding performances (LeCun et al., 2015), thanks to the availability of large labeled datasets, such as ImageNet (Deng et al., 2009), and HPC. Similar to other machine learning models, deep learning can be utilized to complete tasks in classification, clustering, prediction, and so forth. Particularly, two types of DNN, convolutional neural network (CNN) and recurrent neural network (RNN), have received a lot of attention from the geography community. CNN is especially suitable for processing images by extracting and representing abstract features through a cascade of neuron layers and using convolutional filters (Li et al., 2017a; Maggiori et al., 2017; Gong et al., 2018). RNN is suitable for processing sequence data, such as movement trajectories (which can be modeled as a sequence of locations) (Kulkarni & Garbinato, 2017), by memorizing some of the previous states and establishing links between the current and previous states. While many studies applied existing models to geographical problems, researchers also developed new DNN models specifically for handling geographic data. Marcos et al. (2018) proposed Rotation Equivariant Vector Field Network (RotEqNet) for land cover mapping based on remote sensing images. RotEqNet encodes rotation equivariance in a CNN, and can recognize the rotated versions of the same object from remote sensing images while reducing the number of parameters required by a traditional CNN. Srivastava et al. (2018) proposed a Variable Input Siamese Convolutional Neural Network (VIS-CNN) model for classifying urban-object level land use types. Their VIS-CNN model can accept a variable number of Google Street View images for an urban object and aggregate them to learn the land use type in an end-to-end manner.

4. Applications of AI in Geography

There exist a considerable number of applications of AI in the domain of geography. This section summarizes some of these applications.

4.1 Automatic recognition of natural terrain features from remote sensing

imagery

Natural terrain features, such as craters, volcanos, and sand dunes, are important indicators of the Earth geological process. Detecting where they are and extracting their geomorphological properties are of great importance to geographers and geologists in understanding the formation process of different terrain features, differentiating similar landscapes, as well as enriching our geospatial knowledge. While this task has been done previously primarily using Object-Based Image Analysis (OBIA), it can hardly achieve automation of the processing steps. Parameters, such as scale factor, and the merging of segmented super pixels (a cluster of pixels with similar values) are often done manually or semi-automatically. With the proliferation of spatial big data, such as massive remote sensing imagery, as well as the fast evolution of deep learning techniques, it becomes possible for implementing an automatic learning process to detect and characterize different terrain features.

Li et al. (2017b) extended a Faster-RCNN (Faster Region-based CNN) model to support the automatic detection of natural terrain features from remote sensing imagery. Different from the classification tasks using CNN, object detection not only needs to tell what is inside of an image, but also the location of the object, depicted by a bounding box (BBOX). This requires additional location labels in the training data as well. Several challenges were addressed in applying deep learning to such spatial problems. This includes a lack of training database, the ambiguous boundary of natural features in comparison to man-made features, such as building or golf courses, and the inter-category similarity of different features. For instance, there is usually a crater type of feature on top of a volcano, called “volcanic crater”, which looks similar to impact craters. Making such distinctions may require the deep learning network to learn features that is unique to each (sub)category, such as emitting ashes from active volcano and flatter terrain that an impact crater usually appears in. Based on a test database of over 10,000 images, the deep learning network achieves over 90% of mean average precision (mAP) for the task of detecting eight terrain categories (Li & Hsu, 2018). The derived information, both category and BBOX, can complement the limited spatial information which is only a point for representing features of any extent in known gazetteers, such as the USGS (United States Geological Survey) GNIS (Geographic Name Information System), hence providing better support for feature allocation and landscape interpretation. Figure 2 demonstrates detection results for hill, impact crater, meander, and volcano.

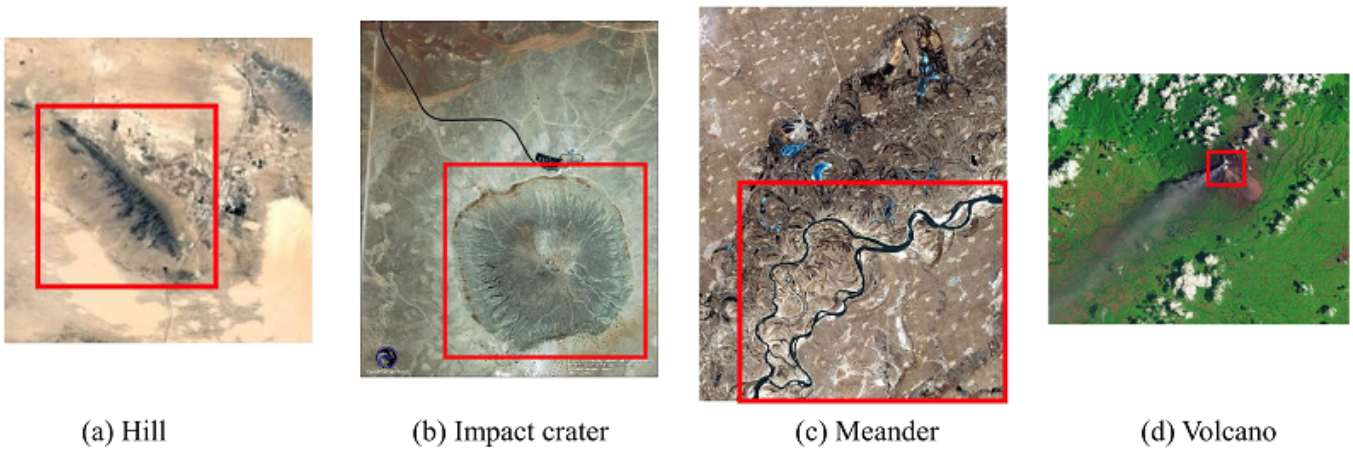


Figure 2. An illustration of terrain feature detection results of hill (a), impact crater (b),

meander (c), and volcano (d) from remote sensing imagery.

4.2 Land cover classification for conservation

Another excellent example of AI application in geography is the work of the Chesapeake Conservancy, a non-profit organization based in Annapolis, MD, and a pioneer in the field of precision conservation to monitor, protect, and restore the natural environment. Their "mission field" covers the entire Chesapeake Bay watershed, supports more than 3,600 species of plants and animals, and is home to 17 million people. Despite the importance of the Chesapeake Bay watershed, its health over decades has been marred by sewage overflows and runoff of animal waste and chemicals. The future of the watershed depends upon smart conservation that is informed by data about where growth is least harmful and where interventions can be most helpful. This requires detailed and highly accurate land cover maps to understand how land is used and to identify environmental issues, such as pollutants entering the bay through agricultural runoff. While Landsat imagery offer a standard 30-meter resolution, they are not sufficient for achieving full understanding of all the possible scenarios of 21st-century land use and associated disruptions of nature. Land cover data at 1-m resolution are not yet available everywhere across the US and is often difficult and expensive to obtain, much less to process and interpret, especially for less affluent communities, smaller nonprofits, and the like. A driving motivation of the Chesapeake Conservancy is to empower natural resource managers and conservationists with access to such a high-resolution data resource.

To achieve this goal, the Chesapeake Conservancy took a deep learning approach to create a semantic segmentation (pixel-level classification) model which predicts high-resolution land cover from aerial imagery (Allenby et al., 2018). The imagery was obtained from the National Agriculture Imagery Program (NAIP), at a 1-m or 60-cm resolution. Figure 3 illustrates this approach based on a location within the Chesapeake watershed. In the upper left is NAIP imagery from the Esri Living Atlas of the World (Kensok, 2017); in the upper right is the original Chesapeake land cover map; and then immediately below that in the lower right is the map as produced by a deep residual neural network executed on that NAIP imagery, on the fly, classifying forests, fields, water, and impervious surfaces such as roads and houses; in the lower left is the same region, except it's a mixture of probabilities from the model, across all those categories, currently at 91% accuracy as compared to ground observations. This was extremely useful as it communicates the uncertainty in the detections and allows one to see how the algorithm was "thinking" and what was needed to improve it. The real power of this approach is that data scientists can use this same algorithm to classify land cover in places that it has never seen before. The system can classify at a rate significantly faster than manual methods, limited only by the available hardware. The time needed to correctly classify the entire Chesapeake Bay watershed was reduced from 2500 hours to 150. The goal is to take this AI approach to scale across all the other watersheds of the US, and to empower more organizations in more places to sustainably manage their lands.

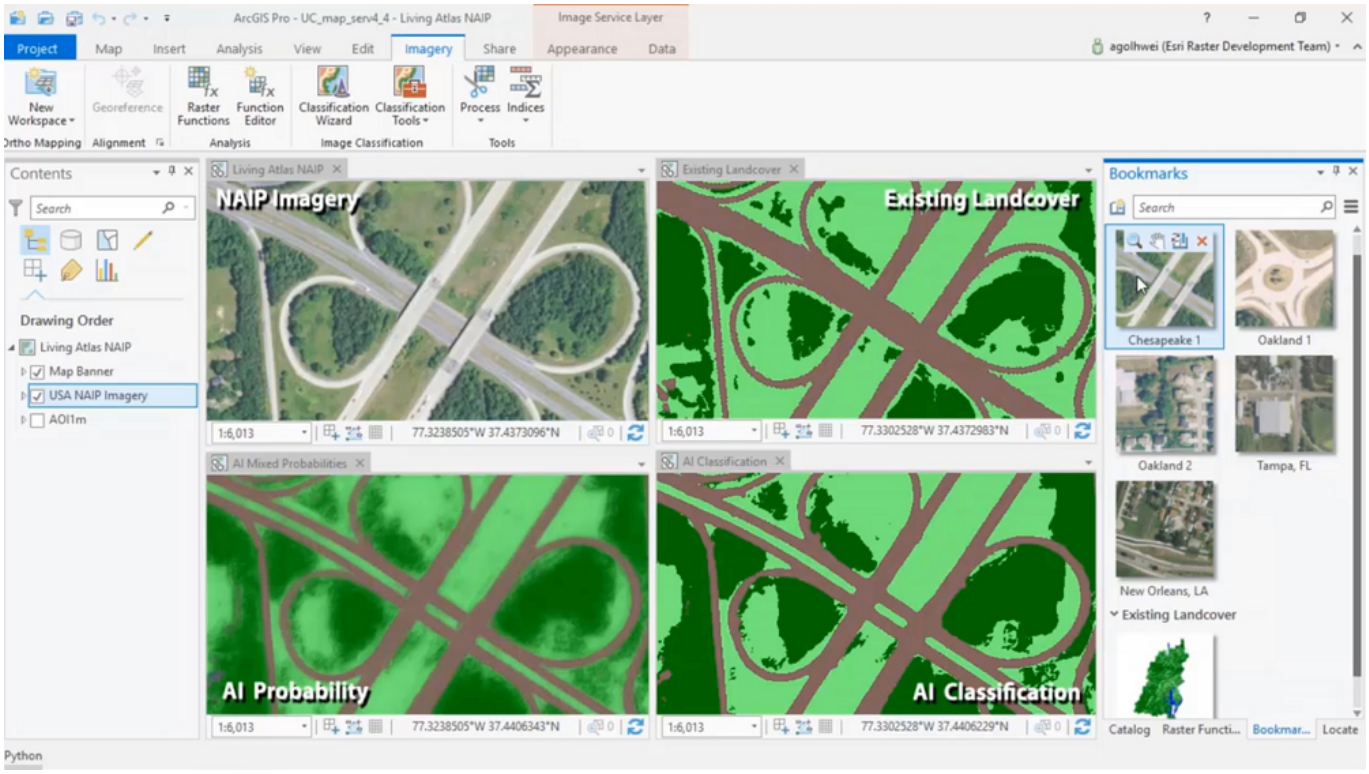


Figure 3. High-resolution land cover classification of Chesapeake Bay watershed using a deep residual neural network implemented in ArcGIS Pro.

4.3 Modelling seagrass habitats in space and time

In this AI application, GIS and machine learning methods are integrated to model the relation between seagrass habitats and ocean conditions. Seagrasses are marine plants that can quickly sequester vast amounts of CO2, up to 100 times more and 12 times faster compared to tropical forests (Parry et al., 2007; Pidgeon, 2009). Only limited amounts of data on global seagrass habitats are available, and existing data are often spatially sparse. Therefore, developing a seagrass habitat model based on existing data can help quantitatively understand the ocean conditions that favor the growth of seagrass. In addition, such a model can also predict future seagrass habitats based on changing ocean conditions. Multiple data sources, including the seagrass data from MarineCadastre.gov and the recently available Ecological Marine Units (EMU) dataset (Wright et al., 2017) were used in this analysis. These data were employed to train a model based on the seagrass occurrence along the U.S. coast to predict seagrass habitats globally up to a depth of 90 meters. In situ oceans data are interpolated using Empirical Bayesian Kriging (EBK) to produce global-scale predictions. A random forest framework is then used to relate the occurrence and absence of seagrass to a set of variables including temperature, salinity, phosphate, silicate, nitrate, and dissolved oxygen. Figure 4(a) shows the global-scale predictions on seagrass habitats using U.S. coast data for training, which is juxtaposed against a map in Figure 4(b) by Short et al. (2007) of seagrass observations.

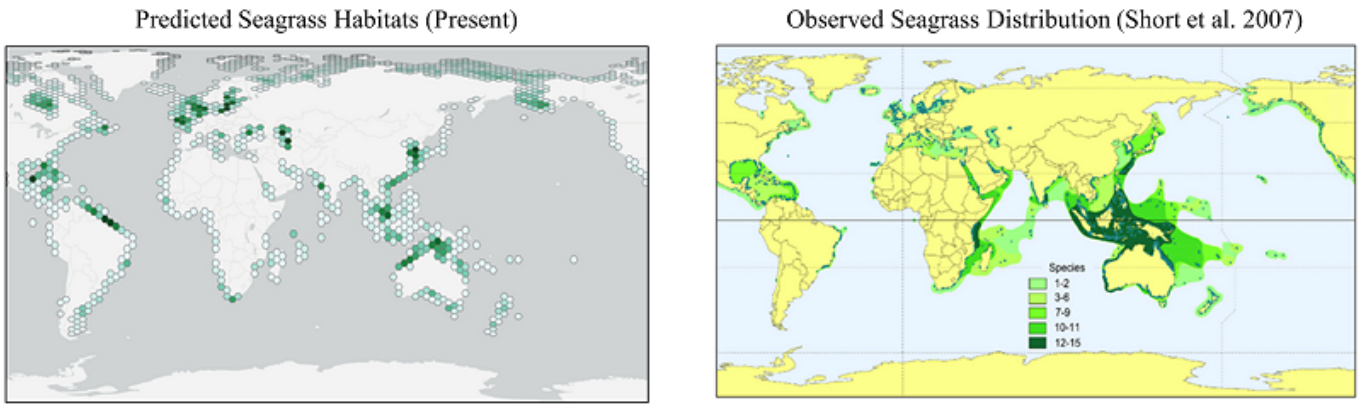


Figure 4. (a) Predicted seagrass habitats using a random forest model (dark green indicates high modelled abundance); (b) reported seagrass occurrence data by Short et al. (2007).

The random forest model trained on existing seagrass habitats is then applied to different ocean conditions to forecast the possible scenarios of seagrass habitats due to the warming of the oceans. An increase of 2 ℃ in average ocean temperature is simulated with an increment of 0.2 ℃. For every snapshot of ocean condition, seagrass habitats are predicted using the trained random forest model. Getis-Ord Gi* statistic (Ord & Getis, 1995) is used for detecting the clusters of seagrass abundance over space-time bins. Assessing changes in the intensity of abundance at individual locations over time using the Mann-Kendall statistic yields the emerging hot spot map in Figure 5.

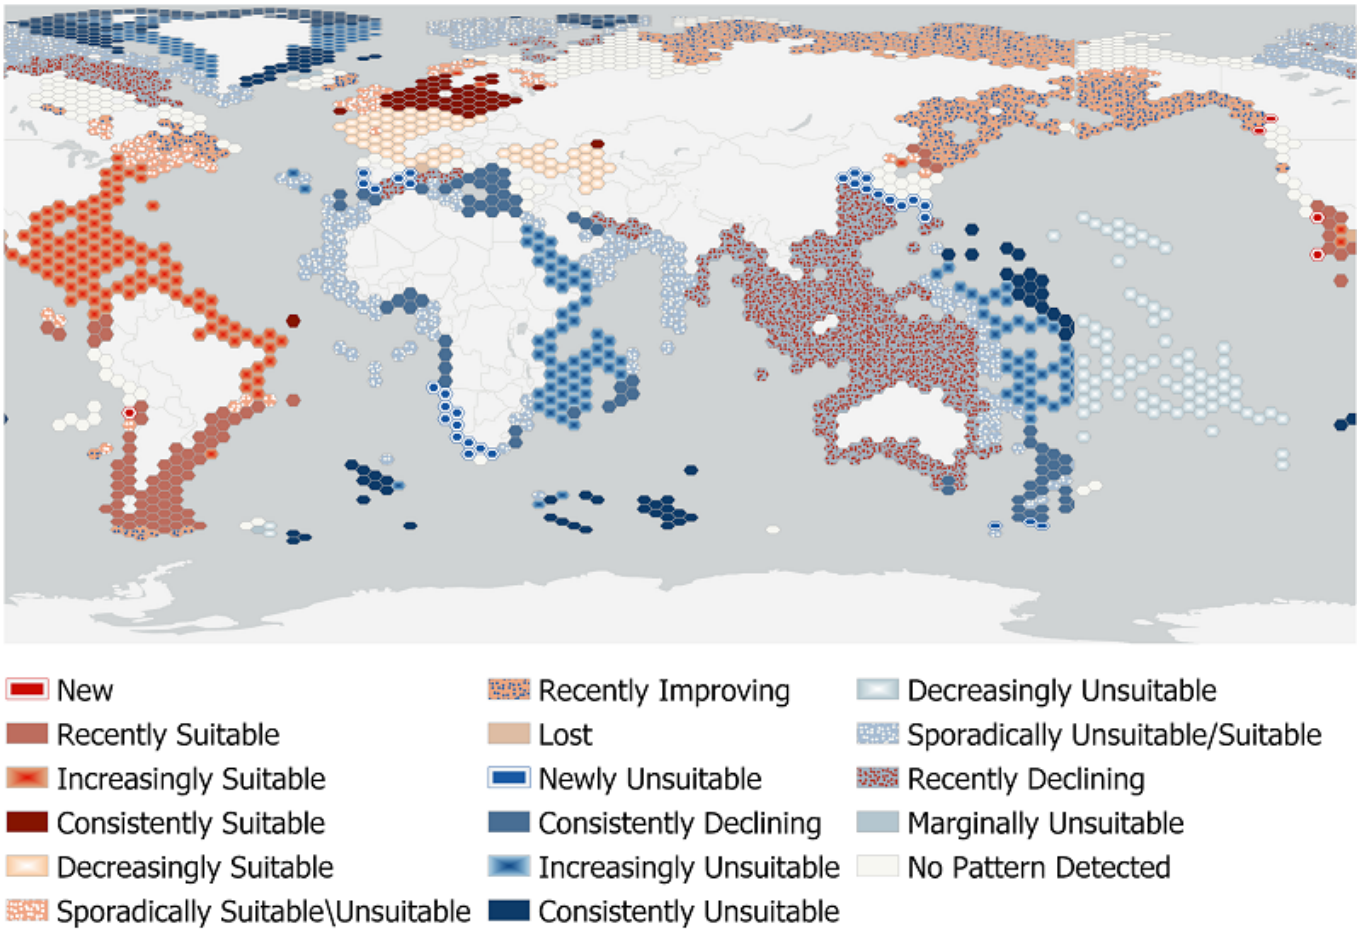


Figure 5. Emerging hot spot map for seagrass habitats under increasing ocean temperature.

As shown in the figure, Australia could lose its seagrasses under changing ocean conditions, whereas Siberia coast may improve its suitability for seagrass habitats. The use of random forest in this AI application enables a data-driven model of seagrass habitats for world's coasts. A tutorial demonstrating the prediction of seagrass with machine learning methods can be found in Additional Resource 5.

4.4 Summary and some other applications of AI in geography

The three studies discussed above are summarized in Table 1.

Table 1. A summary of the three studies discussed from 3.1 to 3.3.

Study	Task Type	Learning Type	Model(s)
Automatic recognition of natural terrain features	Classification	Supervised learning	Deep learning: faster region-based CNN
High-resolution land cover classification	Classification	Supervised learning	Deep learning: deep residual neural network
Global-scale seagrass habitat prediction	Prediction	Supervised learning	Machine learning: EBK and random forest

As shown in Table 1, the first study addresses a classification task (more specifically, object detection) which leverages deep learning to identify the category of a terrain feature and its location in the image. The second study is also a classification task (semantic segmentation), in which a pixel-level land cover classification is achieved. The third study completes a prediction task which predicts the growth of seagrass in simulated situations. All three studies employ supervised learning in which the models are first trained using labeled training data. While supervised learning is indeed very common, many studies also employ unsupervised learning, such as spatial and spatio-temporal clustering approaches, to extract geographic patterns from data (Anbaroglu et al., 2014; Hu et al., 2015).

There exist other applications that integrate AI techniques with geographic research. CNNs and their variants, in addition to their outstanding performances in processing satellite imagery (Collins et al., 2017) and Google Street View photos (Law et al., 2017), are also utilized for aligning vector data to raster historical map images (Duan et al., 2017), detecting the types of online maps (Zhou et al., 2018), classifying geographic texts (Adams & McKenzie, 2018), and so forth. RNNs, such as Long Short-Term Memory (LSTM), are utilized to handle time series data, such as predicting the next locations of trajectories (Li et al., 2018) and examining the temporal patterns of crops (Sun et al., 2018). RNNs are also used for analyzing geotagged tweets and other natural language texts containing geographic information (Mao et al., 2018b; Sit et al., 2019; Santos et al., 2018). Machine learning models, such as hidden Markov model, are integrated with a variety of geospatial applications, such as indoor navigation (Li et al., 2017a) and location prediction of financial services (McKenzie & Slind, 2019). While many applications already exist, new ideas and methods are being constantly developed.

5. Possible Future Directions

GeoAI is a rapidly growing field with many possible directions to be pursued in the near future. Here, we list a few of these directions. First, most AI methods are currently applied to pre-defined and highly specific spatial analysis tasks. Can there be a general GeoAI assistant, similar to Amazon's Alexa or Apple's Siri? Such a GeoAI assistant may be able to understand the needs of a GIS practitioner, automatically formalize and define the tasks, and identify candidate tools from a large GIS toolbox. Second, most of today's AI models are developed based on a training dataset, and as a result, they naturally inherit the potential bias in the data. In geographic research, training data are often collected from a certain geographic area, and consequently, it can be difficult for a model trained using the data from one geographic area to perform well on the data from other areas. An important direction, therefore, is to improve model architectures (or the training process) so that the obtained GeoAI models can be transferred across different geographic areas. Third, many existing GeoAI studies only apply AI methods to geographic problems rather than improving or inventing methods. While this is fine from a perspective of problem solving, it is critical for geographers to not only import methods from outside disciplines but also export geographic knowledge to other fields. Geographically-informed or spatially-explicit AI models can be developed to capture the uniqueness of geographic problems.

References

Adams, B. and McKenzie, G. (2018). Crowdsourcing the character of a place: Character-level convolutional networks for multilingual geographic text classification. Transactions in GIS, 22(2):394–408.

Allenby, J., Joppa, L., & Jojic, N. (2018). Precision conservation. In Wright, D. and Harder, C. (Eds.) GIS for Science: How Mapping and Spatial Analytics Advance Research. Redlands, California: Esri Press.

Anbaroglu, B., Heydecker, B., & Cheng, T. (2014). Spatio-temporal clustering for non-recurrent traffic congestion detection on urban road networks. Transportation Research Part C: Emerging Technologies, 48, 47-65.

Bennett, L. (2018). Machine learning in ArcGIS. ArcUser, 21(2), 8-9.

Blaha, M., Vogel, C., Richard, A., Wegner, J. D., Pock, T., & Schindler, K. (2016). Large-scale semantic 3d reconstruction: an adaptive multi-resolution model for multi-class volumetric labeling. ed. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 3176-3184.

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

Collins, C. B., Beck, J. M., Bridges, S. M., Rushing, J. A., & Graves, S. J. (2017). Deep learning for multisensor image resolution enhancement. Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery, 37-44.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database," 2009 IEEE Conference on Computer Vision and Pattern

Recognition, Miami, FL, USA. pp. 248-255.

Duan, W., Chiang, Y.-Y., Knoblock, C. A., Jain, V., Feldman, D., Uhl, J. H., & Leyk, S. (2017). Automatic alignment of geographic features in contemporary vector data and historical maps. Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery, 45-54.

Flach, P. (2012). Machine Learning: the Art and Science of Algorithms that Make Sense of Data. Cambridge University Press.

Fotheringham, A. S., Brunsdon, C., & Charlton, M. (2003). Geographically Weighted Regression: The Analysis of Spatially Varying Relationships. John Wiley & Sons.

Gong, F.-Y., Zeng, Z.-C., Zhang, F., Li, X., Ng, E., & Norford, L. K. (2018). Mapping sky, tree, and building view factors of street canyons in a high-density urban environment. Building and Environment, 134, 155-167.

Hu, Y., Gao, S., Janowicz, K., Yu, B., Li, W., & Prasad, S. (2015). Extracting and understanding urban areas of interest using geotagged photos. Computers, Environment and Urban Systems, 54, 240-254.

Jombart, T., Devillard, S., Dufour, A., & Pontier, D. (2008). Revealing cryptic spatial patterns in genetic variability by a new multivariate method. Heredity, 101(1), 92.

Kensok, D. (2017). NAIP 2016 Image Layer is Now Available. ArcGIS Blog. Accessed August 16, 2018.

Krivoruchko, K. (2012). Empirical Bayesian Kriging. ArcUser, 6-10.

Kulkarni, V., & Garbinato, B. (2017). Generating synthetic mobility traffic using RNNs. GeoAI '17: Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery, 1-4.

Law, S., Shen, Y., & Seresinhe, C. (2017). An application of convolutional neural network in street image classification: the case study of London. Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery, 5-9.

LeCun, Y. L., Bengio, Y., & Hinton, G. (2015). Deep Learning. Nature, 521(7553), 436-444.

Li, M., Lu, F., Zhang, H., & Chen, J. (2018). Predicting future locations of moving objects with deep fuzzy-LSTM networks. Transportmetrica A: Transport Science, 1-18.

Li, Q., Zhu, J., Liu, T., Garibaldi, J., Li, Q., & Qiu, G. (2017a). Visual landmark sequence-based indoor localization. Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery, 14-23.

Li, W., & Hsu, C.-Y. (2018). Automated terrain feature identification from remote sensing imagery: a deep learning approach. International Journal of Geographical Information

Science, 1-24.

Li, W., Zhou, B., Hsu, C.-Y., Li, Y., & Ren, F. (2017b). Recognizing terrain features on terrestrial surface using a deep learning model: an example with crater detection. GeoAI '17: Proceedings of the 1st Workshop on Artificial Intelligence and Deep Learning for Geographic Knowledge Discovery, 33-36.

Maggiori, E., Tarabalka, Y., Charpiat, G., & Alliez, P. (2017). Convolutional Neural Networks for Large-Scale Remote-Sensing Image Classification. IEEE Transactions on geoscience and remote sensing, 55(2), 645-657.

Mao, H., Hu, Y., Kar, B., Gao, S., & Mckenzie, G. (2018). 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.

Mao, H., Thakur, G., Sparks, K., Sanyal, J., & Bhaduri, B. (2018a). Mapping near-real-time power outages from social media. International Journal of Digital Earth, 12(11):1285-1299.

Marcos, D., Volpi, M., Kellenberger, B., & Tuia, D. (2018b). 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.

McCarthy, J., Minsky, M. L., Rochester, N., and Shannon, C. E. (1956). The Dartmouth Summer Research Project on Artificial Intelligence. Dartmouth College, New Hampshire.

McKenzie, G., & Slind, R. T. (2019). A user-generated data based approach to enhancing location prediction of financial services in sub-Saharan Africa. Applied Geography, 105, 25-36.

Ord, J. K., & Getis, A. (1995). Local Spatial Autocorrelation Statistics: Distributional Issues and an Application. Geographical Analysis, 27(4), 286-306.

Parry, M., Canziani, O., Palutikof, J., Van Der Linden, P., & Hanson, C. (editors). (2007). Climate Change 2007- Impacts, Adaptation and Vulnerability. Working Group II Contribution to the Fourth Assessment Report of the IPCC. Cambridge University Press.

Pidgeon, E. (2009). Carbon Sequestration by Coastal Marine Habitats: Important Missing Sinks. In D. G. Laffoley, The Management of Natural Coastal Carbon Sinks. IUCN, Gland.

Russell, S. J., Norvig, P., Canny, J. F., Malik, J. M., & Edwards, D. D. (2003). Artificial Intelligence: a Modern Approach (1st Edition). Prentice Hall, Upper Saddle River.

Santos, R., Murrieta-Flores, P., Calado, P., & Martins, B. (2018). Toponym matching through

deep neural networks. International Journal of Geographical Information Science, 32(2), 324-348.

Short, F., Carruthers, T., Dennison, W., & Waycott, M. (2007). Global seagrass distribution and diversity: a bioregional model. Journal of Experimental Marine Biology and Ecology, 350(1-2), 3-20.

Sit, M. A., Koylu, C., & Demir, I. (2019). Identifying disaster-related tweets and their semantic, spatial and temporal context using deep learning, natural language processing and spatial analysis: a case study of Hurricane Irma. International Journal of Digital Earth, 1-25.

Srivastava, S., Vargas Muñoz, J. E., Lobry, S., & Tuia, D. (2018). Fine-grained landuse characterization using ground-based pictures: a deep learning solution based on globally available data. International Journal of Geographical Information Science, 1-20.

Sun, Z., Di, L., & Fang, H. (2018). Using long short-term memory recurrent neural network in land cover classification on Landsat and Cropland data layer time series. International Journal of Remote Sensing, 40(2):593-614.

Tuia, D., Ratle, F., Pacifici, F., Kanevski, M. F., & Emery, W. J. (2009). Active Learning Methods for Remote Sensing Image Classification. IEEE Transactions on Geoscience and Remote Sensing, 47(7), 2218-2232.

Turing, A. M. (1950). Computing Machinery and Intelligence. Mind, 59, 433-460.

Valiant, L. G. (1984). A theory of the learnable. Communications of the ACM, 27(11), 1134-1142.

Wang, S. (2010). A CyberGIS Framework for the Synthesis of Cyberinfrastructure, GIS, and Spatial Analysis. Annals of the Association of American Geographers, 100(3): 535–557.

Wright, D., Sayre, R., Breyer, S., Butler, K., Vangraafeiland, K., Goodin, K., Kavanaugh, M., Costello, M., Cressie, N., & Basher, Z. (2017). The Ecological Marine Units Project as a Framework for Collaborative Data Exploration, Distribution, and Knowledge Building. EGU General Assembly Conference Abstracts, 6051.

Zhang, F., Du, B., & Zhang, L. (2015). Scene classification via a gradient boosting random convolutional network framework. IEEE Transactions on Geoscience and Remote Sensing, 54(3), 1793-1802.

Zhou, X., Li, W., Arundel, S. T., & Liu, J. (2018). Deep Convolutional Neural Networks for Map-Type Classification. In Freundschuh, S. M. and Sinton, D. (Eds.), Proceedings of the 2018 UCGIS Symposium, 147-155.